Neural circuitry underlying contrast gain control in auditory cortex

James E. Cooke
Worcester College, University of Oxford

Supervisors:
Benjamin D.B. Willmore, Jan W. Schnupp, Andrew J. King

Thesis submitted for the degree of
Doctor of Philosophy
at the University of Oxford
Hilary, 2016

”As long as our brain is a mystery, the universe - the reflection of the structure of the brain - will also be a mystery.”

Santiago Ramon y Cajal

Acknowledgements

First and foremost, I want to thank every member of the Auditory Neuroscience Group for making the lab such a wonderful environment to spend 3 years in. I'm especially grateful to Mari, Christian and Nicol for making our office a place I looked forward to coming to every day, and to Vani for letting me distract her every time I visited the coffee machine. A special mention goes to Jonathan Webb for helping me with my initial experimental setup and training me as a (mouse) brain surgeon. Above all I'm immensely grateful for the supervision offered by Andy, Jan and Ben. The combination of support and intellectual freedom created an ideal space for me to explore during the first few years of my career. I'm also grateful to Ed Mann, Louise Upton, Tommas Ellender, Alexandru Calin, Colin Ackerman and Luiz Guidi for their eagerness to collaborate and generosity with their equipment, time and mice. Special thanks needs to go to Martin Kahn, whose enthusiasm and generosity made an entire chapter of this thesis possible through our collaboration on the whole-cell recording experiments. Martin gets a second mention, along with the rest of the 2011 MSc in Neuroscience cohort, for offering constant friendship and support during our DPhils. All of you have contributed to making this period of my life one that I am sad to leave behind.

Statement of Originality

Intracellular recording experiments were conducted in collaboration with Martin Kahn, who established the whole-cell recordings. The remainder of this thesis is entirely my own work.

Abstract

While sensory environments can vary dramatically in their statistics, neurons have a limited dynamic range with which they can encode sensory information. In sensory cortex, this problem is in part resolved by the systematic adjustment of neural gain in accordance with the contrast of sensory input. This computation can produce contrast invariant representations in cortex that are also more resilient to static background noise. In visual cortex shunting inhibition by parvalbumin (PV) expressing interneurons and contrast-dependent membrane potential variance have been shown to contribute to contrast gain control (CGC), but whether these mechanisms underlie CGC in auditory cortex (AC) is currently unknown. I aimed to investigate the contributions of these mechanisms to CGC in mouse AC, as optogenetic methods for the circuit level interrogation of the nervous system are well established in that species. First, I characterised CGC in the anaesthetised mouse by performing large scale extracellular recordings across all layers of A1. I found that CGC is present in units recorded in all layers of mouse AC and is marginally stronger in deep layers, implicating intracortical mechanisms in the generation of this CGC. In order to investigate whether PV interneuron activity is capable of modulating the gain of sensory responses, I performed extracellular recordings of sensory evoked multi-unit responses in AC while I manipulated the activity of the PV interneurons optogenetically using Channelrhodospsin (ChR2) or Archaerhodopsin (Arch). PV interneuron activation with ChR2 did not alter spectrotemporal tuning of neuronal responses in AC but reduced both their gain and baseline activity. PV interneuron suppression with Arch left receptive field structure similarly unchanged. The strongest effect of PV suppression was an increase in the gain of sensory evoked responses. Thus, PV interneuron activity does appear capable of modulating the gain of AC auditory responses. However, I found that the activity of PV interneurons did not increase systematically with increasing stimulus contrast, indicating that these neurons may not implement CGC. Finally, I performed whole cell recordings from AC neurons in anaesthetised mice in order to assess the contribution of shunting inhibition and membrane potential variance to CGC. I found that CGC exists at the level of the membrane potential responses in AC, but neither input conductance nor membrane potential variance appeared to contribute to this. This canonical computation therefore appears to be implemented by non-canonical mechanisms in different cortical areas.

Contents

List of Figures

Glossary

2-pi two-photon.

5HT3aR serotonin 3a receptor.

A1 primary auditory cortex.

AC auditory cortex.

Arch archearodopsin.

BF best frequency.

CGC contrast gain control.

ChR2 channelrhodopsin.

Cl^- chloride.

CR calretinin.

CSD Current source density.

DRC dynamic random chord.

E-I excitation-inhibition.

EYFP enhanced yellow fluorescent protein.

FS fast-spiking.

G conductance.

Halo halorhodopsin.

I-O input-output.

IC inferior colliculus.

IT inferior temporal cortex.

L1 Layer 1.

LFP local field potential.

LGN lateral geniculate nucleus.

MGB medial geniculate body.

MSTd dorsal medial superior temporal cortex.

MT middle temporal visual area.

MUA multi-unit activity.

NPY neuropeptide Y.

PINP photo-identification of neuronal populations.

PSP post-synaptic potenital.

PSTH peri-stimulus time histogram.

PV parvalbumin.

RMP resting membrane potential.

RMS root mean square.

SPE signal power explained.

SST somatostatin.

STD short-term depression.

STF short-term facilitation.

STRF spectro-temporal receptive field.

V1 primary visual cortex.

V4 visual area 4.

VIP vasoactive intestinal peptide.

Chapter 1

Introduction

1.1 Neural adaptation to stimulus statistics

The world around us is astoundingly complex. Behaviourally relevant sensory signals vary over a large range of intensities - ambient light levels cover 9 orders of magnitude (Heckaman and Fairchild, 2009; Xiao et al., 2002) while the human auditory system is sensitive to a 12 fold range of sound intensities (Christopher Kirk and Smith, 2003). Sensory neurons, however, have a limited dynamic range over which they can vary their signalling. Stimulus intensity is routinely conveyed by the mean rate of action potential firing, yet physical limitations in the process of action potential generation result in neurons varying their firing rates within 2 orders of magnitude. Sensory systems therefore face the challenge of transmitting highly variable information along channels with limited capacity. In order to achieve this, they must attempt to encode these signals with the highest fidelity possible while minimising the cost of doing so - the code should be optimally efficient.

Efficiency in the neural code was first proposed as a method for avoiding the encoding of redundant sensory information (Barlow, 1961) but is now viewed as a core principle of sensory processing (Atick, 2015; Simoncelli and Olshausen, 2001;

Laughlin and Sejnowski, 2003). Information theoretic analysis has confirmed that in order for sensory codes to be efficient, they must take into account the statistical structure of the sensory input (Shannon, 1948; MacKay, 2003; Attneave, 1954; Brenner et al., 2000; Fairhall et al., 2001) and there is now considerable evidence that individual neurons adapt their responses to such statistics at every stage of sensory processing (Shapley and Enroth-Cugell, 1984; Dunn et al., 2007; Shapley and Victor, 1978, 1981; Benardete et al., 1992; Chander and Chichilnisky, 2001; Baccus and Meister, 2002; Zaghloul et al., 2005; Kaplan et al., 1987; Cheng et al., 1995; Sclar, 1987; Sclar et al., 1990; Smirnakis et al., 1997; Brown and Masland, 2001; Kim and Rieke, 2001; Solomon et al., 2004; Manookin and Demb, 2006; Sharpee et al., 2006; Lewicki, 2002; Rieke et al., 1995; Klein et al., 2003)

1.1.1 Adaptation to mean stimulus intensity

Two fundamental statistical properties of a given sensory stimulus are its mean and variance. Mean stimulus level corresponds to the overall intensity of the sensory input while the variance in stimulus level corresponds to the contrast of the stimulus. Over the course of a day, the visual system is confronted with mean luminance values that vary over approximately 9 orders of magnitude. The ability of the human visual system to successfully operate across this entire range is due predominantly to adaptation to local mean intensity in the retina, a process known as light adaptation (Shapley and Enroth-Cugell, 1984; Dunn et al., 2007). Light adaptation occurs as a result of the bleaching of rhodopsin, the G-protein receptor responsible for phototransduction, in high light levels and its regeneration in low light levels. The sensitivity of rhodopsin to bleaching varies between the two classes of photoreceptor in the retina, the rods and cones. Rods are more sensitive than cones and therefore are responsible for encoding low light levels, while cones are capable of encoding high light levels. The retina therefore employs a division

of labour strategy in which different parts of the input range are processed along separate channels.

Within a channel, adaptation is also crucial to the encoding of a large range of input intensities. The result of light adaptation is that the range of intensities that the neuron responds to is shifted to match the mean intensity of the stimulus and thereby ensures that the output of the sensory epithelium to the brain is invariant to average intensity of the sensory input. By no longer encoding mean intensity, individual neurons can make available coding capacity that can be used in order to signal more behaviourally relevant aspects of the sensory input.

The same problem occurs in the auditory system. The sensitivity of human hearing ranges from 0 dB sound pressure level (SPL) at threshold to 120 dB SPL, sounds intensities that are 12 orders of magnitude larger. While quiet natural environments can have mean intensities of several dB, the presence of vocalising animal groups can result in environments with mean intensities more than 10^8-fold higher (Christopher Kirk and Smith, 2003; Simmons et al., 1971; Aubin, 2004). Despite this enormous variation in input, human listeners can discriminate intensity changes of approximately 1 dB across the majority of this range (Viemeister, 1974; Houtsma and Braida, 1980; Viemeister, 1983).

As with the retina, the cochlea also shows adaptation to stimulus statistics, resulting in adaptation to stimulus mean in auditory nerve firing rates (Wen et al., 2009). When mean intensity is low, the outer hair cells of the cochlea actively amplify the movements of the basilar membrane in order to amplify soft sounds and compress loud sounds (Ruggero and Rich, 1991; Yates, 1995). This results in adaption to low mean sound levels by shifting the intensities that the inner hair cells respond to, depending on the statistics of the input. Adaptation to high mean intensities also occurs in the cochlea but via a combination of mechanisms. Adaptation to transient high mean intensities occurs through suppression of outer hair

cell amplification via the olivocochlear reflex (Dallas, 1992) and as a result of the middle-ear-muscle reflex (Liberman and Guinan, 1998; Borg E, 1989). The latter reflex involves contractions of the stapedius and tensor tympani muscles of the middle ear and results in attenuation of transmission of vibrations through the ossicles, especially in response to high intensity low frequency sounds (Nuttall, 1974). Auditory nerve firing rates also exhibit spike-frequency adaptation to sustained stimuli which is also thought to contribute to high mean adaptation (Westerman and Smith, 1984; Yates, 1985). Spike-frequency adaptation in the auditory nerve is thought to occur as a result of both synaptic depression via pre-synaptic neurotransmitter depletion (Furukawa et al., 1978; Moser and Beutner, 2000; Spassova et al., 2004) and the post-synaptic desensitisation of receptors (Goutman and Glowatzki, 2007) at the inner hair cell synapse.

These adaptive mechanisms fail to enable individual neurons to encode intensities across the entire operating range of the system. A complementary division-of-labour strategy is also employed in the auditory periphery in order to solve the mean intensity problem. A minority of auditory nerve fibres have high thresholds and responses that do not saturate (Colburn et al., 2003), while most tend to have low thresholds and responses that saturate in the low to middle area of the range of hearing. The properties of low threshold and saturation result in the majority of auditory nerve fibres being capable of encoding a ~30-40 dB range of sound levels (Evans and Palmer, 1980; Sachs and Abbas, 1974). It has been shown that the psychophysical ability of human listeners to discriminate different intensities far outperforms predictions based on the properties of these auditory nerve fibres (Viemeister, 1988). This confirms that a dynamic range problem does indeed exist in audition and that central processing is crucial in solving this problem.

After the auditory nerve, sensory signals ascend the auditory pathway, passing sequentially through the cochlea nucleus, the superior olivary complex in the

brainstem, the lateral lemniscus, the inferior colliculus (IC) in the midbrain, the medial geniculate body (MGB) of the thalamus and the primary auditory cortex (A1). The majority of outputs from the dorsal cochlea nucleus bypass the brainstem and synapse directly in the IC (Schweizer, 1981; Ryugo DK, 1981; Winer and Schreiner, 2005). Recordings from the IC of anaesthetised guinea pigs show that the mean intensity problem appears to be resolved by this level of the auditory pathway, as individual IC neurons adjust their sensitivity to sound intensity based on the mean intensity of recent acoustic stimulation (Dean et al., 2005). While this effect undoubtedly builds upon previous stages of mean adaptation, it is not entirely attributable to the adaptive mechanisms of the inner ear. The middle-ear-muscle reflex is unlikely to play a significant role as this reflex is weak in guinea pigs (Avan et al., 1992), is inactive under conditions of anaesthesia (Carmel and Starr, 1963) and is unlikely to be elicited by the stimuli of the intensities and frequencies used in the study (Nuttall, 1974). Spike-rate adaptation in the auditory nerve, as well as further along the ascending auditory pathway (Ingham and McAlpine, 2004), may contribute to this effect as its time constant of 100-400 ms (Wen et al., 2012) fits with the time constant of adaptation in the IC, which has a mean of 160 ms (Dean et al., 2008). Suppression of outer-hair cell amplification via the olivocochlear reflex may also contribute to adaptation in the IC as it has a time constant of approximately 100 ms (Boyev et al., 2002).

1.1.2 Adaptation to stimulus contrast

While adaptation to mean stimulus intensity enables sensory systems to function over a wide range of stimulus levels, the variance of stimulus intensity, or contrast, also poses a problem for encoding. The relationship between sensory input and the firing rate output of neurons typically shows a sigmoidal relationship as a result of intrinsic threshold and saturation nonlinearities in neurons (Figure 1.1). Mean

adaptation produces horizontal shifts of this curve along x-axis, positioning the centre of this curve over the centre of the distribution of input intensities (Figure 1.1A). Encoding in this region of the sigmoid is relatively linear but, as a result of neuronal nonlinearities, its dynamic range is limited. Failing to adapt to different levels of variability in the sensory input would result in the quality of encoding being highly dependent on how closely the range of input intensities matches the dynamic range of the neuron in question. In vision, local contrast can vary dramatically from one part of a natural scene to another, typically varying over 2-3 orders of magnitude but occasionally approaching a difference of 10^6:1 (Heckaman and Fairchild, 2009; Xiao, 2002; Mante et al., 2005). This level of variability in stimulus contrast poses a major problem for sensory processing.

This problem is resolved in a similar manner to the mean intensity problem, by adapting responses of individual neurons in the retina to stimulus contrast (Shapley and Victor, 1978, 1981; Benardete et al., 1992; Chander and Chichilnisky, 2001; Baccus and Meister, 2002; Zaghloul et al., 2005). This process, known as contrast gain control (CGC), does not involve additive or subtractive shifts in neuronal responsiveness as is seen with mean adaptation. Instead, multiplicative or divisive changes in responsiveness to sensory input occur, which result in a change of the slope of the neuron's response (Figure 1.1B). This reflects a change in gain that maintains the overall firing rate of the neuron. This adjustment of sensitivity to stimulus intensity occurs as a function of the root-mean-square (RMS) contrast in the surrounding region, the standard deviation of the input intensities divided by the mean (Sclar et al., 1990; Bonin et al., 2006). This adaptation changes the gain of the neurons responses, so that under low contrast stimulation gain is increased and the dynamic range of the neuron is narrowed to match the distribution of stimulus vales, while under high contrast stimulation gain is reduced and the dynamic range extended. This process continues in the lateral geniculate nu-

cleus (LGN) of the thalamus (Kaplan et al., 1987; Cheng et al., 1995; Sclar, 1987; Sclar et al., 1990) before visual input reaches primary visual cortex (V1). Contrast adaptation throughout the visual pathway occurs on two timescales; fast adaptation occurs on the order of tens to hundreds of milliseconds while slow contrast adaptation takes several seconds (Maffei et al., 1973; Movshon and Lennie, 1979; Smirnakis et al., 1997; Carandini and Ferster, 1997; Brown and Masland, 2001; Kim and Rieke, 2001; Baccus and Meister, 2002; Solomon et al., 2004; Manookin and Demb, 2006).

Neuronal gain is modulated further by stimulus contrast in V1. Computations in the retina result in contrast itself driving responses in the visual pathway but by the level of V1, these responses are scaled by the level of contrast in a wider surround region of visual space. This process results in contrast being represented relative to the contrast of the surrounding area. This form of adaptation, known as contrast normalization, is well described by the normalization model (Heeger, 1992; Heeger et al., 1996; Carandini et al., 1997; Carandini and Heeger, 1994) in which the strength of the sensory input, weighted by the receptive field of the neuron, is divided by a constant plus the summed activity of another more broadly tuned population of neurons. Contrast normalization solves the contrast problem in vision by generating representations of local contrast that are invariant to contrast variations at a less local scale (Albrecht, 1991; Heeger, 1992; Busse et al., 2009; Ringach, 2010).

In the auditory system, the contrast problem is solved by a very similar form of adaptation. In the primary auditory cortex (A1) of anaesthetised ferrets, individual neurons adapt to increases in stimulus contrast by decreasing their gain over a timescale of 100-200msec (Rabinowitz et al., 2011a). This adaptation compensates for changes in stimulus contrast in much the same way as in the visual system and, in keeping with this similarity, these gain changes being well described by

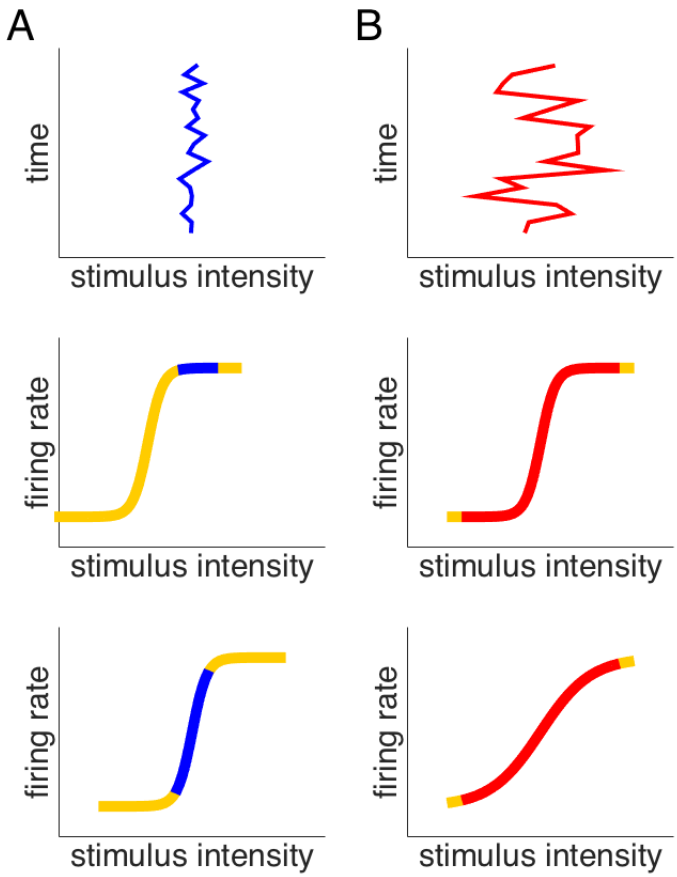

Figure 1.1: A. Upper Panel: Variation of stimulus intensity over time. Middle panel: Sigmoidal stimulus-response function of a sensory neuron before adaption to mean or variance. The region shaded in blue indicates the portion of the curve that is being used to encode the stimulus. As a result of its high mean, the distribution of inputs falls in the saturation range of the curve. As a result, variations in input intensity fail to be encoded as variations in firing rates. Bottom panel: Following mean adaptation, the curve is shifted horizontally so that the distribution of sensory inputs now falls over the linear portion of the stimulus-response function. B. Upper Panel: Large variation of stimulus intensity over time for a high contrast stimulus. Middle panel: Stimulus-response function after mean adaptation only. The high variance in this stimulus results in much of the distribution of stimulus intensities falling outside the linear portion of the curve, as indicated by the region shaded in red. Bottom panel: Contrast gain control results in the slope of the curve decreasing in order to compensate for this change in contrast. The full range of inputs can now be successfully encoded by the neuron.

the normalization equation (Rabinowitz et al., 2011a). Unlike in the visual system however, these gain changes are not entirely compensatory and therefore do not generate responses that are completely contrast invariant. This form of adaptation

is therefore termed contrast gain control (CGC) rather than contrast normalization (Willmore et al., 2014).

In the IC, neurons show heterogeneous changes in their stimulus-response functions in response to changes in stimulus contrast, but many units show weak contrast-dependent gain changes (Rabinowitz et al., 2013). Adaptation to contrast has also been observed at the population level in IC (Dean et al., 2005). CGC is significantly stronger and more pervasive in A1 however, suggesting that adaptation in the IC alone cannot account for CGC in cortex (Rabinowitz et al., 2013). In combination, mean adaptation in the IC and CGC in A1 have been shown to improve the encoding of complex auditory stimuli and increase the tolerance of these representations to background noise as stimulus information ascends the auditory pathway (Rabinowitz et al., 2013).

At the level of sensory cortex, CGC appears to be a widespread phenomenon. It has been observed in V1 of a variety of species, including mice (Atallah et al., 2012), cats (Heeger, 1992; Busse et al., 2009), monkeys (Carandini and Heeger, 1994; Carandini et al., 1997) and humans (Busse et al., 2009). As well as A1 of ferrets (Rabinowitz et al., 2011a), compensatory gain changes have been observed in rat somatosensory cortex in response to changes in the variance of whisker deflection (Garcia-Lazaro et al., 2007). It has even been observed in the avian forebrain (Nagel and Doupe, 2006; Baccus, 2006). The prevalence of this computation across species and across sensory cortical areas indicates that contrast gain control may be a fundamental computation performed by the cortex (Carandini and Heeger, 2012).

Beyond contrast normalization in primary sensory cortical areas, the normalization model captures certain nonlinear aspects of responses in higher visual cortical areas including MT (Heeger et al., 1996; Simoncelli and Heeger, 1998; Rust et al., 2005), V4 (Reynolds and Desimone, 2003) and inferior temporal cortex (IT) (Zoccolan et al., 2005). It also successfully captures aspects of multisensory integra-

tion in the superior colliculus and dorsal medial superior temporal (MSTd) cortex (Morgan et al., 2008; Ohshiro et al., 2011) and the effects of attention in visual (Reynolds and Heeger, 2009) and auditory cortex (Fritz et al., 2007; Atiani et al., 2009; Yin et al., 2014). Given its prevalence and the relative homogeneity of cortical organisation (Mountcastle, 1997), elucidating the biophysical basis of cortical gain control may provide fundamental insight into the mechanistic basis of cortical processing.

1.2 The biophysical basis of gain control

Given the importance of gain control as a neural computation (Carandini and Heeger, 2012), much work has gone into identifying the biophysical basis of this process. While several elegantly simple mechanisms for gain control exist, on-going work is revealing a picture in which gain changes are induced via a variety of mechanisms. Variation across systems and species as well as the presence of multiple redundant mechanisms in the same system may reflect the fundamental importance of this computation.

1.2.1 Short-term synaptic depression

Many synapses dynamically modulate their strength in an activity-dependent manner, resulting in either short-term facilitation (STF) or short-term depression (STD) (Eccles et al., 1941; Feng, 1941). STD can occur as a result of depletion of neurotransmitter stores in the presynaptic inputs to a neuron (Betz, 1970; Saviane and Silver, 2006; Elmqvist and Quastel, 1965) or post-synaptic desensitisation of receptors (Saviane and Silver, 2006; Trussell and Fischbach, 1989). STD results in a reduction in synaptic strength with a time constant that can vary from a few milliseconds (msec) (Carandini et al., 2002) to several seconds (Varela et al., 1997;

Dittman et al., 2000). By reducing the probability of neurotransmitter release during high-frequency stimulation, fast-acting STD introduces saturation into the input-output (I-O) relationship of the synapse. This enables the synapse to act as a low-pass filter and to signal dynamic changes in on-going rate-coded inputs rather than the absolute rate of the inputs (Tsodyks and Markram, 1997). The strength of STD however is proportional to the absolute rate of the inputs. Together, these features allow gain control to occur at the level of the individual synapse, as they signal dynamic changes in the input, scaled by the absolute rate of the input (Abbott et al., 1997; Abbott and Regehr, 2004). This allows STD to provide stimulus-driven gain control when the normalization factor is encoded in the absolute firing rates of the inputs.

STD has been identified at the thalamocortical synapse in cat V1 (Stratford et al., 1996; Boudreau and Ferster, 2005), raising the possibility that it may underlie some forms of stimulus-driven gain control observed there. The synapse-specific nature of STD is thought to enable the divisive changes in the tuning properties of V1 neurons that produce contrast invariance. A scaling down of the neuron's receptive field is possible through synaptic depression as a result of the strength of STD being proportional to the strength of stimulus drive at each synapse (Banitt et al., 2007; Carandini et al., 2002). Given the input-specific nature of these gain changes, synaptic depression can only explain forms of gain control that are driven by inputs that fall within the excitatory receptive field of the neuron. This feature is consistent with STD underlying contrast normalization and the suppressive gain changes observed when orthogonal grating are superimposed, known as cross-orientation suppression. The input-specific nature of STD, however, prevents it from accounting for gain changes produced by stimuli that fall outside of the neurons classical receptive field, as is the case for surround suppression (Carandini et al., 2002).

In the auditory system STD contributes to the spike-frequency adaptation seen

in the auditory nerve (Westerman and Smith, 1984; Yates, 1985). It has also been observed in the IC (Ingham and McAlpine, 2004), making it a potential candidate mechanism for the mean intensity adaptation observed there (Dean et al., 2005). STD has been observed at the thalamocortical synapse in mouse A1 (Rose and Metherate, 2005), raising the possibility that the first synapse onto layer 4 cortical neurons is the site of CGC in A1. This mechanism fails to account for several aspects of CGC however, making the extent to which it may contribute to CGC in this area unclear.

Analysis of the spectral features that drive CGC in A1 has revealed that contrast within the receptive field drives gain changes with over twice the strength of those produced by contrast in frequency bands outside the receptive field (Rabinowitz et al., 2011a). This is the case whether the frequencies fall in the excitatory or inhibitory regions of the receptive field (Rabinowitz et al., 2012). Despite the dominance of local contrast, this analysis shows that CGC also has a significant global component (Rabinowitz et al., 2011a). Given the input specificity of STD, this mechanism cannot be used to explain gain changes produced by contrast outside the receptive field of the neuron. While a single mechanism may be able to account for both the local and global components of CGC, it is entirely possible that these two aspects of CGC rely on separate mechanisms. If this is the case, any role for STD would be restricted to local CGC alone.

STD also faces problems in accounting for certain aspects of local CGC. Neurons in layer 4 have been shown to inherit the spectral range of their excitatory tuning from a combination of thalamic inputs (Liu et al., 2007) while inhibitory tuning has an intracortical origin and is more broadly tuned than excitation (Wu et al., 2008). If STD at the thalamocortical synapse underlies local CGC, only contrast within the excitatory region of the receptive field should modulate gain. CGC has been shown to be driven by contrast within both excitatory and inhibitory regions of the

receptive field (Rabinowitz et al., 2012), indicating that STD at the thalamocortical synapse alone is not a sufficient explanation for CGC. STD would also have to occur at inhibitory intracortical synapses for this mechanism to account for the effects of CGC.

Another feature of CGC that poses a problem for this explanation is the finding that changing contrast in a restricted frequency region changes gain across the entire receptive field of the neuron (Rabinowitz et al., 2011a). If CGC were implemented though STD, this manipulation would produce gain changes in the inputs representing the restricted frequency region only, leaving the gain of the remaining inputs unchanged. STD at the thalamocortical synapse would be capable of accounting for this finding only if the spectral tuning properties of layer 4 neurons in A1 are completely inherited from individual neurons in the MGB. While the width of their spectral sensitivity has been shown to be determined by their thalamic inputs (Liu et al., 2007), these inputs originate from several subdivisions of the MGB and so most likely represent the integration of many differently tuned inputs (Huang and Winer, 2000; Winer et al., 2005). Furthermore, thalamocortical STD would only alter gain in the excitatory regions of the receptive field but changing contrast within a restricted frequency region has been shown to alter gain in inhibitory as well as excitatory parts of the receptive field (Rabinowitz et al., 2012). The finding that gain changes occur across the entire receptive field of the neuron indicates that the mechanism responsible for CGC is capable of altering the sensitivity of the entire neuron and not just specific inputs. Furthermore, the nature of spectral integration in CGC indicates that the signals driving this neuron-level change in gain originate at the network level, not at the level of specific inputs. For these reasons, STD alone appears unlikely to account for CGC in AC.

1.2.2 Intrinsic cellular mechanisms for gain control

In the visual system, intrinsic cellular mechanisms are thought to contribute to contrast invariance, specifically underlying the slow component of contrast adaptation that take place over several seconds in the retina (Smirnakis et al., 1997; Kim and Rieke, 2001; Brown and Masland, 2001; Baccus and Meister, 2002; Manookin and Demb, 2006), LGN (Solomon et al., 2004) and V1 (Maffei et al., 1973; Movshon and Lennie, 1979; Carandini and Ferster, 1997). At each stage in the visual pathway, slow contrast adaptation is accompanied by a tonic hyperpolarisation of the membrane potential. In V1, this is produced by a potassium conductance triggered by sodium influx from action-potential firing (Sanchez-Vives et al., 2000a,b). This potassium conductance scales down excitatory synaptic responses via a negative feedback mechanism whereby greater influx of sodium produces a greater potassium conductance (Nanou et al., 2008). A hyperpolarising potassium conductance that is evoked in a stimulus-dependent manner has also been found in rat A1 neurons in vitro (Abolafia et al., 2011). However, the kinetics of this effect are too slow to account for CGC in A1. Adaptation to high contrasts takes approximately 90msec while adaptation to low contrasts is slower, taking 160msec (Rabinowitz et al., 2011a). This potassium conductance occurs on the order of seconds and even minutes (Schwindt et al., 1989) and does not show an asymmetry of time course in induction and recovery (Schwindt et al., 1989). Furthermore, cellular mechanisms of this kind are only capable of accounting for the local component of CGC, leaving the issue of what mechanism underlies global CGC unanswered.

1.2.3 Intracortical excitation

The finding that contrast outside of a neuron's receptive field modulates gain in A1 necessitates the presence of a network level mechanism (Rabinowitz et al., 2012).

If a single mechanism underlies both local and global components of CGC, it must be implemented at the network-level. Most research into network-level gain control has focused on the recruitment of inhibition but it is possible to modulate neural gain through changes in excitation alone. Intracortical connections provide the vast majority of inputs to cortical neurons (LeVay and Gilbert, 1976) and are known to have a role in shaping response properties of A1 neurons, such as the sharpness of spectral tuning (Liu et al., 2007; White and Keller, 1989; Peters and Payne, 1993; Peters et al., 1994). Recurrent activity in these excitatory circuits has been shown to amplify thalamic inputs in a multiplicative manner in A1 (Li et al., 2013a) and V1 (Douglas and Martin, 1991a; Douglas et al., 1995; Li et al., 2013b; Lien and Scanziani, 2013). This property is likely to play some role in gain control, if only by producing a 'default' high-gain state that can be modulated by other mechanisms.

Variability in the synaptic input to a neuron, generated by this excitation, has also been shown to modulate neuronal gain. The I-O function of a neuron recorded in slice preparations follows a threshold-linear relationship or, if the linear portion is steep enough and the neuron reaches saturation, a step-function. In vivo, however, these neurons are exposed to constant barrages of synaptic inputs that may or may not relate directly to the driving input to the neuron (Ringach, 2009). Under conditions of membrane potential depolarisation above threshold this synaptic noise has the effect of occasionally boosting weak inputs above threshold and strong inputs below threshold, making the relationship between membrane potential and firing rate probabilistic (Figure 1.2A). This variability has the effect of smoothing a step-like I-O function into a sigmoidal function with a shallower slope, reflecting a reduction in gain (Chance et al., 2002).

Trial-to-trial variability in the sub-threshold membrane potential when the neuron is not routinely depolarised above threshold can actually increase gain. Under these conditions, the effect of Gaussian distributed variability is to smooth

the threshold-linear I-O function, producing a power-law function as observed in vivo (Miller and Troyer, 2002). The main effect of further increasing variability under these conditions is boosting weak inputs above threshold, increasing gain (Figure 1.2B). This occurs as the membrane potential is typically below threshold where the suppressive half of the variability distribution has no effect on firing rate. Trial-to-trial variability in membrane potential responses of this kind are believed to contribute to contrast invariance in V1, as low contrast stimuli have been shown to produce larger levels of variability, leading to an increase in gain (Troyer et al., 1998; Hansel and van Vreeswijk, 2002; Carandini, 2004; Ringach and Malone, 2007; Finn et al., 2007). In V1 however, this variability is believed to originate from a combination of feedforward thalamocortical mechanisms rather than from intracortical excitation (Sadagopan and Ferster, 2012). Depending on how depolarised neurons are during acoustic stimulation opposite predictions can be made regarding the contrast-dependent nature of membrane potential variability. Increased sub-threshold trial-to-trial membrane potential variability under low contrast stimulation in A1 may contribute mechanistically to CGC as this form of variability increases gain. It is however also possible that increased variability, or 'synaptic noise', during periods of depolarisation evoked by high contrast stimuli may also contribute to CGC, as this form of variability decreases gain. Increased membrane potential variability either under low or high contrast stimulation may therefore contribute to CGC in cortical neurons, depending on the operating regime of the neuron in question.

1.2.4 Intracortical inhibition

Inhibition has been seen as a very plausible candidate mechanism for gain control. Pharmacologically blocking GABA transmission using GABA antagonists has been used in order to test the involvement of inhibition in different forms of gain

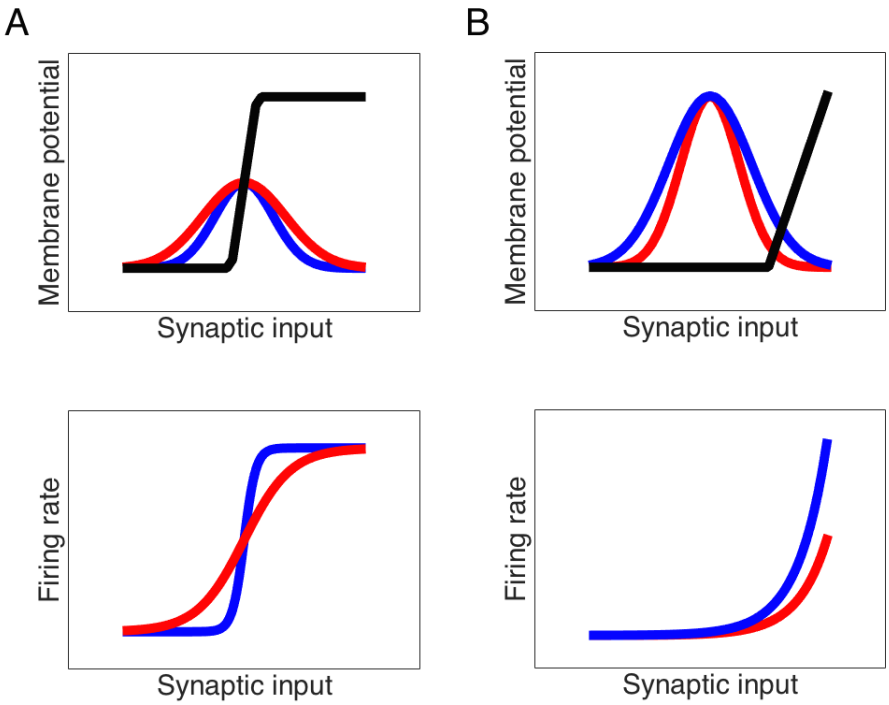

Figure 1.2: A. Top panel: Step-like I-O function of a neuron during depolarisation above spiking threshold (black). Levels of synaptic noise predicted to be produced by low (blue) and high (red) contrast stimuli. Bottom panel: I-O function following smoothing by low (blue) and high (red) levels of synaptic noise, predicted to be produced by low (blue) and high (red) contrast stimuli respectively. B. Top panel: Threshold-linear I-O function under conditions where threshold crossing is less common (black). Levels of trial-to-trial variability predicted to be produced by low (blue) and high (red) contrast stimuli. Under these conditions, increased variability is predicted to be associated with low (blue) rather than high (red) contrast levels. Bottom panel: I-O function following smoothing by low (red) and high (blue) levels of membrane potential variability, corresponding to high (red) and low (blue) contrast levels respectively.

control. Iontophoresis of the GABAa receptor blocker gabazine has been found to increase the gain of sensory evoked responses in V1, but leaves the contrast sensitivity of the neurons unaffected (Katzner et al., 2011). Cross orientation suppression is similarly been found to be unaffected by this manipulation (Katzner et al., 2011). This indicates that, in V1, inhibition may not underlie gain changes that are driven by stimulus contrast within the receptive field of the neuron. Inhibition may play a role in mediating gain changes driven by stimulus contrast outside the receptive field, however, as in the case of surround suppression (Haider et al., 2010),

although this remains controversial (Ozeki et al., 2009). In the auditory system, GABA antagonists have been shown to affect the gain of neurons in the IC (Ingham and McAlpine, 2005). Inhibition has even been shown to underlie gain changes produced by mask odorants in the antennal lobe of the fruit fly (Olsen and Wilson, 2008; Olsen et al., 2010).

One particular form of inhibition has received particular attention in the literature as a mechanism for gain control, shunting inhibition. Classical inhibition occurs when channels are opened that produce a net flow of current out of the neuron, hyperpolarising the membrane potential of the neuron. Shunting inhibition occurs when there is no net flow of current in or out of the cell, and therefore no change in membrane potential, but the opening of channels changes the conductance of the membrane (Borg-Graham et al., 1998a). This can be achieved either through the opening of channels with reversal potentials close to the resting potential of the neuron (Reichardt et al., 1983), or via balanced excitation and inhibition (Capaday, 2002). Given that its effects are not directly observable as changes in current or voltage, shunting inhibition is also known as silent inhibition.

The divisive effect of conductance changes produced by shunting inhibition can be understood through Ohm's law (Coombs et al., 1955; Fatt and Katz, 1953). Ohm's law tells us that voltage (V) is equal to current (I), multiplied by resistance (R).

$$V = IR \tag{1.1}$$

Conductance (G) is the inverse of the resistance (R):

$$G = 1/R \tag{1.2}$$

Ohm's law can the be rewritten as:

$$V = I/G \tag{1.3}$$

This demonstrates that the conductance of the neuron has a scaling effect on the translation of synaptic currents into changes in membrane potential. Modulating this conductance via shunting inhibition offers an elegantly simple way of modulating gain.

While it is clear that changes in conductance scale membrane potential responses, this does not simply translate into a scaling of firing rates. Some modeling studies have suggested that the effect of increased somatic conductance on firing rates is actually subtractive rather than divisive (Gabbiani et al., 1994; Holt and Koch, 1997; Capaday, 2002) and this has been confirmed in in vitro cortical slice preparations (Mitchell and Silver, 2003). The reason conductance had a subtractive effect in these studies was that the neurons were not exposed to the background synaptic input that they would typically be exposed to in vivo (Destexhe and Paré, 1999; Ulrich, 2003; Destexhe et al., 2003; Ringach, 2009). Without synaptic noise to couple membrane potential to firing rate in a probabilistic manner, shunting inhibition simply increases the amount of synaptic input necessary for the neuron to reach threshold, manifesting as a rightward shift in the linear-threshold or step-like I-O function. Once the I-O function of a neuron has been smoothed by synaptic noise, increases in conductance successfully reduce neuronal gain (Mitchell and Silver, 2003; Chance et al., 2002).

Shunting inhibition has long been considered as a mechanism for changing neuronal gain in V1 (Carandini and Heeger, 1994; Carandini et al., 1997) and visual stimuli have been shown to evoke strong changes in the conductance of V1 neurons (Borg-Graham et al., 1998b; Anderson et al., 2000; Monier et al., 2003; Frégnac et al., 2003). Conductance has been shown to increase with stimulus contrast in V1 but these conductance changes are not orientation invariant as predicted by the

hypothesis that these conductance changes only contribute to contrast normalisation (Anderson et al., 2000; Monier et al., 2003; Frégnac et al., 2003). It remains to be seen however, whether contrast-dependent conductance changes can account for CGC in A1.

Classical inhibition has traditionally been thought to have a subtractive effect on membrane potential but, in certain circumstances, has also been shown to be capable of modulating neuronal gain. While it does not have a shunting effect by itself, if balanced by excitation the effects of inhibitory currents on membrane potential can be cancelled by the effects of excitatory currents. As a result, balanced excitatory and inhibitory inputs do not produce a net flow of current in or out of the neuron but the opening of channels in the membrane results in an increase in conductance that has a functional shunting effect (Carandini and Heeger, 1994; Chance et al., 2002). Balanced excitation and inhibition has been observed in both A1 (Wehr and Zador, 2003; Sun et al., 2010; Dorrn et al., 2010; Zhou et al., 2014a) and V1 (Haider et al., 2006; Okun and Lampl, 2008) making this a plausible mechanism for gain control in cortex (Shapley and Xing, 2013). Inhibition may also modulate gain by dampening the multiplicative effect of amplification by recurrent intracortical connections (Douglas and Martin, 1991a; Douglas et al., 1995; Li et al., 2013a,b; Lien and Scanziani, 2013).

In some instances, hyperpolarising inhibition at the cellular level may translate into divisive effects when viewed with respect to the sensory input, rather than the synaptic input to the neuron (Murphy and Miller, 2003). Subtractive inhibition has the effect of shifting the range of synaptic inputs that encode the sensory input to the left along the x-axis. As a result of the power-law relationship between synaptic input and spiking activity at the neuronal level, the mapping from sensory input to firing rate is now positioned over a portion of the neuron's I-O function that has a shallower slope. While the gain of the relationship between synaptic

input and spiking remains the same for the neuron, the gain of the relationship between sensory input and spiking has been reduced (Murphy and Miller, 2003). This mechanism has been shown to account for gain changes produced by a tonic reduction in cortical inhibition in V1 (Atallah et al., 2012) and may contribute to a variety of forms of gain control in sensory systems. As an explanation of CGC, this mechanism faces one problem in particular. It should introduce a significant y-axis offset in the systems level I-O function as a result of moving along the power law function of the neuron but this is not observed in CGC in ferret A1 (Rabinowitz et al., 2011a).

1.2.5 The interaction of mechanisms for gain control

It is becoming increasingly clear that the brain employs a variety of mechanisms for gain control that often interact in order to generate changes in gain. We have already seen in the case of shunting inhibition that it is essential to take into account its interaction with another mechanism of gain control, membrane potential variability, in order to understand its arithmetic effects on neuronal responses (Mitchell and Silver, 2003; Chance et al., 2002). Under low-noise conditions, the effectiveness of gain control can be improved further if conductance changes are introduced via balanced excitation and inhibition (Chance et al., 2002). As well as producing a change in slope, low levels of synaptic noise can result in a leftward shift in the I-O function of a neuron by enabling weak inputs to cross the spiking threshold while having little effect on strong inputs that reliably cross threshold. The balance of excitation and inhibition produce a change in conductance in the neuron that, under low-noise conditions, introduces a subtractive rightward shift in the I-O function. By introducing synaptic noise and increasing conductance simultaneously, balanced excitation causes the leftward and rightward shifts introduced by these two mechanisms to cancel, resulting in a pure gain change (Chance et al., 2002).

STD has also been found to contribute to the gain modulation by membrane potential variability observed in V1. As long as correlated contrast-dependent variability exists between neurons in the thalamus, the compressive nonlinearity introduced by STD results in the variability of large depolarisations being reduced (Sadagopan and Ferster, 2012). This combination of gain control mechanisms allows contrast-dependent membrane potential variability to modulate gain in a purely feed-forward manner, without intracortical excitation contributing to this variability (Sadagopan and Ferster, 2012).

Experimental evidence from dynamic-clamp recordings in morphologically simple cerebellar granule cells has revealed that that STD and classical inhibition can also interact in order to produce changes in neuronal gain (Rothman et al., 2009). STD in the excitatory input introduces a nonlinearity into the relationship between pre-synaptic firing rate and post-synaptic conductance - the synaptic I-O function. It changes this previously linear relationship into saturating exponential. Subtractive inhibition results in a rightward shift in the relationship between post-synaptic current and firing rate, the neuronal I-O function (Figure 1.3A). This offset results in the level of current necessary to drive the spiking responses being moved towards the saturating region of the synaptic I-O function (Figure 1.3B). The I-O function that maps pre-synaptic firing rate to post-synaptic firing rate is essentially a convolution of these two functions (Figure 1.3C). With the added nonlinearity introduced by STD, subtractive inhibition results in a reduction in gain by moving the neuronal I-O function into the more compressive region of the synaptic I-O function. This mechanism operates without synaptic noise and can produce gain changes that are four times as strong as those produced by noise alone (Rothman et al., 2009). The strength of gain control using this combined mechanism relies on both the strength of STD and the strength of inhibition, providing two channels by which to modulate gain.

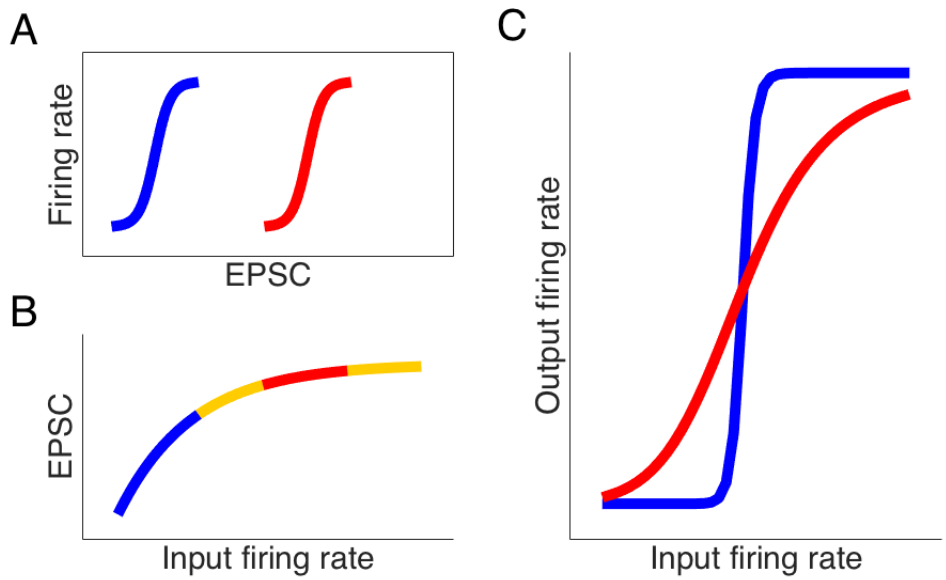

Figure 1.3: A. The neuronal I-O function maps excitatory post-synaptic current (EPSC) onto firing rate. Adding tonic inhibition produces a rightward shift in this relationship (Blue curve: before inhibition. Red curve: after inhibition). B. This synaptic I-O function shows the nonlinear mapping between pre-synaptic firing rate and post-synaptic excitatory current in a neuron with STD. Without STD this function would be linear. The offset in the neuronal I-O function produced by inhibition results in more current being necessary in order to achieve the same level of spiking, increasing the level of pre-synaptic spiking necessary to achieve this current. This places the synaptic input in the saturating range of the synaptic I-O function (red portion of curve. The blue portion of the curve indicates input range prior to inhibition). C. The mapping from pre-synaptic firing rate to post-synaptic firing rate involves passing through both of these nonlinearities. The offset in the neuronal I-O function produced with inhibition results in this function being convolved with a part of the synaptic I-O function with a shallow slope, reducing the gain of the overall I-O function (Blue curve: before inhibition. Red curve: after inhibition).

As an explanation of CGC in cortical neurons, these experimental findings are limited by the drastically different morphology of the neurons used in the study. The authors addressed this issue, however, through theoretical work. They demonstrated that this mechanism successfully modulates gain in simulations of morphologically complex cortical pyramidal neurons in which asynchronous inhibitory inputs are distributed over the basal dendrites of the neuron (Rothman et al., 2009). Furthermore, this combination of mechanisms introduces a subtractive rightward shift along the x-axis that is also observed in CGC in A1 (Rabinowitz et al., 2011a). The involvement of inhibition in this mechanism compensates for the input-specific

nature of STD, as inhibition is capable of modulating responses across the entire receptive field of the neuron if it acts following dendritic integration of excitatory inputs. Combined STD and classical inhibition results in compression of the spiking output dynamic range of the neuron while leaving the input dynamic range the same (Rothman et al., 2009). CGC in A1 is typically viewed as a problem of input dynamic range compression but the extent to which CGC reflects input vs output gain control has not been addressed.

1.3 The source of inhibition for gain control

1.3.1 Subtypes of cortical interneuron

Inhibitory mechanisms appear to be the most plausible candidates for explaining CGC in A1. Shunting inhibition, under conditions of synaptic noise observed in vivo (Mitchell and Silver, 2003; Chance et al., 2002), is a very plausible mechanism for CGC. Non-shunting classical inhibition could also plausibly underlie CGC, if combined with STD (Rothman et al., 2009), balanced by excitation or as a result of the mapping sensory input onto firing rate via a power-law I-O function (Murphy and Miller, 2003). A major strength of intracortical inhibition as a mechanism for CGC is that it is capable of accounting for both the local and global components of CGC in a single mechanism. It satisfies the criteria for explaining local CGC that it operates at a cellular level. It is also capable of accounting for the global component of CGC as sources of intracortical inhibition sample more broadly from the cortical network than pyramidal cells (Atencio and Schreiner, 2008; Moore and Wehr, 2013; Liu et al., 2011).

If inhibition underlies CGC in A1, where does this inhibition originate? Inhibitory interneurons comprise approximately 20% of neurons in the cortex, yet they exhibit extensive diversity and consist of an unknown number of subtypes

(Markram et al., 2004; Klausberger and Somogyi, 2008). Given their diverse properties, it is thought that different subtypes serve different functional roles in the cortex, raising the possibility that a particular subtype of cortical interneuron may be specialized for stimulus-driven gain control.

Historically, attempts to identify neuronal subtypes have been limited by the particular methodology available. Initial attempts depended on Golgi staining methods and microscopy. The reliance on these methods led to cellular morphology and anatomical differences in the targeting of subcellular compartments being used as the basis for classification (Somogyi et al., 1998; Markram et al., 2004). The development of single-cell electrophysiological techniques led researchers to classify neurons based on the physiological properties of these neurons, such as firing pattern (McCormick et al., 1985; Connors and Gutnick, 1990). While anatomical and electrophysiological methods of classification show some correlation, neurons of the same morphological subtype can demonstrate a variety of firing patterns (Kawaguchi, 1993, 1995; Kawaguchi and Kubota, 1997; Gupta et al., 2000; Wang et al., 2002).

In recent years, advances in molecular biology have resulted in gene expression patterns also being used as criteria for classification (Rudy et al., 2011). While attempts have been made to integrate these three approaches into a single classification system (Ascoli et al., 2008), the methodological advantages in targeting subtypes based on gene expression have led to this becoming a widely accepted method for separating neuronal subtypes in sensory neuroscience (Taniguchi et al., 2011). While no simple mapping exists between gene expression patterns, firing patterns and morphology, strong correlations do exist between these features. This is most likely a result of the molecular phenotype contributing to the generation of the anatomical and physiological phenotypes (Nelson et al., 2006; Sugino et al., 2006; Taniguchi et al., 2011).

Three classes of interneurons have been found to account for almost 100 % of interneurons in the cortex (Rudy et al., 2011) (Figure 1.4). One class expresses the calcium-binding protein parvalbumin (PV) and comprises approximately 40 % of cortical interneurons. Another class expresses the neuropeptide somatostatin (SST) and makes up approximately 30 % of the interneuron population. The final 30 % of interneurons express the ionotropic serotonin 3a receptor (5HT3aR). Each of these groups are themselves comprised of multiple subtypes that vary anatomically, physiologically and in their molecular profile (Rudy et al., 2011; Taniguchi et al., 2011). While this may be seen as a relatively crude method classification, variation in these phenotypes is far greater between these three groups than within them, giving this approach a level of face validity.

The 5HT3aR group represents the largest group of interneurons in layers 1-3 of the cortex and can be subdivided into two major subtypes, based on whether or not they express vasoactive intestinal peptide (VIP). The majority of 5HT3aR neurons, around 60 %, are non-VIP expressing cells and this category includes over 90 % of layer 1 interneurons (L1) (Rudy et al., 2011). These non-VIP 5HT3aR interneurons tend to express neuropeptide Y (NPY) (Vucurovic et al., 2010) and reelin (Rudy et al., 2011), have a multipolar neurogliaform morphology (Kawaguchi and Kubota, 1997; Oláh et al., 2007) and a late spiking firing pattern (Rudy et al., 2011). VIP interneurons make up the remaining 40 % of the 5HT3aR group and are primarily located in layer 2/3 (Rudy et al., 2011; Taniguchi et al., 2011). These neurons often coexpress the neuropeptide cholecystokinin (CCK) (Férézou et al., 2002; Inta et al., 2008; Varga et al., 2009; Vucurovic et al., 2010), they have a bipolar/bitufted morphology and can show either an irregular spiking or fast adapting firing pattern. VIP interneurons have been shown to primarily target SST interneurons, innervating pyramidal cells weakly with a low probability of $\sim$12.5 % (Pfeffer et al., 2013).

SST interneurons preferentially innervate the dendrites of pyramidal cells and

all classes of interneuron in cortex (Pfeffer et al., 2013). The vast majority of SST cells have been classified morphologically as Martinotti cells. These interneurons target the distal dendritic tufts of pyramidal cells in layer 1, while having their cell bodies in all other layers, predominantly in layer 5. SST interneurons of layers 5 in particular tend to show a low-threshold intrinsic bursting firing pattern while the remaining SST interneurons tend to show a regular adapting firing pattern (Kawaguchi and Kubota, 1997; Wang et al., 2004; Uematsu et al., 2008). Martinotti cells can be themselves subdivided into two subtypes depending on whether or not they coexpress calretinin (CR). SST/CR interneurons differ from CR negative SST interneurons in their a spiking properties (Xu et al., 2006), connectivity (Xu and Callaway, 2009) and developmental origins (Fogarty et al., 2007; Sousa et al., 2009), making the case for their being independent subtypes convincing. A subpopulation of non-Martinotti SST interneurons also exists in cortex. These neurons were identified in the X94 transgenic mouse line, generated through random insertion of transgene that expresses GFP under a GAD67 promoter (Ma et al., 2006). X94 interneurons innervate thalamorecipient layer 4 and have their cell bodies in layers 4 and 5. Unlike Martinotti cells, they have low input resistance and fire with a high frequency stuttering firing pattern that also exhibits spike adaptation (Helmstaedter et al., 2009a).

PV interneurons consist of two major subtypes that show distinct anatomical profiles. By far the most common subtype is the basket cell, which targets the soma and proximal dendritic region of pyramidal cells. The other major subtype of PV interneuron is the chandelier cell, which targets the axon initial segment of pyramidal cells. Both of these subtypes show a distinctive fast-spiking (FS) phenotype, although they do show some differences in their electrophysiological properties (Woodruff et al., 2009). The FS phenotype is characterized by high firing rates, narrow action potential width and is sustained with little spike frequency adapta-

tion and has repeatedly been shown to correlate strongly with PV expression (Cauli et al., 1997; Kawaguchi and Kubota, 1997; Gibson et al., 1999; Ascoli et al., 2008; Xu and Callaway, 2009). As the most numerous group of interneuron, basket cells may themselves consist of several subtypes (Markram et al., 2004). Large, small and nest basket cells have been identified based on the extent and pattern of their arborisations (Markram et al., 2004; Uematsu et al., 2008; Helmstaedter et al., 2009b; Hellwig et al., 2012). Basket cells also show variability in their molecular profiles (Kawaguchi and Kubota, 1997; Chow et al., 2010; Vruwink et al., 2001; Saganich et al., 2001) and in their firing properties (Goldberg et al., 2008; Puig et al., 2008; Hellwig et al., 2012) but there is currently no consensus as to whether this variation reflects diversity of subtypes or variability within a single interneuron class (Golomb et al., 2007).

1.3.2 Interneuron subtypes and gain control

Are any of these interneuron subtypes plausible candidates for the source of inhibition that may underlie CGC in A1? Certain differences between these subtypes indicate that they may have significantly different functional specialisations. 5HT3AR expressing interneurons appear to be specialised for top-down disinhibition, via inhibition of other interneuron subtypes (Pfeffer et al., 2013; Pi et al., 2013; Fu et al., 2014; Zhou et al., 2014b). VIP interneurons in V1 have been shown to primarily target SST interneurons but also innervate pyramidal cells weakly and with a low $\sim$12.5 % probability (Pfeffer et al., 2013). Innervation of other interneuron subtypes by VIP interneurons is marginally stronger than innervation of pyramidal cells, but is still weak.

Locomotion has been found to result in activation of VIP interneurons in V1, S1 and A1 (Fu et al., 2014). In V1, this activation appears to be mediated by cholinergic inputs from the basal forebrain and results in an increase in the gain of

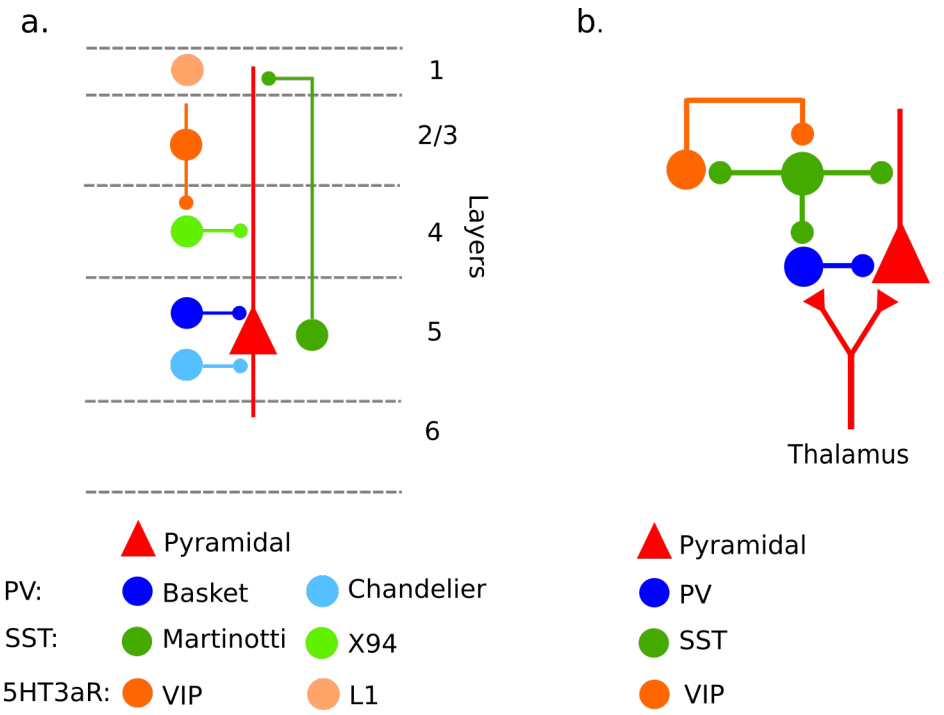

Figure 1.4: A. PV expressing interneurons can be divided into basket (dark blue) and chandelier cells (light blue) and target the soma and AIS of pyramidal cells respectively. The majority of SST expressing interneurons are Martinotti cells (dark green) which tend to be concentrated in deep layers and innervate the distal dendritic tufts of pyramidal cells in layer 1. These interneurons may also be divisible into CR+ and CR- subtypes. X94 interneurons represent another, non-Martinotti, SST subtype (light green). Almost half of 5HT3aR neurons (40 %) express VIP. These interneurons tend to be concentrated in layers 2/3, have a bipolar/bitufted morphology and primarily innervate SST interneurons. The majority of non-VIP expressing 5HT3aR neurons are L1 interneurons that express NPY and reelin. B. A simplified classification scheme based on 3 molecular markers alone accounts for over 80 % of interneurons in cortex (Pfeffer et al., 2013). PV interneurons receive inputs from thalamus and innverate pyramidal cells on and near their cell bodies, as well as inhibiting other PV interneurons. SST interneurons target the dendrites of pyramidal cells and inhibit all other classes of interneuron. VIP interneurons primarily target SST interneurons.

pyramidal cell responses (Polack et al., 2013; Fu et al., 2014). Artificial activation of VIP interneurons in this area while the animal is stationary has been found to increase the gain of pyramidal cell responses (Fu et al., 2014). Similarly, VIP interneurons in mouse A1 have been found to be activated by reinforcement signals, resulting in an increase in the gain of pyramidal cell responses via inhibition of SST and a subset of PV interneurons (Pi et al., 2013). While in V1 locomotion results in

an increase in gain and an increase in the signal-to-noise ratio of visual responses (Bennett et al., 2013), locomotion has been found to result in a reduction in gain in A1 (Zhou et al., 2014b). This occurs as a result of L1 interneuron activation and suppression of layer 2/3 excitatory responses and PV interneuron responses that results in gain modulation through a scaling down of balanced excitation and inhibition (Zhou et al., 2014b). These findings indicate that, while VIP interneurons are unlikely to implement gain control of pyramidal cell responses directly, they are capable of regulating gain via the inhibition of interneuron subtypes that are mechanistically involved in gain control.

The targeting of SST interneurons by VIP interneurons in V1 has implicated SST interneurons in gain control (Fu et al., 2014). The effects of SST interneuron activity on pyramidal cell responses in V1 has been investigated through the use of optogenetic techniques. Optogenetic modulation of neuronal subtypes involves targeting light-sensitive constructs to genetically defined groups of neurons through either viral or transgenic methods. Depending on the function of the protein encoded by the construct, exposure to specific wavelengths of light can result in either activation, as with channelrhodopsin (ChR2) (Boyden et al., 2005) or inhibition, as with Halorhodopsin (Halo) (Zhang et al., 2007) or archearodopsin (Arch) (Chow et al., 2010). Lee et al (Lee et al., 2012) found that ChR2-mediated transient activation of SST interneurons at the onset of visual stimulation results in a subtractive narrowing of tuning width (Wilson et al., 2012; Lee et al., 2014). Tonic activation of SST interneurons during visual stimulation however has been found to scale down the responses of pyramidal cells without affecting tuning width, reflecting a change in gain. Unlike PV interneurons, SST interneurons are not routinely co-active with pyramidal cells, but tonic optogenetic stimulation of this kind forces them to be. SST interneuron mediated gain control may therefore result from an artificial E-I balance created through this manipulation (El-Boustani et al., 2014), although this

issue remains controversial (Lee et al., 2014; El-Boustani et al., 2014).

PV interneurons have been found to play many functional roles in cortex. They are targeted by thalamic input in all layers of cortex and provide fast and powerful feed-forward inhibition to excitatory neurons (Ji et al., 2015). Their fast synaptic responses are the result of their low input resistance and fast membrane time constant (Connors and Gutnick, 1990; Cauli et al., 1997; Gibson et al., 1999; Kawaguchi and Kubota, 1997; Ascoli et al., 2008; Goldberg et al., 2008). These features make them crucial for increasing the precision of sensory evoked cortical responses (Pinto et al., 2000; Miller et al., 2001; Pouille and Scanziani, 2001; Pinto et al., 2003; Lawrence and McBain, 2003; Gabernet et al., 2005) and the prevention of propagating epileptiform activity (Sessolo et al., 2015). This feed-forward configuration, combined with mechanisms of synaptic plasticity that are dependent on post-synaptic activity levels, result in PV interneurons underlying the balance of excitation and inhibition in cortex (Xue et al., 2014; Hasenstaub et al., 2005; Haider and McCormick, 2009). They have also been shown to be involved in the generation of gamma-frequency oscillatory activity (Traub et al., 2004; Bartos et al., 2007; Cardin et al., 2009) and the regulation of critical-period plasticity in cortex (Hensch, 2005).

In addition to these many roles, there is also compelling evidence for PV interneurons underlying inhibition-mediated gain control in cortex. Feed-forward recruitment of putative PV interneurons, identified through their FS phenotype, has been shown to modulate the dynamic range of post-synaptic neurons in both hippocampus and cortex in vitro (Pouille et al., 2009). In vivo, optogenetic techniques have been used in order to test the arithmetic effects of PV interneuron activity on pyramidal cell responses in V1. Light-driven activation of PV interneurons with ChR2 has been found to reduce the gain of visually evoked pyramidal cell spiking responses while light-driven suppression with Arch increased gain (Atallah et al., 2012; Wilson et al., 2012). Importantly, these manipulations did not affect the

tuning widths of these cells but scaled them, reflecting a gain change. The effects of PV interneuron activity appeared to be attributable to broadly tuned subtractive inhibition combined with a power law output function (Atallah et al., 2012). The possibility that these effects were mediated through shunting inhibition was not directly tested however. The divisive nature of PV cell activation was found both when populations of these cells were activated and when individual PV interneurons were activated (Wilson et al., 2012).

When activated more strongly, PV interneurons can have a subtractive effect on pyramidal cell responses, resulting in a sharpening of orientation tuning (Lee et al., 2012, 2014). This occurs when pyramidal cell firing rates are reduced by more than half, at which point much of the excitatory drive to the neuron is pushed below threshold producing a subtractive 'iceberg effect' (Lee et al., 2014; Atallah et al., 2014). Confirmation of the divisive effect of PV interneuron activity came from the finding that strong suppression of PV interneurons results in a multiplicative scaling of pyramidal cell responses, as threshold does not effect the tuning width under these conditions (Atallah et al., 2012, 2014). These findings indicate that, under typical physiological conditions, the primary effect of PV interneuron activity on pyramidal cell responses is to modulate gain.

In keeping with their potential role in gain control, PV interneurons appear to be ideally suited to providing shunting inhibition. The inhibition provided by these cells occurs primarily via GABAa receptors (Klausberger et al. 2002) which are mainly permeable to chloride (Cl^-). The reversal potential for chloride determines the effect of inhibitory currents on the post-synaptic neuron, and the reversal potential is determined by the Cl^- concentration across the cell membrane. If the Cl^- reversal potential is below the resting membrane potential (RMP) of the neuron, opening of these channels will result in an outward flow of current, producing a hyperpolarising effect. If the reversal potential is close to the RMP and below the spike

initiation threshold, GABAa receptor transmission will have a shunting effect. The reversal potential for chloride at the pyramidal cell soma in vivo is thought to be approximately -70mV, very close to the RMP. This appears to allow PV interneurons to provide shunting conductances to cortical pyramidal cells. The somatic innervation of pyramidal cells by PV interneurons also appears to be suited to gain modulation. Innervation of the soma also enables these cells to modulate the spiking responses of pyramidal neurons without affecting their tuning (Isaacson and Scanziani, 2011; Fino et al., 2013). PV interneurons may also produce shunting conductance changes as a result of their balanced configuration with excitatory input (Xue et al., 2014).

1.3.3 The contrast signal

In order for CGC to be implemented in cortex, thalamocortical inputs must carry a contrast-dependent signal that can drive the mechanism that underlies contrast gain control. Since an increase in contrast reflects an increase in the spectral and temporal variance in the strength of the sensory input, this variability may be reflected in pre-synaptic firing rates. In rodents, the majority of neurons in A1 respond to increasing intensity with a monotonic increase in firing rate (Bakin and Weinberger, 1990; LeDoux and Semple; Polley et al., 2004) allowing them to encode temporal contrast using the variance of the neuron's firing rate over time (Rees and Møller, 1983; Joris and Yin, 1992; Eggermont, 1994; Liang et al., 2002; Nelson and Carney, 2007; Malone et al., 2007). Given the rectification of the input signal that occurs during spike generation, increased contrast is also likely to be encoded as an increase in mean firing rate in neurons that do not show CGC, as has been observed in the IC (Escabí et al., 2003). Contrast-dependent modulation of the mean and variance of presynaptic firing rates would allow feedforward recruitment of PV interneurons by thalamic inputs (Ji et al., 2015) to modulate gain in cortical neurons. In A1 PV interneurons have been found to show more linear response to stimulus

intensity than pyramidal cells, indicating that they are well-suited to translating increased mean drive into inhibitory output that may modulate gain (Moore and Wehr, 2013).

Spectral contrast may produce a similar increase in mean firing rate as PV responses are believed to indiscriminately pool excitatory inputs from the local population (Yoshimura and Callaway, 2005; Hofer et al., 2011; Wilson et al., 2012). The variability of PV interneuron responses may be unaffected by spectral contrast as the instantaneous stimulus intensity in neighbouring frequency channels should cancel out. This should only occur if stimulus level at each time point is uncorrelated across frequency channels, however, as it is in stimuli known to elicit CGC (Rabinowitz et al., 2011b). Integration over the local population may have the effect of elevating their membrane potential near to threshold under high contrast stimulation, allowing these neurons to fire rapid non-adapting trains of action potentials. This may allow PV interneurons to modulate gain in a highly dynamic manner. Whether or not PV interneurons pool local pyramidal cell responses in this manner can be assessed from their spectral tuning properties. Evidence for the pooling hypothesis comes from V1 where PV interneurons are well tuned in species that have show a columnar organisation of orientation tuning (Hirsch et al., 2003; Cardin et al., 2007; Nowak et al., 2008), but show broad tuning in mouse V1 (Ma et al., 2010; Kuhlman et al., 2011; Zariwala et al., 2011; Wilson et al., 2012) where local populations show a range of orientation preferences (Kerlin et al., 2010; Hofer et al., 2011).

In mouse A1, there is debate over both the extent of tonotopic frequency organization at the local level (Guo et al., 2012; Bandyopadhyay et al., 2010; Rothschild et al., 2010) and over the width of PV interneuron tuning (Moore and Wehr, 2013; Li et al., 2014). Tonotopic organisation is observed in mouse A1 at the scale of hundreds of microns (Guo et al., 2012) but 2-pi imaging studies have revealed het-

erogeneous tuning at the local scale in layers 2/3, with neighbouring neurons showing strikingly different tuning preferences (Bandyopadhyay et al., 2010; Rothschild et al., 2010; Issa et al., 2014). This finding appears to be the result of layer specific differences in intracortical circuitry, as layer 4 has been found to be tonotopic at the local scale (Winkowski and Kanold, 2013).

If PV interneurons sample the responses of local excitatory cells in an indiscriminate manner, one might expect PV interneurons in layer 4 to show frequency tuning similar to pyramidal cells, and PV interneurons in layers 2/3 to show broader tuning. Intracellular recordings from fast-spiking putative PV interneurons in layer 4 of rat A1 have shown that the width of subthreshold tuning is indeed similar to that of pyramidal cells, although as a result of a lower intensity threshold, the tuning of their spiking response is in fact broader than that of pyramidal cells (Wu et al., 2008). Broader tuning of fast-spiking putative PV interneurons have been observed across all cortical layers in cat A1(Atencio and Schreiner, 2008).

Moore and Wehr (2013) assessed tuning properties of PV interneurons in mouse A1 by performing tungsten electrode recordings in transgenic mice that expressed ChR2 in PV interneurons. This method, known as photo-identification of neuronal populations (PINP) (Lima et al., 2009) allowed them to identify these neurons from their responses to light and assess their tuning properties. They found that fast-spiking neurons did indeed show broader tuning than pyramidal cells but when this analysis was restricted to light-responsive PV interneurons, tuning was as narrow as in pyramidal cells both within and outside layer 4 (Moore and Wehr, 2013). Using 2-pi guided targeting of cell attached recordings to layer 2/3 PV and SST interneurons, Li et al (Li et al., 2014) found that tuning is in fact broader for PV interneurons than both pyramidal cells and SST interneurons. Cell attached recording is the gold-standard technique for single unit isolation and indicates that the finding of similar tuning between PV interneurons and pyramidal neurons may be

due to contamination in the PV recordings by pyramidal cell spiking activity. One study that used this approach to assess PV interneuron tuning in mouse V1 found that these neurons are in fact well tuned (Runyan et al., 2010), contrary to many other reports (Ma et al., 2010; Kuhlman et al., 2011; Zariwala et al., 2011; Wilson et al., 2012). Further work will be necessary in order to identify the causes of these variations between studies.

While there is on-going debate as to the breadth of tuning of PV interneurons in A1, there is a consensus that PV interneurons are tuned for frequency to some extent. This property allows them to account for local CGC, in which the range of contrasts that drive CGC are similar to the tuning of the excitatory cell undergoing gain control (Rabinowitz et al., 2011a). This feature may also be attributable to the somatic innervation patterns of pyramidal cells by PV interneurons. By modulating the sensitivity of the neuron at the soma, following dendritic integration of synaptic inputs, gain can be changed equally across all inputs while tuning is left unaffected. Even if PV interneurons show pyramidal-like tuning, these interneurons show extensive connections through gap junctions (Galarreta and Hestrin, 1999; Gibson et al., 1999; Tamás et al., 2000; Beierlein et al., 2000; Deans et al., 2001; Galarreta and Hestrin, 2001; Meyer et al., 2002; Galarreta and Hestrin, 2002; Amitai et al., 2002; Gibson et al., 2005; Hestrin and Galarreta, 2005) allowing them to respond as a population. This feature may underlie the global aspect of CGC, in which contrast outside the receptive field of the excitatory neuron can modulate gain (Rabinowitz et al., 2011a).

If thalamocortical inputs provide the contrast-dependent signal that drives CGC, this computation may be implemented through feedforward PV-mediated inhibition in layer 4, where innervation of PV interneurons by thalamic input is strongest (Figure 1.5 A). A contrast-dependent increase in mean spiking may also allow STD to contribute to CGC in layer 4 by acting at the thalamocortical synapse (Stratford

et al., 1996; Boudreau and Ferster, 2005). STD may also serve to increase contrast-dependent variability in the input, a mechanism of gain control that has been shown to contribute to contrast-invariance in V1 (Sadagopan and Ferster, 2012). Implementing gain control at this earliest possible stage of cortical processing in this manner has the effect of normalising cortical input so that it can be modified in later stages of cortical processing outside of layer 4.

If thalamocortical circuitry alone underlies CGC, these gain changes may simply be inherited outside of layer 4 (Douglas and Martin, 1991b), resulting in gain changes across the depth of cortex. PV interneurons have been shown to receive thalamic input across the entire depth of the cortex, potentially enabling CGC to be implemented via feedforward inhibition across all cortical layers simultaneously (Figure 1.5 B). Another possibility is that CGC is initially implemented by thalamocortical circuitry and is subsequently boosted by intracortical mechanisms outside of layer 4, resulting in observable differences in the strength of CGC across cortical layers (Figure 1.5 C). Such differences may be even stronger if intracortical mechanisms alone are involved in the generation of gain control. Layer 4 responses may vary with contrast, allowing these neurons to drive CGC in other layers via intracortical excitation, STD or feedforward intracortical inhibition (Barbour and Callaway, 2008; Adesnik et al., 2012).

In mouse V1, corticothalamic layer 6 excitatory neurons have been found to modulate the gain of pyramidal cells in all other cortical layers (Olsen et al., 2012) through the recruitment of a translaminar-projecting, PV interneuron subtype, whose cell bodies also reside within layer 6 (Bortone et al., 2014). This deepest cortical layer receives direct thalamic input, which may drive gain changes across the cortical column simultaneously via this intracortical circuit (Figure 1.5 D). Unlike excitatory inputs to PV interneurons in other cortical layers which tend to show STD, layer 6 cortico-thalamic pyramidal cells preferentially innervate interneurons

via facilitating synapses (Gil et al., 1999; Beierlein and Connors, 2002; Thomson and Lamy, 2007) indicating that this may represent a highly specialized circuit for gain control. If this circuit underlies CGC in A1, CGC should be absent from layer 6 neurons and contrast dependent modulation of firing rates or variability should be observed in this layer.

These different circuit hypotheses lead to different predictions regarding the strength of CGC across the layers of the cortex. Knowledge of the laminar organisation of CGC may therefore serve to constrain the possible mechanisms and columnar circuitry that may underlie CGC.

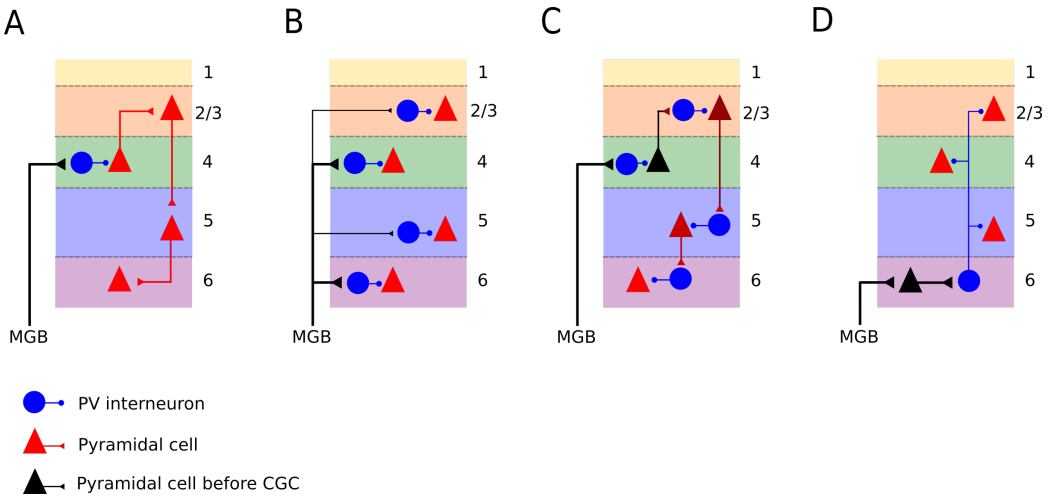

Figure 1.5: A. CGC may be implemented within layer 4 via feedforward PV interneuron inhibition, synaptic depression, synaptic noise or a combination of these mechanisms. Gain changes may be inherited subsequently by pyramidal cells in other cortical layers. B. Feedforward PV interneuron inhibition may occur simultaneously across all cortical layers. C. Alternatively, PV interneurons may be recruited sequentially in each layer by intracortical excitatory connections. This could lead to a gradual increase in the strength of gain control, indicated here by the transition from black to red in the pyramidal cells. D. A subpopulation of thalamorecipient layer 6 excitatory neurons that project to thalamus may drive gain changes in all other cortical layers via feedforward inhibition, mediated by a translaminar projecting PV interneuron (adapted from (Bortone et al., 2014)). In this case, CGC should be absent from layer 6. MGB: medial geniculate body.

1.4 Organisation of this thesis

The overall aim of this work was to investigate the mechanistic basis of CGC in AC. The tools required in order to do this are not available in the ferret, where CGC was first identified, and so the first aim of this work was to establish the mouse as a model organism for the investigation of mechanisms of CGC. This involved characterising CGC in mouse AC through extracellular recordings. I also examined the laminar organisation of CGC in mouse AC in order to constrain circuit hypotheses regarding the mechanisms underlying CGC. This information was also important in allowing us to identify whether certain neuronal populations in specific layers needed to be targeted in future experiments. The second aim was to investigate the impact of PV interneuron activity on responses in AC, in order to test the hypothesis that this cell type drives reductions in gain in response to increases in stimulus contrast. Finally, I aimed to test the involvement of two possible biophysical mechanisms in CGC through intracellular recordings, shunting inhibition and membrane potential variability.

In Chapter 2, I outline the experimental methodology employed for both extracellular and intracellular electrophysiological experiments in anaesthetised mice and cover the optogenetic methods used in a subset of experiments. I also outline the stimuli and analytical approaches that were used in this work. These include Dynamic Random Chord stimuli and linear and nonlinear models that were used in order to estimate the stimulus-response function of extracellular recorded AC neurons.

In Chapter 3, I examine whether CGC occurs in mouse AC and whether any interspecies differences exist between the mouse and the ferret. I also examine how CGC is organised across cortical layers in mouse AC.

In Chapter 4, I test the effects of PV interneuron inhibition on extracelluar recorded unit responses through optogenetic activation and suppression of PV in-

terneuron activity. I also examine the effect of contrast on putative PV interneuron activity.

In Chapter 5, I examine the biophysical basis of CGC though intracellular recordings. I assess the contributions of shunting inhibition and membrane potential variance to CGC in AC neurons.

Finally, in Chapter 6, I discuss what these experimental results tell us about the mechanistic basis of CGC in AC and more generally in cortex. I also identify areas for future investigation.

Chapter 2

Methods

2.1 Experimental Methods

2.1.1 Surgical methods

All extracellular and intracellular recordings reported in this thesis were made from the auditory cortex (AC) of anaesthetised mice, all of which were between 8 and 12 weeks old. Induction of anaesthesia was achieved through intraperitoneal injection of 0.1 mg/kg ketamine (Narketan, Chassot AG, Belp, Bern, Switzerland) and 0.25 mg/kg medetomidine (Domitor, Pfizer, Walton Oaks, Surrey, UK). 0.1 mg atropine (Atrocare, Animal Care ltd, Nether Poppleton, York) and 0.01 mg dexamethasone (Dexadreson, MSD Animal Health, Milton Keynes, Buckinghamshire) per animal were delivered subcutaneously in order to reduce cerebral edema and lung secretions, respectively. Anaesthesisa was maintained via intraperitoneal infusion of 50 mg/kg/hr ketamine and 0.06 mg/kg/h medetomidine. The scalp was anaesthetized locally using 0.1 μg marcain (AstraZenica, London, UK) before being removed in order to expose the left temporal and parietal bones of the skull. The left temporal muscle was anaesthetized locally using 0.1 mg marcain and subsequently retracted, using a diathermy in order to cauterize any bleeding. The skull was fixed to a head-

plate using a combination of tissue adhesive (Vetbond, Santa Cruz Animal Health, Paso Robles, CA, United States) and dental acrylic.

A craniotomy was performed over AC, the location of which was identified using cranial landmarks. Craniotomies extended from the lambdoid suture at the most caudal extent to 1 mm rostral of the point at which the squamosal suture crosses the temporal ridge. In the dorsal-ventral axis, the craniotomy extended from 2mm dorsal of the temporal line to the squamosal suture at the most ventral extent. Local vasculature was also used in order to localise AC. Sites close to the largest vessel in this area oriented along the dorsal-ventral axis were consistently responsive to auditory stimuli. It was not possible to map specific auditory cortical fields in these experiments and so AC recordings likely include responses from secondary as well as primary cortical areas. Frequency tuning of auditory responses was also used to confirm that recordings were localised to AC.

Following extraction of multi-unit activity on each channel, 56% of recordings showed noise-evoked responses, defined as a doubling of baseline firing in response to noise onset. For intracellular recordings the dura was removed and saline was applied to the exposed cortex in order to avoid dehydration. Silver-plated reference and ground wires were inserted into frontal and parietal cortices, respectively, through two small craniotomies. These wires were fixed in place using dental acrylic. Core body temperature was monitored throughout the experiment and maintained at 37.5 $\pm$0.5 degrees C. Anaesthesia levels were assessed regularly via the pedal withdrawal (paw pinch) reflex and any signs of whisker movement. All procedures received approval from the Oxford University Animal Care and Ethical Review Committee and were licensed by the UK Home Office.

2.1.2 Extracellular recordings

Extracellular recordings were performed using multi-channel silicon probes (NeuroNexus, Ann Arbor, United States). Single shank probes with 32 recording sites spaced 50 μm apart and arranged in a linear configuration (1x32) were used in order to record neural responses from all cortical layers simultaneously. Single units could not be reliably extracted from the data obtained using linear probes, due to the geometry of their recording sites and so the small amplitude of recorded spiking activity. Multi-unit activity (MUA) was instead extracted from recordings. MUA is less precise a measure than single-unit spiking and is more likely to suffer from false positives, with non-spiking currents and electrical noise being classified as spiking activity. It is routinely used however and often acts as reasonable proxy for single-unit activity (Guo et al., 2012).

An analog measure of MUA recorded on each channel was used instead of traditional threshold-crossing methods for MUA extraction. The analog MUA extraction method used here involved measuring the voltage signal power in the frequency range occupied by spiking activity by bandpass filtering the recorded signals. This method has been used in several previous studies in order to extract MUA (Chung et al., 1987; King and Carlile, 1994; Schroeder et al., 1998; Choi et al., 2010; Schnupp et al., 2015). For each channel, the recorded voltage signal was filtered between 300 and 6000 Hz. The absolute values comprising the signal were taken and the resulting signal was subsequently low-pass filtered below 6000 Hz and downsampled to a sample rate of 12000 Hz. Local field potentials (LFPs) were extracted by low-pass filtering the recorded signals below 300Hz using a digital 8^{th} order Chebyshev Type I filter. "Buzsaki" style NeuroNexus probes were used for single unit recordings. These probes had 8 shanks, each with 8 sites arranged in a staggered configuration (8x8) designed in order to isolate single units. For the Buzsaki probes, 136 units were isolated across 6 penetrations, resulting in a yield

of ∼23 well-isolated single units per penetration. Extracellular recordings were performed in a sound-insulated chamber. A TDT PZ2-64 multichannel preamplifier (Tucker Davis Technologies, Alachua, FL) was used in order to amplify extracellular recordings and a TDT RZ2 digital signal processor was used to digitize the data at a sample rate of ∼25 kHz.

Single units were isolated in some experiments by spike sorting. Automated spike sorting was performed using KlustaKwik (github.com/klusta-team/klustakwik), which fits principal components in the data with a mixtures-of Gaussians model using an Expectation Maximisation (EM) algorithm. After automated clustering, clusters were visually inspected and classified as noise, single units or multi-units (MUs) on the basis of spike waveform, distribution over channels and autocorrelation structure. Signals from single units were required to have a waveform that was restricted to a small number of channels and to show a 1 ms refractory period in their autocorrelation. Clusters that did not meet these criteria were classified as multi-units, unless their waveforms indicated that they reflected non-biological electrical activity, in which case they were classified as noise.

2.1.3 Stimulus presentation

Stimuli were presented using an Ultrasonic Dynamic Speaker (Avisoft Bioacoustics, Glienicke, Germany), modified for monaural in-ear delivery. The speaker was driven using a TDT RX6 multifunction processor at a sample rate of ∼200 kHz. Stimuli were calibrated across the frequency range of 1-64 kHz. A Bruel and Kjaer (Naerum, Denmark) Type 4138 $1/8^{th}$ inch pressure-field microphone was used in order to assess the response of the speaker across this range. An inverse filter was created based on this response and was used in order to produce a flat frequency response in this frequency range, within ±3 dB. Stimulus presentation and data acquisition was controlled using Benware (github.com\beniamino38\benware). Sub-

sequent analysis was performed in Matlab (The MathWorks, Natick, MA).

2.2 Stimuli

2.2.1 DRC stimuli

Throughout this thesis, the main stimulus used is the dynamic random chord sequence (DRC). This stimulus consists of a series of chords that themselves consist of a superposition of pure tones (Figure 2.1 A and B). The frequencies of these tones are fixed, as is the duration of each chord. The intensity of each tone at any time point, however, is drawn from a distribution of possible values (Figure 2.1 C), leading to fluctuations in intensity over time for each tone (Figure 2.1 E). This method also generates variations in intensity between neighbouring tones (Figure 2.1 D). Changing the range of possible intensities therefore controls both the temporal and spectral contrast of the stimulus.

All DRCs in this study consist of 25 pure tones with frequencies log-spaced between 1 and 64 kHz, inclusive. The frequency range of the stimulus covered 6 octaves and tone frequencies were spaced 1/4 octave apart. Chord duration was fixed at 25 ms and 5 ms linear ramps were included between sequential chords in order to reduce “spectral splatter”.

In order to investigate CGC using these stimuli, the contrast of each DRC had to be controlled. Traditionally, contrast, c, has been defined as the standard deviation of the stimulus intensity, σ_I, over the mean intensity, μ_I:

$$c = \sigma_I / \mu_I \tag{2.1}$$

It is therefore possible to define auditory contrast as the standard deviation of the root mean square (RMS) sound pressure, σ_p, over the mean sound pressure, μ_p:

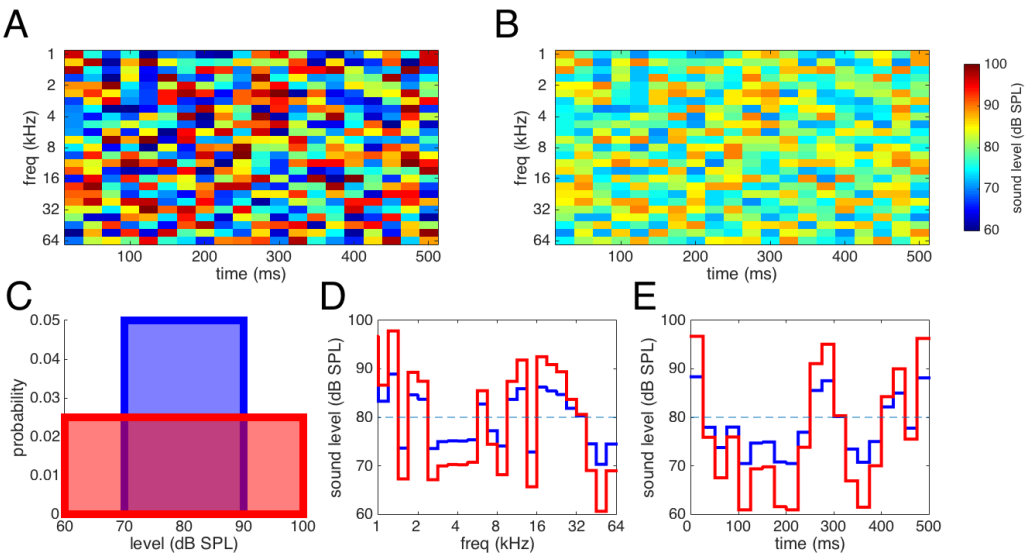

Figure 2.1: A. High contrast DRC stimulus. DRCs are comprised of pure tones of specific frequencies that vary in their level over time, indicated by colour on this plot. Levels for each tone at each time point are drawn randomly from a tone-level distribution. For this high contrast DRC, the range of possible intensities is ±20 dB. B. Low contrast dynamic random chord (DRC). The range of possible intensities for this low contrast DRC is ±10 dB. C. Tone level distributions used in order to generate DRCs in A and B. The mean level for both distributions are the same, but the widths of the distributions, and therefore the standard deviations, are different. The increased width of the red distribution results in the generation of a higher contrast DRC (A) than that generated by the blue distribution (B). D. Cross-section though the DRC across frequencies for a given time point. At each time point, the width of the tone-level distribution also controls the variation in level across frequency. E. Cross-section though the DRC within a single frequency band. For a given frequency, the intensity fluctuates randomly over time. These fluctuations are larger for high contrast DRCs generated using a tone-level distribution with a larger width (red). This increase in level variance results in an increase in the temporal contrast of the stimulus. This method of DRC generation leads to stimuli with matched spectral and temporal contrast.

$$c = \sigma_p / \mu_p : \tag{2.2}$$

In auditory physiology or psychoacoustics, however, sound level is not usually quantified as pressure however, but as sound pressure level (SPL), L:

$$L = 20log_{10}(p/p_0) \tag{2.3}$$

where p_0 = 20 μPa, the reference RMS sound pressure used when calculating SPL.

This logarithmic measurement allows the very large range of pressure levels that the auditory system is exposed to to be expressed in a more convenient manner than is possible with a linear measurement. This log transform also conveniently results in SPL standard deviation being a very close approximation of sound pressure contrast.

$$c = \frac{\sigma_p}{\mu_p} \approx a.\sigma_L + b \tag{2.4}$$

where a and b are constants.

The standard deviation of the stimulus SPL, and therefore the contrast, can simply be controlled by changing the range of tone intensities that comprise the DRC (Figure 2.1 B,C). This method of controlling contrast changes both the variance in intensity across frequency bands (spectral contrast) (Figure 2.1 D) and the variance over time within frequency bands (temporal contrast) (Figure 2.1 E). Throughout this thesis I use DRCs with one of two contrast levels. Low contrast DRCs have an intensity range of ±10 dB (σ_L ~6.2 dB, c = σ_P/ μP = 0.68). High contrast DRCs have the same mean but the range was doubled to ±20 dB (σ_L ~11.97 dB, c = σ_P/ μ_P = 1.2). For these ranges, there is approximately a doubling in contrast from the low contrast to high contrast stimulus. Tone levels were generated using a variety of random seeds in order to generate different DRCs. Each DRC sequence consisted of 1,600 25 ms chords and lasted 40 seconds in total. Stimulus presentation was randomised (Figure 2.2).

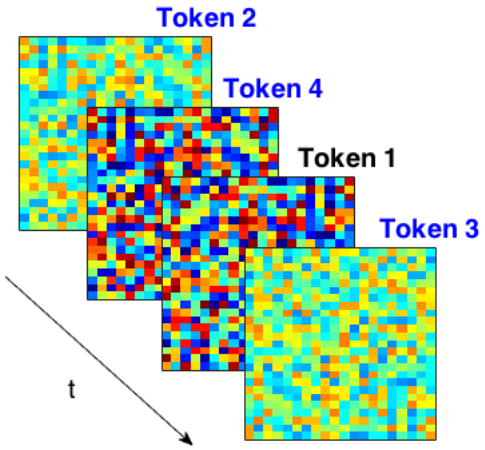

Figure 2.2: Schematic of DRC stimulus presentation. DRCs could be of high or low contrast with random sound levels determined by one of four seed tokens. Each DRC was presented 3 times with optogenetic stimulation and 3 times without optogenetic stimulation. Stimuli were presented in a random order. Experiment time (t) is indicated here by the arrow. This schematic shows four 40 second DRCs being presented sequentially with the token indicated above and the optogenetic condition indicated by the font color. Blue indicates a "light on" optogenetic condition while black indicates a "light off" control condition. Here a low contrast DRC seeded with token 2 during an optogenetic trial is followed by a high contrast DRC seeded with token 4 during a second optogenetic trial. This is followed by another contrast DRC seeded with token 1 during an control trial and finally a low contrast DRC seeded with token 3 during an optogenetic trial. Control data was used in chapter 3 while both data sets were used in chapter 4.

2.2.2 Noise stimuli

Bursts of white noise lasting 50 ms in duration were used in order to assess whether neurons responded to auditory stimuli. Evoked LFP responses to these noise bursts were also used in order to identify cortical layers. These noise bursts were also embedded in DRCs in order to probe evoked auditory responses during intracellular recordings.

2.3 Modelling

2.3.1 Linear models

DRCs have been used extensively in order to assess the spectrotemporal selectivity of auditory neurons (deCharms et al., 1998; Schnupp et al., 2001; Rutkowski et al., 2002; Sahani and Linden, 2003; Linden et al., 2003; Ahrens et al., 2008; Christianson et al., 2008; Bitterman et al., 2008) as has the spectrotemporal receptive field (STRF). In most mammals, neurons in the auditory system respond to frequencies within a restricted range as well as to temporal modulations, captured by the envelope of the sound. STRFs are models of neuronal responses that are capable of capturing both of these features (Shamma et al., 1993; Kowalski et al., 1995; Bizley et al., 2005). Fluctuations in stimulus energy within frequency bands over time can be represented as the spectrogram of an auditory stimulus, *X(f,t)* (Figure 2.4A). The STRF, *k*, is a linear filter that defines frequencies that drive or suppress neuronal responses over time, y_t (Figure 2.4B). The STRF captures spectral preferences of the neuron by assigning different weights to different frequency bins. At any given time in the stimulus, t, the STRF also weights the stimulus history, h, over a set number of time steps. This linear weighting in frequency and time yields a linear prediction of the neurons firing rate, $\hat{y}_t$. The hat operator denotes an estimate of the parameter in question.

$$\hat{y}_t = k_0 + \sum_{f,h}[k(f,h).X(f,t-h)] \tag{2.5}$$

As temporal tuning is modelled by weighting recent stimulus history at each time step, the addition of this history dimension results in the stimulus being represented

as a three-dimensional tensor, X_{tfh} (Figure 2.3):

$$X_{tfh} = X(f, t-h) \tag{2.6}$$

The STRF, k_{fh}, captures features in the two dimensional frequency-history space that drive responses over time. Linear filtering of the stimulus, X_{tfh}, by the STRF k_{fh}, yields the predicted firing rate, $\hat{y}_t$.

$$\hat{y}_t = k_0 + \sum_{f,h} X_{tfh}.k_{fh} \tag{2.7}$$

2.3.2 Linear-nonlinear models

The linear STRF is capable of capturing important aspects of neuronal responses to auditory stimuli. These models are limited, however, by their linearity. The relationship between the spectrogram of a sound and neuronal responses is not entirely linear and this must be incorporated into models in order to improve their accuracy. One approach that can be used in order to address this issue is to add a nonlinear stage to the model, following the linear STRF. Instead of using the STRF to predict firing rates directly, the output of the STRF can be treated as a variable within the model, z_t (Figure 2.4).

$$z_t = \sum_{f,h} X_{tfh}.k_{fh} \tag{2.8}$$

The output of the STRF, z_t, can then be passed through a static sigmoidal output non-linearity (Figure 2.4C):

$$F\big[z_t\big] = a + \frac{b}{1 + e^{-\big(z_t - c\big)d}} \tag{2.9}$$

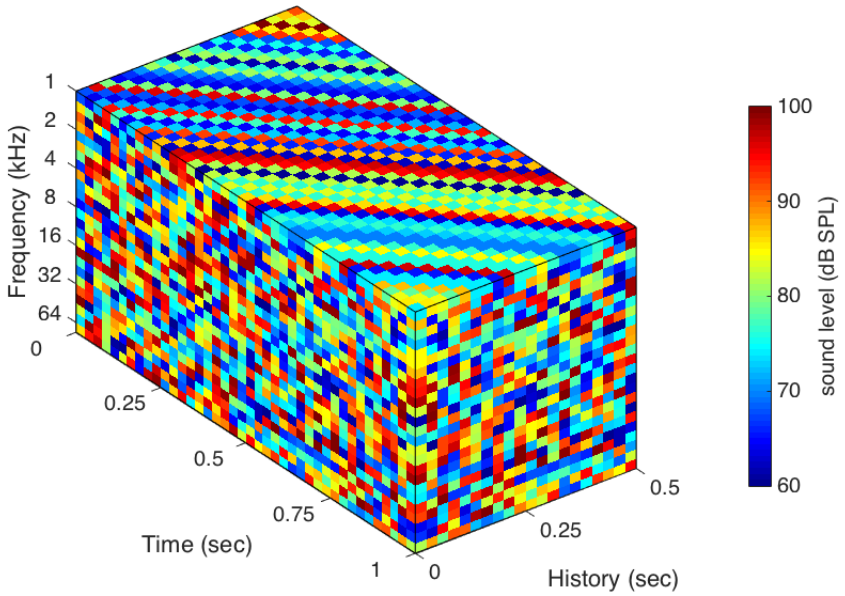

Figure 2.3: Three-dimensional DRC representation, X_{tfh}. The pure tones that make up the DRC vary in their intensity over both frequency, f, and time, t, as can be seen on the left face of the stimulus schematic. The STRF model produces an estimate of the unit's firing rate at each time point, t, based on a linear combination of stimulus frequency and recent stimulus history over a set number of time steps. Given that the length of the stimulus history vector being considered within each frequency channel is of a different length to the entire time vector in each channel, it is convenient to represent recent stimulus history at each time step along a third dimension, h. The STRF is fit to the frequency-history matrix at each time point, X_{fh}. This matrix for $t = 1$ is visible on the right-hand face of the schematic. With each increase in t, stimulus history also shifts by one time-step, as can be seen from the top face of the schematic.

This nonlinearity that maps STRF output , z_t, onto predicted firing rate, $\hat{y}_t$:

$$\hat{y}_t = F[z_t] \tag{2.10}$$

The sigmoidal shape of this nonlinearity results in rectification of the signal, preventing predictions of negative firing rates, and captures two major neuronal non-linearities; spike threshold and saturation. Adding this non-linearity has been shown to improve predictions of neural responses in AC (Rabinowitz et al., 2011a).

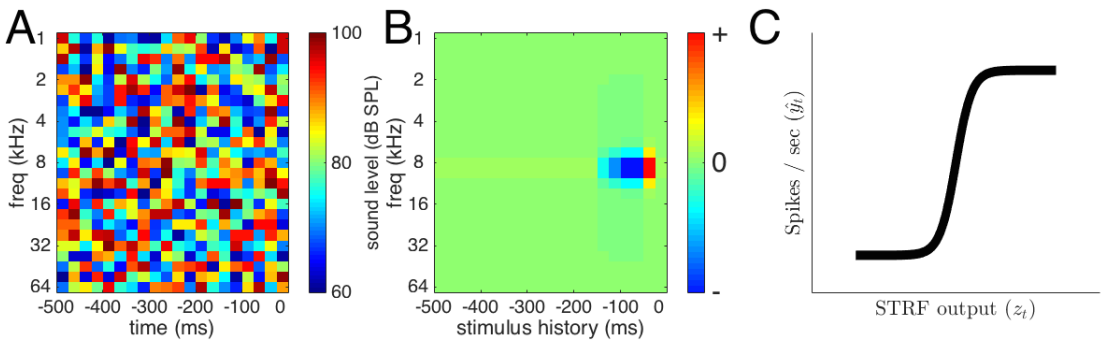

Figure 2.4: A. The aim of our modelling is to predict neuronal responses from the spectrogram of the DRC. B. This spectrogram is filtered through the STRF resulting in a predicted firing rates that are linear with respect to the spectrogram. C. This linear output is then passed through a sigmoidal output nonlinearity that maps STRF output z_t onto the firing rate of the unit y_t.

2.3.3 Model fitting

Responses during the first second of auditory stimulation were excluded from model fitting in order to avoid fitting to onset responses. Linear STRFs and output non-linearities were fitted to the data by using gradient decent to minimise a sum-of-squares error term:

$$E = \sum_t (y_t - \hat{y}_t)^2 \tag{2.11}$$

Cortical responses tend to be very noisy, however, and when fitting STRFs this method regularly results in overfitting to the training set (Sahani and Linden, 2003; Ahrens et al., 2008). An effective way to avoid such over-fitting is by including a regularisation term. Regularisation involves making prior assumptions regarding the structure of the model that can then be used to constrain the fitting of the model. Smoothness and sparseness (Sahani and Linden, 2003; Linden et al., 2003; Ahrens et al., 2008) are two routinely used regularisation constraints (Machens et al., 2004). Regularisation parameters can be chosen manually or via automated Bayesian approaches, such as Automatic Relevance Determination and Automatic Smoothness Determination (Sahani and Linden, 2003; Linden et al., 2003; Ahrens et al., 2008).

Throughout this thesis, I reduce model overfitting by assuming that the tuning of the MUs are separable in both space and time. This involves fitting a frequency kernel, k_f, and a history kernel, k_f, and computing the outer product of these two kernels, k_{fh}:

$$k_{\mathrm{fh}} = k_{\mathrm{f}} \otimes k_{\mathrm{h}} \tag{2.12}$$

Each kernel is fit by least-squares linear regression while the other kernel is fixed. The error term, E, captures the degree by which the predicted firing rate produced by the model differs from the actual firing rate of the MU, y_t.

$$E = \sum_{t} \left(y_t - \sum_{f,h} [k_0 + X_{tfh}.(k_f \otimes k_h)] \right)^2 \tag{2.13}$$

The coefficients of the kernel being fitted are then updated in order to improve this prediction. This procedure is repeated for each kernel in an iterative manner until a local minimum is found in the error term. While separable STRFs cannot model sensitivity to stimulus features that covary in frequency and time, this is compensated for by the reduction in over-fitting observed when using these models - separable STRFs tend to perform as well or better than inseparable STRFs in predicting responses to unseen stimuli (Linden et al., 2003; Simon et al., 2007; Ahrens et al., 2008).

The bandwidth of tuning and the frequency preference of each STRF was measured from the two kernels. The frequency kernel was rectified by squaring the coefficients before the best frequency (BF) was measured (Figure 2.5 B). The BF was defined as the frequency bin with the largest coefficient, following rectification. The bandwidth was measured as the width of the tuning curve around this peak response at half of its amplitude. The bandwidth of temporal tuning was measured as the width of the tuning curve surrounding the peak coefficient observed in the first 100 ms of the history kernel (Figure 2.5 C). The measurement was taken

at half of the amplitude of the peak coefficient.

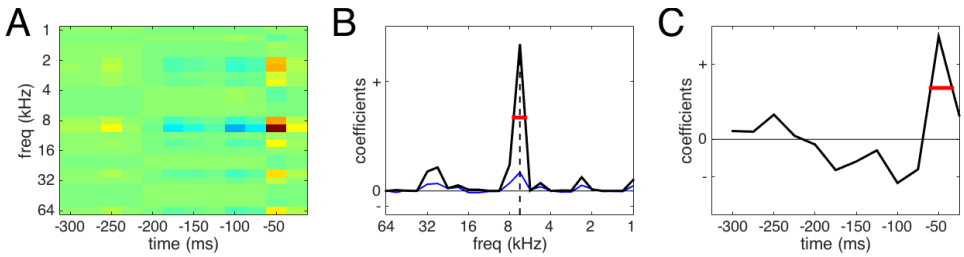

Figure 2.5: A. Complete STRF, the product of frequency and history kernels in B and C. B. Frequency kernels (blue line) were smoothed by squaring the coefficients (solid black line) before bandwidths were measured at half of the amplitude of the largest coefficient. The frequency associated with the largest coefficient was taken as the BF of the MU. BF is indicated here by the broken black line and bandwidth is indicated by the red line. C. History kernel coefficients (solid black line). The largest coefficient between 0 and -100 ms was identified and temporal bandwidth was measured at half of the amplitude of this coefficient. Bandwidth is indicated by the red line.

2.3.4 Current source density analysis

Bursts of white noise lasting 50 ms in duration were used in order to identify cortical areas that responded to auditory stimuli. The broadband nature of white noise makes it convenient as a search stimulus but one caveat of using a simple stimulus of this kind is that neurons with nonlinear auditory responses may go undetected. A low contrast stimulus of this kind may also bias unit detection towards neurons that respond preferentially to low contrast stimuli. This stimulus was used to identify penetrations in auditory cortical areas in which many neurons were recorded simultaneously. It was therefore not necessary for every recorded neuron to be noise responsive for the data from the whole penetration to be included in the analysis. This approach results in activity from potentially non-noise responsive auditory neurons being included in the analysis, avoiding the caveats associated with using white noise to identify auditory neurons on a unit by unit basis.

In recordings made with linear probes, evoked local field potential (LFP) responses to these noise bursts were also used in order to identify cortical layers. The

electrode spanned all cortical layers and was inserted orthogonally to the cortical surface. LFP responses to noise bursts underwent current source density (CSD) analysis in order to identify the pattern of current density flowing to and from the extracellular space, across depth and time. Current sinks are associated with synaptic activity, as this results in a net flow of current in to neurons and away from the extracellular medium, while current sources reflect the flow of current in to the extracellular medium (Nicholson and Freeman, 1975). The pattern of current sinks and sources can be used to identify the location of recording channels in different cortical layers.

In the standard method, homogeneity of activity orthogonal to the electrode axis is assumed and the CSD is estimated for each channel, $\hat{C}_z$, from the recorded potential, Φ, taking into account electrode spacing, h, and extracellular conductivity, σ (Pitts, 1952; Nicholson and Freeman, 1975). Here, σ is used to indicate extracellular conductivity, not standard deviation as used previously in this chapter:

$$\hat{C}_z = -\sigma \frac{\Phi(z+h) - 2\Phi(z) - \Phi(z-h)}{h^2} \tag{2.14}$$

An alternative method is the inverse or iCSD method (Pettersen et al., 2006). The aim of CSD analysis is to estimate the CSD from recorded potentials. The iCSD method solves this by first doing the inverse of this, calculating potentials for a known distribution of current sources. This method provides a matrix of the electrostatic forward solutions, *F*, that can be inverted in order to calculate current sources from recorded potentials:

$$\hat{C} = F^{-1}\Phi \tag{2.15}$$

Several variants of the iCSD method exist; here I use the δ-source method. This method assumes that current sources originate from within infinitesimally thin disks

around each electrode contact. If each disk is assumed to have an infinite radius, this method is equivalent to the standard method. Assuming a more plausible radius for each disk improves CSD calculations using this method. This radius corresponds to the area of cortex that is assumed to contribute to stimulus-evoked changes in current and voltage. Here I assumed a diameter of 250 μm. While assuming different values for this parameter can change CSD estimates, this has been found to not impact negatively on the use of CSDs for layer estimation (Szymanski et al., 2009).

The conductivity of the extracellular matrix, σ, must be factored into CSD calculations. Known anisotropic variation in extracellular conductivity can also be incorporated into this calculation, but I assumed a homogenous, isotropic conductivity of 0.3 S/m. Changing the value of σ will not alter the spatial or temporal pattern of CSDs used here for laminar identification, as it will scale the CSD across all channels equally.

The most prominent feature of both the LFP and CSD is the reversal that occurs at the border of layers 1 and 2 (Figure 2.6 A and B). This feature was used in order to align recordings so as to achieve consistent cortical depth measurements across animals. Layers were assigned based on experimentally obtained measurements of cortical thickness (Ji et al., 2015; Anderson et al., 2009). Layers 2 and 3 were combined into a single layer 2/3, as is commonly done. Recordings between 0 and 225 μm below this border were classified as layer 2/3, between 225 and 425 μm as layer 4, between 425 and 675 μm as layer 5 and below 675 μm as layer 6. Layer assignments using this method were validated using CSD features, such as a prominent early sink corresponding to thalamic input in layer 6 and at the border of layers 2/3 and 4 (Figure 2.6 B and C).

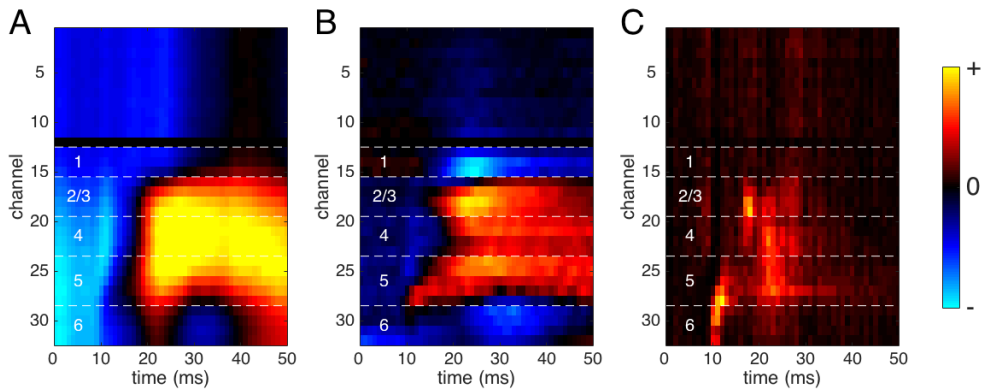

Figure 2.6: A. Mean local field potential (LFP) responses to noise burst stimuli on all 32 multi-electrode channels on the linear probe, inserted orthogonally to the cortical surface. Channel 1 is furthest outside the brain while channel 32 is deepest in the brain. All channels are separated by 50 μm. Poor spatial localisation of the LFP signal makes it difficult to identify cortical layers from the LFP response. B. Mean current source density (CSD) responses to noise burst stimuli. The prominent reversal in polarity on channel 15 indicates the border of layer 1 and 2/3. All other layer boundaries (indicated by white broken lines) were assigned based on distance from this reference point. Distances between layer boundaries were obtained from published measurements of layer thickness (Ji et al., 2015; Anderson et al., 2009). White numbers indicate assignment of cortical layers. The early activity at the border of layer 2/3 and 4, as well as at the border of layer 5 and 6, corresponds to the known anatomy of thalamic input to mouse AC. C. Mean spiking responses to noise across all channels further validates layer assignment. Strong feed forward activity driven by thalamic input can be seen in layer 6 and lower layer 2/3. Spiking responses stop above layer 2/3, in keeping with the very sparse number of cell bodies found in layer 1. Warmer colours indicate increased activity (depolarisation of LFP in A, current sinks in B and an increase in firing rate in C) while colder colours indicate reduced activity (Hyperpolarisation of LFP in A and current sources in B. Negative firing rates are not possible in C). Channel 12 on this probe was excluded on the basis of an extreme impedance value, indicating damage to the channel.

2.3.5 MU selection criteria

Recordings were only performed during penetrations in which MUs were responsive to noise burst stimuli, indicating that neurons in that area were acoustically driven. In order to perform our analysis, MUs had to have STRFs that could successfully predict neural responses to novel DRC stimuli. In order to quantify model performance, data from a single condition were divided into fitting and test sets. I fitted models to 9/10ths of the data and then used the model to predict responses to the remaining 1/10th of the data. Performance was quantified as the correlation

between predicted and actual responses in the test set. Cross-condition performance was measured by training the model on the data from one condition and then using the model to predict responses to the stimulus from the other condition. Performance was again quantified using the correlation between the predicted and observed responses to the test stimuli.

2.3.6 Statistical analysis

I used standard non-parametric statistical tests in order to ascertain the significance of effects reported throughout this thesis. Non-parametric tests were used in order to avoid assumptions regarding the distribution of the data being analysed. Wilcoxen signed-rank tests were used in order to compare two related samples and to test whether the median of a single sample was significantly different from 0, or 1 where the data are ratios. The Kruskal-Wallis one-way analysis of variance was used in order to compare multiple samples. Where this test showed a significant effect post-hoc pairwise multiple comparisons were used in order identify which samples were significantly different from each other. The threshold for significance I used throughout this thesis was $p<0.05$.

Chapter 3

Contrast gain control in mouse auditory cortex

3.1 Introduction

Understanding the brain as an information processing system will require an understanding of the computations performed by different neural populations and the mechanisms underpinning these computations (Poggio, 2012). The most efficient progress in this area may be achieved by investigating the mechanisms underlying "canonical computations". It has been argued that certain computations are canonical in the sense that they are employed by different populations of neurons in different areas of the brain, and even in different species (Carandini and Heeger, 2012). Contrast gain control (CGC), a compensatory reduction in the gain of sensory evoked activity in response to increased stimulus contrast, may represent such a computation. CGC has been observed in individual neurons throughout sensory systems in a variety of species, including mice (Atallah et al., 2012), rats (Garcia-Lazaro et al., 2007), rabbits (Brown and Masland, 2001) guinea pig (Zaghloul et al., 2005; Manookin and Demb, 2006), cats (Shapley and Victor, 1978, 1981; Heeger,

1992; Busse et al., 2009; Kaplan et al., 1987; Cheng et al., 1995; Sharpee et al., 2006; Sclar, 1987), ferrets (Rabinowitz et al., 2011a), monkeys (Carandini and Heeger, 1994; Carandini et al., 1997; Dunn et al., 2007; Benardete et al., 1992; Chander and Chichilnisky, 2001; Kaplan et al., 1987; Solomon et al., 2004; Sclar et al., 1990) and humans (Busse et al., 2009), as well as in non-mammalian species such as songbirds (Nagel and Doupe, 2006; Baccus, 2006), salamanders (Baccus and Meister, 2002; Chander and Chichilnisky, 2001; Kim and Rieke, 2001; Smirnakis et al., 1997) and bullfrogs (Rieke et al., 1995). In the mammalian cortex, it been observed in a variety of sensory areas, including primary visual (V1) (Heeger, 1992; Heeger et al., 1996; Carandini et al., 1997; Carandini and Heeger, 1994) and auditory (A1) (Rabinowitz et al., 2011a) and somatosensory cortex (S1) (Garcia-Lazaro et al., 2007), as well as in higher visual cortical areas MT (Heeger et al., 1996; Simoncelli and Heeger, 1998; Rust et al., 2005), V4 (Reynolds and Desimone, 2003) and inferior temporal cortex (IT) (Zoccolan et al., 2005). Given the widespread nature of this computation, it has been proposed that this form of gain control may represent a canonical neural computation (Carandini and Heeger, 2012).

The possibility that stimulus-driven gain control may be in some sense canonical as a computation does not imply the existence of a common mechanism across modalities (Carandini and Heeger, 2012). A wide variety of biophysical mechanisms for gain control have been identified and it is certainly true that differences exist between mammalian and non-mammalian species in how gain control is implemented (Olsen et al., 2010; Sadagopan and Ferster, 2012). Within mammals however, the cortex has long been thought to comprise of canonical circuits that have evolved in order to implement equivalent computations across a variety of modalities (Douglas and Martin, 1991b). CGC appears to be one of the best existing candidates for a canonical cortical computation. A common mechanism, or

combination of mechanisms, may therefore underlie cortical CGC across sensory systems and species.

In order to investigate the mechanisms underpinning computations in the mammalian brain, investigators are increasingly turning to the mouse as a model organism (Wu et al., 2011; Atallah et al., 2012; Wilson et al., 2012; Olsen et al., 2012). This is largely due to the abundance of genetic manipulations that are possible in this species (Boyden et al., 2005; Zhang et al., 2007; Chow et al., 2010; Sternson and Roth, 2014; Huang and Zeng, 2013; Akerboom et al., 2013) as well its suitability to other circuit-level techniques such as head-fixed whole-cell recording (Wu et al., 2011; Zhou et al., 2014a) and two-photon imaging (Akerboom et al., 2013; Fu et al., 2014). By exploiting these techniques in order to investigate CGC in mouse auditory cortex (AC), it may be possible to not only increase our understanding of the neural basis of this this computation in the auditory system but may also inform the debate regarding the canonical nature of cortical circuitry (Douglas and Martin, 1991b). Testing for the presence of this computation in the mouse allows the canonical nature of CGC to be tested experimentally. If it is truly canonical it should be observed across species as well as across cortical areas, and should therefore be found in mouse AC.

A crucial piece of information that can constrain circuit hypotheses is whether, and if so how, the strength of CGC varies across cortical layers. If thalamocortical inputs drive CGC, gain changes should be observed in layer 4, where thalamic inputs are strongest. Furthermore, gain changes should be consistent across the depth of cortex. This could arise either by thalamic inputs driving gain changes in all layers independently (Ji et al., 2015) or via CGC being implemented in layer 4, and subsequently inherited in other cortical layers. If layer specific intracortical mechanisms underlie CGC however, variation in the strength of CGC across different layers should be observed. Such an intracortical circuit for gain control has been

identified in primary visual cortex (V1), where thalamic input drives corticothalamic projecting layer 6 neurons, which in turn drive a translaminar projecting PV interneuron subtype via facilitating synapses (Olsen et al., 2012; Bortone et al., 2014). Inhibition via this route has been found to reduce gain across all other cortical layers but not layer 6. If this circuit underlies CGC in AC, CGC should be absent from layer 6. Knowledge of the laminar organisation of CGC in mouse AC would serve to distinguish the likelihood of these different circuit hypotheses.

Here I characterised CGC in mouse A1 in order to lay a foundation for circuit-level dissection of the mechanisms involved in this computation. This line of investigation will allow for direct comparison with findings in visual and other sensory modalities, allowing the question of whether a canonical gain control mechanism exists to be addressed experimentally. Finally, I examined the laminar organisation of CGC in order to constrain possible circuit hypotheses regarding the mechanistic basis of CGC in AC.

3.2 Methods

Extracellular recordings were made from 815 sites in the left AC of 24 anaesthetised mice, while auditory stimuli were presented to the contralateral ear. Each penetration showed noise-evoked responses on multiple channels. Of the MUs recorded, 56% showed noise evoked responses, defined as a doubling of baseline firing in response to noise onset. Noise responsiveness was not used as a selection criterion for further analysis however. Data was collected from transgenic mice expressing cre-recombinase under the PV promoter as these animals were also being used in other experiments which are presented in chapter 4. One third of the mice used (N = 8) were crosses between homozygous PV ^{cre} (Jax No: 008069) (Hippenmeyer et al., 2005) and homozygous Ai35D mice (Jax No: 012735) (Madisen

et al., 2012), which express ARCH and GFP in a cre-dependent manner. The other two thirds (N = 16) were PV ^{cre} mice that had previously undergone intracranial injection of an adenoassociated virus that enables cre-dependent expression of channelrhodopsin (ChR2) and enhanced yellow fluorescent protein (EYFP) (AAV-EF1a-DIO-hChR2(H134R)-EYFP-WPRE-pA serotype 2) (UNC vector core, Chapel Hill, NC, United States), several weeks before these experiments were performed.

Anaesthesia and surgery were performed as described in section 2.1.1. Linear 1x32 multi-electrodes (NeuroNexus, Ann Arbor, United States) were used in order to record activity from all cortical layers simultaneously. In 5 experiments 64 channel 8 shank Buzsaki-style multi-electrodes were used in order to improve single-unit isolation, at the expense of the laminar information provided when using linear probes. Recordings were band pass filtered an analog measure of multi-unit activity (MUA) on each channel. The same raw recordings were low-pass filtered in order to extract local field potentials (LFPs). LFPs then underwent Current source density (CSD) analysis as described in section 2.3.4.

Bursts of white noise lasting 50 ms were delivered in order to test acoustic responses and to identify cortical layers based on the laminar pattern of evoked LFP responses. Dynamic random chord (DRC) stimuli of either low or high contrast were also delivered. In the low contrast condition, the animal was presented with DRCs that had a mean of 80 dB and a range of ±10 dB (σ_L ~6.2 dB, c = σ_P/μP = 0.68). In the high contrast condition, the DRCs had the same mean but the range was doubled ±20 dB (σ_L ~11.97 dB, c = σ_P/μ_P = 1.2).

3.3 Results

3.3.1 Stimulus contrast has negligible effects on STRF tuning

First, the effects of altering stimulus contrast on receptive field tuning in AC neurons were assessed. If stimulus contrast produces a pure gain change, tuning in these neurons should remain constant between contrast conditions. MUs were required to have a predictive spectro-temporal receptive field (STRF) in order to be included in this analysis. For each MU, the responses to both low and high contrast DRCs were used in order to fit the STRF. The STRF was fitted to 90% of this data and was then used in order to predict the response of the MUA to the remaining 10% of the data. MUAs were deemed predictive if the correlation coefficient (CC) between the predicted and actual response was larger than 0.04. This threshold was determined from the distribution of CCs recorded across MUs. There was a clear population of recordings that had CCs centered on 0, indicating that these recordings were dominated by noise (Figure 3.1). The threshold was positioned in order to exclude these recordings. 160 MUs passed this criterion and were used in order to assess the effect of contrast on receptive field shape in mouse AC.

In order to quantify the effect of contrast on receptive field parameters, a single spectro-temporal receptive field (STRF) was estimated for each contrast condition, yielding separate high and low contrast STRFs (Figure 3.2). This made it possible to asses whether tuning remains constant as stimulus contrast changes, as would be the case if MUs showed a pure gain change in response to a change in contrast.

The best frequency (BF) of the STRFs, defined as the frequency associated with the peak coefficient in the frequency kernel, did not change systematically with stimulus contrast (medians: Low contrast; 9.5kHz, High contrast; 9.5kHz, sign-rank test; $p=0.571$) (Figure 3.2 A). Some individual MUs did show large changes in BF however. 34% of MUs with predictive STRFs maintained their best frequency

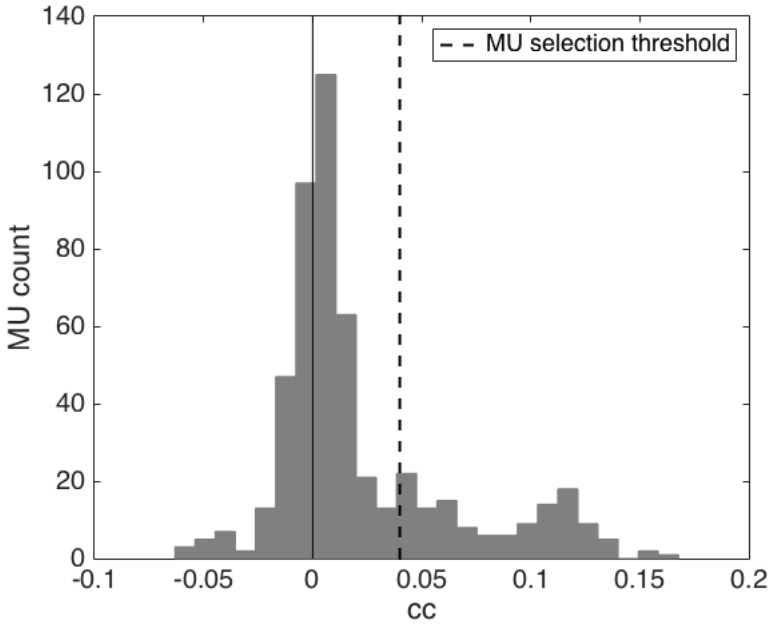

Figure 3.1: Histogram of CCs between actual MU responses and responses predicted by the STRF for each MU. It was necessary to exclude the population of recordings with CCs near 0. A selection threshold of 0.04 was used in order to exclude these recordings. All MUs to the right of this threshold were included in subsequent analyses.

across contrast conditions while 30% maintained their best frequency within 1/2 an octave (35% up, 65% down). 22% showed a reduction in their BF more than 1/2 an octave and of these 51% shifted their BF by more than an octave. 14% showed a increase in their BF more than 1/2 an octave and of these 50% shifted their BF by more than an octave.

Spectral bandwidth showed a small but significant increase during high contrast stimulation (medians: Low contrast; 1.17 oct, High contrast; 1.35 oct, sign-rank test; $p=0.007$) (Figure 3.2 B). No systematic changes in bandwidth of temporal tuning of STRFs were found (medians: Low contrast; 8.11 ms, High contrast; 5.17 ms, sign-rank test; $p=0.177$) (Figure 3.2 C).

The most striking change that occurred between STRFs from each condition was a suppressive effect on tuning during high contrast stimulation. 159 of the 160 MUs studied showed a reduction in the largest STRF coefficient under high compared

with low contrast stimulation (medians: Low contrast; 0.5μV/dB/s, High contrast; 0.29μV/dB/s, sign-rank test; $p<0.001$) (Figure 3.2 D). The effects of contrast on STRF shape appeared to be far less significant than this suppressive effect.

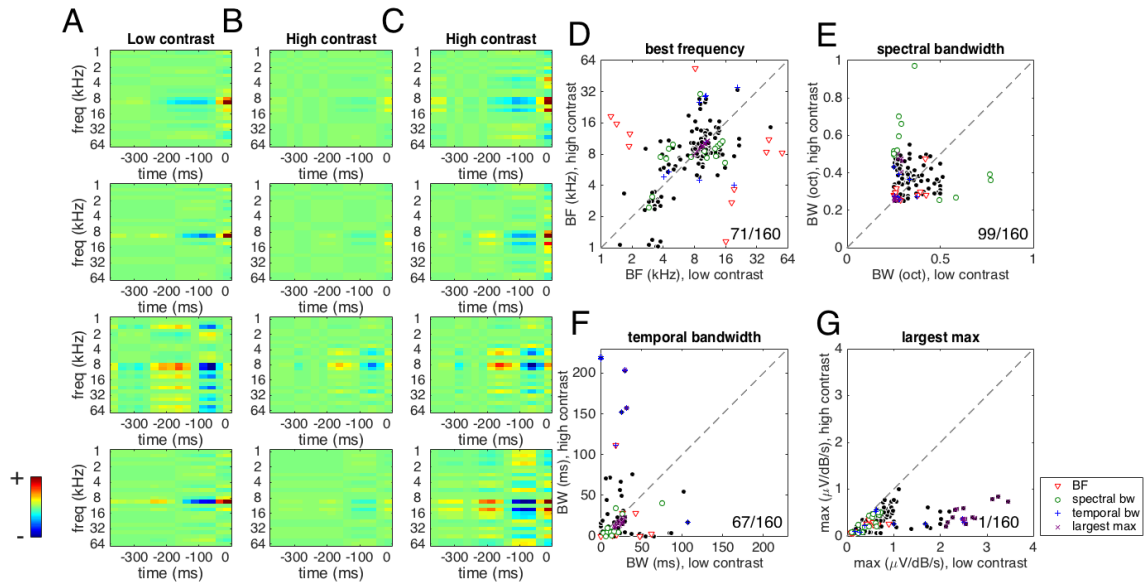

Figure 3.2: A. Example STRFs estimated under low contrast stimulation for 4 MUs. Red areas of the STRF indicate spectrotemporal features of the stimulus that increase the amplitude of the MUA while blue areas indicate features that suppress MUA. B. STRFs estimated under high contrast stimulation for the same MUs revealing a strong suppression of the STRF coefficients. Color scaling is the same as in A. High contrast STRFs, as shown in B, but with independent colour scaling, revealing that spectrotemporal selectivity is broadly similar between contrast conditions (Panels A and C). D. Scatter plot showing the effect of contrast on best frequency (BF). The BF of each MU did not vary systematically with stimulus contrast. For individual MUs Estimates of best frequency did vary between conditions but this was primarily attributable to the diffuse nature of some STRFs which lacked a clear peak in their spectral tuning. E. The bandwidth of tuning broadens slightly as stimulus contrast increases. F. Temporal tuning showed no systematic changes between contrast conditions. G. By far the clearest effect of altering contrast is a reduction in stimulus sensitivity under high contrast stimulation, as indicated by a reduction in the largest STRF coefficient. For each parameter, outliers were identified as values that fell outside 1.96 standard deviations of the distribution of these values. They are shown here in each plot using different markers and colors as indicated in the legend. The outlier parameters did not tend to come from the same population of MUs, as indicated by the lack of overlap in outliers identified from the different parameters. Ratios in the corner of each plot indicate the number of MUs above the identity line, y=x, over the total number of MUs.

A surprising amount of variability was observed for changes in these parame-

ters across the population of MUs. It is possible that the MUs showing the greatest changes in BF, spectral bandwidth and temporal bandwidth across contrast conditions may have simply had poor STRFs and, as a result, gave poor estimates of these values. In order to asses whether these outliers come from a single subpopulation of MUs with poor STRFs, I identified outliers for each of these parameters independently and examined the overlap of these outliers. If the outliers for each parameter come from a single population, the MUs defined as outliers based on the different parameters should be largely overlapping. Outliers were defined as MUs that showed parameter changes between contrast conditions that fell outside 1.96 standard deviation of the distribution of changes, as these values would have a $<$5% probability of occurring in such a distribution (i.e. $p < 0.05$). BF outliers are indicated by the inverted red triangle symbol, spectral bandwidth outliers are indicated by the green circle temporal bandwidth outliers are indicated by the blue plus symbol while largest max coefficient outliers are indicated by the purple cross (Figure 3.2 D-G). Only 3 out of 50 outlier MUs were identified as outliers based on changes in more than one parameter, indicating that they do not come from a single population with poor STRFs.

The STRFs of the outliers identified for each condition were examined in order to investigate the cause of the variability in these parameters. MUs that showed large changes in BF between contrast conditions had poor STRFs which made estimating the BF of the unit between contrast conditions unreliable (Figure 3.3A). The STRFs of MU outliers identified through large changes in spectral (Figure 3.3B) and temporal bandwidth (Figure 3.3C) had reasonable STRFs in one but not both of the contrast conditions. This resulted is noisy estimates of spectral bandwidth in one of the two contrast conditions. This indicates that erroneous estimates of bandwidth changes occured as a result of comparing a reasonable estimate of in one condition with a noisy estimates in the other condition. Changes in the largest STRF coeffi-

cient between contrast conditions were predicted. The outliers identified through large changes in this parameter were consistent with our prediction that suppression would occur in the high contrast condition and showed good STRFs under low contrast stimulation, indicating that these are not erroneous estimates but are instead the MUs that show the strongest effect (Figure 3.3D).

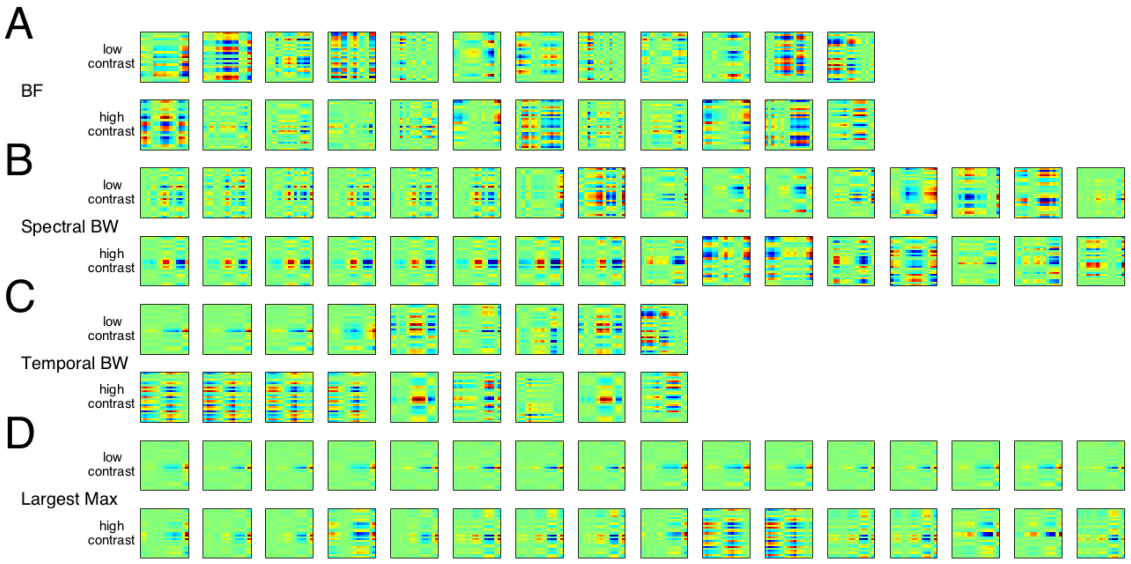

Figure 3.3: Low and high contrast STRFs of outliers identified from changes in each STRF parameter between contrast conditions. Outliers were defined as values that fell outside 1.96 standard deviations of the distribution of these values, corresponding to a p value of <0.05. A. STRFs for MU outliers identified through large changes in BF. These MUs have poor STRFs, resulting is noisy estimates of BF. B. STRFs for MU outliers identified through large changes in spectral bandwidth. These MUs had reasonable STRFs in one but not both of the contrast conditions, resulting is noisy estimates of spectral bandwidth in a single condition and noisy estimates of changes in bandwidth between conditions. C. STRFs for MU outliers identified through large changes in temporal bandwidth. As was the case for spectral bandwidth outliers, these MUs had reasonable STRFs in one but not both of the contrast conditions, resulting is noisy estimates of changes in bandwidth between conditions. D. STRFs for MU outliers identified through large changes in the largest STRF coefficient. Such changes were predicted and so these outliers merely represent the MUs showing the largest effect. STRF coefficients are independently scaled for plotting, resulting in STRFs with low coefficients looking noisy, as in the bottom row of STRFs. Vertical and horizontal stripes in STRFs are the result of the separable nature of the kernels.

I next investigated whether this variability is the result of the change in contrast

of the stimulus or of noisy STRF estimate. I split the data from each contrast condition in half and fit a STRF to each half of the data. For the low contrast data, this produced STRFs that looked largely similar for each MU but were noisier than the estimates when the full dataset was used (Figure 3.4A). The noisiness of the estimates resulted in large changes in STRF parameters between the two estimates for the same MU under low contrast stimulation (Figure 3.4C-D).

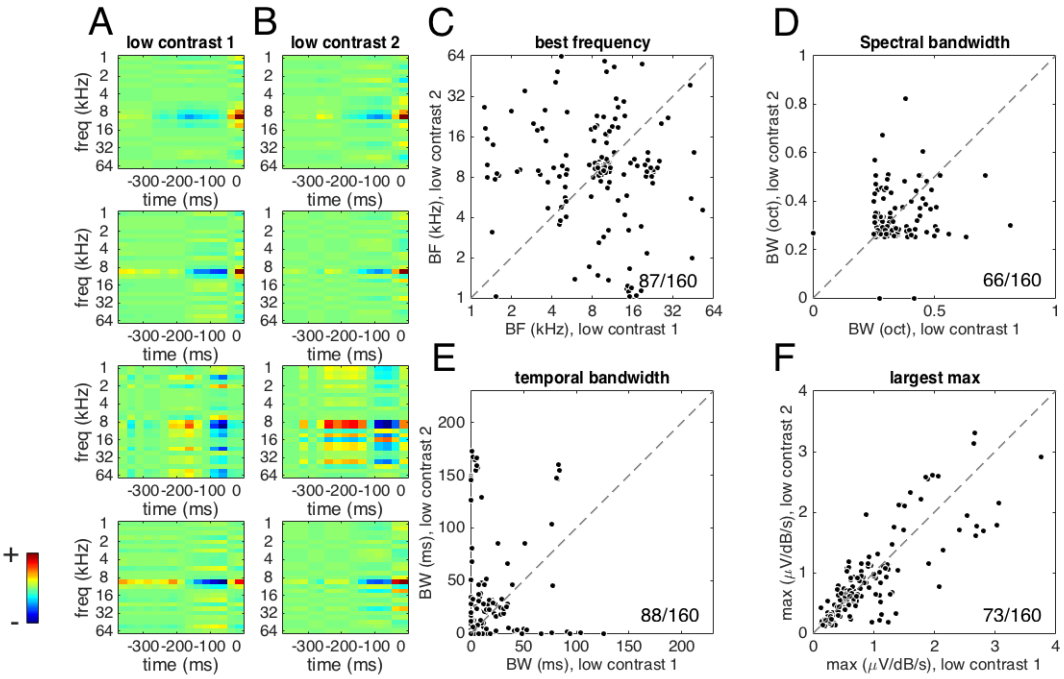

Figure 3.4: A. Example STRFs estimated using half of the low contrast stimulation data for 4 MUs. Red areas of the STRF indicate spectrotemporal features of the stimulus that increase the amplitude of the MUA while blue areas indicate features that suppress MUA. B. STRFs estimated under low contrast stimulation for the same MUs as in A, using the remaining half of the data. Color scaling is the same as in A. Some STRF shapes were consistent between estimates while other estimates were very noisy, due to being fitted to half of the available data. C-F. Scatter plots comparing estimates of best frequency (BF) (C), spectral bandwidth (D), temporal bandwidth (E) and larges coefficient (F) from the two independent STRFs during low contrast stimulation. The noisiness of the estimates resulted in many units showing large changes in these parameters. Ratios in the corner of each plot indicate the number of MUs above the identity line, y=x, over the total number of MUs.

For the high contrast data, the STRF estimates were equally noisy (Figure 3.5A). This also resulted in large changes in STRF parameters between the two estimates

for the same MU under high contrast stimulation (Figure 3.5C-D). This indicates that large changes in STRF parameters between low and high contrast STRFs may be largely due to noisy STRF estimates, rather than to genuine changes in tuning between contrast conditions.

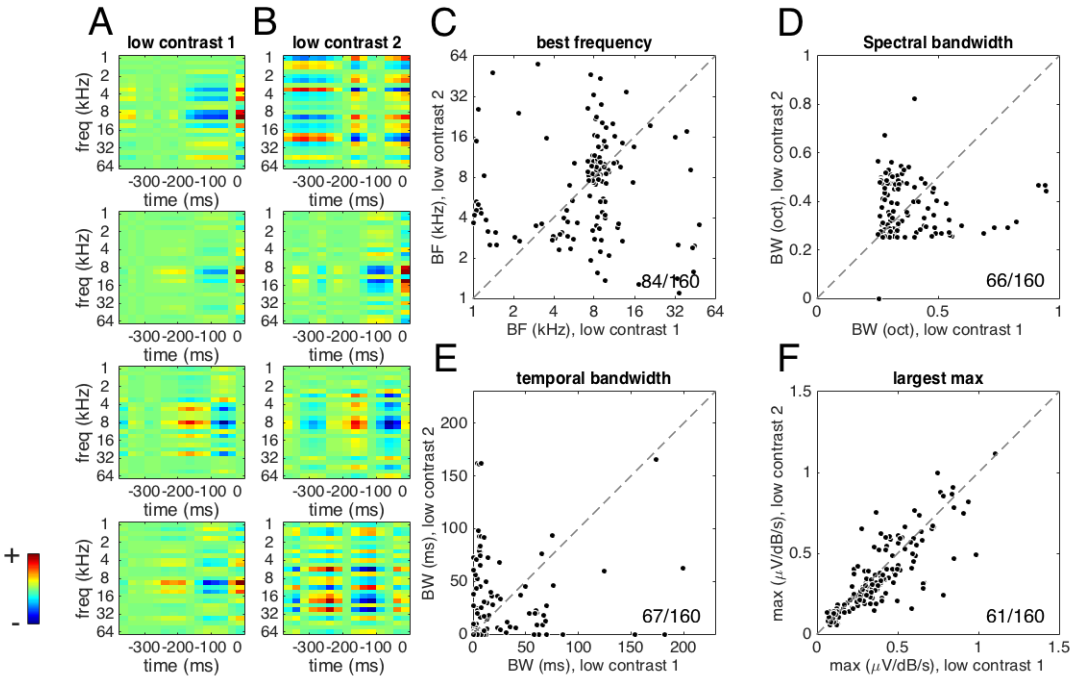

Figure 3.5: A. Example STRFs estimated using half of the high contrast stimulation data for 4 MUs. Red areas of the STRF indicate spectrotemporal features of the stimulus that increase the amplitude of the MUA while blue areas indicate features that suppress MUA. B. STRFs estimated under high contrast stimulation for the same MUs as in A, using the remaining half of the data. Color scaling is the same as in A. Some STRF shapes were consistent between estimates while other estimates were very noisy, due to being fitted to half of the available data. C-F. Scatter plots comparing estimates of best frequency (BF) (C), spectral bandwidth (D), temporal bandwidth (E) and larges coefficient (F) from the two independent STRFs during high contrast stimulation. For the low contrast condition, the noisiness of the estimates resulted in many units showing large changes in these parameters. Ratios in the corner of each plot indicate the number of MUs above the identity line, y=x, over the total number of MUs.

To test the importance of changes in STRFs tuning with stimulus contrast, I compared the 2D correlation of the two STRFs estimated from low contrast responses to the mean 2D correlation between low contrast and high contrast STRF estimates (Figure 3.6 A). I found that across contrast condition STRF CCs were not

significantly different to within low contrast condition STRF CCs (sign-rank test; p=0.665). I also compared the 2D correlation of the two STRFs estimated from high contrast responses to the mean 2D correlation between low contrast and high contrast STRF estimates (Figure 3.6 B). I found that across contrast condition STRF CCs were significantly lower than within high contrast condition STRF CCs (sign-rank test; $p<0.001$). Many STRF estimates has near 0 CCs however when STRFs were compared within a single contrast condition. This indicates that the STRF estimates when fitted to half the available data are too poor to act as a meaningful comparison. This brings into question the utility of comparing this poor within-condition performance with across-condition performance as a method of assessing the significance of changes in tuning.

In order to avoid this issue a single STRF was fitted to 80% of both high and low contrast data and was used to predict the remaining 10% from each contrast condition separately. The predictive power of the STRF was quantified as the CC between the predicted and actual response for each contrast condition. Using this approach, it was possible to examine whether a single STRF can be used to model responses to both high and low contrast stimuli in an unbiased manner, a necessary criterion for modelling CGC using linear-nonlinear models. These STRFs predicted low and high contrast responses equally well (Median CCs: High contrast predicted: 0.09, Low contrast predicted: 0.08, sign-rank test; p=0.139)(Figure 3.6 C). STRFs were therefore not significantly biased towards either contrast condition in their performance, indicating a single STRF can be used to model potential changes in gain between contrast conditions.

3.3.2 Gain control compensates for changes in stimulus contrast

I next attempted to assess gain changes more accurately by incorporating the data of both the high and the low contrast conditions into a single model. This model

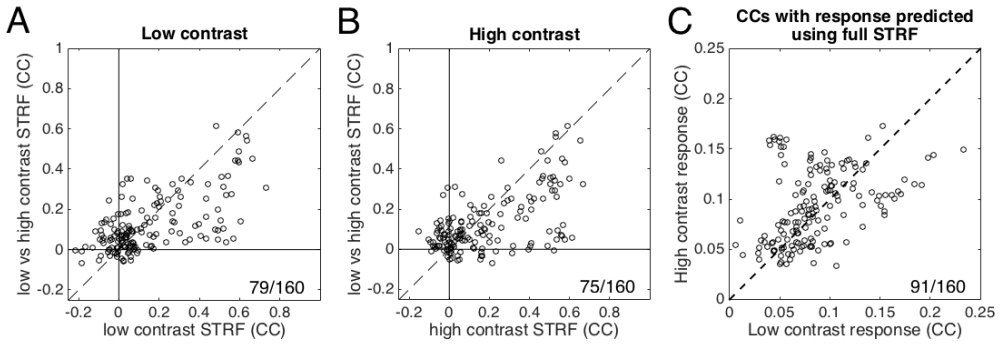

Figure 3.6: For each MU, low contrast responses were split in half and two STRFs were estimated. This was also done for high contrast responses. 2D correlation values (CC) were compared within and between contrast conditions in order to investigate whether STRF shapes changes with stimulus contrast. A. Scatter plot of CCs between the two low contrast STRFs compared to the mean CC between low and high contrast STRFs. CC values were not significantly different between within (low vs low) and across condition STRFs (low vs high). B. Scatter plot of CCs between the two high contrast STRFs compared to the mean CC between low and high contrast STRFs. CC values were significantly higher for within (high vs high) compared to across condition STRFs (low vs high). C. Cross predictions were used also used in order to test the functional significance of changes in STRF tuning between contrast conditions. A single STRF was fitted to 90% of both the low and high contrast DRC responses for each MU. This STRF was then used to predict the remaining 10% of the low and high contrast responses that the model had not been trained on. The predictive power of the STRF, quantified as the CC between the predicted and actual response, was compared between these conditions. No significant difference was observed between predictions of high and low contrast responses. This indicates that, despite small changes in tuning, a single STRF can be used to model contrast-dependent gain changes for these MUs. Ratios in the corner of each plot indicate the number of MUs above the identity line, $y=x$, over the total number of MUs.

consisted of a single linear STRF, fitted to the data from both high and low contrast conditions (Figure 3.7). A sigmoidal output nonlinearity (Equation 2.9) that could vary between contrast conditions was fitted to the output of this STRF, as described in section 2.3.3.

Nonlinearities were defined by four parameters which reflect the y-axis offset (a), y-axis range (b), x-axis offset (c) and the slope (d) of the curve. Y-axis offset reflects baseline activity, which has previously been shown to be unaffected by changes in stimulus contrast (Rabinowitz et al., 2011a). This parameter of the nonlinearity was therefore kept constant between contrast conditions in order to

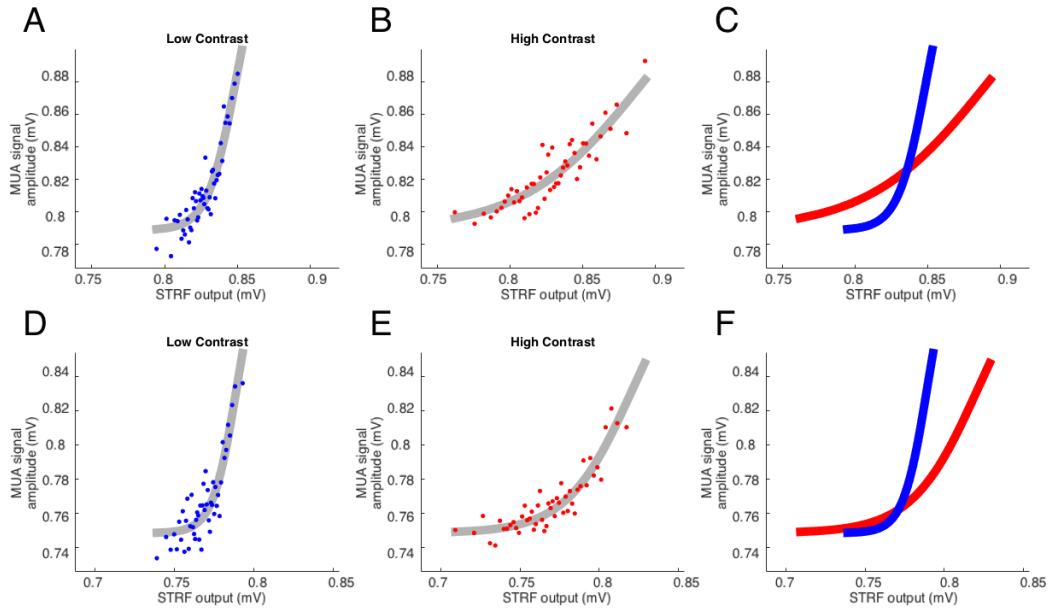

Figure 3.7: In order to quantify gain changes, a single STRF was fitted to the data from both contrast conditions for each MU. An output nonlinearity that varied between contrast conditions was fitted to the output of this STRF. A: Sigmoidal output nonlinearity fitted under low contrast stimulation. Blue dots are binned responses to low contrast stimulation. B: Nonlinearity fitted under high contrast stimulation for the same MUs. Red dots are binned responses to high contrast stimulation. C: Overlay of nonlinearities for both conditions. D-F, same as in A-C but for a second MU. The main effect of increasing contrast is to reduce the slope of this nonlinearity, capturing a change in the gain of the MUA.

more accurately quantify changes in the remaining parameters. In order to validate this assumption, baseline activity was compared under low (Figure 3.8 A) and high contrast stimulation (Figure 3.8 B). Changes in the offset of baseline activity (R_{offset}) was quantified by taking difference in the 5th percentile of the MUA under low (R_{min}^{low}) and high (R_{min}^{high}) contrast stimulation. This quantity was then normalised by the full range of responses (R_{max} - R_{min}) for each neuron:

$$R_{\text{offset}} = \frac{R_{\text{min}}^{\text{low}} - R_{\text{min}}^{\text{high}}}{R_{\text{max}} - R_{\text{min}}} \tag{3.1}$$

Across the population, a negligible increase was observed under high contrast stimulation (Median increase: 1.29%, sign-rank test; $p<0.001$) Figure 3.8 C).

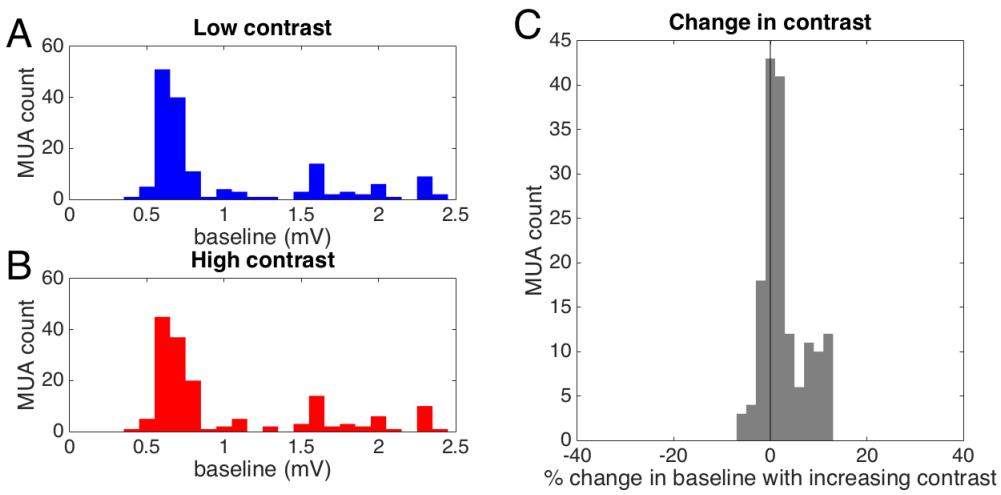

Figure 3.8: Changes in baseline activity between contrast conditions. A. Histogram of baseline activity values across all MUs during low contrast stimulation. B. Histogram of baseline activity values across all MUs during high contrast stimulation. C. Histogram of the % change in baseline activity under high compared with low contrast stimulation. Positive values indicate an increase in baseline activity during high contrast stimulation. The solid line indicates 0% change. Across the population MUs showed a small but significant increase in baseline activity during high contrast stimulation

Nonlinearities could vary in their slope and their offset along the x-axis. An increase in the slope of the nonlinearity indicates an increase in gain while an increase in the x-offset reflects an increase in the threshold of the MUA. In theory, it would be possible to measure changes in the unit's saturation point using a sigmoidal output nonlinearity, however, as with the ferret, DRCs did not drive responses to saturation. This resulted in it not being possible to estimate this value in practice. An increase in stimulus contrast was found to result in a compensatory reduction in gain, from low (Figure 3.9 A) to high contrast conditions (Figure 3.9 B). The ~2 fold increase in contrast decreased gain by median value of 55.2% (sign-rank test: $p<0.001$) (Figure 3.9 C).

The suppressive effects of increasing contrast observed in the STRF coefficients might also be explained by a reduction in threshold, in combination with the observed gain change. A negligible change in threshold from low (Figure 3.10 A) to high contrast (Figure 3.10 B) was observed (median threshold change: 0.6%,

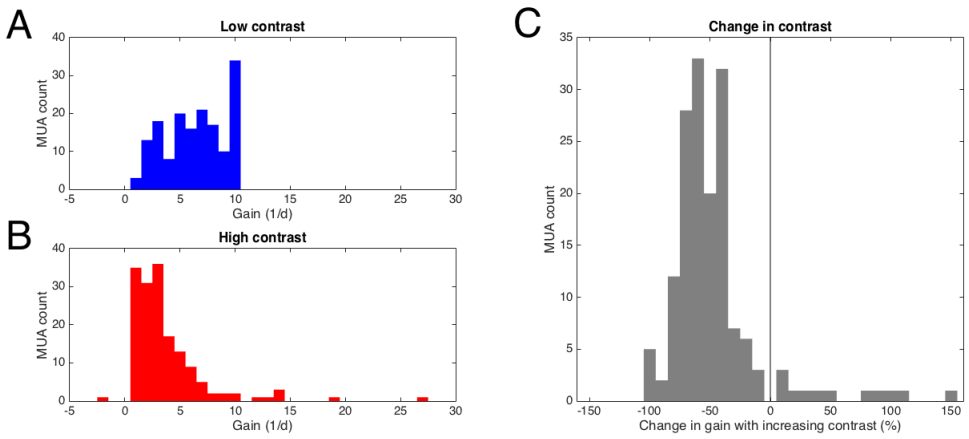

Figure 3.9: Changes in the output nonlinearity gain parameter (1/d) between contrast conditions. A. Histogram of gain values across all MUs during high contrast stimulation. B. Histogram of gain values across all MUs during low contrast stimulation. C. Histogram of the % change in gain under high compared with low contrast stimulation. Negative values indicate a reduction in gain. The solid line indicates 0% change. The vast majority of MUs showed a compensatory reduction in gain during high contrast stimulation.

sign-rank test: $p<0.001$) (Figure 3.10 C).

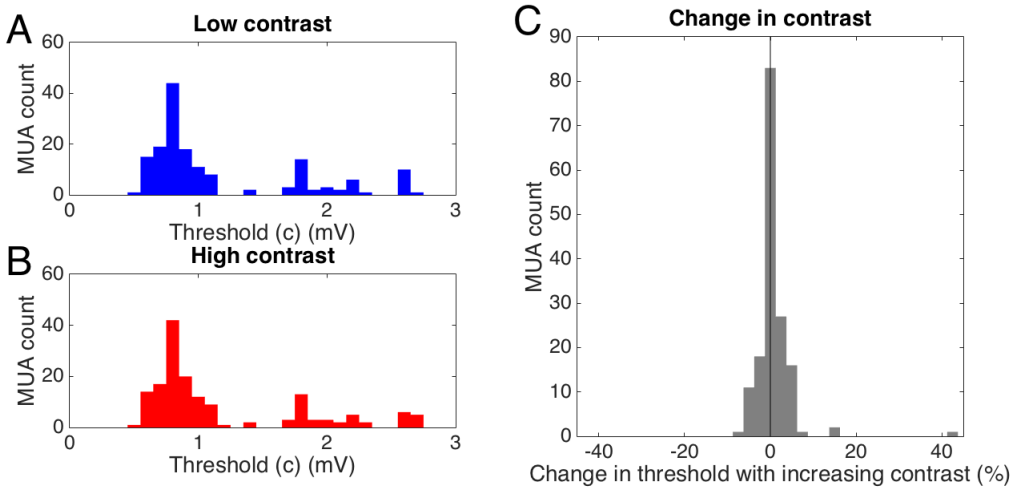

Figure 3.10: Changes in the output nonlinearity threshold parameter (c) between contrast conditions. A. Histogram of threshold values across all MUs during low contrast stimulation. B. Histogram of threshold values across all MUs during high contrast stimulation. C. Histogram of the % change in threshold under high compared with low contrast stimulation. Negative values indicate a reduction in threshold, a leftward shift along the x-axis. The solid line indicates 0% change. MUs showed a negligible increase in threshold during high contrast stimulation.

3.3.3 Gain control is strongest in deep cortical layers

Previous investigations of CGC in AC did not attempt to control for the layers in which the recordings were made (Rabinowitz et al., 2011a). As described in section 2.3.4 I used linear probes that allowed us to identify cortical layers using current source density (CSD) analysis. A reversal in the CSD profile is observed at the border of layer 1 and 2/3 and this feature was used in order to align recordings. After alignment of our recordings, MUs were assigned to cortical layers based on their cortical depth from the layer 1-2/3 border (Figure 3.11 A), using experimentally obtained measurements of layer thickness (Ji et al., 2015; Anderson et al., 2009). Layer estimates were inspected visually for features such as an early current sink in layer 6 and were compared to the depth distribution of spiking responses in order to confirm layer estimates. Linear probes were used in 19 of the 24 experiments, resulting in it being possible to estimate cortical depth for 143 of 160 MUs. Median gain changes with increasing contrast showed significant laminar variation (Layer 2/3: -41.4%, N = 23, Layer 4: -51.9%, N = 42, Layer 5: -60.6%, N = 50, Layer 6: -65.1%, N = 24. Kruskal-Wallis test: $H(3) = 15.46$, $p=0.002$). Post-hoc multiple comparisons of mean ranks indicated that CGC is significantly stronger in layers 5 and 6 compared to layer 2/3 (Figure 3.11 B). I performed this analysis excluding the 2 outliers in layer 2/3 (Figure 3.11 B). I found that this pattern of laminar differences was remained significant (Kruskal-Wallis test: $H(3) = 12.61$, $p=0.006$). This indicates that this effect was not due to these outliers in this layer.

I next examined whether this laminar pattern of gain changes was attributable to laminar variations in the noise of responses. Cortical neurons show highly variable responses to multiple presentations of the same stimulus, the average variance across stimulus presentations is known as the “total power” (Sahani and Linden, 2003). It is possible to quantify the proportion of this variance that is reliably locked to the stimulus, known as the “signal power”. The remaining variance is not

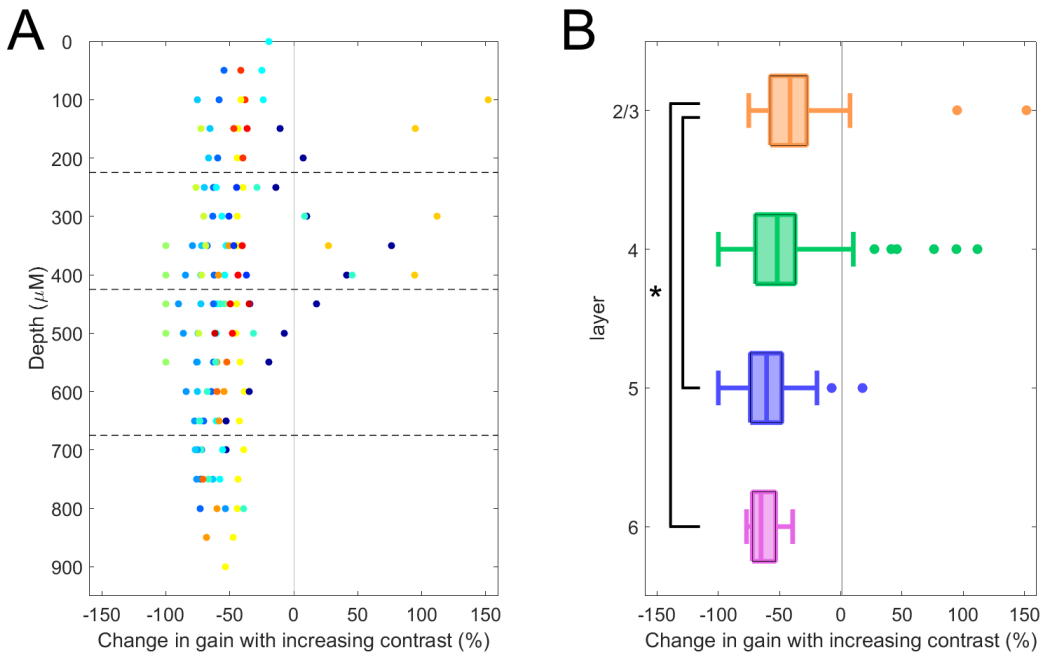

Figure 3.11: A. Strength of gain change for each MU plotted as a function of cortical depth from the layer 1-2/3 border. Layer boundaries are indicated by lines and data points are color-coded by penetration. B. Box plots of gain control within each cortical layer. Gain changes were significantly stronger in layers 5 and 6 compared to layer 2/3.

locked to the stimulus and is known as the "noise power":

$$totalpower = signalpower + noisepower \tag{3.2}$$

The noise in individual MU responses can be quantified by the "noise ratio", ratio of noise power to signal power:

$$noiseratio = \frac{noisepower}{signalpower} \tag{3.3}$$

A lower noise ratio indicates a more reliable response. If responses were identical on every trial there would be 0 noise power and the noise ratio would also be 0.

I measured the noise ratio for each MU in order to investigate whether lami-

nar differences exist in the noise level of MU responses (Figure 3.12 A). I found that NRs decreased significantly with cortical depth (median NR: Layer 2/3=22.63, Layer 4=15.53, Layer 5=12.44, Layer 6=10.91, Kruskal-Wallis test: $H(3) = 10.97$, $p=0.012$). Post-hoc multiple comparisons of mean ranks indicated that layer 2/3 MUs displayed significantly higher noise ratios than layers 5 and 6 (Figure 3.12 B). This mirrors the pattern of laminar variation observed for CGC 3.11 B) and indicates that weaker CGC in layer 2/3 may be due to noisier responses in this layer.

If an increase in the NR of MUA results in weaker CGC, these measures should be significantly correlated. In order to investigate the relationship between NR and CGC I ran a correlation analysis and found that these measures were not significantly correlated when all MUs were included in the analysis ($r=0.094$, $p=0.237$) (Figure 3.12 C). I ran the correlation analysis on the data from each layer separately and found that in layer 6 NR and CGC were significantly correlated, with weaker gain changes observed for higher NR values ($r=0.456$, $p=0.025$) (Figure 3.12 C). Significant correlations were not observed for data from other cortical layers (Layer 2/3; $r=-0.12$, $p=0.58$, Layer 4; $r=-0.04$, $p=0.8$, Layer 5; $r=0.056$, $p=0.698$). This may have been due to these responses being less reliable, as indicated by their higher NR values.

Finally, I examined the laminar organisation of CGC after controlling for NRs. I normalised all NRs by the largest NR and multiplied the resulting values by the % change in gain for each MU (Figure 3.13 A). I found that noise-controlled gain changes showed no significant variation with cortical depth (median NR: Layer 2/3=-8.91, Layer 4=-7.02, Layer 5=-6.12, Layer 6=-6.91, Kruskal-Wallis test: $H(3) = 2.27$, $p=0.518$). This indicates that laminar variation in CGC is primarily attributable to laminar variation in the signal to noise ratio of MU responses.

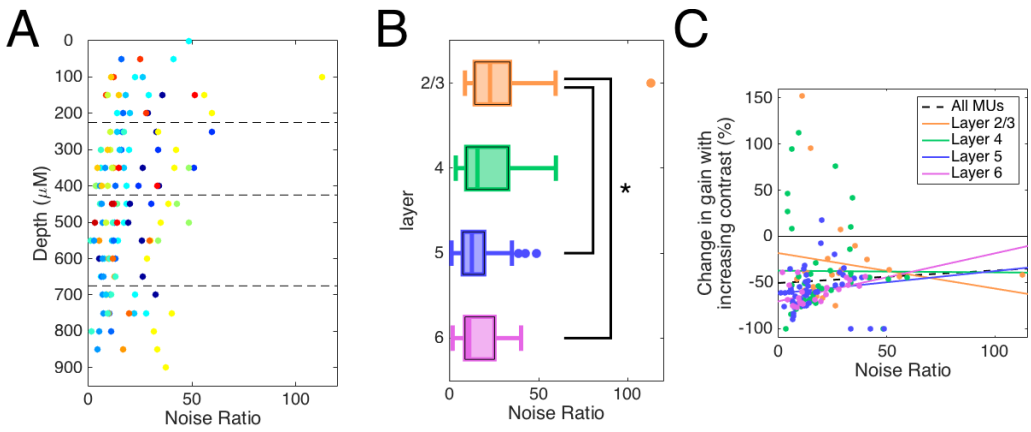

Figure 3.12: A. NRs for each MU plotted as a function of cortical depth from the layer 1-2/3 border. Layer boundaries are indicated by lines and data points are color-coded by penetration. B. Box plots of NRs within each cortical layer. NRs were significantly larger in layer 2/3 compared to layers 5 and 6. C. Scatter plot of NRs against the strength of gain changes in response to an increase in stimulus contrast for each MU. MUs are color coded by cortical layer, as are regression lines. The dashed black line indicates the regression line for all MUs.

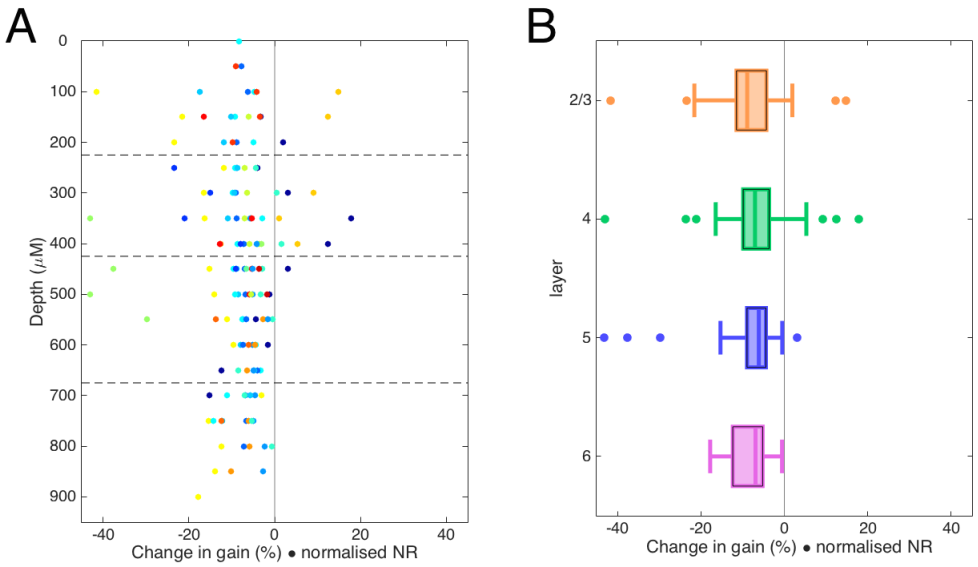

Figure 3.13: A. Strength of gain change after controlling for noise for each MU plotted as a function of cortical depth from the layer 1-2/3 border. Layer boundaries are indicated by lines and data points are color-coded by penetration. B. Box plots of gain control after controlling for noise within each cortical layer. Gain changes showed no significant variation across cortical layers after controlling for noise.

3.4 Discussion

Here I attempted to characterize CGC in the mouse as a test of the canonical nature of this computation and to lay a foundation for investigations into the mechanisms underpinning it. I first investigated the effects of stimulus contrast on the spectrotemporal tuning of MU responses in mouse AC. If responses in AC show a pure gain change in response to a change in stimulus contrast, tuning properties should remain the same. In the ferret, 89% of MUs maintained their best frequency within 1/6 of an octave when contrast was changed by a factor of 3 (Rabinowitz et al., 2011a). In the mouse, however, only 64% of MUs with predictive STRFs maintained their best frequency across contrast conditions, within 1/2 an octave when contrast was changed by a factor of 2, indicating that a greater proportion of units showed changes in BF with stimulus contrast. As in the ferret, however, there was no systematic effect of stimulus contrast on BF.

In the ferret, a 3 fold increase in contrast was found to result in slightly narrower bandwidth of tuning under high contrast stimulation (Rabinowitz et al., 2011a). I found that, in the mouse, the opposite pattern occurs. For a $\sim$2 fold increase in contrast, there is a small but significant tendency for STRFs to increase their bandwidth as contrast increases. STRF estimates for the mouse data appeared noisier than for the ferret data however. This may account for these observed differences as noisy STRF estimates will provide a less reliable assessment of STRF parameters. As in the ferret (Rabinowitz et al., 2011a), no systematic changes in the temporal tuning of STRFs was observed and the most striking change that occurred between STRFs from each condition was a suppressive effect on the peak of tuning under high contrast stimulation. Every MU recorded apart from one showed a reduction in the largest STRF coefficient under high compared with low contrast stimulation. Changes in shape were small enough that a single STRF could be used to successfully predict responses to both low and high contrast stimuli.

These results indicated that, as in the ferret, sensory responses in mouse AC may undergo a reduction in gain in response to an increase in stimulus contrast. I quantified any contrast dependent changes in gain by fitting a single STRF for each MU with an output nonlinearity that could vary between contrast conditions. This allowed me to quantify the effect of contrast on gain and x-axis offset for each MU. As was found to be the case when using these stimuli in the ferret (Rabinowitz et al., 2011a), it was not possible to estimate the saturation point of MU responses as DRCs did not drive responses to saturation. I found that 92.5% of MUs in mouse AC showed a reduction in gain in response to high contrast compared with low contrast stimulation. The median strength of gain control across the population was -55.2% in response to an approximate doubling of stimulus contrast, indicating that in mouse AC contrast gain control may be completely compensatory. This was not found to be the case in ferret AC, where a 3 fold increase in contrast was found to produce a 50% reduction in gain, instead of the 66% reduction that would be entirely compensatory (Rabinowitz et al., 2011a). These findings indicate that unlike in ferret AC, CGC in mouse AC may approximate complete normalisation, as is thought to be the case in V1 (Heeger, 1992).

In the ferret, a contrast dependent increase in threshold that correlated with the strength of gain control was also observed (Rabinowitz et al., 2011a). No such change in threshold was observed here however, indicating that, in the mouse, adaptation to increased stimulus contrast may be more purely divisive than in the ferret where there is also a subtractive component. The lack of a subtractive component in the changes in output nonlinearities is also consistent with the broadening of STRF tuning under high contrast stimulation observed here, as broadening would reflect a reduction, not an increase, in threshold.

The finding that CGC is prevalent in mouse AC allows for investigations into the mechanisms underlying this computation in this particularly useful model organ-

ism. Great progress has been made in understanding the mechanisms underlying CGC in mouse visual cortex using optogenetics, where PV interneurons have been implicated (Atallah et al., 2012; Wilson et al., 2012). A similar line of investigation in mouse AC will not only increase our understanding of information processing mechanisms in the auditory system but will also allow for direct comparisons with findings in visual and other sensory modalities, enabling the question of whether a canonical gain control mechanism exists to be addressed experimentally. The findings presented here indicate that, while CGC is present in mouse AC, interesting differences between species and modalities may exist. This may be seen as evidence for a "serial homology" hypothesis of functional cortical organization (Harris and Shepherd, 2015), in which different cortical areas employ modified versions of common computations and mechanisms that are tailored to the system in question.

As a first step towards understanding the circuit basis of CGC in mouse AC I tested predictions made from different circuit hypotheses regarding the laminar organisation of CGC. I found that the strength of CGC exhibits modest laminar variation in mouse AC, with CGC being strongest in deep cortical layer 5 and 6. In V1, thalamic input has been shown to drive corticothalamic projecting layer 6 neurons which in turn drive a translaminar projecting PV interneuron subtype via facilitating synapses (Olsen et al., 2012; Bortone et al., 2014). Inhibition via this route has been found to reduce gain across all other cortical layers in V1, but not in layer 6. The presence of strong CGC in layer 6 indicates that it may not be implemented through translaminar projecting layer 6 PV interneurons (Olsen et al., 2012; Bortone et al., 2014), as this mechanism would require layer 6 responses to not be invariant to stimulus contrast. Layer 6 corticothalamic projecting neurons may contribute to CGC, however, via interactions with the thalamus. Simultaneous recordings in these two structures combined with the temporal evolution of gain control would be necessary in order to investigate this possibility.

The presence of strong CGC in layer 4 is consistent with gain control being primarily implemented in this layer as it is thought to represent the earliest stage of cortical processing (Douglas and Martin, 1991b). In keeping with this, I found evidence that laminar variation may be primarily attributable to variation in the noise of MU responses across cortical layers, as this variation mirrored that of CGC. This indicates that intracortical processing outside of layer 4 may contribute less to the CGC compared with mechanisms in layer 4. Thalamic inputs to AC have been found to specifically innervate PV interneurons in all cortical layers (Ji et al., 2015) however. This raises the alternative possibility that feedforward inhibition across all cortical layers may be the primary mechanism of CGC, rather than it being simply inherited outside of layer 4. Circuit level experiments, made possible by the work presented in this chapter, will be necessary in order to test these circuit models experimentally. This is the topic of the next chapter, where I consider the role of cortical PV interneurons in CGC.

Chapter 4

Parvalbumin-positive interneurons and gain control

Having characterised contrast gain control (CGC) in mouse auditory cortex (AC), I next investigated whether parvalbumin (PV) expressing interneurons in cortex contribute to this computation. Multiple lines of evidence have implicated PV interneurons in gain control in non-auditory cortical areas. The earliest indication that this subtype may play a role in gain control came from *in vitro* evidence that fast-spiking interneurons in hippocampus and cortex to modulate the dynamic range of excitatory neurons via feed-forward inhibition (Pouille et al., 2009). More recently, molecular techniques have allowed researchers to target the expression of optogenetic constructs such as channelrhodopsin (ChR2) (Boyden et al., 2005) and archaerhodopsin (Arch) (Chow et al., 2010) to PV interneurons *in vivo*, allowing the activity of these neurons to be increased or decreased through exposure to specific wavelengths of light.

Using this approach, moderate increases in PV interneuron firing have been found to reduce the gain of sensory-evoked spiking responses in primary visual cortex (V1), while reducing PV firing has been found to increase gain (Atallah

et al., 2012; Wilson et al., 2012). These manipulations left the bandwidth of tuning and the orientation preferences of the pyramidal cells unchanged, indicating that the only effect of PV interneuron activity was to modulate neuronal gain. If PV interneurons are activated strongly, to the point where they suppress pyramidal spiking by more than 50%, they are capable of reducing tuning bandwidth by suppressing weak responses below spike threshold (Lee et al., 2012, 2014). This level of activation is unlikely to routinely occur under natural sensory stimulation, however, indicating that the primary computational function of PV activity in V1 is to control pyramidal cell gain, in a manner that is inversely related to PV firing rates. Increases in PV interneuron activity have also been shown to be capable of increasing pyramidal cell gain in cortex through the generation of gamma rhythms, raising the possibility that PV interneurons may modulate gain in a variety of ways (Sohal et al., 2009a).

PV interneuron inhibition is also a plausible candidate mechanism for CGC in AC. The responses of AC PV interneurons are more linear with respect to stimulus intensity than pyramidal cell responses, a characteristic that has been suggested to make them well-suited for implementing stimulus driven gain control in AC (Moore and Wehr, 2013). PV interneurons in AC show frequency tuning that may be matched to the local population (Moore and Wehr, 2013), a feature that is consistent with the finding that CGC in A1 predominantly depends on contrast within the receptive field of the neuron (Rabinowitz et al., 2011b). PV interneuron inhibition may therefore underlie CGC in auditory as well as visual cortex and may represent a canonical mechanism for gain control across sensory cortical areas.

Here, I investigated whether PV interneuron inhibition contributes to CGC in AC by asking whether PV interneuron activity modulates the gain of sensory responses in AC and, if so, whether this activity is appropriately driven by stimulus contrast. I hypothesised that PV interneurons increase their firing rate during high contrast

stimulation (Figure 4.1A,B) and that increases in PV interneuron firing rates produce a reduction in the gain of sensory responses in AC (Figure 4.1C). I tested the latter hypothesis by optogeneticially manipulating PV interneuron activity in mouse AC while simultaneously performing extracellular recordings during acoustic stimulation in order to asses the impact of PV interneuron activity on MU gain. I predicted that optogenetically increasing PV interneuron activity should decrease the gain of AC responses while optogenetically reducing PV interneuron activity should result in an increase in gain. I subsequently tested the hypothesis that PV interneurons show an increase in firing rate under high contrast stimulation by performing single unit recordings from putative PV interneurons during stimulation with high and low contrast dynamic random chord (DRC) stimuli.

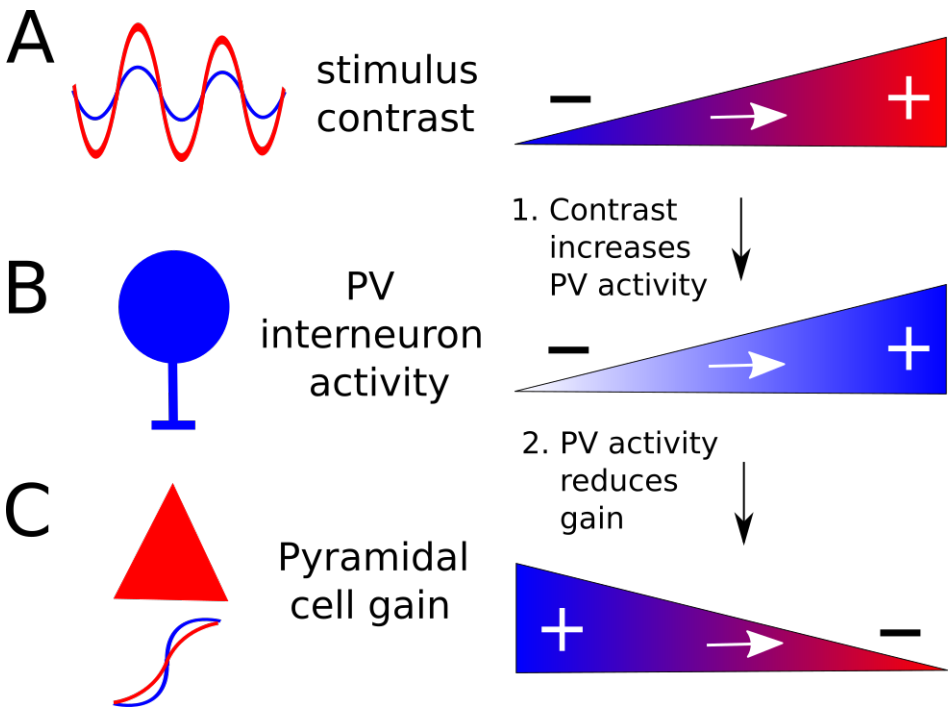

Figure 4.1: Predictions regarding PV interneuron activity and gain control A. Increasing stimulus contrast is known to result in a decrease in the gain of sensory responses in AC (C). One mechanism that could mediate this is PV interneuron inhibition (B). Two predictions arise from this hypothesis; 1: Increasing stimulus contrast results in an increase in PV interneuron activity. 2: Increasing PV interneuron activity reduces gain and decreasing PV interneuron activity increases gain.

4.1 Methods

I made extracellular recordings from the AC of 24 anaesthetised mice between 8 and 12 weeks old using silicon multi-electrodes (NeuroNexus, Ann Arbor, United States). The data presented in this chapter was collected from the same animals as in chapter 3. All mice expressed cre-recombinase under the PV promoter, allowing for targeting of optogenetic constructs to PV interneurons. In order to suppress PV interneuron activity, I crossed homozygous PV ^{cre} (Jax No: 008069) (Hippenmeyer et al., 2005) with homozygous Ai35D mice (Jax No: 012735) (Madisen et al., 2012), which express Arch (Chow et al., 2010) and green fluorescent protein (GFP) in a cre-dependent manner. PV ^{cre} x Ai35D offspring therefore express Arch and GFP in PV interneurons. MUs were recorded in PV ^{cre} x Ai35D mice (N = 8) between the ages of 8 and 12 weeks while auditory stimuli were presented to the contralateral ear. On half of the trials, I suppressed PV cell activity by exposing AC to amber light, delivered from an LED (Doric Lenses Inc, Quebec, Canada) coupled to a 0.55mm diameter fibre-optic cable. I positioned the fibre-optic cable over the recording site using a micromanipulator, with the tip $<$1mm above the surface of the cortex. The maximum output power of the amber LED was $\sim$70mW/mm^2, although the light level was calibrated for each penetration in order to achieve a physiologically defined level of disinhibition of sensory evoked responses, as discussed below.

I targeted expression of ChR2 (Boyden et al., 2005) and enhanced yellow fluorescent protein (eYFP) to PV interneurons in order to increase PV interneuron activity during recordings. PV ^{cre} mice were anaesthetised using 0.05 mg/kg fentanyl (Sublimaze, Janssen Ltd, High Wycombe, Bucks, UK), 0.5 mg/kg medetomidine (Domitor, Pfizer, Walton Oaks, Surrey, UK) and 5 mg/kg midazolam (Hypnovel, Roche Products Limited, Welwyn Garden City Hertfordshire, UK), delivered intraperitoneally. Animals subsequently received stereotaxic injections of $\sim$500nL

AAV-EF1a-DIO-hChR2(H134R)-EYFP-WPRE-pA serotype 2 (UNC Vector Core, Chapel Hill, NC, United States) in AC. I identified AC using cranial landmarks as described in section 2.1.1. I performed injections 3-5 weeks prior to terminal recordings, in order to allow for expression of ChR2 and EYFP in PV interneurons. I recorded MUA in 14 of these mice using 1x32 linear multi-electrodes and performed single unit recordings in 6 more of these mice using 64 channel Buzsaki probes (NeuroNexus, Ann Arbor, United States). I identified single units manually following spike sorting using KlustaKwik, as described in section 2.1.2. Animals were between 8 and 12 weeks of age at the time of recording. I increased PV cell activity by exposure to blue light, delivered from an LED (Doric Lenses Inc, Quebec, Canada) coupled to a fibre-optic cable. As with the Arch experiments, I positioned the fibre-optic cable over the recording site using a micromanipulator, with the tip $<$1mm above the surface of the cortex. The maximum output power of the blue LED was $\sim$200mW/mm^2, although the light level delivered was typically lower than this. Light level was calibrated for each penetration in order to achieve a physiologically defined level of inhibition.

In all animals, I made recordings in the left AC and delivered stimuli to the contralateral ear. All stimuli had a mean sound level of 80 dB. DRC stimuli that were either low (σ_L $\sim$6.2 dB, c = σ_P/μP = 0.68) or high in contrast (σ_L $\sim$11.97 dB, c = σ_P/μ_P = 1.2) were delivered. Bursts of white noise that were 50 ms in duration were also delivered in order to test acoustic responses of MUs and the effect of optogenetic activity on these responses. Each penetration showed noise-evoked responses on multiple channels. For the ChR2 expressing mice, 85% of MUs recorded showed noise evoked responses, defined as a doubling of baseline firing in response to noise onset. For the Arch expressing mice, 40% of MUs recorded showed noise evoked responses. Noise responsiveness was only used as a selection criterion when assessing the strength of optogentic effects. For both noise and DRC

stimuli, LED illumination of the cortex occurred on half of trials. "Light on" and "light off" trials were presented in a pseudorandom order. For noise stimuli, light onset occurred 250 ms before sound onset and lasted for 750ms. For DRC stimuli, light onset was coincident with sound onset and lasted for the entire duration of the stimulus.

I constructed a peri-stimulus time histogram (PSTH) of MU responses to noise burst stimuli by binning spike times after stimulus onset into 25 ms time bins. The first two bins following light onset were removed in order to avoid light artifacts. For both "light on" and "light off" conditions, the bin within 150 ms of noise onset with the highest MUA signal amplitude was taken as the peak response. I used the difference in peak firing rate between "light on" and "light off" conditions in order to quantify the effect of the optogenetic manipulations. I calibrated the strength of the optogenetic manipulations by adjusting light intensity in order to produce a $\sim$5% difference in the peak evoked firing rate to noise burst stimuli. I performed post-hoc quantification of the strength of the optogenetic manipulations only for MUs that were included in subsequent analysis, after meeting the selection criterion of having a predictive STRF. I assigned cortical layers using CSD profiles of noise responses, as described in section 2.3.4.

Following completion of terminal experiments, I removed the brain from the skull and immersed it in 4% paraformaldehyde for a minimum of 24 hours. After fixation was complete, 50 μm sections were cut using a cryotome. I identified PV interneurons by immunohistochemical labeling with a fluorescent marker. I rinsed sections with 0.25% Triton in phosphate buffered saline (PBS) before incubating them in a 10% donkey serum solution for one hour in order to block background labeling. I then incubated the sections in a 1/4000 PV goat antibody PBS solution with 3% donkey serum for 48 hours at 4°C. I then rinsed the sections in 0.25% PBS Triton before incubating them in a secondary red fluorescent antibody in 0.25%

PBS Triton (1/1000 Donkey anti-goat IgG Cy3). Finally I rinsed sections in PBS before mounting them on glass slides and cover-slipping with VectaShield HardSet mounting medium with DAPI (VectorLabs, Burlingame, CA, United States).

4.2 Results

4.2.1 Validation of optogenetic manipulations

First, I validated my optogenetic manipulations via two methods. The AAV-EF1a-DIO-hChR2(H134R)-EYFP-WPRE-pA construct used in order to target expression of ChR2 to PV interneurons also encodes EYFP so that it is co-expressed with ChR2. This allowed me to validate the specificity of ChR2 expression by immunohistochemically staining sections of AC with a PV antibody and identifying neurons which co-expressed PV. I found expression of ChR2-eYFP in ∼90% of PV positive neurons (Figure 4.2A). Expression of ChR2 in non-PV interneurons was rare, in keeping with previous reports that found <10% misexpression in cortex using this mouse line (Sohal et al., 2009b; Cardin et al., 2009, 2010; Kerlin et al., 2010; Taniguchi et al., 2011; Adesnik et al., 2012; Pfeffer et al., 2013; Kvitsiani et al., 2013; Seybold et al., 2015). I used the same approach in order to validate endogenous Arch expression, as the Ai35D mouse line co-expresses Arch and GFP. I found that ∼95% of PV positive neurons expressed Arch-GFP (Figure 4.2B) while <10% of PV negative neurons showed Arch-GFP expression. Endogenous Arch-GFP expression was weaker than virally mediated ChR2-eYFP expression, as expected.

The second validation method I employed was to examine the effect of these optogenetic manipulations on noise responses in AC. This allowed me to validate the efficacy of these manipulations for each MU presented in subsequent analyses (Figure 4.2C,D). In order to be included in all further analysis, I required MUs to meet the selection criterion of having a predictive spectro-temporal receptive

field (STRF). I fitted STRFs to 90% of responses to DRCs during "light off" control trials and examined their predictive performance by using the STRFs to predict responses to the remaining 10% of DRCs. I deemed STRFs predictive if the correlation coefficient (CC) between the predicted and actual response was larger than 0.04. For the ChR2 experiments, 28 MUs passed this criterion. For these sites, optogenetic activation of PV interneurons significantly inhibited peak MU responses to noise (median change: -1.15%; sign-rank test: $p<0.001$)(Figure 4.2C,E). For the Arch experiments 121 MUs passed this criterion. For these sites, optogenetic suppression of PV interneuron activity resulted in significant disinhibition of peak MU responses (median change in spike probability: 2.44%; sign-rank test: $p<0.001$) (Figure 4.2D,E).

Next, I investigated whether these effects depended on cortical depth. I positioned the LED over the cortical surface, which may have resulted in weaker effects in deep cortical layers. I found that for optogenetic activation of PV interneurons using ChR2, the inhibition of noise evoked responses was not depth-dependent (Kruskal Wallis test: Chi-square(3) = 6.65, $p=0.919$) (Figure 4.2F). I found that the disinhibition of evoked responses produced by optogenetic suppression of PV interneuron activity was depth dependent. The disinhibition was significantly stronger in layer 4 than in layer 6 (median change: Layer 2/3 = 2.17%, Layer 4 = 2.93%, Layer 5 = 2.97%, Layer 6 = 1.29%) (Kruskal Wallis test: Chi-square(3) = 11.89, $p=0.008$) (Figure 4.2F). This result is most likely attributable to the reduced light penetration in deep cortical layers and the weaker expression observed for the endogenous Arch, compared with the virally-mediated expression of ChR2. The strength of virally-mediated expression for ChR2 and the increased power output of the blue LED used for this manipulation made it possible to inhibit responses even in the deep layers of the cortex. The lower wavelength of amber light, used for the Arch manipulation, should in theory enable deeper light penetra-

tion. In practice, however, the weaker expression levels achieved through endogenous expression, coupled with the lower power output of the amber LED, may have resulted in the lack of disinhibition observed in deep cortical layers, where the light level will inevitably be weaker.

I also analysed the effect of these manipulations on baseline activity in the pre-noise stimulus period. Arch-mediated suppression of PV interneuron activity resulted in significant disinhibition of baseline activity (median change: 2.74%; sign-rank test: $p<0.001$) whereas ChR2-mediated activation of PV interneuron activity left baseline activity unaffected (median change: -0.15%; sign-rank test: $p=0.601$) (Figure 4.2G). The disinhibition resulting from PV suppression was not depth dependant (median change: Layer 2/3 = 2.07%, Layer 4 = 2.93%, Layer 5 = 2.51%, Layer 6 = 2.92%) (Kruskal Wallis test: Chi-square(3) = 1.05, $p=0.788$) (Figure 4.2H). PV activation similarly had no layer specific effects on baseline activity (median change: Layer 2/3 = -0.15%, Layer 4 = -0.22%, Layer 5 = 0.05%, Layer 6 = -0.55%) (Kruskal Wallis test: Chi-square(3) = 2.51, $p=0.474$) (Figure 4.2H).

4.2.2 Bidirectional modulation of PV interneuron activity has negligible effects on STRF tuning

I next investigated whether changes in PV interneuron activity alter the spectrotemporal response properties of AC MUA. If PV interneuron activity controls only the gain of neuronal responses, tuning properties should remain largely unchanged. I examined the effects of PV interneuron activation by fitting STRFs to MU responses under “light off” (Figure 4.3A) and “light on” (Figure 4.3B) conditions. The same approach was used in order to examine the effects of PV interneuron suppression (Figure 4.5). I examined the effects of these optogenetic manipulations on 4 STRF parameters, the best frequency (BF), bandwidth of the frequency kernel, bandwidth of the temporal kernel and the largest coefficient in the STRF. Stimulus contrast has

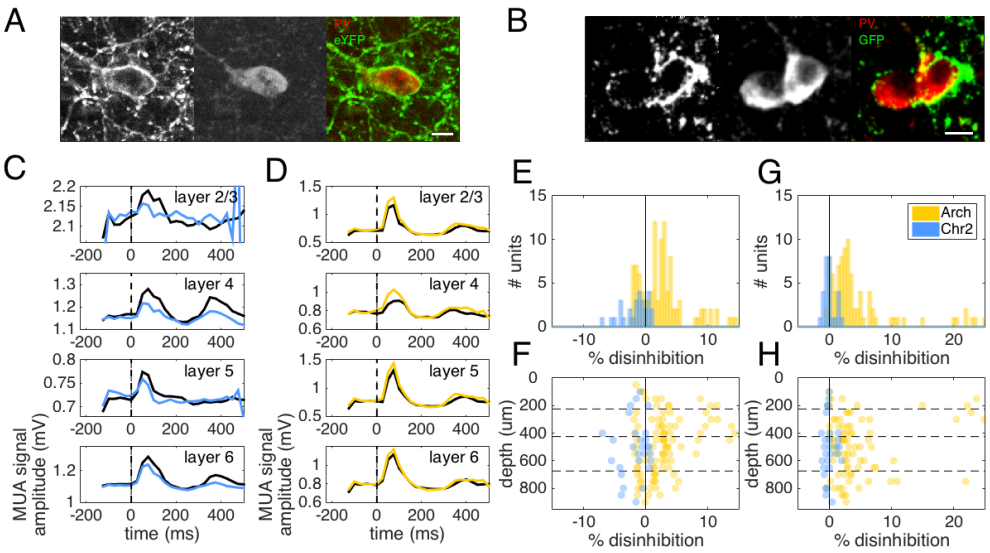

Figure 4.2: A. Left Panel: Confocal micrograph of a ChR2-eYFP positive neuron in a section of mouse AC. Middle Panel: I treated slices with a PV antibody and subsequently counterstained them with a red fluorescent dye in order to visualise PV expression. Right Panel: Merged image of green (ChR2-eYFP) and red (PV) channels showing co-expression of ChR2-eYFP and PV in the same neuron. Scale bar: 10μm. B. Same as in A but for co-expression of Arch-GFP and PV in two neurons. C. PSTHs of responses to noise burst stimuli during optogenetic activation of PV interneurons (blue traces). I removed the first two bins of the PSTH in order to exclude photoelectric artifacts from the analysis. Control traces with no optogenetic stimulation are shown in black. All PSTHs show suppression of peak responses under conditions of increased PV activity (blue) but no suppression of baseline activity. D. PSTHs of noise responses during optogenetic suppression of PV interneurons (amber traces). As in C, control traces with no optogenetic stimulation are shown in black. E. Histograms showing the strength of the effect of PV interneuron activation (blue) and suppression (amber) on peak firing rate in the PSTH for MUs that showed noise evoked responses. PV activation resulted in significant inhibition of evoked responses, while suppression resulted in significant disinhibition of responses to noise stimuli. F. Effects of optogenetic manipulations on evoked neural activity across the depth of cortex. Amber circles correspond to individual MUs under PV suppression while blue circles correspond to individual MUs under PV activation. PV activation suppressed noise responses across all cortical layers and did not vary significantly over cortical depth. The strength of the disinhibition of noise-evoked responses observed during PV suppression was found to decrease significantly with cortical depth. G. Histograms showing the strength of the effect of PV interneuron activation (blue) and suppression (amber) on baseline activity PSTH for the same MUs as in E and F. PV activation had no effect on baseline activity, while suppression resulted in significant disinhibition of baseline responses. H. Effects of optogenetic manipulations on baseline activity across the depth of cortex for the same MUs as in E and F. Effects on baseline activity were consistent across the depth of cortex.

been shown to leave these parameters largely unchanged, apart from the largest STRF coefficient which shows a compensatory reduction in response to increasing contrast (Rabinowitz et al., 2011b) (Chapter 3). If PV interneuron activity drives contrast induced gain changes, optogenetic manipulations of their activity should similarly leave these STRF parameters largely unchanged. I assessed the significance of changes in STRF parameters produced by PV interneuron activation by using a non-parametric paired-sample Wilcoxon signed-rank test to test for differences in median values between "light on" and "light off" conditions. I did not perform comparisons for MUs with STRFs that failed to predict responses in either "light on" or "light off" conditions.

First, I considered the effect of ChR2-mediated PV interneuron activation on these parameters. I found that, for PV interneuron activation, STRF shapes were similar under "light off" (Figure 4.3A) and "light on" conditions (Figure 4.3B). This manipulation had no significant effect on BF (sign-rank test; $p=0.074$) (Figure 4.3C), bandwidth of the frequency kernel (sign-rank test; $p=0.338$) (Figure 4.3D) or the bandwidth of the temporal kernel (sign-rank test; $p=0.121$) (Figure 4.3E). The magnitude of the largest coefficient in the STRF however was significantly reduced under "light on" conditions in these MUs ("light off" median; 0.16μV/dB/s, "light on" medians: 0.14μV/dB/s, sign-rank test; $p<0.001$) (Figure 4.3F).

For each STRF parameter a small population of MUs that showed large changes with PV activation was observed. In order to asses whether these outliers come from a single subpopulation of MUs I identified outliers for each of these parameters independently and examined the overlap of these outliers. If the outliers for each parameter come from a single population, the MUs defined as outliers based on the different parameters should be largely overlapping. Outliers were defined as units that showed parameter changes between light conditions that fell outside 1.65 standard deviation of the distribution of changes, as these values would have

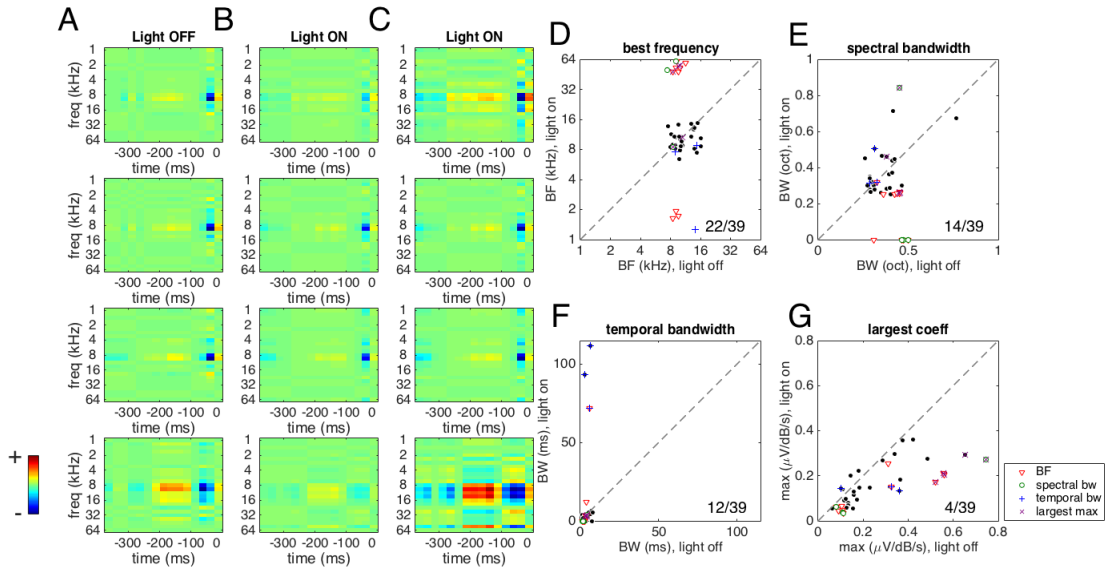

Figure 4.3: A. Example control STRFs estimated for the "light off" condition. Spectrotemporal features that increase the MUA amplitude are shown in warmer colours while suppressive features are shown in cooler colours. B. STRFs estimated for the same MUs as in A but under "light on" conditions, during ChR2-mediated PV interneuron activation. Color scaling is the same as in A. This manipulation results in supression of STRF coefficients. C. "Light on" STRFs, as shown in B, but with independent colour scaling. STRFs are largely the same under both "light off" and "light on" conditions (Panels A and C). D-G. Comparison of STRF parameters under "light off" (abscissa) vs. "light on" (ordinate) conditions. I observed no significant changes in best frequency (D), frequency tuning bandwidth (E) or temporal bandwidth (F). The largest STRF coefficients were significantly reduced during PV interneuron activation (G). For each parameter, outliers were identified as values that fell outside 1.645 standard deviations of the distribution of these values. They are shown here in each plot using different markers and colors as indicated in the legend. Several MUs were identified as outliers on the basis of changes in multiple parameters, as indicated by the overlap in outlier symbols. Ratios in the corner of each plot indicate the number of MUs above the identity line, y=x, over the total number of MUs.

a <10% probability of occurring in such a distribution (i.e. p <0.1). A threshold of p <0.05 is used elsewhere in this thesis for outlier detection. This threshold failed to capture outliers in this particular dataset however, due to differences in the distribution of parameter changes. As a result of this the threshold was increased to p <0.1. BF outliers are indicated by the inverted red triangle symbol, spectral bandwidth outliers are indicated by the green circle temporal bandwidth outliers are indicated by the blue plus symbol while largest max coefficient outliers are indi-

cated by the purple cross (Figure 4.3 D-G). 8 out of 15 outlier MUs were identified as outliers based on changes in more than one parameter, indicating that a small population of MUs that show changes in multiple STRF parameters does exits. In keeping with this, all of these outlier MUs were recorded from 2 penetrations, each from a different animal.

The STRFs of the outliers identified for each condition were examined in order to investigate the cause of the variability in these parameters. The STRFs of outliers identified through large changes in BF (Figure 4.4A), spectral bandwidth (Figure 4.4B), temporal bandwidth (Figure 4.4C) and the largest coefficient (Figure 4.4D) all showed strong suppression during PV interneuron activation. These large changes are therefore likely a result of inaccurate STRF parameter estimates in the "light on" condition for this sub-population, rather than genuine changes in tuning as a result of PV interneuron activation.

I examined the effect of Arch-mediated PV interneuron suppression on STRF parameters using the same approach as described above. As with PV activation, STRF shapes were similar under "light off" (Figure 4.5A) and "light on" conditions (Figure 4.5B). PV suppression did not produce a significant change in the bandwidth of frequency (sign-rank test; $p=0.59$) (Figure 4.5D) or temporal tuning (sign-rank test; $p=0.528$) (Figure 4.5E). I observed a significant change in best frequency although median values remained unchanged between "light off" and "light on" conditions ("light off" median; 9.51kHz, "light on" median; 9.51kHz; "light off" mean; 7.61kHz, "light on" mean; 7.81kHz; sign-rank test; $p=0.047$) (Figure 4.5C). I also observed a significant increase in the largest STRF coefficient in these MUs during PV suppression ("light off" median; 0.39μV/dB/s, "light on" median: 0.4μV/dB/s, sign-rank test; $p<0.001$) (Figure 4.5F).

Few MUs showed drastic changes in tuning as a result of PV interneuron suppression. I identified outliers for each of these parameters independently. Outliers

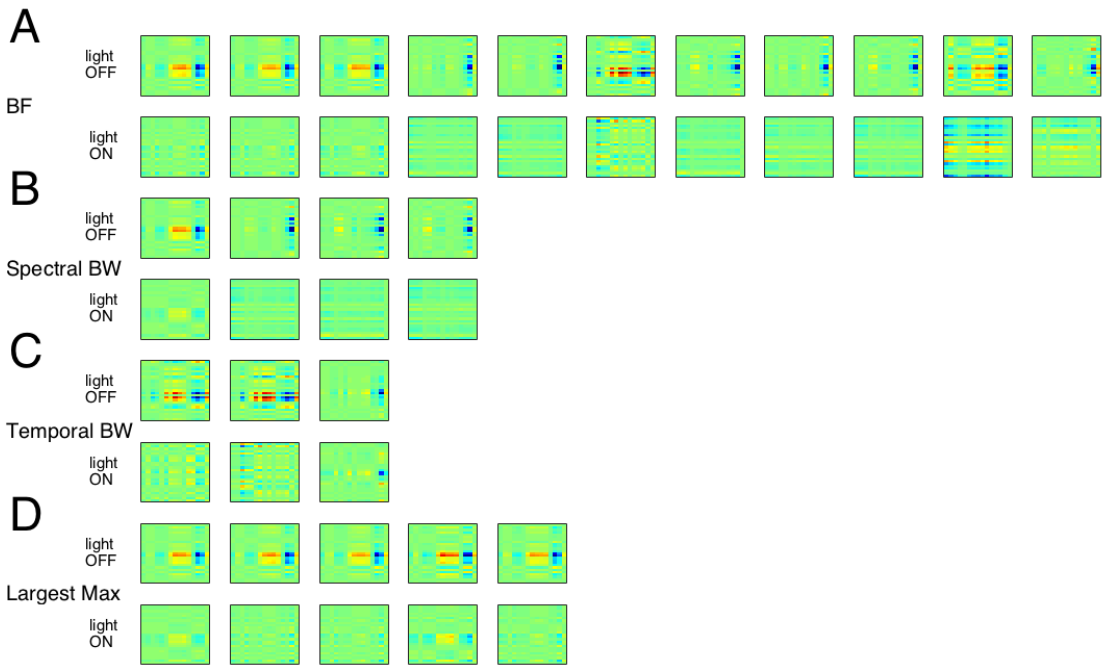

Figure 4.4: STRFs of outliers identified from changes in each STRF parameter as a result of PV activation. Outliers were defined as values that fell outside 1.645 standard deviations of the distribution of these values, corresponding to a p value of <0.1. A. STRFs for MU outliers identified through large changes in BF. B. STRFs for MU outliers identified through large changes in spectral bandwidth. C. STRFs for MU outliers identified through large changes in temporal bandwidth. D. STRFs for MU outliers identified through large changes in the largest STRF coefficient. Estimates under "light off" and "light on" conditions are shown with the same color scaling and demonstrate that these MUs undergo strong suppression in the "light on" condition, resulting in erroneous estimates of STRF parameters under this condition.

were required to show changes between light conditions that fell outside 1.96 standard deviation of the distribution of changes, as these values would have a <5% probability of occurring in such a distribution (i.e. p <0.05) (Figure 4.5 D-G). Only 3 of 22 outlier MUs were identified as outliers based on changes in more than one parameter. All of these outlier MUs were recorded from 3 penetrations, each from a different animal.

The STRFs of the outliers identified for each condition were examined in order to investigate the cause of the variability in these parameters. The STRFs of outliers

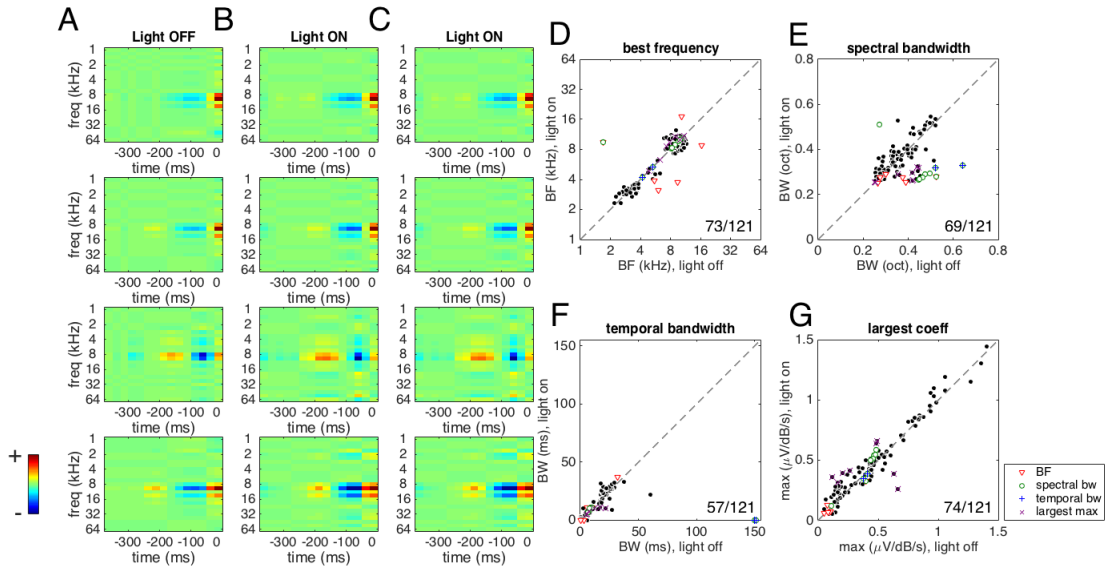

Figure 4.5: A. Example control STRFs estimated for the "light off" condition. Spectrotemporal features that increase the MUA amplitude are shown in warmer colours while suppressive features are shown in cooler colours. B. STRFs estimated for the same MUAs as in A but under "light on" conditions, during ARCH-mediated PV interneuron suppression. Color scaling is the same as in A. This manipulation results in a small increase in peak STRF coefficient. C. "Light on" STRFs, as shown in B, but with independent colour scaling. STRFs had similar shapes under light off" and "light on" conditions (Panels A and C). D-G. Comparison of STRF parameters under "light off" (abscissa) vs. "light on" (ordinate) conditions. A small but significant change was observed in best frequency (D) under "light on" conditions. The bandwidth of frequency (E) and temporal tuning (F) remained unaffected. The largest STRF coefficient decreased during PV interneuron suppression. (G). For each parameter, outliers were identified as values that fell outside 1.96 standard deviations of the distribution of these values. They are shown here in each plot using different markers and colors as indicated in the legend. The outlier parameters did not tend to come from the same population of MUs, as indicated by the lack of overlap in outliers identified from the different parameters. Ratios in the corner of each plot indicate the number of MUs above the identity line, y=x, over the total number of MUs.

identified through large changes in BF showed relatively noisy STRFs, indicating that BF estimates may have been innacurate in both "light off' and "light on" conditions (Figure 4.4A). Outliers identified through large changes in spectral bandwidth had reasonable STRFs and appeared to show a genuine narrowing of tuning with PV suppression (Figure 4.4B). Outliers identified through large changes in temporal bandwidth had large bandwidth under "light off' conditions and near 0 ms bandwidth under"light on" conditions. This likely reflects an inaccurate estimate

of temporal bandwidth as a result of the lack of a clear short latency excitatory response in these STRFs (Figure 4.4C). As for the spectral bandwidth outlier MUs, outliers identified through large changes in the largest STRF coefficient had reasonable STRFs, indicating a genuine change in this parameter with PV suppression for these MUs (Figure 4.4D).

In order to assess the significance of the observed changes in STRF shape, as well as the importance of unmeasured changes, I predicted MU responses to stimuli across light conditions. This involved using "light off" STRFs, scaled by a MU-specific gain factor, to predict responses to "light on" responses. For ChR2-mediated PV interneuron activation, STRFs that were trained and tested on "light off" data predicted responses under "light on" and "light off" conditions equally well ("light off" and "light on" median CC: 0.071/0.061, sign-rank test: $p=0.665$) (Figure 4.7A). For Arch mediated PV interneuron suppression, CCs were marginally lower under "light on" conditions compared to "light off" conditions ("light off" median; 0.081%, "light on" median; 0.079%; sign-rank test: $p<0.001$) (Figure 4.7B). This decrement in performance was only of 1.85% however, indicating that STRF performance is largely the same under "light off" and "light on" conditions.

4.2.3 PV interneuron activity is capable of modulating the gain of cortical responses

I next investigated whether PV interneuron inhibition is capable of modulating the gain of neural responses in AC, a necessary criteria for them to be capable of implementing CGC. I fitted a linear-nonlinear model to each MU in order to quantify gain changes induced by the optogenetic manipulations, in the same way as I did for contrast-induced gain changes in Chapter 3. To do this, I fitted a single STRF to the data from both "light off" and "light on" conditions for each MU. The output of the STRF was then passed through a sigmoidal output nonlinearity, as in equation 2.9,

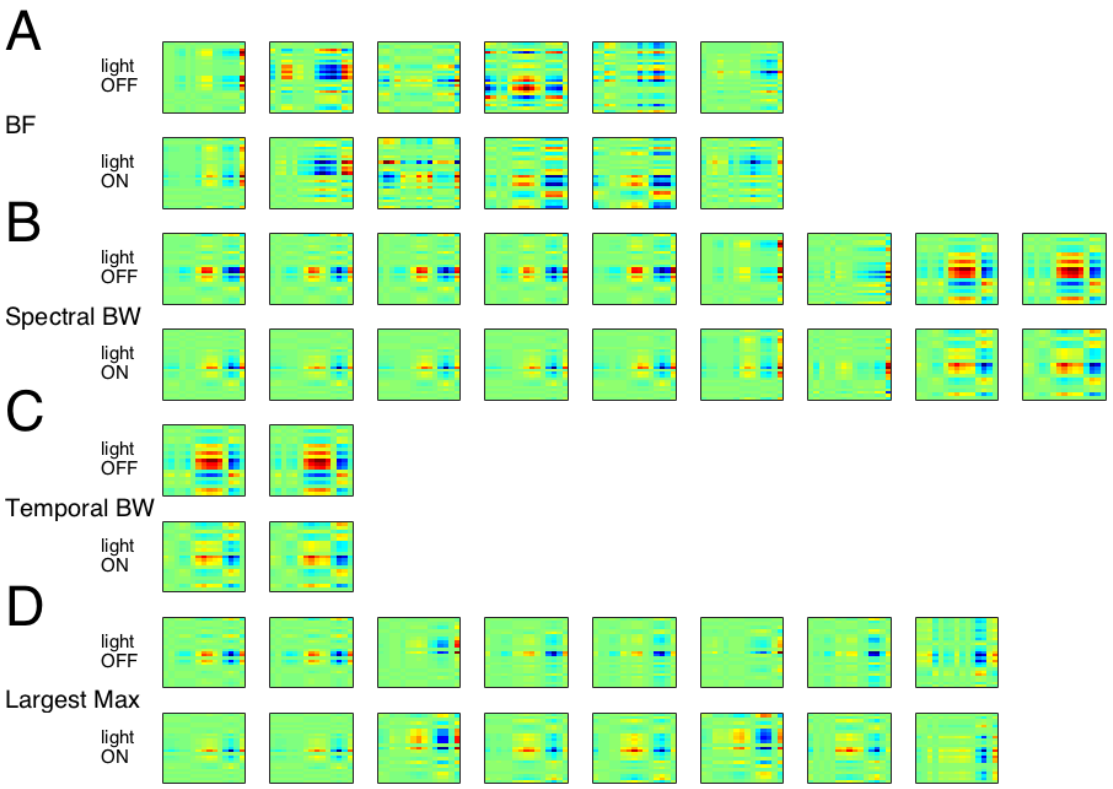

Figure 4.6: STRFs of outliers identified from changes in each STRF parameter as a result of PV suppression. Outliers were defined as values that fell outside 1.96 standard deviations of the distribution of these values, corresponding to a p value of <0.05. A. STRFs for MU outliers identified through large changes in BF. These STRF estimates are relatively poor which may have resulted in erroneous measurements of BF in both "light off" and "light on" conditions. B. STRFs for MU outliers identified through large changes in spectral bandwidth. A narrowing in tuning can be observed in the "light on" condition. C. STRFs for MU outliers identified through large changes in temporal bandwidth. These STRFs do not show a short latency excitatory response which may have resulted in an inaccurate estimate of changes in temporal bandwidth. D. STRFs for MU outliers identified through large changes in the largest STRF coefficient. These STRF estimates are reasonable, indicating a genuine change in this parameter as opposed to incorrect measurements across conditions. STRF coefficients are independently scaled for plotting.

that could vary between light conditions. This allowed me to capture the effects of optogenetic manipulations on offsets along y (parameter a) and x (parameter c) axes and the slope (parameter d) of the output nonlinearity. The saturation point (parameter b) was fixed between light conditions as responses did not typically sat-

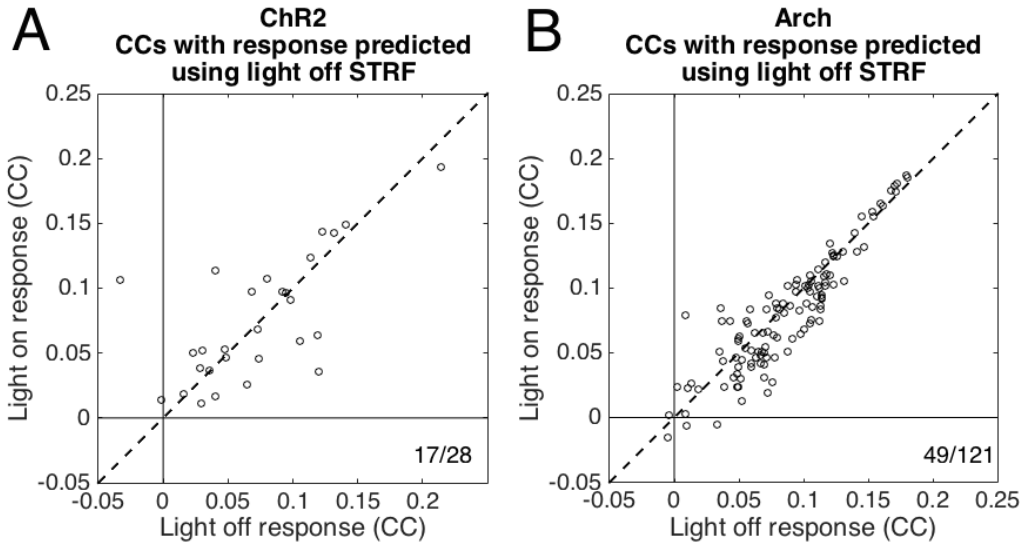

Figure 4.7: Correlation coefficients between actual responses and responses predicted by ChR2 (A) and Arch (B) STRFs. A. For PV interneuron activation, STRFs fitted under "light off" conditions predicted responses under "light on" and "light off" conditions equally well. B. A slight decrement in performance in STRFs fitted under "light off" conditions was observed during "light on" conditions for the PV interneuron suppression manipulation. This decrement was of <2% however. These findings indicate that tuning was largely unaffected by the manipulation of PV interneuron activity. Ratios in the corner of each plot indicate the number of MUs above the identity line, y=x, over the total number of MUs.

urate. Parameters a, c and d correspond to the baseline activity, threshold and gain of the MUs respectively.

I first examined the effects of ChR2-mediated PV interneuron activation on output nonlinearities. Under "light off" control conditions, MUs responded to DRCs as expected (Figure 4.8 A,D). Under "light on" conditions when PV interneurons were activated, responses were suppressed (Figure 4.8 B,E). Many MUs showed a reduction in gain but the largest effect appeared to be a subtractive reduction in baseline activity (Figure 4.8 C,F).

I quantified the effect of PV interneuron activation on sensory responses by comparing the output nonlinearity parameters under "light off" and "light on" conditions (Figure 4.9A-C). ChR2-mediated PV interneuron activation significantly reduced baseline activity (Median change: -3.793%, sign-rank test: $p<0.001$) (Fig-

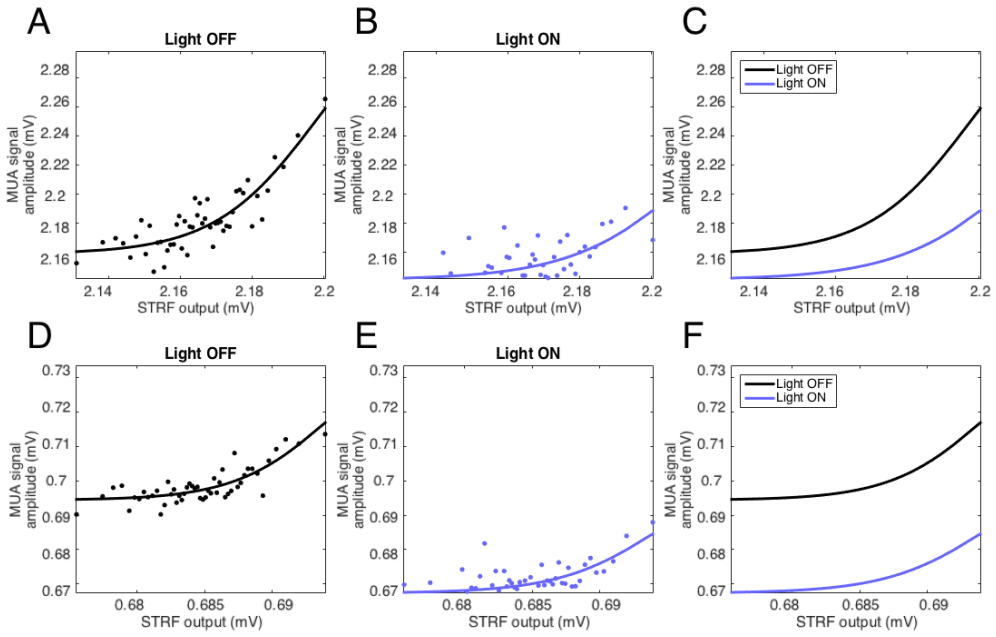

Figure 4.8: Gain changes produced by PV activation were quantified by fitting a single STRF to the data from each light condition for each MU. An output non-linearity that could vary between contrast conditions was fitted to the output of this STRF. A: Sigmoidal output nonlinearities fitted to one representative MU under "light off" (black) conditions. Open circles are binned responses. B. Nonlinearity fitted to the same MU under "light on" (blue) conditions. C. Overlay of output nonlinearities for this MU. The clearest effect of ChR2-mediated PV interneuron activation was a change in offset as well as a small reduction in gain. D-F. Same as in A-C for a second representative MU.

ure 4.9D) but left thresholds unaffected (Median change: -2.554%, sign-rank test: p=0.439) (Figure 4.9E). Gain was also significantly reduced during PV interneuron activation (Median change: -39.573%, sign-rank test: p=0.004)(Figure 4.9F).

I next examined the effects of Arch-mediated PV interneuron suppression on output nonlinearities. Under "light off" control conditions, MUs responded to DRCs as expected (Figure 4.10 A,D). Under "light on" conditions when PV interneurons were suppressed, responses were disinhibited (Figure 4.10 B,E). The majority of MUs appeared to show a clear increase in gain (Figure 4.10 C,F).

As for ChR2-mediated PV activation, I quantified the effect of Arch-mediated PV interneuron suppression on sensory responses by comparing the output non-

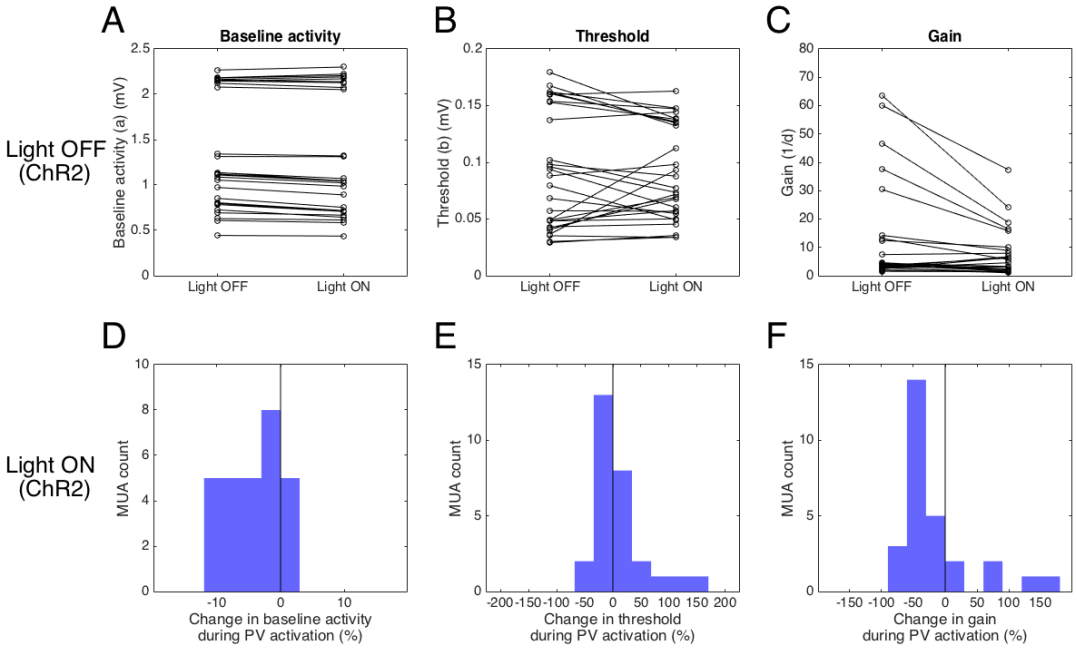

Figure 4.9: A. Effect of ChR2-mediated PV interneuron activation (Light ON) on parameter a, reflecting baseline activity, for each MU. B. Effect of PV interneuron activation (Light ON) on parameter c, reflecting threshold. C. Same as in A and B for parameter 1/d, reflecting gain. MUs with the largest gain values under "light off" conditions showed the largest decrease in gain in absolute terms. D-F. Histograms of % change in these parameters with PV interneuron activation. D. Baseline activity showed a slight reduction during PV interneuron activation. E. PV interneuron PV activation had no significant effect on thresholds. F. This manipulation significantly reduced the gain of MU responses, as predicted. Despite MUs with the largest gain values under "light off" conditions showing the largest decrease in gain in absolute terms (C), the consistent shape of the distribution of gain changes indicates that the relative reduction in gain was consistent across MUs.

linearity parameters under "light off" and "light on" conditions (Figure 4.11A-C). PV suppression had no effect on baseline activity (Median change: 0.008%, sign-rank test: $p=0.724$) (Figure 4.11D), thresholds (Median change: 3.046%, sign-rank test: $p=0.525$) (Figure 4.9E) or gain (Median change: -1.111%, sign-rank test: $p=0.245$) (Figure 4.11F).

The lack of a clear increase in gain across the population was surprising as the majority of MUs appears qualitatively to show a multiplicative disinihibtion, reflecting an increase in gain. The lack of any consistent effect on gain may be attributable to a combination of the minimal nature of the opotogenetic manipula-

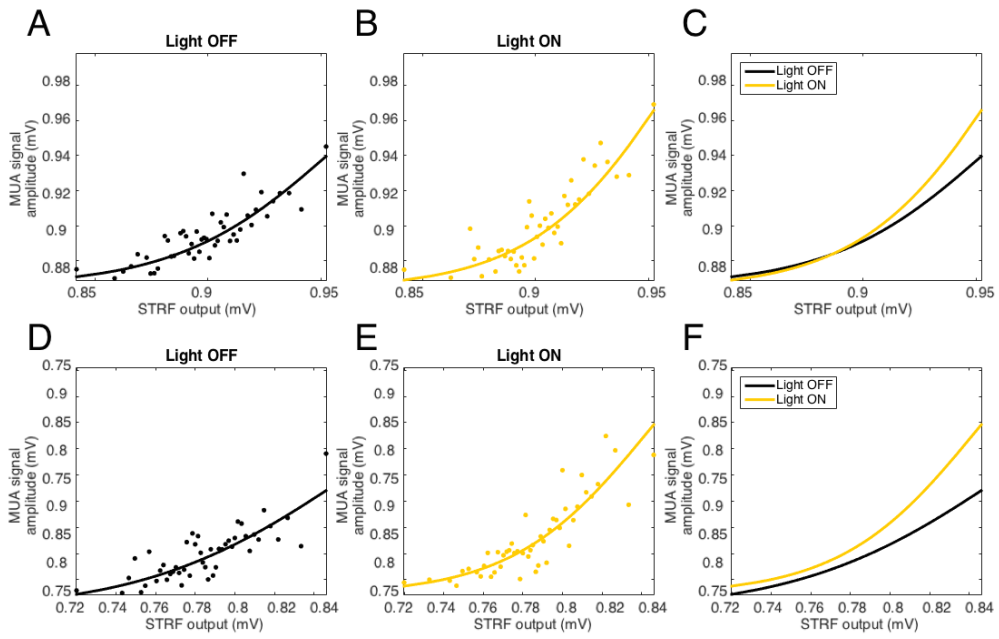

Figure 4.10: As with PV activation, gain changes produced by PV suppression were quantified by fitting a single STRF to the data from each light condition for each MU. An output nonlinearity that could vary between light conditions was fitted to the output of this STRF. A-C: Sigmoidal output nonlinearities fitted to one representative MU under "light off" (black) conditions. B. A nonlinearity fitted to the same MU under "light on" (amber) conditions. C. Overlay of output nonlinearities for this MU The clearest effect of PV interneuron suppression was an increase in the gain of MU responses. D-F. Same a A-C for a second representative MU.

tion and the fact the output nonlinearities can vary in 3 separate parameters. Small changes in any one of these parameters could potentially be modelled by combinations of changes in the remaining parameters, making it difficult to accurately quantify changes in individual parameters using this approach.

In order to test the divisive vs subtractive nature of PV interneuron inhibition more directly I used a more sensitive non-parametric approach. This was possible for this data set as responses to identical stimuli were recorded during "light on" and "light off" conditions, allowing for direct comparisons between MUA recorded under these two conditions. For each MUA, I quantified baseline activity under each condition as the 5th percentile of the response of the MUA (R_{min}) and the maximum response as the 95th percentile (R_{max}). It was possible to measure changes in R_{max}

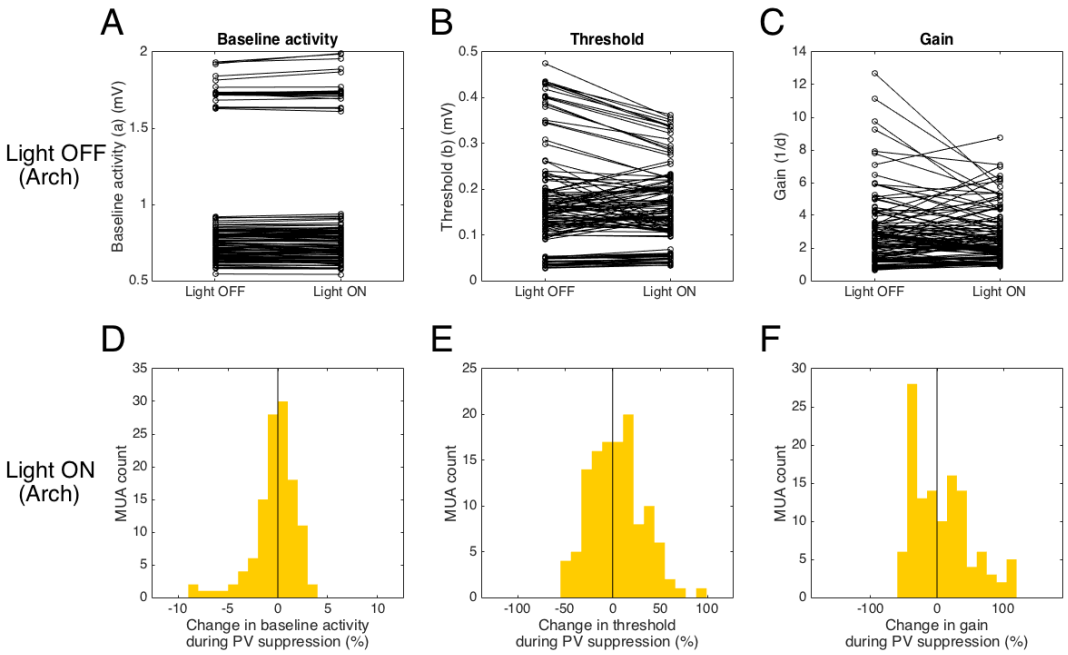

Figure 4.11: A. Effect of Arch-mediated PV suppression (Light ON) on parameter a, reflecting baseline activity, for each MU. B. Effect of PV interneuron suppression (Light ON) on parameter c, reflecting threshold. C. Same as in A and B for parameter 1/d, reflecting gain. D-F. Histograms of % change in these parameters with PV interneuron suppression. D. Baseline activity showed a slight increase during PV interneuron suppression. E. PV interneuron PV suppression had no significant effect on thresholds. F. This manipulation had no consistent effect on the gain of MU responses across the population.

and R_{min} by taking the ratio of each of these quantities under "light on" (R^{on}_{max}) and "light off" (R^{off}_{min}) conditions. I quantified additive vs subtractive changes in the offset of responses along the y-axis (R_{offset}) produced by changes in PV interneuron activity by taking the difference between R_{min} under "light on" (R^{on}_{min}) and "light off" conditions (R^{off}_{min}) and normalising this by the full range of responses:

$$R_{\text{offset}} = \frac{R^{\text{on}}_{\text{min}} - R^{\text{off}}_{\text{min}}}{R_{\text{max}} - R_{\text{min}}} \tag{4.1}$$

I quantified gain as the difference between R_{max} and R_{min} for both contrast and light conditions (high contrast, light on; $R^{on}_{min,high}$, $R^{on}_{max,high}$, low contrast, light on; $R^{on}_{min,low}$, $R^{on}_{max,low}$, high contrast, light off; $R^{off}_{min,high}$, $R^{off}_{max,high}$, low contrast, light off;

$R^{off}_{min,low}$, $R^{off}_{max,low}$):

$$G^{on}_{low} = R^{on}_{max,low} - R^{on}_{min,low} \tag{4.2}$$

$$G^{off}_{low} = R^{off}_{max,low} - R^{off}_{min,low} \tag{4.3}$$

$$G^{on}_{high} = R^{on}_{max,high} - R^{on}_{min,high} \tag{4.4}$$

$$G^{off}_{high} = R^{off}_{max,high} - R^{off}_{min,high} \tag{4.5}$$

I quantified changes in gain produced by optogenetic manipulations as the ratio of "light on" and "light off" gain, for high ($G^{on}_{high}/G^{off}_{high}$) and low contrast ($G^{on}_{low}/G^{off}_{low}$) stimulation separately. In addition to investigating the effects of PV interneuron activity on the gain of sensory responses for a given stimulus, the use of both high and low contrast DRCs permitted the effect of manipulating PV interneuron activity on CGC to be investigated. If PV interneurons are involved in CGC, artificially altering their activity may be expected to affect the magnitude of contrast-dependent gain changes.

The most direct way to test whether these manipulations alter the strength of CGC would be to measure contrast dependent gain changes nonparametrically under "light on" ($G^{on}_{low}/G^{on}_{high}$) and "light off" conditions ($G^{off}_{low}/G^{off}_{high}$). The ratio of these two quantities then provides a measure of the relative strength of gain changes between light conditions ($G_{relative}$), reflecting how the strength of CGC is affected by the optogenetic manipulations.

$$G_{relative} = \frac{G^{on}_{low}/G^{on}_{high}}{G^{off}_{low}/G^{off}_{high}} \tag{4.6}$$

It is not possible to calculate CGC nonparametrically however as the high and low contrast stimuli are not identical. It is possible, and mathematically equivalent, to instead calculate the magnitude of gain changes between "light off" and "light on" conditions for low ($G^{on}_{low}/G^{off}_{low}$) and high ($G^{on}_{high}/G^{off}_{high}$) contrast conditions separately, and then take the ratio of these two quantities. This provides a measure of how the strength of CGC is affected by manipulating PV interneuron activity:

$$G_{\text{relative}} = \frac{G^{\text{on}}_{\text{low}}/G^{\text{on}}_{\text{high}}}{G^{\text{off}}_{\text{low}}/G^{\text{off}}_{\text{high}}} = \frac{G^{\text{on}}_{\text{low}}/G^{\text{off}}_{\text{low}}}{G^{\text{on}}_{\text{high}}/G^{\text{off}}_{\text{high}}} \tag{4.7}$$

For low contrast stimulation, ChR2-mediated PV interneuron activation produced a significant reduction in baseline (R_{min} median; 0.97, sign-rank test; $p<0.001$) (Figure 4.12 A) and maximum activity (R_{max} median; 0.967, sign-rank test; $p=0.013$) (Figure 4.12 B). The comparable size of the reduction in minimum and maximum response resulted in a significant reduction in offset (R_{offset} median; -0.196, sign-rank test; $p<0.001$) (Figure 4.12 C), in addition to a small reduction in gain ($G^{on}_{low}/G^{off}_{low}$ median; 0.958, sign-rank test; $p=0.036$) (Figure 4.12 D). Under high contrast stimulation, similar comparable reductions in baseline (R_{min} median; 0.967, sign-rank test; $p=0.004$) (Figure 4.12 E), and maximum activity (R_{max} median; 0.967, sign-rank test; $p<0.001$) (Figure 4.12 F) were observed. This resulted in the same reduction in offset (R_{offset} median; -0.19, sign-rank test; $p<0.001$) (Figure 4.12 G) and gain ($G^{on}_{high}/G^{off}_{high}$ median; 0.946, sign-rank test; $p<0.001$) (Figure 4.12 H) observed for low contrast stimulation. Despite the similarity of the changes in gain induced by PV activation under low (Figure 4.12 D) and high contrast stimulation, ChR2-mediated PV interneuron activation resulted in an increase in the strength of CGC ($G_{relative}$ median; 1.013, sign-rank test; $p=0.013$) (Figure 4.12 I). This indicates that it's effects were contrast dependent.

For low contrast stimulation, Arch-mediated PV interneuron suppression produced a reduction in baseline activity (R_{min} median; 0.997, sign-rank test; $p=0.006$)

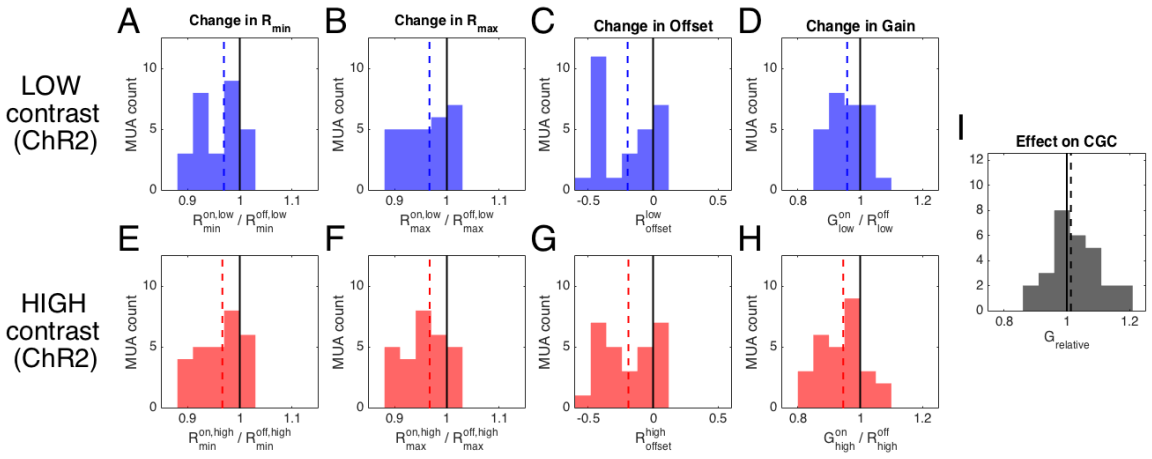

Figure 4.12: A. Reduction in baseline responses (R_{min}) produced by PV interneuron activation during low contrast stimulation. B. A comparable reduction in the maximum response (R_{max}) was also observed. C. This comparable reduction resulted in the effect of PV interneuron activation being primarily captured by a change in offset. D. A reduction in gain was also observed during PV interneuron activation, however. Similar reductions in (R_{min}) (R_{max}), offset and gain were observed during high contrast stimulation (E-H). I. PV interneuron activation increased the strength of contrast gain control.

(Figure 4.13 A) and an increase in maximum activity (R_{max} median; 1.014, sign-rank test; $p<0.001$) (Figure 4.13 B), in keeping with the predicted disinhibitory action of suppressing inhibitory neural activity. This pattern of changes corresponded to a small reduction in offset (R_{offset} median; -0.007, sign-rank test; $p<0.001$) (Figure 4.13 C) and a significant increase in gain ($G^{on}_{low}/G^{off}_{low}$ median; 1.062, sign-rank test; $p<0.001$) (Figure 4.13 D). Under high contrast stimulation, Arch-mediated PV interneuron suppression produced a reduction in baseline activity (R_{min} median; 0.998, sign-rank test; $p=0.046$) (Figure 4.13 E) and a disinhibitory increase in maximum activity, similar to that observed under low contrast stimulation (R_{max} median; 1.012, sign-rank test; $p<0.001$) (Figure 4.13 F). This produced the same small change in offset (R_{offset} median; 0.016, sign-rank test; $p<0.001$) (Figure 4.13 G) and increase in gain ($G^{on}_{high}/G^{off}_{high}$ median; 1.053, sign-rank test; $p<0.001$) (Figure 4.13 H) as observed under low contrast stimulation. As was observed for PV interneuron activation, PV interneuron suppression left the strength of CGC unaf-

fected ($G_{relative}$ median; 1.012, sign-rank test; p=0.146) (Figure 4.13 I). This is in keeping with the similarity of the effects of PV interneuron suppression observed under low (Figure 4.13 A-D) and high (Figure 4.13 E-H) contrast stimulation.

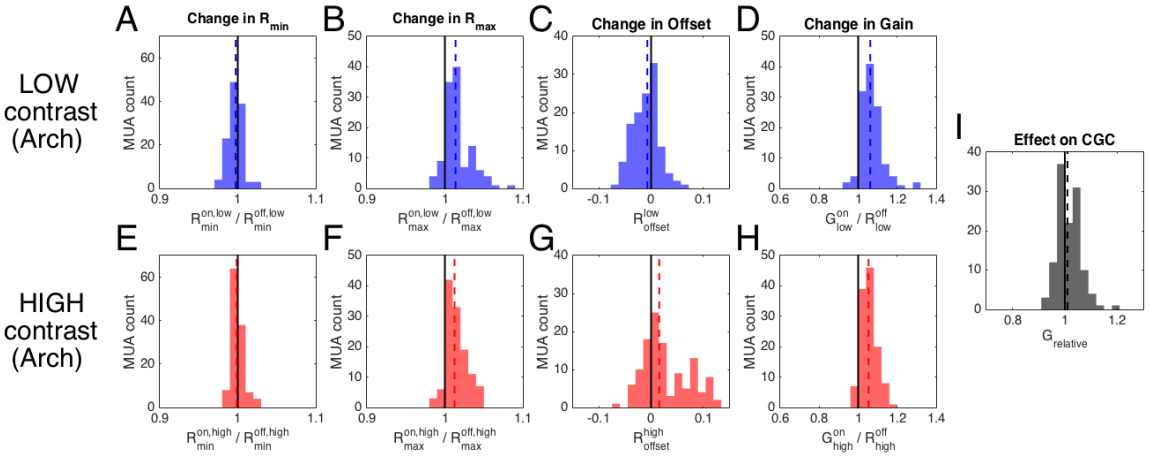

Figure 4.13: A. Negligible reduction in baseline responses (R_{min}) produced by PV interneuron suppression during low contrast stimulation. B. A more substantial increase in the maximum response (R_{max}) was also observed. C. The small change in R_{min} led to a small subtractive change in offset. D. The asymmetrical effects on R_{min} and R_{max} led to an increase in gain during PV interneuron activation, however. E. A similar reduction in (R_{min}) was observed under high contrast stimulation. F. A similar increase in (R_{max}) was also observed under high contrast stimulation. G. Unlike under low contrast stimulation, PV interneuron suppression produced a small increase in offset during high contrast stimulation. H. The robust increase in gain observed under low contrast stimulation was also observed during high contrast stimulation. I. The similarity of the effects on gain under low and high contrast stimulation resulted in this manipulation leaving the strength of contrast gain control unaffected.

4.2.4 PV interneuron activity is not modulated by stimulus contrast

I next examined whether PV interneuron activity increases during high contrast stimulation. In order to do this I performed recordings in AC using Buzsaki style multi-electrodes in 6 mice (8-12 weeks old) expressing ChR2 in PV interneurons. This allowed me to isolate single units and to identify putative PV interneurons by examining the extent to which the activity of single units was modulated by expo-

sure to blue light (Figure 4.14A). The latency of the first spike following light onset is often used for photo-identification of neurons. This approach is compatible with tungsten electrode recordings but is not compatible with silicon probes recordings, as used here, as large photoelectric artifacts are produced by light onset. These artifacts prevent spike extraction for a period of 10s of ms following light onset. This precluded the use of spike-onset latency as a criterion for PV interneuron identification and so different selection criteria were employed here.

Putative PV interneurons were required to show a significant increase in their firing rate under "light on" conditions, as assessed by a paired sample t-test, and to have spikewidths of less than 0.25ms (Cardin et al., 2009; Letzkus et al., 2011; Pi et al., 2013). This second criterion was included in order to exclude pyramidal cells that may be recruited via disinhibition during light stimulation. Of the 136 single units recorded, 34 passed these criteria and were classified as putative PV interneurons (Figure 4.14B). The remaining 102 putative non-PV units were not included in following analysis (Figure 4.14C).

I played both high and low contrast DRC stimuli to anaesthetised mice under "light off" conditions in order to investigate whether PV interneurons respond more strongly during high contrast stimulation. I assessed the contrast response of putative PV neurons by taking the ratio of firing rates during high contrast over low contrast conditions. As a population, putative PV units did not show a significant increase in firing rate with increased contrast (high/low contrast median: 1.18, sign-rank test: $p=0.14$) (Figure 4.15). The significance of changes in firing rate with contrast was assessed on a unit-by-unit basis using paired t-tests. Twelve of the 36 units showed a significant increase in firing rate with contrast, with two of these units only responding during high contrast stimulation(Figure 4.15). Eight units showed a significant decrease in firing rates during high contrast stimulation (Figure 4.15). These results demonstrate that, as a population, PV interneurons in

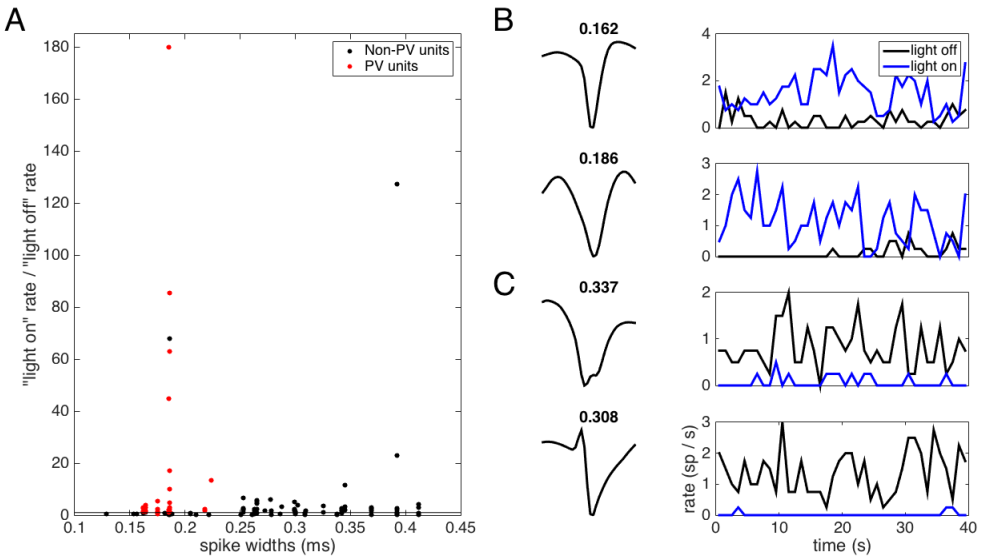

Figure 4.14: A. Scatter plot of spike widths (abscissa) against the strength of the optogenetic response (ordinate), measured as the ratio of firing rates during "light on" over "light off" conditions. In order to be classified as a putative PV neuron (red), units were required to show a significant increase in firing rate during light stimulation and to have spikewidths of <0.25 ms. B. Spike waveforms and PSTHs of spiking responses of putative PV units to high contrast DRC stimuli with "light off" (black) and "light on" (blue). Spike widths are shown above the waveforms. C. Examples of putative non-PV units that did not meet the selection criteria.

AC are not driven by stimulus contrast.

4.3 Discussion

I set out to test the hypothesis that PV interneuron activity is capable of reducing the gain of sensory evoked responses in AC and that such activity is recruited during high contrast stimulation, in order to mediate the effects of contrast gain control (CGC).

I first investigated whether PV interneurons are capable of modulating the gain of sensory evoked responses in AC, as they have found to be in V1 (Atallah et al., 2012; Wilson et al., 2012). I did this by optogenetically increasing or decreasing their activity during sensory stimulation while I recorded MUA from the AC of

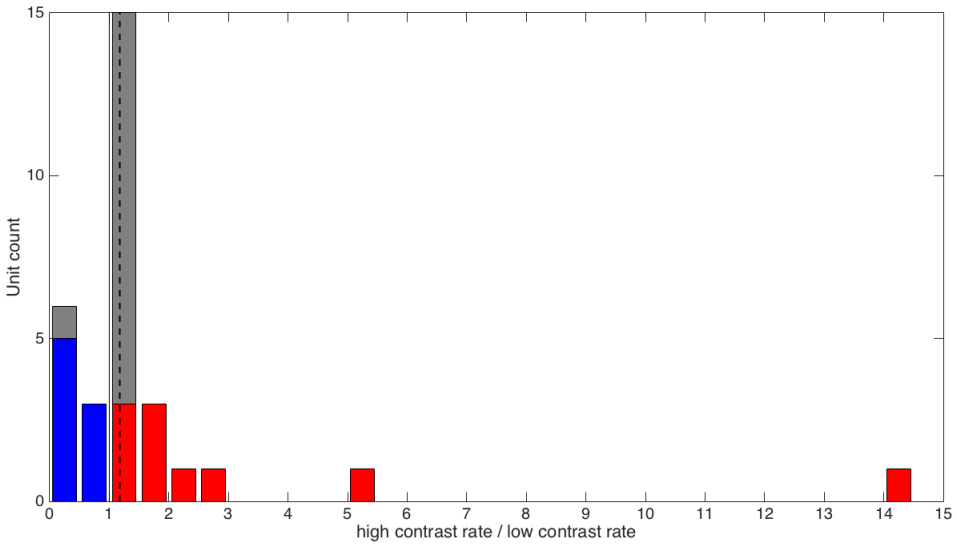

Figure 4.15: Histogram of the change in firing rates of putative PV units in response to a doubling of stimulus contrast. Units that showed a significant increase in firing rate are shown in red and units that showed a significant decrease are shown in blue. No systematic effect of contrast on firing rates was observed.

anaesthetised mice. If these neurons only control gain, the tuning properties of AC neurons should remain unchanged. For PV interneuron activation and suppression, this was found to be the case as no changes in spectral or temporal tuning bandwidths of STRFs were observed. A small but significant change in the distribution of best frequencies across the population during PV suppression was observed however. The main effect on the STRFs of altering PV activity was to change the largest STRF coefficient. PV interneuron activation was found to reduce the largest coefficient while PV interneuron suppression increased it. Taken with the finding that spectral tuning bandwidths remained unchanged for both of these manipulations, these findings indicate that PV interneuron activity controls the gain of MU responses in mouse AC.

In order to test the significance of any unobserved changes in tuning, I used STRFs fitted under control conditions in order to predict responses to stimuli presented during PV activation or suppression. STRFs predicted spiking responses

during optogenetic manipulations as well as they predicted responses during the control condition, indicating that STRF tuning is not significantly altered by optogenetic PV interneuron activation or suppression.

The similarity of STRFs under these two conditions allowed me to fit a single STRF to responses from both of these conditions and to model the effects of PV interneuron activity on gain, baseline firing and threshold. I did this by fitting a nonlinearity to the output of the STRF that was allowed to vary in these three parameters between control and optogenetic conditions.

Activation of PV interneurons in mouse AC with ChR2 has recently been found to produce a mixture of divisive and subtractive effects on sensory evoked responses (Seybold et al., 2015). In keeping with this finding I found that while PV interneuron activity reduces the gain, it also produces a large reduction in offset. Previous conflicting results regarding the subtractive or divisive nature of PV interneuron activity in V1 has been found to be attributable to unnatural levels of interneuron activation artificially suppressing responses below threshold, changing divisive changes in gain into subtractive changes (Atallah et al., 2012; Lee et al., 2012; Wilson et al., 2012; Atallah et al., 2014; Lee et al., 2014). PV interneuron activation here was extremely mild however, and so should not have had such an effect. This raises the possibility that the effects of PV interneuron activation may differ between these cortical areas. Another explanation is that firing rates in AC in response to DRCs are typically much lower than V1 responses to drifting grating stimuli, typically used in these experiments. Smaller levels of suppression may therefore be required to produce subtractive effects due to interaction with spike threshold when studying sensory cortex during a low firing rate regime. If this is the case, PV interneuron activation should have a subtractive effect in V1 during stimulation with natural scenes, which typically result in lower firing rates compared with gratings (Pecka et al., 2014).

Seybold et al. (2015) argue that the true effects of inhibition provided by distinct subpopulations could be masked when these neurons were activated artificially, due to recurrent connectivity of cortical networks and to interactions with spike threshold. In keeping with this, optogenetic activation is thought to introduce highly artificial synchronous activity into the activity patterns of the stimulated population, which may significantly disrupt normal network operation. Controlled optogenetic suppression may disrupt typical network operation less by leaving the statistical structure of spiking activity largely unchanged, as it does not introduce new spikes with unnatural correlations. Based on these considerations, I also tested whether PV interneuron activity modulates the gain of sensory evoked responses by suppressing PV interneuron activity as mildly as experimentally possible. The used of an LED to deliver light stimulation instead of a laser, as was used in some V1 studies (Lee et al., 2012; Atallah et al., 2014), made weak optogenetic manipulations possible. This is due to the lower power output of LEDs and allowed for more physiologically plausible changes in PV interneuron activity to be achieved. Arch-mediated suppression of PV interneuron activity resulted in a significant increase in gain in addition to a small subtractive change in offset. Furthermore, ChR2-mediated activation was found to increase the strength of CGC. These findings indicate that PV interneuron activity is capable of modulating the gain of sensory evoked responses in AC.

I used a measure of MUA to quantify sensory evoked responses in the present study, whereas previous studies have often used single unit recordings when investigating the role of inhibition in sensory processing (Atallah et al., 2012; Wilson et al., 2012; Lee et al., 2012, 2014; Seybold et al., 2015). The finding of Seybold et al. (2015) that optogenetic manipulations can have dramatic effects on network activity indicates that using a measure of population activity, such as MUA, may result in heterogeneous effects of such a manipulation being masked. An inves-

tigation of the effects of PV interneuron inhibition on single-unit activity will be necessary in order to achieve a more detailed understanding of the role of this class of interneuron in AC.

A fibre-optic with relatively large diameter (0.55 nm) was used to ensure that optotgenetic manipulations were consistent across AC, so that AC function was not disrupted by spatially uneven local disruption of interneuron activity. The fibre-optic that delivered LED illumination was designed in order to deliver a relatively consistent power output across the width of the fibre, making spatially broad optogenetic manipulations possible. In the mouse, however, it is not possible to accurately localise primary auditory cortex (A1) without functional mapping of tuning curves across the whole of AC. This was not possible here and, as a result, optogenetic manipulations may have activated secondary auditory cortical areas as well as A1. Craniotomies were very limited in their spatial extent however which should have prevented inadvertent manipulations of PV interneurons in cortical areas surrounding AC. V1 is easily targeted in the mouse and previous studies have often used local optogenetic manipulations (Wilson et al., 2012). These experimental differences may account for some of the differences between findings in V1 and those reported here.

Given the ability of PV interneurons to modulate the gain of sensory responses in AC, I next investigated whether they are suitably recruited during stimulation with high contrast stimuli, in a manner that would allow them to mediate CGC in AC. I did this by performing extracellular single unit recordings in the AC of anaesthetised mice during acoustic stimulation with DRCs of different contrasts. I used virally-mediated expression of ChR2, targeted to PV interneurons, in order to optically tag PV interneurons.

Unlike in V1 (Atallah et al., 2012) 65% of PV interneurons in AC did not significantly increase their activity during high contrast stimulation. This indicates that

the contribution of PV interneurons to this computation may vary between sensory areas. This dataset does not definitively demonstrate that PV interneurons in AC are not driven by stimulus contrast however. Recordings were approximately targeted to deep cortical layers, based on recording depth, as neurons here are larger and more densely packed (Sakata and Harris, 2009). This was done in order to maximise the number of units recorded. These deep layers are considered the output layers of the cortex and so it is entirely plausible that PV interneurons in layers 4 and 2/3 are driven by stimulus contrast and may implement CGC before the deep layers, where it may simply be inherited by excitatory neurons. The finding that PV interneurons in V1 are driven by stimulus contrast was based on results from superficial cortical layers (Atallah et al., 2012), limiting the ability to perform a direct comparison with these results. Recordings from PV interneurons in layers 4 and 2/3 during stimulation with stimuli of varying contrasts will be necessary in order to confirm whether PV interneurons in AC are driven by stimulus contrast.

PV interneurons in AC appear to be capable of modulating the gain of sensory evoked responses. The PV interneurons recorded here, however, did not increase their activity in response to increasing stimulus contrast. If this is the case for all subpopulations of PV interneurons in AC, PV interneuron inhibition cannot contribute to CGC in AC, but may instead be specialised for other forms of gain control.

Chapter 5

Biophysical basis of CGC

5.1 Introduction

What biophysical mechanisms might contribute to contrast gain control (CGC) in auditory cortex (AC)? Multiple mechanisms for gain control have been identified both between different species and systems but also within specific cortical areas. In primary visual cortex (V1), gain control has been attributed to a combination of synaptic depression (Carandini et al., 2002), shunting inhibition (Carandini and Heeger, 1994; Carandini et al., 1997) and fluctuations in membrane potential (Finn et al., 2007). Each of these mechanisms may also contribute to CGC in AC.

The results from the previous chapter raise the possibility that PV interneuron inhibition may not contribute to CGC in AC as it has been found to in V1 (Atallah et al., 2012; Wilson et al., 2012). If this is the case, non-inhibitory mechanisms of gain control may contribute significantly to CGC. One such candidate mechanism is contrast-dependent fluctuations in membrane potential. With no synaptic noise present, the relationship between membrane potential and the firing rate of cortical pyramidal cells is captured by a linear threshold function (Figure 5.1A). *In vivo*, membrane potential variability smooths this relationship so that it approximates

a power law function (Miller and Troyer, 2002) and increases gain by increasing the probability of spiking for a given mean membrane potential. In V1, reductions in this variability with increasing contrast have been shown to reduce the gain of cortical responses, contributing to CGC in V1 (Troyer et al., 1998; Hansel and van Vreeswijk, 2002; Carandini, 2004; Ringach and Malone, 2007; Finn et al., 2007) (Figure 5.1B). It is unknown whether a similar mechanism contributes to CGC in AC.

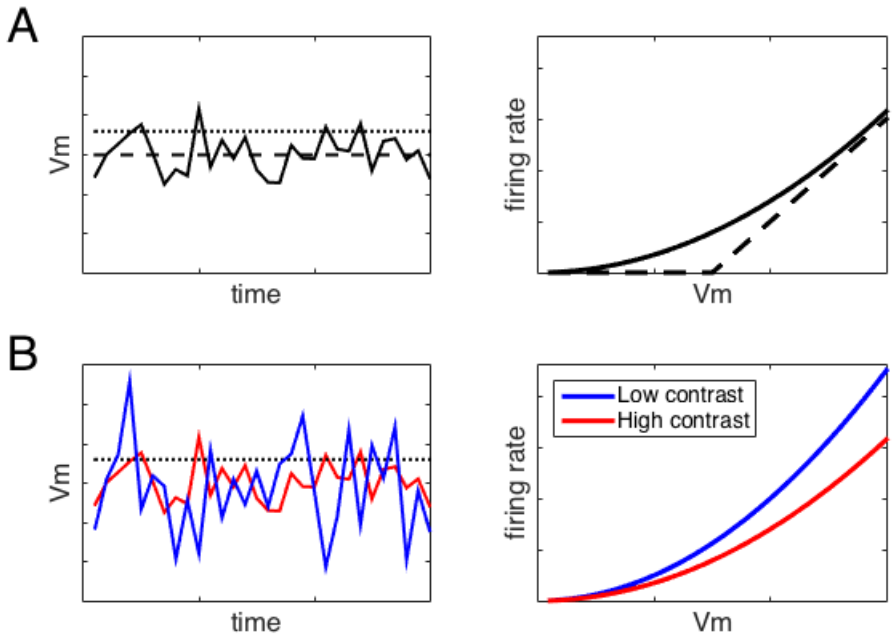

Figure 5.1: A. Under conditions of no synaptic noise (left panel, dashed line) the relationship between membrane potential and firing rate in neurons is a linear threshold function (right panel, dashed line). Membrane potential variability that is present under in vivo conditions transforms this relationship into a power law function (right panel, solid line). This occurs as this variability boosts the membrane potential above threshold (left panel, dotted line) on some, but not all trials, changing the relationship between mean membrane potential and mean firing rate into a probabilistic relationship. B. Increasing membrane potential variability during low contrast stimulation (left panel, blue trace) should result in an increase in the gain of this relationship (right panel, blue curve). Gain would increase as a result of the increased probability of spiking responses being observed for a given mean membrane potential, due to larger variability boosting the membrane potential above threshold (left panel, dotted line) more often.

Inhibitory mechanisms may still contribute to CGC in AC however. Shunting in-

hibition in particular offers an elegant way of controlling neuronal gain through the modulation of input conductance (Borg-Graham et al., 1998a). This has a scaling effect on the amplitude of post-synaptic potentials (PSPs), with increases in conductance reducing PSP size. In keeping with shunting inhibition contributing to gain control in V1, the input conductance of these neurons has been found to increase by up to 300% during stimulation with gratings of 100% contrast (Borg-Graham et al., 1998b; Anderson et al., 2000; Monier et al., 2003; Frégnac et al., 2003). While both excitatory and inhibitory conductances contribute to changes in the global input conductance of the neuron, estimates of conductance based on somatic intracellular recordings are biased to detect large GABAa-mediated somatic changes in conductance as opposed to smaller distal excitatory conductance changes that primarily occur on dendrites (Williams and Mitchell, 2008). Furthermore, theoretical studies have shown that changes in global conductance at the soma must be on the order of a 100-200% increase compared with no sensory stimulation in order to effectively scale membrane potential responses and control gain (Koch et al., 1990). Furthermore, as excitatory and inhibitory conductances both increase global input conductance as opposed to cancelling out, it is not necessary to separate these components when investigating whether inhibitory shunting conductance changes between stimulus conditions. Changes in global input conductance have therefore been used extensively in order to detect the presence of shunting inhibition in cortical neurons (Borg-Graham et al., 1998b; Anderson et al., 2000; Monier et al., 2003; Frégnac et al., 2003). If shunting inhibition contributes to CGC in AC, an increase in the input conductance of AC neurons should be observed during high contrast stimulation. This would scale down the amplitude of sensory evoked membrane potential responses and would result in a reduction in gain.

Here, I address the question of whether shunting inhibition and fluctuations in membrane potential contribute mechanistically to CGC in A1. First, I investi-

gate whether PSP amplitude is modulated by stimulus contrast, as this effect is not predicted by the membrane potential variance hypothesis but is predicted by the shunting inhibition hypothesis. Next, I asked whether membrane potential variance is reduced during high contrast stimulation, as it is in V1, and, if so, whether this results in gain changes in the relationship between membrane potential and firing rate in these neurons. Finally, I estimated the input conductance of AC pyramidal cells during low and high contrast auditory stimulation in order to investigate the involvement of shunting inhibition in CGC in AC.

5.2 Methods

5.2.1 Intracellular recordings

Blind whole-cell patch recordings were made from the AC of anaesthetised C57BL/6 Mice, between the ages of 4 and 10 weeks old. These animals were not used in any other experiments presented in this thesis. For some intracellular recordings, a durotomy was performed before recording took place. Electrodes were oriented 45 degrees to the cortical surface and were advanced to depths of between 300 and 700 μm in order to record from the middle layers of the cortex. Conventional, low-resistance (4—7 MΩ) patch pipettes were made from glass (World Precision Instruments, Sarasofa, Florida, USA). The intracellular solution consisted of (in mM): 110 K-gluconate, 40 HEPES, 2 ATP-Mg, 0.3 GTP, 4 NaCl and 4mg/ml biocytin, pH 7.2-7.3. A Multiclamp 200b amplifier (Molecular Devices, CA, USA) was used in order to amplify intracellular recordings. As with extracellular recordings, a TDT RZ2 digital signal processor was used to digitize the data. For intracellular recordings, I performed the surgery while Martin Kahn prepared the pipettes and established the patch recordings.

Unlike extracellular recordings, which can be stable over many hours, whole

cell recordings can typically only be maintained for tens of minutes. This limitation prevented me from collecting enough data to fit receptive field models as has been done with the extracellular data. As a result, the spectrotemporal tuning properties of the neurons that were recorded from remained unknown. In order to probe response properties in neurons with unknown tuning properties, I embedded 50ms bursts of frozen noise halfway through each DRC (Figure 5.2). This enabled me to drive responses in all neurons, irrespective of their individual tuning properties. A similar approach has successfully been used previously in order to assess the time course of CGC in ferret AC (Rabinowitz et al., 2011a). Stimuli consisted of a 1 second DRC that was repeated 11 times. The contrast varied from low (σ_L ~6.2 dB, c = σ_P/μP = 0.68) to high (σ_L ~11.97 dB, c = σ_P/μ_P = 1.2) with each repetition while the spectrotemporal pattern of the DRC remained the same. The 50 ms noise bursts occurred 500 ms into each DRC. The two 25 ms chords that would otherwise have occurred at this time were omitted during the noise burst. Noise bursts were spectrally frozen so that each burst in a stimulus was identical. DRCs and noise bursts were constructed using 4 different tokens in order to generate the random numbers that determined pure tone levels for the DRCs and the spectral structure of the noise burst. This allowed me to ensure that our results generalised across stimuli and are not specific to the particular pattern present in a particular DRC or noise burst. The first low contrast DRC was omitted from the analysis in order to avoid onset responses.

5.2.2 Access resistance compensation

I attempted to measure the input conductance, *G*, of recorded AC neurons by injecting different levels of current and measuring the effect on membrane potential, V_m. The extent to which V_m changes in response to injected current,I^{inj}, depends on *G*, as is captured in Ohm's law:

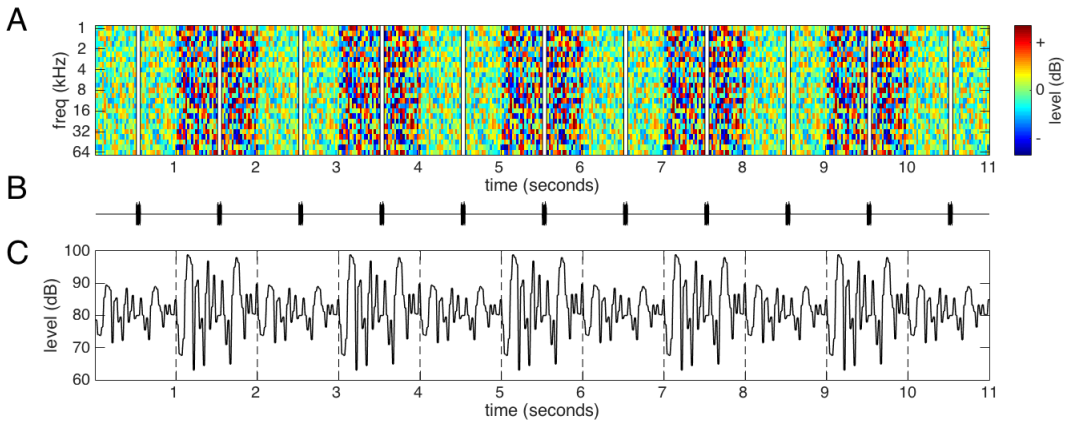

Figure 5.2: A. For whole cell recordings, I created DRC stimuli consisting of one second of DRC that repeated 11 times, changing between two contrast levels with every repetition. B. A 50 ms frozen noise burst was embedded 500 ms into each DRC. This noise stimulus allowed me to investigate evoked responses under different contrast conditions without requiring me to know the tuning of each neuron. C. Temporal cross section of the DRC showing the variation in sound level at 1kHz.

$$V_{\mathrm{m}} = \frac{I^{\mathrm{inj}}}{G} \tag{5.1}$$

This allowed me to estimate G from the recorded change in V_m in response to different levels of I^{inj}. The inverse of G is the input resistance of the neuron R^{input}. The resistance of the electrode after the patch is formed, also known as the series of access resistance and here referred to as R^{access}, also contributes to the relationship between V_m and I^{inj}:

$$V_m = I^{inj} * (R^{\mathrm{access}} + R^{\mathrm{input}}) \tag{5.2}$$

In order to accurately estimate stimulus induced changes in G over time, it was first necessary to compensate for R^{access} in our recordings. Before each stimulus presentation current with a square waveform (20Hz, 20 cycles) was injected into each cell. The recorded change in membrane potential in response to I^{inj} consists of two components, a fast change due to R^{access} and a slow exponential change due

to R^{input} (Anderson et al., 2000). In order to separate these two components, I fit a double exponential model to the mean recorded change in V_m over time, *V(t)*:

$$V(t)/I^{\text{inj}} = R^{\text{access}}[1 - exp(-t/\tau^{\text{access}})] + R^{\text{input}}[1 - exp(-t/\tau^{\text{input}})] \quad (5.3)$$

where τ^{input} corresponds to the time constant of the neuron and τ^{access} corresponds to the time constant of the patch electrode. The mean response to I^{inj} was calculated from V_m responses that did not include an "up state". Up states are periods of spontaneous depolarisation (Destexhe et al., 2003) that can mask the effect of I^{inj} on V_m. As this results in a bimodal distribution of V_m values, up states could be detected by threshold crossing where the threshold is the mean V_m recorded during both up and down states (Figure 5.3A). Sweeps including up states were excluded from further analysis. The model was fit to the mean V_m response (Figure 5.3B). A clear minimum in the mean squared error of the fit for different values of electrode resistance and time constant was consistently observed (Figure 5.3C). Subtracting the voltage component corresponding to R^{access} allowed me to estimate the true V_m response to I^{inj}.

In order to validate this method of R^{access} compensation, I examined the effect of current injection on spike threshold both before and after compensation for R^{access}. In keeping with previous studies (Anderson et al., 2000) I assumed that spike threshold would not change dramatically with I^{inj}. Any change in spike threshold with different levels of I^{inj} can therefore be attributed to changes in V_m resulting from R^{access} (Figure 5.3C). I detected spikes by crossing of a threshold positioned 7 standard deviations of the V_m distribution above the mean V_m. Spike threshold was detected by finding the inflection point in the spike waveform within 1 ms prior to threshold crossing. After identifying the mean spike threshold for each I^{inj}, I fit a

linear model, with a slope that corresponded to R^{access} that had not been successfully compensated for (Figure 5.3D). The slope of this fit was consistently reduced to near 0 following R^{access} compensation.

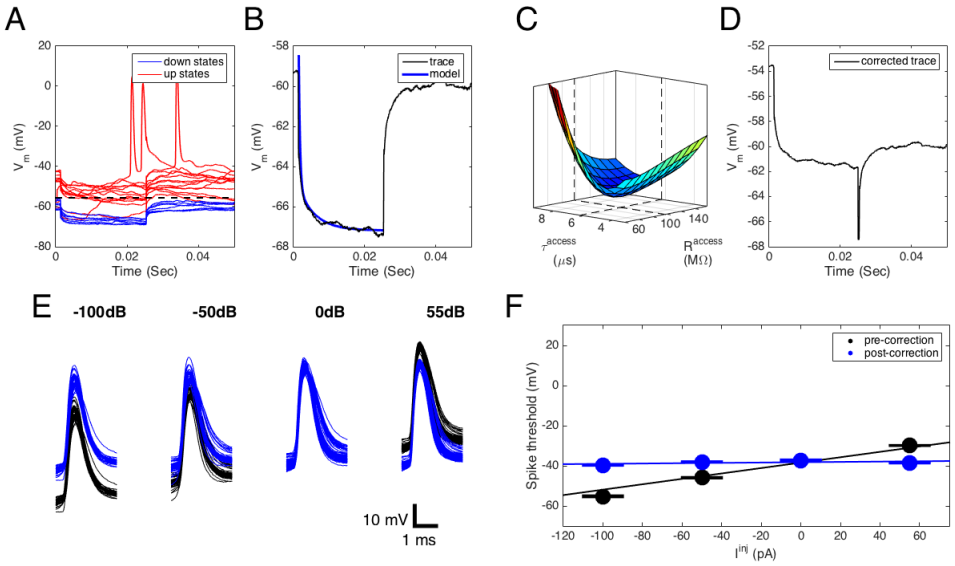

Figure 5.3: R^{access} compensation. A. I injected a series of 40pA current pulses into each neuron before each sweep. This allowed me to estimate both R^{access} and R^{input} from the change in V_m. Under anaesthesia, cortical neurons undergo periods of spontaneous depolarisation, known as up states. During these periods, the impact of I^{inj} on V_m is difficult to quantify. Cycles of current injection that included an up state (red traces) were identified by crossing a threshold corresponding to the mean of the V_m distribution. These traces were excluded and the remaining traces (blue) were used to compute the mean V_m response to the current pulses. B. The mean V_m (blue trace) consists of fast and slow exponential components that correspond to R^{access} and R^{input}, respectively. A double exponential model was fit to the mean V_m response in order to quantify the contribution of R^{access} so that it could subsequently be compensated for. C. Each fit showed a clear minimum in the error function for a range of resistance, R^{access}, and time constant, τ^{access}, values. D. Estimate of the true V_m in response to I^{inj} following compensation for R^{access}. E. An estimate of the remaining R^{access} following compensation was also obtained for each neuron by quantifying the change in spike threshold under different levels of I^{inj}. Spikes before (black) and after (blue) correction for R^{access} using these two methods. F. A linear fit (black line) to mean spike threshold for different levels of I^{inj} (black circles, standard errors indicated by black lines) can be used to quantify R^{access}. Error bars appear as a single line due to the variation in spike threshold between spikes being minimal. After correction for R^{access}, spike threshold is impacted far less by I^{inj} (blue), demonstrating that R^{access} has successfully been compensated for.

5.3 Results

5.3.1 Membrane potential responses undergo contrast gain control

If contrast-dependent membrane potential variance alone underlies CGC at the level of spiking responses in AC, the amplitude of subthreshold V_m responses should remain unchanged between contrast conditions. If inhibition contributes to CGC of spiking responses, subthreshold V_m responses should also be modulated by stimulus contrast. Current clamp recordings were performed from 17 pyramidal cells in the AC of anesthestised mice. No net current was injected into the neuron (0pA holding current), allowing V_m responses to DRCs of different contrasts to be recorded (Figure 5.4A-F). This approach allowed the effect of stimulus contrast on evoked V_m responses in AC neurons to be examined.

I first examined whether the mean V_m responses to DRCs changed with stimulus contrast. I found a small but significant depolarisation during high contrast compared to low contrast stimulation (low contrast V_m: -52.62 mV, high contrast V_m: -52.39 mV; sign-rank test: p=0.009) (Figure 5.4E). I next examined whether noise-evoked PSP amplitudes were modulated by DRC contrast. Noise bursts lasting 50 ms were embedded 500 ms into a 1 second DRC of either low or high contrast. I found that noise-evoked PSP amplitude was suppressed under high contrast stimulation (low contrast: 4.67 mV, high contrast: 3.77 mV; sign-rank test: p=0.042) (Figure 5.4F).

In order to investigate whether this suppression reflects divisive gain control, I varied the level of the noise probe stimuli between 70 and 90 dB in 5 dB steps and recorded the amplitude of PSPs evoked by these stimuli when embedded in a high or low contrast DRC (Figure 5.5A,B). If the PSP suppression observed during high contrast stimulation reflects divisive gain control, PSP amplitude should be

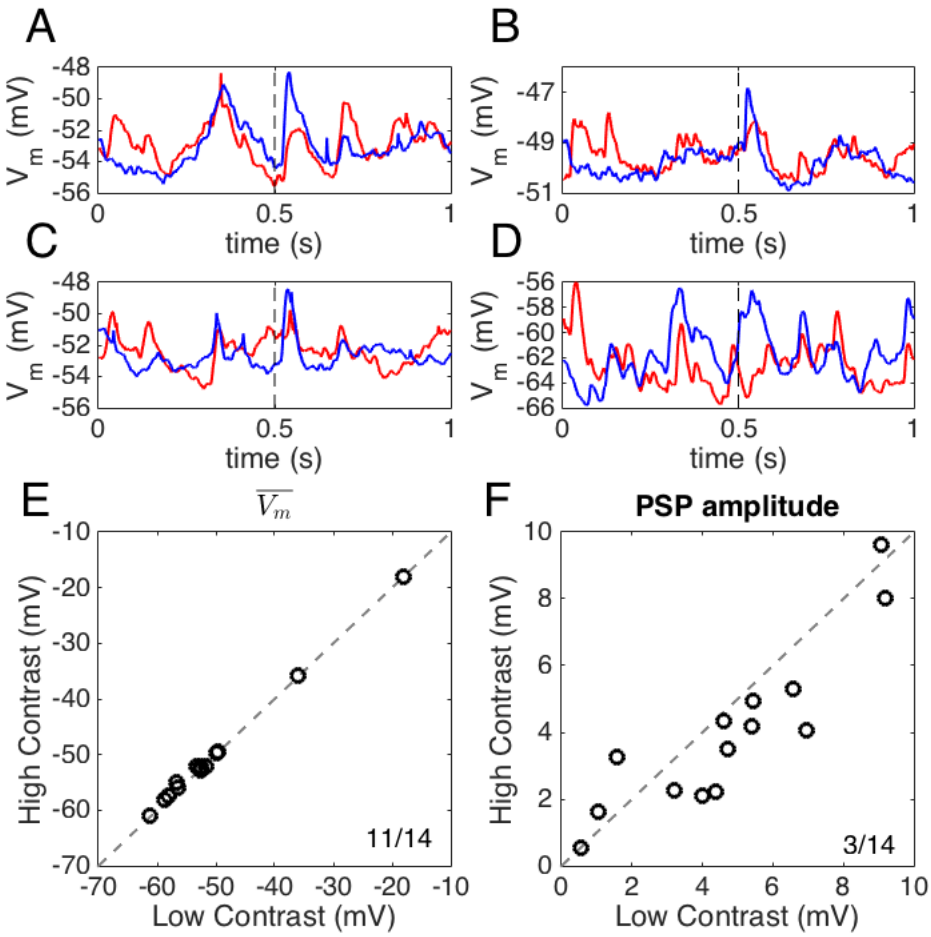

Figure 5.4: A-D. mean V_m responses to 1 ms high contrast (red) and low contrast (blue) DRCs for four different AC neurons. Broken lines indicate the onset of a 50 ms 80 dB noise burst. E. Mean V_m is not significantly modulated by stimulus contrast. F. The amplitude of the noise evoked PSP was increased under low contrast stimulation, relative to high contrast stimulation. Ratios in the corner of plots indicate the number of neurons above the identity line, y=x, over the total number of neurons.

scaled down with increasing stimulus level. I fit a linear model to the relationship between noise probe level and PSP amplitude for each recorded neuron in order to quantify the extent to which this suppression was divisive or subtractive (Figure 5.5C). Six out of 7 neurons showed a significant divisive reduction in the slope of this relationship in response to increased stimulus contrast (Figure 5.5D). Recorded neurons did not show any systematic effect on x-axis intercept with stimulus contrast (Figure 5.5E). CGC therefore appears to exist at the level of the membrane

potential response as well as in the spiking response of AC neurons.

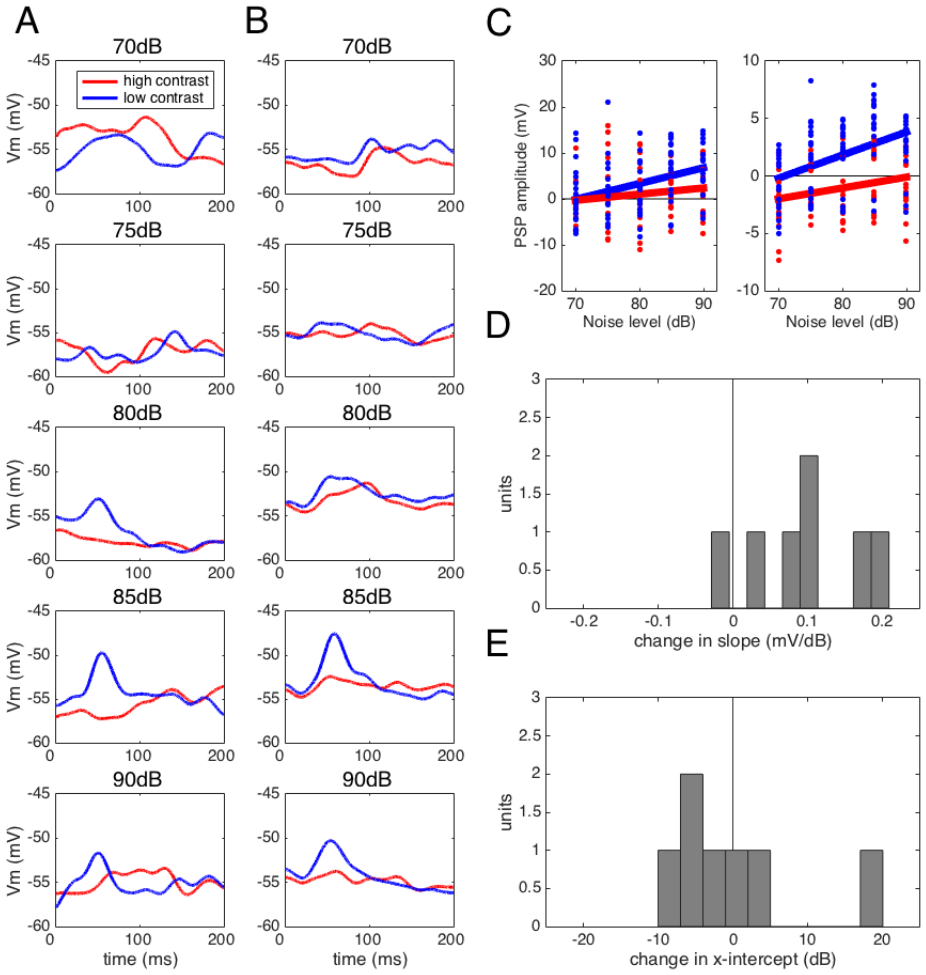

Figure 5.5: In order to investigate whether the observed contrast dependent suppression of PSP amplitude was predominantly suppressive or divisive in nature, I varied the amplitude of noise burst stimuli between 70 and 90 dB in 5 dB steps. Noise burst stimuli were embedded in high and low contrast DRCs. In all plots blue indicates low contrast while red indicates high contrast. A. PSPs recorded in response to noise probe stimuli of different sound levels for a single representative neuron. B. Same as in A for a second neuron. C. Linear fits to PSP amplitudes recorded in response to noise probe stimuli of different levels for the same neurons in A and B. Dots indicate mean PSP amplitude for a given stimulus token. D. Histogram of slope coefficients of the linear fit for each recorded neuron. The majority of neurons showed a reduction in slope indicative of a gain change. E. Histogram of x-intercepts of the linear fit for each recorded neuron. The neurons did not show a systematic change in x-intercept with contrast.

5.3.2 Membrane potential variability in auditory cortex is not contrast dependent

Trial-to-trial V_m variability alters gain by changing the probability that a mean value of V_m will result in a spiking response but leaves PSP amplitude unaffected (Troyer et al., 1998; Hansel and van Vreeswijk, 2002; Carandini, 2004; Ringach and Malone, 2007; Finn et al., 2007). This mechanism therefore cannot account for the observed reduction in PSP amplitude under high contrast stimulation but may still contribute to CGC at the level of spiking responses. If this is the case, responses should be more variable under low contrast stimulation, resulting in an increase in the gain of this relationship. I fitted a power law function to the mean V_m, $\overline{V_m}$, and mean spiking response across trials, $\overline{S}$, in order to investigate whether the gain of this relationship is indeed increased under low contrast stimulation.

$$\overline{S} = k[\overline{V_{\mathrm{m}}} - V_{\mathrm{rest}}]^{\mathrm{p}}{}_{+} \tag{5.4}$$

V_{rest} refers to the resting membrane potential of the neuron and the subscript "+" indicates half-wave rectification. The exponent p was fixed between contrast conditions but the gain factor k was allowed to vary. Some neurons showed an increase in gain under high contrast stimulation (Figure 5.6A) while others showed no change (Figure 5.6B). No neurons showed a reduction in the gain of this relationship under high contrast stimulation as would be predicted if changes in membrane potential variance contributed to contrast gain control (Figure 5.6C). As fluctuations in V_m are thought to underlie changes in this relationship, this finding indicates that V_m variability is not contrast dependent. I addressed this issue by quantifying both the variance of V_m responses to high and low contrast DRCs and the variability in noise-evoked PSP peak amplitude across trials. As predicted, no contrast dependent changes in the standard deviation of V_m, σ_{V_m} were observed (low contrast: 4.34 mV,

high contrast: 4.19 mV; sign-rank test: p=0.426) (Figure 5.6E). No contrast dependent changes were observed in the standard deviation of peak response amplitude to the embedded noise stimuli either (low contrast: 2.96 mV, high contrast: 2.91 mV; sign-rank test: p=0.426) (Figure 5.6F).

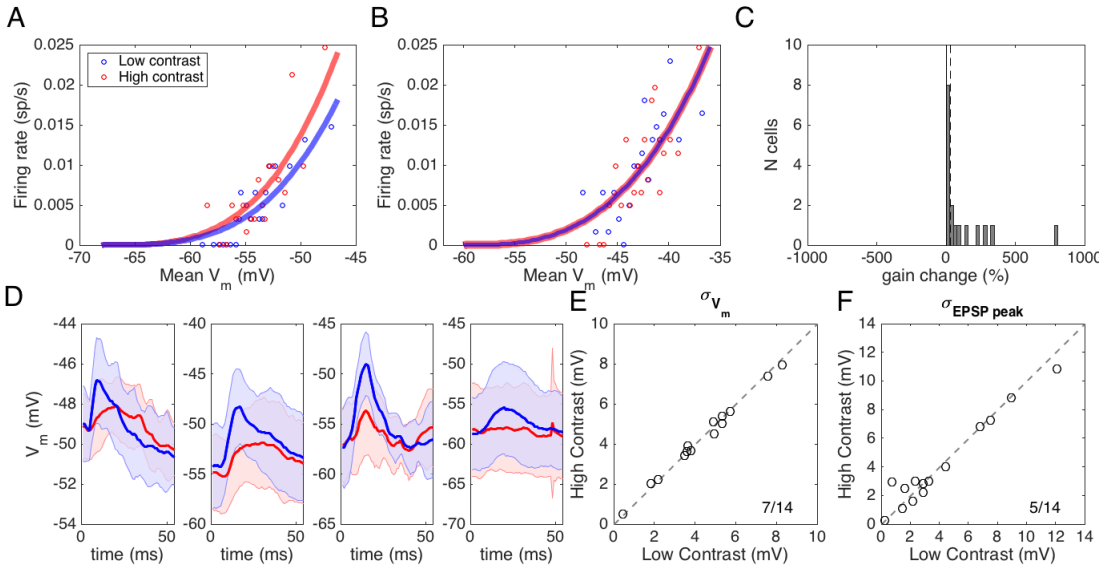

Figure 5.6: A. A power law function was used to map V_m to firing rate for each unit. Some units, such as this one, showed an increase in the gain of this relationship under high contrast stimulation. B. Other units showed no change in this relationship with stimulus contrast. C. Histogram of gain changes resulting from increased contrast. No units showed the predicted reduction in gain. The broken line corresponds to the median increase in gain with increased contrast. D. Four representative units showing PSPs evoked by noise stimuli during high (red) and low (blue) contrast stimulation. Shaded areas indicate the standard deviation of V_m responses across trials. F. Membrane potential standard deviation during DRC stimulation did not change with stimulus contrast. G. Similarly, the standard deviation of PSP responses did not show a systematic change with stimulus contrast. Ratios in the corner of plots indicate the number of neurons above the identity line, y=x, over the total number of neurons.

5.3.3 Contrast dependent changes in *G* do not explain gain control

I next investigated whether the input conductance (G) of AC neurons is modulated by stimulus contrast, as predicted by the hypothesis that shunting inhibition con-

tributes to CGC. G was estimated by injecting different levels of current into the neuron during sensory stimulation (Figure 5.7A) as this allowed me to estimate G by examining the relationship between injected current and recorded membrane potential. In order to do this, I fit a linear model at each time point of the stimulus to the recorded membrane potential responses (Figure 5.7B). The inverse of slope of the model was used as the estimate for G at each time point (Figure 5.7C). The estimates of G over time could then be used to reconstruct the recorded V_m in order to validate the estimates of G (Figure 5.7D).

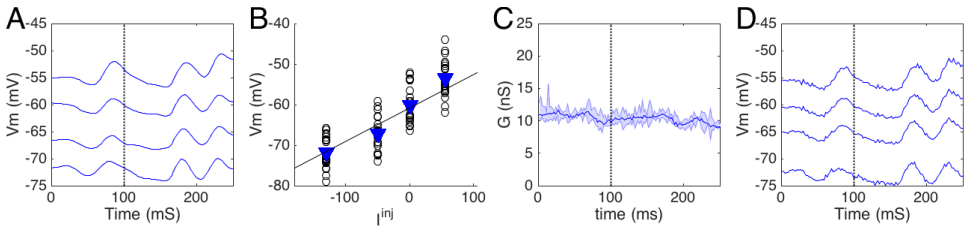

Figure 5.7: A. Mean V_m recorded for four different levels of current injection during DRC and noise stimulation. G was estimated at each time point in the stimulus. The broken line indicates a single time point 100 ms into the stimulus for which I estimated a single value of G. B. Open circles indicate the V_m recorded at this time point for 25 stimulus repetitions for the four levels of current injection. Blue triangles indicate mean V_m values plotted in A. G could be estimated at this time point from these responses by examining the relationship between injected current and recorded V_m response. This relationship was modelled by fitting a line to this data as Ohm's law predicts that this relationship should be linear. C. The inverse of the slope of the linear fit was used as the estimate for G for this time point. D. Using Ohm's law, the V_m could be reconstructed from this estimate for G and the known levels of current injection, in order to validate the estimates. This was necessary as neurons are not perfectly linear devices.

In order to asses the reliability of the linear model over the recorded range of membrane potentials, I compared the estimated (Figure 5.7D) V_m to the recorded V_m (Figure 5.7A). This was performed for the entire data set (Figure 5.8A). Predicted values of V_m captured 97% of the variance in the membrane potential response indicating that the estimates of G were successful and that the relationship between injected current and recorded membrane potential is highly linear (Figure 5.7B).

I next asked whether contrast dependent changes in G could account for the

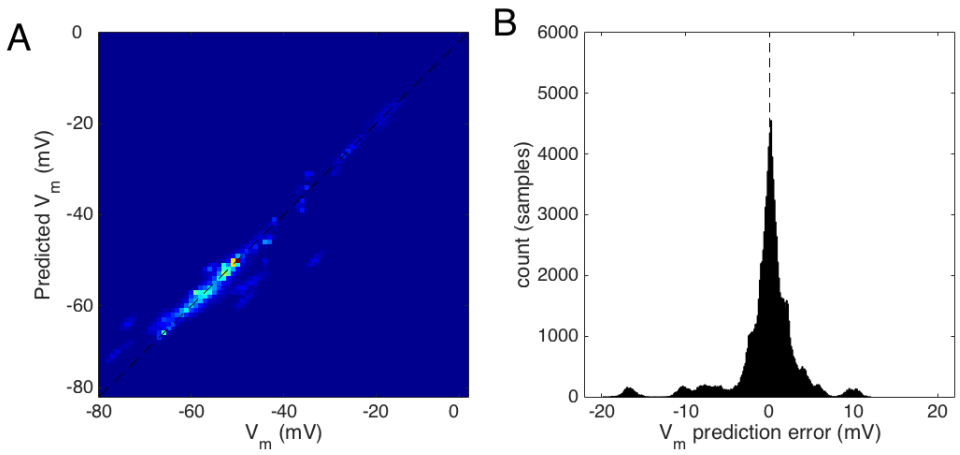

Figure 5.8: A. Density plot of predicted vs recorded membrane potential responses for the entire data set. B. Histogram of the errors in the membrane potential prediction. Predicted recorded membrane potential responses were highly similar to recorded membrane potential responses, validating the use of the linear model to estimate input conductance.

reduction in the amplitude of the V_m response under high contrast stimulation. In order to assess the contrast dependency of input conductance in the recorded neurons, G was estimated during stimulation with both high and low contrast DRCs. Increasing contrast by $\sim$2 had no systematic effect on G (Median change in G 2.37%; sign-rank test: p=0.296). In order to account for the contrast dependent scaling of PSPs, changes in G should be correlated with the observed changes in PSP amplitude. I found that contrast dependent changes in G were unrelated to the contrast dependent changes the size of PSPs (Figure 5.9H). This was primarily attributable to the small magnitude of changes in G between contrast conditions, indicating that input conductance does not change with stimulus contrast.

5.4 Discussion

I aimed to investigate the biophysical mechanisms underlying CGC in AC by examining the involvement of two mechanisms, shunting inhibition and membrane potential variability. My approach was to perform whole-cell current clamp record-

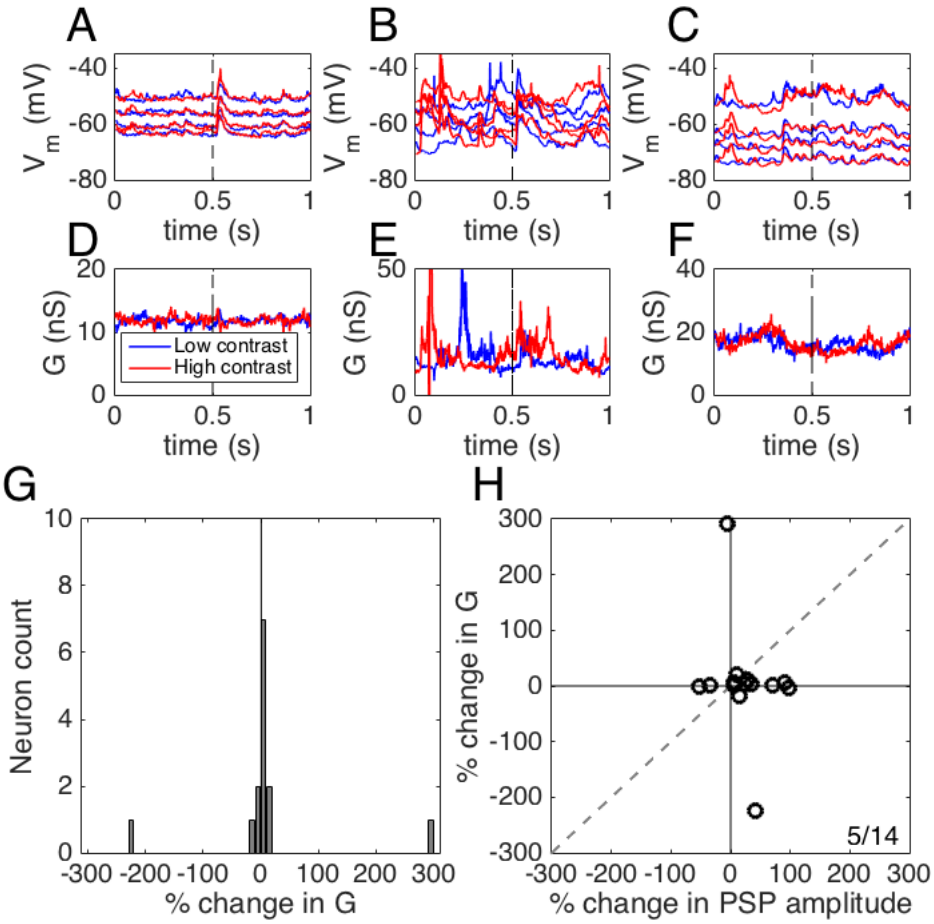

Figure 5.9: A-C. Membrane potential responses recorded under 4 different levels of current injection for 3 representative units, during high (red) and low (blue) contrast stimulation. Broken lines indicate the timing of noise stimulus presentation. D-F. G estimated for these same 3 units under high (red) and low (blue) contrast stimulation. G. Histogram of % change in G in response to an approximate doubling of stimulus contrast. G showed no systematic variation with stimulus contrast. H. Scatter plot of contrast dependent changes in PSP amplitude (abscissa) and G (ordinate). The change in G between contrast conditions was far smaller than changes that would be required to account for the contrast dependent scaling of the evoked V_m response. The ratio in the corner indicates the number of neurons above the identity line, $y=x$, over the total number of neurons.

ings from the AC of anaesthetised mice while they were presented with auditory stimuli of high or low contrast. This allowed me to examine the effect of stimulus contrast on responses at the level of the membrane potential, to quantify variability in the membrane potential and to estimate the input conductance of each neuron.

CGC in AC has been observed at the level of spiking responses (Rabinowitz et al., 2011a) but it has previously been unknown whether or not it exists at the level of the membrane potential. According to the hypothesis that shunting inhibition predominantly underlies CGC, gain changes observed in spiking responses occur as a result of a divisive scaling of PSP amplitude. This hypothesis therefore predicts that CGC should also be observed at the level of the membrane potential. The hypothesis that contrast-dependent fluctuations in membrane potential responses underlie CGC in the spiking response predicts that CGC should not exist at the level of the membrane potential responses. According to this hypothesis, gain changes occur as a result of interactions between membrane potential variability and spike threshold but the size of evoked PSPs remains unchanged. In order to test these predictions, I recorded membrane potential responses to high and low contrast DRC stimuli containing noise probe stimuli. The PSPs generated by these noise stimuli showed a divisive scaling in their amplitude during high contrast compared to low contrast stimulation. This is consistent with the shunting inhibition hypothesis but not the membrane potential variance hypothesis. This finding indicates that contrast-dependent membrane potential variance alone cannot account for CGC in AC but it does not preclude it from contributing to this computation in combination with another mechanism that can account for the contrast-dependent reduction in PSP amplitude. In order to investigate whether membrane potential variability contributes at all to CGC in AC, I compared the standard deviation of evoked PSP amplitudes between contrast conditions and found no systematic difference. In keeping with this, the gain of the relationship between mean membrane potential and firing rate was not reduced during high contrast stimulation in AC neurons, unlike in V1 where this is found to be the case Finn et al. (2007). These findings indicate that contrast dependent changes in membrane potential variability do not contribute to CGC in AC, but are consistent with the shunting inhibition

hypothesis.

Shunting inhibition is thought to modulate gain by increasing the input conductance of neurons in order to divisively scale PSP amplitude. In order to examine the contribution of shunting inhibition to CGC in AC, I estimated the input conductance of each neuron during stimulation with high and low contrast DRCs. I did this by injecting multiple levels of current and fitting a series of linear models to the relationship between this injected current and the observed membrane potential response at each time point in the recording. This analysis assumes that this relationship is linear, an assumption which may reasonably be questioned given the number of voltage dependent nonlinearities present in neurons. I primarily injected negative currents into the neurons in order to avoid depolarising them and thereby avoid interactions with voltage dependent nonlinearities. I also validated this modelling approach and tested the assumption of linearity by predicting the membrane potential responses from the model. In keeping with previous reports (Anderson et al., 2000), I found that this relationship is indeed highly linear. This approach generated estimates of input conductance over time during sensory stimulation.

Using this modelling approach, I found that input conductance was not significantly modulated by stimulus contrast. In V1, input conductance can increase up to 300% during the presentation of a high contrast grating with reference to baseline input conductance (Anderson et al., 2000). Here I observed no such effect, indicating that in AC, shunting inhibition does not contribute to CGC. Our recordings were performed across the depth of cortex however, and it is possible that shunting inhibition is recruited by thalamic inputs and may therefore be confined to neurons in layer 4. Targeted recordings of input conductance in layer 4 neurons during stimulation with stimuli of different contrasts will be necessary in order to address this issue.

Chapter 6

Discussion

6.1 Conclusions

In this thesis, I aimed to investigate the mechanisms underlying a fundamental sensory computation, contrast gain control (CGC), in auditory cortex (AC). Little was previously known regarding this issue, in contrast to primary visual cortex where the mechanistic basis of this computation has been studied extensively. The presence of this computation in other sensory cortical areas, such as primary visual cortex (V1), was a strong motivating factor in this research, as investigating the mechanisms underlying CGC in AC would not only add to our understanding of auditory processing but would also contribute to the fundamental issue of whether cortical areas exploit canonical circuit mechanisms in order to implement canonical computations.

The methods required in order to investigate these mechanisms necessitated the use of a different model organism to the ferret, in which CGC had originally been characterised. Mice were used due to their suitability for both optogenetic manipulations of genetically determined neuronal subtypes and for *in-vivo* whole-cell recording. Before these techniques could be used in order to investigate the mech-

anistic basis of CGC in AC, it was necessary to determine whether it was present in mouse AC. I found that CGC was widely observed in spiking responses in AC, allowing for investigations into its underlying mechanisms using the previously mentioned techniques.

I first examined whether PV interneuron inhibition is capable of modulating gain in AC. I found that, as in V1, PV interneuron activity is capable of modulating the gain of sensory evoked responses in AC. I next examined whether PV interneurons increase their firing in response to an increase in stimulus contrast, as they have been found to do in V1. I found that in the neurons I recorded from in AC, this was not the case. This raises the possibility that a subpopulation of PV interneurons, not recorded from in these experiments, may be capable of contributing to contrast driven gain changes in AC. I finally examined contributions of two biophysical mechanisms to CGC in AC; shunting inhibition and membrane potential variability. Unlike in V1, I found no role for either of these mechanisms in CGC in AC, raising the possibility that this canonical cortical computation is implemented by non-canonical circuit mechanisms in different cortical areas. Alternatively, PV interneuron mediated shunting inhibition may underlie CGC but in AC this may be restricted to layer 4.

6.2 CGC in mouse AC

I aimed to characterise CGC in AC both in order to allow research into its mechanistic basis but also in order to examine the canonical nature of this computation (Carandini and Heeger, 2012). CGC has been observed in both the AC of the ferret (Rabinowitz et al., 2011a) and V1 of several species, including the mouse (Atallah et al., 2012). Characterising CGC in mouse AC therefore enabled me to test whether CGC is truly canonical, both in the sense that it is observed both across species and

across cortical areas. I found that the majority of multi-unit (MU) responses in AC showed a reduction in gain during high contrast stimulation, indicative of CGC. I also observed several differences compared with CGC in ferret AC. Unlike in the ferret, the strength of gain control in the mouse appeared to approximate complete compensation. As in the ferret, spectrotemporal tuning was left largely unaffected by changes in stimulus contrast, except for a small change in the bandwidth of spectral tuning. Interestingly, however, whereas bandwidths narrowed slightly with increasing contrast in ferret AC (Rabinowitz et al., 2011a), bandwidths increased slightly under high contrast stimulation in the mouse. When a sigmoidal output nonlinearity was fitted to these receptive field models, no change in threshold was observed for responses in mouse AC during high contrast stimulation, compared with a reduction in threshold in the ferret during high contrast stimulation.

These differences indicate that there may be differences in the way in which this computation is implemented across species. This possibility can only be investigated through the characterisation and comparison of the mechanisms involved across these two species. Such interspecies differences in CGC have not been reported in V1 however. This raises the possibility that the variation in the effect of stimulus contrast on responses in ferret A1 and mouse AC is due to differences in experimental paradigms. Ferret A1 is easily located based on anatomy. This is not the case for mouse A1 which shows variation in its precise location across animals. It was not possible to map auditory cortical fields in order to precisely localise A1 and so the dataset presented here likely includes responses from both primary and secondary cortical areas. Responses in secondary fields may show different contrast-dependencies which could have contributed to the differences observed between species. Another experimental difference is that investigations of CGC in ferret A1 have been based largely on single-unit, not multi-unit data (Rabinowitz et al., 2011a, 2013). This may also have contributed to the observed differences.

6.3 Non-canonical circuit mechanisms

Gain control has been proposed to represent a canonical computation. Taken in combination with suggestions that the cortex consists of canonical circuits for such computations, it appears possible that the same mechanisms may implement this computation across sensory cortical areas and across species. In ferret AC, CGC has been shown to compensate for 75% of changes in stimulus contrast. I found that, in mouse AC, CGC approximates complete compensation for changes in stimulus contrast. Similarly, In V1, CGC is thought to be completely compensatory (Heeger, 1992). These differences raise the possibility that, despite the canonical nature of this computation and the relative homogeneity of cortical architecture, different mechanisms may implement this computation in these different sensory systems and species. Much work has been done in investigating the mechanistic basis of CGC in V1, making experimentally determining whether the same mechanisms contribute to CGC in AC as well the most direct test of the canonical circuit hypothesis.

In V1, CGC is thought to be implemented via a combination of mechanisms. Early research focused on shunting inhibition as a mechanism of contrast dependent gain control (Carandini and Heeger, 1994; Carandini et al., 1997) and, more recently, efforts have turned to characterising the contribution of the most likely source of shunting inhibition, PV interneurons (Atallah et al., 2012; Wilson et al., 2012; Lee et al., 2012, 2014; El-Boustani et al., 2014). In V1, PV interneuron activity in superficial layers increases with stimulus contrast (Atallah et al., 2012) and the effect of their inhibition has been shown to be divisive, i.e. it produces a reduction in the gain of sensory evoked responses. These two findings establish a logical chain of events that strongly implicates PV interneuron inhibition in gain control; as stimulus contrast increases, so does PV interneuron activity and as PV interneuron activity increases, gain is reduced. I aimed to investigate whether a similar circuit

logic exists in AC that may account for CGC in this area.

In order to investigate whether PV interneuron activity is capable of modulating the gain of sensory evoked responses as in V1, I optogenetically modulated PV interneuron activity during sensory stimulation. ChR2 activation of PV interneurons, had no effect on spectro-temporal tuning but produced a small reduction in the gain of sensory responses. The largest effect however was a subtractive reduction in baseline fairing rates. In keeping with this finding, a recent study in which ChR2 was used in order to activate PV interneurons in AC found that this artificial activation led to complex network effects in which the true nature of the inhibitory effect was masked by interactions with the spike threshold of pyramidal cells (Seybold et al., 2015). Another possibility however is that this result is due to activation of non-primary AC areas as, unlike V1, it is not possible to precisely target A1 without functional mapping of auditory cortical fields.

By suppressing PV interneuron activity using Arch, I avoided issues of interactions with spike threshold, which have similarly been found to mask the effect of PV interneuron inhibition in V1 (El-Boustani et al., 2014). This approach also avoids introducing artificial structure into the spiking activity of the population being manipulated. As with PV interneuron activation, PV interneuron suppression left tuning properties largely unchanged. The strongest effect of PV suppression was a multiplicative disinhibition of sensory evoked responses, corresponding to an increase in gain. This demonstrates that PV interneuron activity is capable of modulating gain in AC as it has been shown to do in V1 (Atallah et al., 2014)

I next investigated whether PV interneurons are typically recruited under conditions that would allow them to contribute to CGC. I tested whether PV interneurons increase their firing rates during high contrast stimulation, relative to low contrast stimulation, allowing them to mediate the suppressive reduction in gain observed during high contrast stimulation. In order to do this, I expressed ChR2 in PV in-

terneurons and used their responsiveness to light in order to identify these cells in extracellular recordings (Lima et al., 2009). This allowed me to track the activity of PV interneurons during stimulation with high and low contrast stimuli. In contrast to the findings in V1 (Atallah et al., 2012) I found that the majority of putative PV interneurons I recorded from did not increase their firing with increased stimulus contrast. The neurons recorded from, however, represent a small subset of PV interneurons, precluding any conclusive statements being made regarding the involvement of PV interneurons in CGC. Several methods exist for targeting PV interneurons and these approaches have previously resulted in contradictory results (Moore and Wehr, 2013; Li et al., 2014). Variation in the results of different studies may be in part due to PV interneurons consisting of multiple sub-classes of interneurons which are indiscriminately manipulated or recorded from in these studies. The development of more precise genetic techniques to target these subclasses may reveal a variety of different functional contributions to CGC and other cortical computations.

6.4 Layer-specific circuitry

These experiments attempted to investigate whether a correlative link between stimulus contrast and PV interneuron activity exists and whether a causal link between PV interneuron activity and the gain of sensory evoked responses could be established. One major difference in the approaches taken to investigate these two issues relates to the subpopulations of PV interneurons being examined. In investigating the hypothesised link between PV interneuron activity and response gain in AC, the activity PV interneurons across cortical layers was modulated, particularly those in the most superficial cortical layers. In investigating how PV interneuron activity relates to stimulus contrast, PV interneurons were recorded predominantly

in deep cortical layers, although the precise laminar location of the recordings could not be verified. This raises the possibility that sub-populations of PV interneurons may contribute to CGC in AC. These subpopulations may have been recruited by the optogenetic manipulations but were not recorded from in my attempts to relate stimulus contrast to PV interneuron activity.

In chapter 3 I characterised the laminar organization of CGC in mouse AC. I found that CGC is strongest in deep cortical layers. Despite this laminar variation, CGC was present to a similar degree in all cortical layers, including thalamorecipient layer 4. Furthermore, this laminar pattern was mirrored by the variation in the noise of MU responses, with deep cortical layers showing the least noisy responses. These findings indicate that laminar variation in CGC may not be functionally relevant. It is possible therefore that CGC may be primarily implemented in layer 4 and subsequently inherited outside of layer 4, with relatively little modification. A PV interneuron subpopulation that contributes strongly to CGC might therefore be confined to layer 4, explaining the lack of consistent contrast-driven responses in deep PV interneurons. These neurons would have been targeted by the optogenetic manipulations but we most likely not recorded from in the stimulus contrast experiment. In order to characterise the contribution of PV interneurons in different cortical layers more systematically, newly developed multi-channel silicon electrodes designed for both laminar identification and single unit isolation (Blanche et al., 2005) may be used in combination with optogenetic-mediated PV interneuron identification (Lima et al., 2009), in order to simultaneously record from PV interneuron subtypes in all cortical layers.

Previous work indicates that layer 4 PV interneurons may form a functional subpopulation. These neurons receive the densest thalamic projections of any cortical interneuron subtype, including PV interneurons found in all other cortical layers (Ji et al., 2015). A feedforward thalamocortical circuit may therefore underlie CGC in

AC, with layer 4 PV interneurons providing feedforward inhibition. In the ferret, CGC has a time course of 62-85ms (Rabinowitz et al., 2011a) making it likely that other mechanisms, including recurrent feedback inhibition, contribute to CGC.

6.5 Biophysical mechanisms of CGC

If PV interneuron inhibition underlies CGC, this should be detectable through intracellular recordings. PV interneurons are thought to mediate shunting inhibition, in which large changes in somatic conductance have a divisive scaling effect on membrane potential responses, that can result in a reduction in gain (Borg-Graham et al., 1998b). Such changes in gain are observed in a contrast-dependent manner in V1, and are thought to contribute to CGC. Stimuli of 100% contrast have been shown to produce increases in conductance of up to 300% in V1 pyramidal cells (Borg-Graham et al., 1998b; Anderson et al., 2000; Monier et al., 2003; Frégnac et al., 2003). I found that the input conductance of neurons in AC did not increase with stimulus contrast as would be expected if shunting inhibition contributes to CGC in AC. Again, however, recordings were performed across the depth of cortex. If CGC is implemented via shunting inhibition from layer 4 PV interneurons, recordings confined to layer 4 may find contrast-dependent modulation of input conductance. Intracellular recording *in vivo* could be combined with optogenetic activation of PV interneurons in order to test whether this subtype really does modulate the conductance of pyramidal cells, as predicted by the hypothesis that they signal via shunting inhibition.

If cortical inhibition does contribute to CGC in AC but no contrast-dependent changes in conductance are observed in layer 4, *in vivo* voltage clamp recordings from pyramidal cells during sensory stimulation may be used in order to identify the mechanism by which inhibition modulates gain. One likely candidate mechanism

is balanced excitation and inhibition, which has been observed extensively in AC (Wehr and Zador, 2003; Sun et al., 2010; Dorrn et al., 2010; Zhou et al., 2014a) and other cortical areas (Haider et al., 2006; Okun and Lampl, 2008; Xue et al., 2014; Hasenstaub et al., 2005; Haider and McCormick, 2009).

Inhibition appeared to be the most plausible candidate mechanism based on results from CGC in the ferret. In ferret AC, stimulus contrast in both excitatory and suppressive regions of frequency space are known to contribute to CGC (Rabinowitz et al., 2011a). The contribution of stimulus contrast in suppressive regions of the receptive field of neurons implicated intracortical inhibition as a mechanism of CGC in AC. It is not known whether this finding generalises to the mouse, however, as I did not attempt to characterise the spectral region within which stimulus contrast modulates gain in mouse AC. Suppressive side bands in frequency tuning were rarely observed however, making it possible that this finding may not generalise to the mouse. If this is the case, it increases the possibility that non-inhibitory mechanisms may contribute to CGC in mouse AC.

Pharmacological blockade of GABA transmission offers an approach for testing the involvement of cortical inhibition in CGC in both species, although in V1 this has produced contradictory results (Katzner et al., 2011; Morrone et al., 1987; Ozeki et al., 2009; Haider et al., 2010). If cortical inhibition can be shown to play a causal role in CGC, systematic characterisation of the contrast-response properties of different interneuron subtypes may allow the crucial components to be identified. Calcium-dependent sensors for *in vivo* imaging can be targeted to specific interneuron subtypes in the mouse, allowing the activity of populations of these neurons to be imaged during sensory stimulation (Hippenmeyer et al., 2005; Li et al., 2014). Such imaging techniques are typically restricted to layers 1-4 due to lack of light penetration to deeper cortical layers (Winkowski and Kanold, 2013).

6.6 Non-inhibitory mechanisms of gain control

While it is possible that PV interneurons in layer 4 mediate CGC via shunting inhibition, it is also possible that non-inhibitory mechanisms contribute to CGC in mouse AC. It may be the case that the divisive nature of PV interneuron inhibition has been recruited for other purposes in AC, instead of being used to implement contrast driven gain changes. In keeping with this, PV interneurons in AC have been shown to be recruited by top-down inputs that are known to mediate other forms of gain control (Pi et al., 2013).

One non-inhibitory mechanism of gain control is contrast dependent membrane potential variability. This mechanism is known to contribute to CGC in V1, where the increased variability of membrane potential responses across trials under low contrast stimulation results in an increased probability of spike-threshold crossing (Troyer et al., 1998; Hansel and van Vreeswijk, 2002; Carandini, 2004; Ringach and Malone, 2007; Finn et al., 2007). This results in an increase in the gain of the relationship between mean membrane potential across trials and mean spiking response across trials. I found no contrast-dependent modulation of membrane potential variability in mouse AC, indicating that this mechanism is not canonical, in that it is not preserved across cortical areas.

Short-term synaptic depression (STD) at the thalamocortical synapse is believed to contribute to CGC in V1. The prime reason that it was not considered as a mechanism for CGC in AC was that is could only account for gain changes within restricted excitatory regions of frequency space. If stimulus contrast only contributes to CGC in mouse AC when it is within the excitatory region of frequency tuning, STD may contribute to CGC in mouse AC. STD is known to occur at the auditory thalamocortical synapse in the mouse (Rose and Metherate, 2005) but it is not known whether stimulus contrast outside the receptive field of the neuron can modulate gain in the mouse. Replicating in the mouse the approach used in order to quantify the spec-

tral regions that contribute to CGC in the ferret (Rabinowitz et al., 2011a) will be necessary in order to address this question.

In the ferret, systematic CGC is not observed in the inferior colliculus (IC) but emerges in AC (Rabinowitz et al., 2013). It is not known whether CGC has been implemented by the level of the IC in the mouse. Furthermore, it is not known whether CGC exists in the medial geniculate body of the thalamus (MGB) of either of these species. STD may contribute to CGC in each of these stations of the ascending auditory pathway, sequentially strengthening CGC (Rabinowitz et al., 2013).

6.7 CGC in the thalamus

Systematic characterisation of the strength and time-course of CGC in the ascending auditory pathway of the mouse will be necessary in order to identify the locations of the crucial mechanisms involved in CGC. If, as in the ferret, CGC is less strong and systematic in the IC, the location of the mechanisms underlying CGC observed in cortex can be most likely localized to the thalamocortical system. Given the observation that CGC is present at full strength in thalamorecipient layer 4, the site of the crucial mechanisms is most likely to be found at the thalamocortical synapse or perhaps in the thalamus itself. Given the strong interconnectedness of the thalamus and cortex, simply demonstrating the presence of CGC in the thalamus will most likely not be adequate in order to ascertain whether it is implemented solely in this structure. Careful characterisation of the time course of CGC during simultaneous recordings from MGB and AC may be necessary in order to identify which, if either, of these structures plays the more important role in implementing CGC.

If CGC is found to occur more quickly and at full strength in MGB, and this can be shown to occur in the absence of corticothalamic feedback projections from layer

6 of AC, the site of the mechanisms implementing CGC will more likely be found in thalamus than cortex. These feedback projections can be targeted for optogenetic or chemogenetic inactivation in the mouse using the NTSR1-Cre GN220 mouse line (Olsen et al., 2012; Gong et al., 2007). This approach can also be replicated in other species, such as the ferret, via chromophore-targeted laser photolysis (Bajo et al., 2010) or via dual injections of cre-dependent Arch in AC and a cre-expressing virus in MGB (Znamenskiy and Zador, 2013).

If CGC is found to be weaker in the thalamus than cortex, or to have a slower time course, the most likely site of CGC will be the thalamocortical synapse, or more generally layer 4 of AC. STD at the thalamocortical synapse may be capable of accounting for CGC in the mouse if it can be demonstrated that only the excitatory regions of the receptive field contribute to CGC in AC. Furthermore it would also be necessary to demonstrate that the entirety of a cortical neuron's excitatory tuning can be inherited from individual MGB afferents. This may be achievable by inactivating different portions of MGB while simultaneously recording from AC in order to examine how tuning is affected. Experimental manipulations or recordings of sensory driven STD are not currently possible *in vivo*, so simulations may be necessary in order to examine any role for this mechanism in CGC. Several predictions are made from the STD hypothesis, however, that can be tested experimentally.

6.8 Linking mechanisms to computation and behaviour

The functional importance of any mechanisms that contribute to CGC can ultimately be examined using the same causal techniques employed in this thesis. CGC has been shown to result in cortical representations that are invariant to background noise (Rabinowitz et al., 2013). It may be possible to demonstrate the behavioural importance of specific mechanisms by disrupting these mechanisms in

animals trained to perform an appropriate behavioural task. In the case of noise invariance, an auditory signal-in-noise discrimination task might be used. It may also be possible to causally manipulate these mechanisms in order to strengthen CGC and thereby benefit behaviour. A primary function of CGC is the efficient use of the dynamic range of sensory neurons. A consequence of this is that, when in a low-gain state, discrimination of stimulus intensity should be improved. It may therefore be possible to improve intensity discrimination by blocking mechanisms that typically reduce gain under natural conditions.

Other functions of CGC can be examined at the electrophysiological rather than behavioural level. CGC is thought to produce less redundant and sparser population codes that are capable of winner-take-all dynamics. The importance of specific mechanisms to these proposed functions may similarly be examined through causal disruption of these mechanisms. Such experiments will ultimately allow links to be drawn between mechanistic, computational and behavioural levels of analysis in neuroscience.

Bibliography

Abbott, L. F. and W. G. Regehr

2004. Synaptic computation. *Nature*, 431(7010):796–803.

Abbott, L. F., J. A. Varela, K. Sen, and S. B. Nelson

1997. Synaptic depression and cortical gain control. *Science (New York, N.Y.)*, 275(5297):220–224.

Abolafia, J. M., R. Vergara, M. M. Arnold, R. Reig, and M. V. Sanchez-Vives

2011. Cortical auditory adaptation in the awake rat and the role of potassium currents. *Cerebral cortex (New York, N.Y. : 1991)*, 21(5):977–990.

Adesnik, H., W. Bruns, H. Taniguchi, Z. J. Huang, and M. Scanziani

2012. A neural circuit for spatial summation in visual cortex. *Nature*, 490(7419):226–231.

Ahrens, M. B., J. F. Linden, and M. Sahani

2008. Nonlinearities and contextual influences in auditory cortical responses modeled with multilinear spectrotemporal methods. *The Journal of neuroscience*, 28(8):1929–1942.

Akerboom, J., N. Carreras Calderón, L. Tian, S. Wabnig, M. Prigge, J. Tolö, A. Gordus, M. B. Orger, K. E. Severi, J. J. Macklin, R. Patel, S. R. Pulver, T. J. Wardill, E. Fischer, C. Schüler, T.-W. Chen, K. S. Sarkisyan, J. S. Marvin, C. I. Bargmann,

D. S. Kim, S. Kügler, L. Lagnado, P. Hegemann, A. Gottschalk, E. R. Schreiter, and L. L. Looger

2013. Genetically encoded calcium indicators for multi-color neural activity imaging and combination with optogenetics. *Frontiers in Molecular Neuroscience*, 6(2).

Albrecht, DG, G. W.

1991. Motion sensitivity and the contrast-response function of simple cells in the visual cortex. *Vis. Neurosci.*, 7:531–546.

Amitai, Y., J. R. Gibson, M. Beierlein, S. L. Patrick, A. M. Ho, B. W. Connors, and D. Golomb

2002. The spatial dimensions of electrically coupled networks of interneurons in the neocortex. *The Journal of neuroscience*, 22(10):4142–4152.

Anderson, J. S., M. Carandini, and D. Ferster

2000. Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *Journal of neurophysiology*, 84(2):909–926.

Anderson, L. A., G. B. Christianson, and J. F. Linden

2009. Mouse auditory cortex differs from visual and somatosensory cortices in the laminar distribution of cytochrome oxidase and acetylcholinesterase. *Brain Research*, 1252:130–142.

Ascoli, G. A., L. Alonso-Nanclares, S. A. Anderson, G. Barrionuevo, R. Benavides-Piccione, A. Burkhalter, G. Buzsáki, B. Cauli, J. Defelipe, A. Fairén, D. Feldmeyer, G. Fishell, Y. Frégnac, T. F. Freund, D. Gardner, E. P. Gardner, J. H. Goldberg, M. Helmstaedter, S. Hestrin, F. Karube, Z. F. Kisvárday, B. Lambolez, D. A. Lewis, O. Marin, H. Markram, A. Muñoz, A. Packer, C. C. H. Petersen, K. S. Rockland, J. Rossier, B. Rudy, P. Somogyi, J. F. Staiger, G. Tamás, A. M. Thomson, M. Toledo-

Rodriguez, Y. Wang, D. C. West, and R. Yuste

2008. Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex. *Nature reviews. Neuroscience*, 9(7):557–568.

Atallah, B. V., W. Bruns, M. Carandini, and M. Scanziani

2012. Parvalbumin-Expressing Interneurons Linearly Transform Cortical Responses to Visual Stimuli. *Neuron*, 73(1):159–170.

Atallah, B. V., M. Scanziani, and M. Carandini

2014. Atallah et al. reply. *Nature*, 508(7494):E3.

Atencio, C. A. and C. E. Schreiner

2008. Spectrotemporal Processing Differences between Auditory Cortical Fast-Spiking and Regular-Spiking Neurons. *The Journal of Neuroscience*, 28:3897–3910.

Atiani, S., M. Elhilali, S. V. David, J. B. Fritz, and S. A. Shamma

2009. Task difficulty and performance induce diverse adaptive patterns in gain and shape of primary auditory cortical receptive fields. *Neuron*, 61(3):467–480.

Atick, J. J.

2015. Could information theory provide an ecological theory of sensory processing? *Network: Computation in Neural Systems*, 3:213–251.

Attneave, F.

1954. Some informational aspects of visual perception. *Psychol. Rev.*, 61:183–93.

Aubin, T.

2004. Penguins and their noisy world. *Anais da Academia Brasileira de Ciências*, 76(2):279–283.

Avan, P., D. Loth, C. Menguy, and M. Teyssou

1992. Hypothetical roles of middle ear muscles in the guinea-pig. *Hearing Research*, 59(1):59–69.

Baccus, S. A.

2006. From a whisper to a roar: adaptation to the mean and variance of naturalistic sounds. *Neuron*, 51(6):682–684.

Baccus, S. A. and M. Meister

2002. Fast and slow contrast adaptation in retinal circuitry. *Neuron*, 36(5):909–919.

Bajo, V. M., F. R. Nodal, D. R. Moore, and A. J. King

2010. The descending corticocollicular pathway mediates learning-induced auditory plasticity. *Nature neuroscience*, 13(2):253–260.

Bakin, J. S. and N. M. Weinberger

1990. Classical conditioning induces CS-specific receptive field plasticity in the auditory cortex of the guinea pig. *Brain Research*, 536(1-2):271–286.

Bandyopadhyay, S., S. A. Shamma, and P. O. Kanold

2010. Dichotomy of functional organization in the mouse auditory cortex. *Nature Neuroscience*, 13(3):361–368.

Banitt, Y., K. A. C. Martin, and I. Segev

2007. A biologically realistic model of contrast invariant orientation tuning by thalamocortical synaptic depression. *The Journal of neuroscience*, 27(38):10230–10239.

Barbour, D. L. and E. M. Callaway

2008. Excitatory local connections of superficial neurons in rat auditory cortex. *The Journal of neuroscience*, 28(44):11174–11185.

Barlow, H.

1961. Possible principles underlying the transformation of sensory messages. *Sensory communication*, Pp. 217–234.

Bartos, M., I. Vida, and P. Jonas

2007. Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature reviews. Neuroscience*, 8(1):45–56.

Beierlein, M. and B. W. Connors

2002. Short-term dynamics of thalamocortical and intracortical synapses onto layer 6 neurons in neocortex. *Journal of neurophysiology*, 88(4):1924–1932.

Beierlein, M., J. R. Gibson, and B. W. Connors

2000. A network of electrically coupled interneurons drives synchronized inhibition in neocortex. *Nature neuroscience*, 3(9):904–910.

Benardete, E. A., E. Kaplan, and B. W. Knight

1992. Contrast gain control in the primate retina: P cells are not X-like, some M cells are. *Visual Neuroscience*, 8(05):483–486.

Bennett, C., S. Arroyo, and S. Hestrin

2013. Subthreshold Mechanisms Underlying State-Dependent Modulation of Visual Responses. *Neuron*, 80(2):350–357.

Betz, W. J.

1970. Depression of transmitter release at the neuromuscular junction of the frog. *The Journal of physiology*, 206(3):629–644.

Bitterman, Y., R. Mukamel, R. Malach, I. Fried, and I. Nelken

2008. Ultra-fine frequency tuning revealed in single neurons of human auditory cortex. *Nature*, 451(7175):197–201.

Bizley, J. K., F. R. Nodal, I. Nelken, and A. J. King

2005. Functional organization of ferret auditory cortex. *Cerebral Cortex*, 15(10):1637–1653.

Blanche, T. J., M. A. Spacek, J. F. Hetke, and N. V. Swindale

2005. Polytrodes: High-Density Silicon Electrode Arrays for Large-Scale Multiunit Recording. *Journal of neurophysiology*, 93(5):2987–3000.

Bonin, V., V. Mante, and M. Carandini

2006. The statistical computation underlying contrast gain control. *The Journal of neuroscience*, 26(23):6346–6353.

Borg E, C. S.

1989. The middle-ear muscles. *Sci Am.*, 261:74–80.

Borg-Graham, L. J., C. Monier, and Y. Frégnac

1998a. Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, 393(6683):369–373.

Borg-Graham, L. J., C. Monier, and Y. Frégnac

1998b. Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, 393(6683):369–373.

Bortone, D. S., S. R. Olsen, and M. Scanziani

2014. Translaminar inhibitory cells recruited by layer 6 corticothalamic neurons suppress visual cortex. *Neuron*, 82(2):474–485.

Boudreau, C. E. and D. Ferster

2005. Short-term depression in thalamocortical synapses of cat primary visual cortex. *The Journal of neuroscience*, 25(31):7179–7190.

Boyden, E. S., F. Zhang, E. Bamberg, G. Nagel, and K. Deisseroth

2005. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature neuroscience*, 8(9):1263–1268.

Boyev, K. P., M. C. Liberman, and M. C. Brown

2002. Effects of anesthesia on efferent-mediated adaptation of the DPOAE. *Journal of the Association for Research in Otolaryngology : JARO*, 3(3):362–373.

Brenner, N., W. Bialek, and R. de Ruyter van Steveninck

2000. Adaptive rescaling maximizes information transmission. *Neuron*, 26(3):695–702.

Brown, S. P. and R. H. Masland

2001. Spatial scale and cellular substrate of contrast adaptation by retinal ganglion cells. *Nature neuroscience*, 4(1):44–51.

Busse, L., A. R. Wade, and M. Carandini

2009. Representation of concurrent stimuli by population activity in visual cortex. *Neuron*, 64(6):931–942.

Capaday, C.

2002. A re-examination of the possibility of controlling the firing rate gain of neurons by balancing excitatory and inhibitory conductances. *Experimental Brain Research*, 143(1):67–77.

Carandini, M.

2004. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS biology*, 2(9):E264.

Carandini, M. and D. Ferster

1997. A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science (New York, N.Y.)*, 276(5314):949–952.

Carandini, M. and D. J. Heeger

1994. Summation and division by neurons in primate visual cortex. *Science (New York, N.Y.)*, 264(5163):1333–1336.

Carandini, M. and D. J. Heeger

2012. Normalization as a canonical neural computation. *Nature reviews. Neuroscience*, 13(1):51–62.

Carandini, M., D. J. Heeger, and J. A. Movshon

1997. Linearity and normalization in simple cells of the macaque primary visual cortex. *The Journal of Neuroscience*, 17(21):8621–8644.

Carandini, M., D. J. Heeger, and W. Senn

2002. A synaptic explanation of suppression in visual cortex. *The Journal of neuroscience*, 22(22):10053–10065.

Cardin, J. A., M. Carlén, K. Meletis, U. Knoblich, F. Zhang, K. Deisseroth, L.-H. Tsai, and C. I. Moore

2009. Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, 459(7247):663–667.

Cardin, J. A., M. Carlén, K. Meletis, U. Knoblich, F. Zhang, K. Deisseroth, L.-H. Tsai, and C. I. Moore

2010. Targeted optogenetic stimulation and recording of neurons in vivo using cell-type-specific expression of Channelrhodopsin-2. *Nature protocols*, 5(2):247–254.

Cardin, J. A., L. A. Palmer, and D. Contreras

2007. Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. *The Journal of neuroscience*, 27(39):10333–10344.

Carmel, P. W. and A. Starr

1963. Acoustic and nonacoustic factors modifying middle-ear muscle activity in waking cats. *Journal of neurophysiology*, 26:598–616.

Cauli, B., E. Audinat, B. Lambolez, M. C. Angulo, N. Ropert, K. Tsuzuki, S. Hestrin, and J. Rossier

1997. Molecular and physiological diversity of cortical nonpyramidal cells. *The Journal of Neuroscience*, 17(10):3894–3906.

Chance, F. S., L. F. Abbott, and A. D. Reyes

2002. Gain modulation from background synaptic input. *Neuron*, 35(4):773–782.

Chander, D. and E. J. Chichilnisky

2001. Adaptation to temporal contrast in primate and salamander retina. *The Journal of neuroscience*, 21(24):9904–9916.

Cheng, H., Y. M. Chino, E. L. Smith, J. Hamamoto, and K. Yoshida

1995. Transfer characteristics of lateral geniculate nucleus X neurons in the cat: effects of spatial frequency and contrast. *Journal of neurophysiology*, 74(6):2548–2557.

Choi, Y.-S., M. A. Koenig, X. Jia, and N. V. Thakor

2010. Quantifying time-varying multiunit neural activity using entropy based measures. *IEEE transactions on bio-medical engineering*, 57(11):2771–2777.

Chow, B. Y., X. Han, A. S. Dobry, X. Qian, A. S. Chuong, M. Li, M. A. Henninger, G. M. Belfort, Y. Lin, P. E. Monahan, and E. S. Boyden

2010. High-performance genetically targetable optical neural silencing by light-driven proton pumps. *Nature*, 463(7277):98–102.

Christianson, G. B., M. Sahani, and J. F. Linden

2008. The consequences of response nonlinearities for interpretation of spectrotemporal receptive fields. *The Journal of neuroscience*, 28(2):446–455.

Christopher Kirk, E. and D. W. Smith

2003. Protection from acoustic trauma is not a primary function of the medial olivocochlear efferent system. *Journal of the Association for Research in Otolaryngology : JARO*, 4(4):445–465.

Chung, S. H., L. C. Jones, B. J. Hammond, M. C. King, R. J. Evans, C. Knott, M. J. Keating, and M. Anson

1987. Signal processing technique to extract neuronal activity from noise. *Journal of neuroscience methods*, 19(2):125–139.

Colburn, H. S., L. H. Carney, and M. G. Heinz

2003. Quantifying the information in auditory-nerve responses for level discrimination. *Journal of the Association for Research in Otolaryngology : JARO*, 4(3):294–311.

Connors, B. W. and M. J. Gutnick

1990. Intrinsic firing patterns of diverse neocortical neurons. *Trends in neurosciences*, 13(3):99–104.

Coombs, J. S., J. C. Eccles, and P. Fatt

1955. The electrical properties of the motoneurone membrane. *The Journal of physiology*, 130(2):291–325.

Dallas, P.

1992. The active cochlea. *The Journal of Neuroscience*, 12:4575–4585.

Dean, I., N. S. Harper, and D. McAlpine

2005. Neural population coding of sound level adapts to stimulus statistics. *Nature Neuroscience*, 8(12):1684–1689.

Dean, I., B. L. Robinson, N. S. Harper, and D. McAlpine

2008. Rapid neural adaptation to sound level statistics. *The Journal of neuroscience*, 28(25):6430–6438.

Deans, M. R., J. R. Gibson, C. Sellitto, B. W. Connors, and D. L. Paul

2001. Synchronous activity of inhibitory networks in neocortex requires electrical synapses containing connexin36. *Neuron*, 31(3):477–485.

deCharms, R. C., D. T. Blake, and M. M. Merzenich

1998. Optimizing Sound Features for Cortical Neurons. *Science (New York, N.Y.)*, 280(5368):1439–1444.

Destexhe, A. and D. Paré

1999. Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *Journal of neurophysiology*, 81(4):1531–1547.

Destexhe, A., M. Rudolph, and D. Paré

2003. The high-conductance state of neocortical neurons in vivo. *Nature reviews. Neuroscience*, 4(9):739–751.

Dittman, J. S., A. C. Kreitzer, and W. G. Regehr

2000. Interplay between facilitation, depression, and residual calcium at three presynaptic terminals. *The Journal of Neuroscience*, 20(4):1374–1385.

Dorrn, A. L., K. Yuan, A. J. Barker, C. E. Schreiner, and R. C. Froemke

2010. Developmental sensory experience balances cortical excitation and inhibition. *Nature*, 465(7300):932–936.

Douglas, R. J., C. Koch, M. Mahowald, K. A. Martin, and H. H. Suarez

1995. Recurrent excitation in neocortical circuits. *Science (New York, N.Y.)*, 269(5226):981–985.

Douglas, R. J. and K. A. Martin

1991a. A functional microcircuit for cat visual cortex. *The Journal of physiology*, 440:735–769.

Douglas, R. J. and K. A. Martin

1991b. A functional microcircuit for cat visual cortex. *The Journal of physiology*, 440:735–769.

Dunn, F. A., M. J. Lankheet, and F. Rieke

2007. Light adaptation in cone vision involves switching between receptor and post-receptor sites. *Nature*, 449(7162):603–606.

Eccles, J. C., B. Katz, and S. W. Kuffler

1941. Nature of the "endplate potential" in curarized muscle. *Journal of neurophysiology*.

Eggermont, J. J.

1994. Temporal modulation transfer functions for AM and FM stimuli in cat auditory cortex. Effects of carrier type, modulating waveform and intensity. *Hearing Research*, 74(1-2):51–66.

El-Boustani, S., N. R. Wilson, C. A. Runyan, and M. Sur

2014. El-Boustani et al. reply. *Nature*, 508(7494):E3–4.

Elmqvist, D. and D. M. Quastel

1965. A quantitative study of end-plate potentials in isolated human muscle. *The Journal of physiology*, 178(3):505–529.

Escabí, M. A., L. M. Miller, H. L. Read, and C. E. Schreiner

2003. Naturalistic auditory contrast improves spectrotemporal coding in the cat inferior colliculus. *The Journal of neuroscience*, 23(37):11489–11504.

Evans, E. F. and A. R. Palmer

1980. Relationship between the dynamic range of cochlear nerve fibres and their spontaneous activity. *Experimental Brain Research*, 40(1):115–118.

Fairhall, A. L., G. D. Lewen, W. Bialek, and R. R. de Ruyter Van Steveninck

2001. Efficiency and ambiguity in an adaptive neural code. *Nature*, 412(6849):787–792.

Fatt, P. and B. Katz

1953. The effect of inhibitory nerve impulses on a crustacean muscle fibre. *The Journal of physiology*, 121(2):374–389.

Feng, T. P.

1941. *Studies on the neuromuscular junction. XXVI. The changes of the end-plate potential during and after prolonged stimulation*. Chin. J. Physiol.

Férézou, I., B. Cauli, E. L. Hill, J. Rossier, E. Hamel, and B. Lambolez

2002. 5-HT3 receptors mediate serotonergic fast synaptic excitation of neocortical vasoactive intestinal peptide/cholecystokinin interneurons. *The Journal of neuroscience*, 22(17):7389–7397.

Finn, I. M., N. J. Priebe, and D. Ferster

2007. The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron*, 54(1):137–152.

Fino, E., A. M. Packer, and R. Yuste

2013. The logic of inhibitory connectivity in the neocortex. *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry*, 19(3):228–237.

Fogarty, M., M. Grist, D. Gelman, O. Marin, V. Pachnis, and N. Kessaris

2007. Spatial Genetic Patterning of the Embryonic Neuroepithelium Generates

GABAergic Interneuron Diversity in the Adult Cortex. *The Journal of Neuroscience*, 27(41):10935–10946.

Frégnac, Y., C. Monier, F. Chavane, P. Baudot, and L. Graham

2003. Shunting inhibition, a silent step in visual cortical computation. *Journal of physiology, Paris*, 97(4-6):441–451.

Fritz, J. B., M. Elhilali, and S. A. Shamma

2007. Adaptive changes in cortical receptive fields induced by attention to complex sounds. *Journal of neurophysiology*, 98(4):2337–2346.

Fu, Y., J. M. Tucciarone, J. S. Espinosa, N. Sheng, D. P. Darcy, R. A. Nicoll, Z. J. Huang, and M. P. Stryker

2014. A cortical circuit for gain control by behavioral state. *Cell*, 156(6):1139–1152.

Furukawa, T., Y. Hayashida, and S. Matsuura

1978. Quantal analysis of the size of excitatory post-synaptic potentials at synapses between hair cells and afferent nerve fibres in goldfish. *The Journal of physiology*, 276:211–226.

Gabbiani, F., J. Midtgaard, and T. Knöpfel

1994. Synaptic integration in a model of cerebellar granule cells. *Journal of neurophysiology*, 72(2):999–1009.

Gabernet, L., S. P. Jadhav, D. E. Feldman, M. Carandini, and M. Scanziani

2005. Somatosensory Integration Controlled by Dynamic Thalamocortical Feed-Forward Inhibition. *Neuron*, 48(2):315–327.

Galarreta, M. and S. Hestrin

1999. A network of fast-spiking cells in the neocortex connected by electrical synapses. *Nature*, 402(6757):72–75.

Galarreta, M. and S. Hestrin

2001. Electrical synapses between Gaba-Releasing interneurons. *Nature reviews. Neuroscience*, 2(6):425–433.

Galarreta, M. and S. Hestrin

2002. Electrical and chemical synapses among parvalbumin fast-spiking GABAergic interneurons in adult mouse neocortex. *Proceedings of the National Academy of Sciences of the United States of America*, 99(19):12438–12443.

Garcia-Lazaro, J. A., S. S. M. Ho, A. Nair, and J. W. H. Schnupp

2007. Shifting and scaling adaptation to dynamic stimuli in somatosensory cortex. *The European journal of neuroscience*, 26(8):2359–2368.

Gibson, J. R., M. Beierlein, and B. W. Connors

1999. Two networks of electrically coupled inhibitory neurons in neocortex. *Nature*, 402(6757):75–79.

Gibson, J. R., M. Beierlein, and B. W. Connors

2005. Functional properties of electrical synapses between inhibitory interneurons of neocortical layer 4. *Journal of neurophysiology*, 93(1):467–480.

Gil, Z., B. W. Connors, and Y. Amitai

1999. Efficacy of thalamocortical and intracortical synaptic connections: quanta, innervation, and reliability. *Neuron*, 23(2):385–397.

Goldberg, E. M., B. D. Clark, E. Zagha, M. Nahmani, A. Erisir, and B. Rudy

2008. K+ Channels at the Axon Initial Segment Dampen Near-Threshold Excitability of Neocortical Fast-Spiking GABAergic Interneurons. *Neuron*, 58(3):387–400.

Golomb, D., K. Donner, L. Shacham, D. Shlosberg, Y. Amitai, and D. Hansel

2007. Mechanisms of firing patterns in fast-spiking cortical interneurons. *PLOS Computational Biology*, 3(8):e156.

Gong, S., M. Doughty, C. R. Harbaugh, A. Cummins, M. E. Hatten, N. Heintz, and C. R. Gerfen

2007. Targeting Cre recombinase to specific neuron populations with bacterial artificial chromosome constructs. *The Journal of neuroscience*, 27(37):9817–9823.

Goutman, J. D. and E. Glowatzki

2007. Time course and calcium dependence of transmitter release at a single ribbon synapse. *Proceedings of the National Academy of Sciences of the United States of America*, 104(41):16341–16346.

Guo, W., A. R. Chambers, K. N. Darrow, K. E. Hancock, B. G. Shinn-Cunningham, and D. B. Polley

2012. Robustness of cortical topography across fields, laminae, anesthetic states, and neurophysiological signal types. *The Journal of neuroscience*, 32(27):9159–9172.

Gupta, A., Y. Wang, and H. Markram

2000. Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science (New York, N.Y.)*, 287(5451):273–278.

Haider, B., A. Duque, A. R. Hasenstaub, and D. A. McCormick

2006. Neocortical Network Activity In Vivo Is Generated through a Dynamic Balance of Excitation and Inhibition. *The Journal of Neuroscience*, 26(17):4535–4545.

Haider, B., M. R. Krause, A. Duque, Y. Yu, J. Touryan, J. A. Mazer, and D. A. Mc-

Cormick

2010. Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron*, 65(1):107–121.

Haider, B. and D. A. McCormick

2009. Rapid neocortical dynamics: cellular and network mechanisms. *Neuron*, 62(2):171–189.

Hansel, D. and C. van Vreeswijk

2002. How Noise Contributes to Contrast Invariance of Orientation Tuning in Cat Visual Cortex. *The Journal of Neuroscience*, 22(12):5118–5128.

Harris, K. D. and G. M. G. Shepherd

2015. The neocortical circuit: themes and variations. *Nature Neuroscience*, 18(2):170–181.

Hasenstaub, A., Y. Shu, B. Haider, U. Kraushaar, A. Duque, and D. A. McCormick

2005. Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron*, 47(3):423–435.

Heckaman, R. L. and M. D. Fairchild

2009. Jones and Condit Redux in High Dynamic Range and Color. *Color and Imaging Conference*.

Heeger, D.

1992. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.*, 9:181–197.

Heeger, D. J., E. P. Simoncelli, and J. A. Movshon

1996. Computational models of cortical visual processing. *Proceedings of the National Academy of Sciences of the United States of America*, 93(2):623–627.

Hellwig, S., I. Hack, B. Zucker, B. Brunne, and D. Junghans

2012. Reelin Together with ApoER2 Regulates Interneuron Migration in the Olfactory Bulb. *PloS one*, 7(11):e50646.

Helmstaedter, M., B. Sakmann, and D. Feldmeyer

2009a. Neuronal Correlates of Local, Lateral, and Translaminar Inhibition with Reference to Cortical Columns. *Cerebral Cortex*, 19(4):926–937.

Helmstaedter, M., B. Sakmann, and D. Feldmeyer

2009b. The Relation between Dendritic Geometry, Electrical Excitability, and Axonal Projections of L2/3 Interneurons in Rat Barrel Cortex. *Cerebral Cortex*, 19(4):938–950.

Hensch, T. K.

2005. Critical period plasticity in local cortical circuits. *Nature reviews. Neuroscience*, 6(11):877–888.

Hestrin, S. and M. Galarreta

2005. Electrical synapses define networks of neocortical GABAergic neurons. *Trends in neurosciences*, 28(6):304–309.

Hippenmeyer, S., E. Vrieseling, M. Sigrist, T. Portmann, C. Laengle, D. R. Ladle, and S. Arber

2005. A developmental switch in the response of DRG neurons to ETS transcription factor signaling. *PLoS biology*, 3(5):e159.

Hirsch, J. A., L. M. Martinez, C. Pillai, J.-M. Alonso, Q. Wang, and F. T. Sommer

2003. Functionally distinct inhibitory neurons at the first stage of visual cortical processing. *Nature neuroscience*, 6(12):1300–1308.

Hofer, S. B., H. Ko, B. Pichler, J. Vogelstein, H. Ros, H. Zeng, E. Lein, N. A. Lesica,

and T. D. Mrsic-Flogel

2011. Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nature Neuroscience*, 14(8):1045–1052.

Holt, G. R. and C. Koch

1997. Shunting inhibition does not have a divisive effect on firing rates. *Neural Computation*, 9(5):1001–1013.

Houtsma, AJM, D. N. and L. Braida

1980. Intensity perception. xi. experimental results on the relation of intensity resolution to loudness matching. *J. Acoust. Soc. Am.*, 68:807?813.

Huang, C. L. and J. A. Winer

2000. Auditory thalamocortical projections in the cat: laminar and areal patterns of input. *The Journal of comparative neurology*, 427(2):302–331.

Huang, Z. J. and H. Zeng

2013. Genetic approaches to neural circuits in the mouse. *Annual Review of Neuroscience*, 36:183–215.

Ingham, N. J. and D. McAlpine

2004. Spike-frequency adaptation in the inferior colliculus. *Journal of neurophysiology*, 91(2):632–645.

Ingham, N. J. and D. McAlpine

2005. GABAergic inhibition controls neural gain in inferior colliculus neurons sensitive to interaural time differences. *The Journal of neuroscience*, 25(26):6187–6198.

Inta, D., J. Alfonso, J. von Engelhardt, M. M. Kreuzberg, A. H. Meyer, J. A. van Hooft, and H. Monyer

2008. Neurogenesis and widespread forebrain migration of distinct GABAer-

gic neurons from the postnatal subventricular zone. *Proceedings of the National Academy of Sciences of the United States of America*, 105(52):20994–20999.

Isaacson, J. S. and M. Scanziani

2011. How Inhibition Shapes Cortical Activity. *Neuron*, 72(2):231–243.

Issa, J. B., B. D. Haeffele, A. Agarwal, D. E. Bergles, E. D. Young, and D. T. Yue

2014. Multiscale optical Ca2+ imaging of tonal organization in mouse auditory cortex. *Neuron*, 83(4):944–959.

Ji, X.-y., B. Zingg, L. Mesik, Z. Xiao, L. I. Zhang, and H. W. Tao

2015. Thalamocortical Innervation Pattern in Mouse Auditory and Visual Cortex: Laminar and Cell-Type Specificity. *Cerebral cortex (New York, N.Y. : 1991)*, P. bhv099.

Joris, P. X. and T. C. T. Yin

1992. Responses to amplitude-modulated tones in the auditory nerve of the cat. *The Journal of the Acoustical Society of America*, 91(1):215–232.

Kaplan, E., K. Purpura, and R. M. Shapley

1987. Contrast affects the transmission of visual information through the mammalian lateral geniculate nucleus. *The Journal of physiology*, 391:267–288.

Katzner, S., L. Busse, and M. Carandini

2011. GABAA inhibition controls response gain in visual cortex. *The Journal of neuroscience*, 31(16):5931–5941.

Kawaguchi, Y.

1993. Groupings of nonpyramidal and pyramidal cells with specific physiological and morphological characteristics in rat frontal cortex. *Journal of neurophysiology*, 69(2):416–431.

Kawaguchi, Y.

1995. Physiological subgroups of nonpyramidal cells with specific morphological characteristics in layer II/III of rat frontal cortex. *The Journal of Neuroscience*, 15(4):2638–2655.

Kawaguchi, Y. and Y. Kubota

1997. GABAergic cell subtypes and their synaptic connections in rat frontal cortex. *Cerebral Cortex*, 7(6):476–486.

Kerlin, A. M., M. L. Andermann, V. K. Berezovskii, and R. C. Reid

2010. Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. *Neuron*, 67(5):858–871.

Kim, K. J. and F. Rieke

2001. Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. *The Journal of neuroscience*, 21(1):287–299.

King, A. J. and S. Carlile

1994. Responses of neurons in the ferret superior colliculus to the spatial location of tonal stimuli. *Hearing Research*, 81(1-2):137–149.

Klausberger, T. and P. Somogyi

2008. Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science (New York, N.Y.)*, 321(5885):53–57.

Klein, D. J., P. König, and K. P. Körding

2003. Sparse spectrotemporal coding of sounds. *EURASIP Journal on Applied Signal Processing*, 2003(7):659–667.

Koch, C., R. Douglas, and U. Wehmeier

1990. Visibility of synaptically induced conductance changes: theory and sim-

ulations of anatomically characterized cortical pyramidal cells. *The Journal of Neuroscience*, 10(6):1728–1744.

Kowalski, N., H. Versnel, and S. A. Shamma

1995. Comparison of responses in the anterior and primary auditory fields of the ferret cortex. *Journal of neurophysiology*, 73(4):1513–1523.

Kuhlman, S. J., E. Tring, and J. T. Trachtenberg

2011. Fast-spiking interneurons have an initial orientation bias that is lost with vision. *Nature Neuroscience*, 14(9):1121–1123.

Kvitsiani, D., S. Ranade, B. Hangya, H. Taniguchi, J. Z. Huang, and A. Kepecs

2013. Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature*, 498(7454):363–366.

Laughlin, S. B. and T. J. Sejnowski

2003. Communication in Neuronal Networks. *Science (New York, N.Y.)*, 301(5641):1870–1874.

Lawrence, J. J. and C. J. McBain

2003. Interneuron diversity series: containing the detonation–feedforward inhibition in the CA3 hippocampus. *Trends in neurosciences*, 26(11):631–640.

LeDoux, J. E. and M. N. Semple

. Redefining the tonotopic core of rat auditory cortex: Physiological evidence for a posterior field - Doron - 2002 - Journal of Comparative Neurology - Wiley Online Library. *Journal of Comparative Neurology, year = 2002 volume = 453, pages = 345–360,*.

Lee, S.-H., A. C. Kwan, and Y. Dan

2014. Interneuron subtypes and orientation tuning. *Nature*, 508(7494):E1–E2.

Lee, S.-H., A. C. Kwan, S. Zhang, V. Phoumthipphavong, J. G. Flannery, S. C. Masmanidis, H. Taniguchi, Z. J. Huang, F. Zhang, E. S. Boyden, K. Deisseroth, and Y. Dan

2012. Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature*, 488(7411):379–383.

Letzkus, J. J., S. B. E. Wolff, E. M. M. Meyer, P. Tovote, J. Courtin, C. Herry, and A. Lüthi

2011. A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature*, 480(7377):331–335.

LeVay, S. and C. D. Gilbert

1976. Laminar patterns of geniculocortical projection in the cat. *Brain Research*, 113(1):1–19.

Lewicki, M. S.

2002. Efficient coding of natural sounds. *Nature neuroscience*, 5(4):356–363.

Li, L.-y., Y.-t. Li, M. Zhou, H. W. Tao, and L. I. Zhang

2013a. Intracortical multiplication of thalamocortical signals in mouse auditory cortex. *Nature Neuroscience*, 16(9):1179–1181.

Li, L.-y., X. R. Xiong, L. A. Ibrahim, W. Yuan, H. W. Tao, and L. I. Zhang

2014. Differential Receptive Field Properties of Parvalbumin and Somatostatin Inhibitory Neurons in Mouse Auditory Cortex. *Cerebral Cortex*, P. doi:10.1093/cercor/bht417.

Li, Y.-t., L. A. Ibrahim, B.-h. Liu, L. I. Zhang, and H. W. Tao

2013b. Linear transformation of thalamocortical input by intracortical excitation. *Nature Neuroscience*, 16(9):1324–1330.

Liang, L., T. Lu, and X. Wang

2002. Neural Representations of Sinusoidal Amplitude and Frequency Modulations in the Primary Auditory Cortex of Awake Primates. *Journal of neurophysiology*, 87(5):2237–2261.

Liberman, M. C. and J. J. Guinan

1998. Feedback control of the auditory periphery: anti-masking effects of middle ear muscles vs. olivocochlear efferents. *Journal of communication disorders*, 31(6):471–483.

Lien, A. D. and M. Scanziani

2013. Tuned thalamic excitation is amplified by visual cortical circuits. *Nature Neuroscience*, 16(9):1315–1323.

Lima, S. Q., T. Hromádka, P. Znamenskiy, and A. M. Zador

2009. PINP: a new method of tagging neuronal populations for identification during in vivo electrophysiological recording. *PloS one*, 4(7):e6099.

Linden, J. F., R. C. Liu, M. Sahani, C. E. Schreiner, and M. M. Merzenich

2003. Spectrotemporal Structure of Receptive Fields in Areas AI and AAF of Mouse Auditory Cortex. *Journal of Neurophysiology*.

Liu, B.-h., Y.-t. Li, W.-P. Ma, C.-j. Pan, L. I. Zhang, and H. W. Tao

2011. Broad inhibition sharpens orientation selectivity by expanding input dynamic range in mouse simple cells. *Neuron*, 71(3):542–554.

Liu, B.-h., G. K. Wu, R. Arbuckle, H. W. Tao, and L. I. Zhang

2007. Defining cortical frequency tuning with recurrent excitatory circuitry. *Nature Neuroscience*, 10(12):1594–1600.

Ma, W.-P., B.-h. Liu, Y.-t. Li, Z. J. Huang, L. I. Zhang, and H. W. Tao

2010. Visual representations by cortical somatostatin inhibitory neurons–

selective but with weak and delayed responses. *The Journal of neuroscience*, 30(43):14371–14379.

Ma, Y., H. Hu, A. S. Berrebi, P. H. Mathers, and A. Agmon

2006. Distinct subtypes of somatostatin-containing neocortical interneurons revealed in transgenic mice. *The Journal of neuroscience*, 26(19):5069–5082.

Machens, C. K., M. S. Wehr, and A. M. Zador

2004. Linearity of cortical receptive fields measured with natural sounds. *The Journal of neuroscience*, 24(5):1089–1100.

MacKay, D. J.

2003. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.

Madisen, L., T. Mao, H. Koch, J.-m. Zhuo, A. Berenyi, S. Fujisawa, Y.-W. A. Hsu, A. J. I. Garcia, X. Gu, S. Zanella, J. Kidney, H. Gu, Y. Mao, B. M. Hooks, E. S. Boyden, G. Buzsaki, J. M. Ramirez, A. R. Jones, K. Svoboda, X. Han, E. E. Turner, and H. Zeng

2012. A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nature Neuroscience*, 15(5):793–802.

Maffei, L., A. Fiorentini, and S. Bisti

1973. Neural correlate of perceptual adaptation to gratings. *Science (New York, N.Y.)*, 182(4116):1036–1038.

Malone, B. J., B. H. Scott, and M. N. Semple

2007. Dynamic Amplitude Coding in the Auditory Cortex of Awake Rhesus Macaques. *Journal of neurophysiology*, 98(3):1451–1474.

Manookin, M. B. and J. B. Demb

2006. Presynaptic mechanism for slow contrast adaptation in mammalian retinal ganglion cells. *Neuron*, 50(3):453–464.

Mante, V., R. A. Frazor, V. Bonin, W. S. Geisler, and M. Carandini

2005. Independence of luminance and contrast in natural scenes and in the early visual system. *Nature neuroscience*, 8(12):1690–1697.

Markram, H., M. Toledo-Rodriguez, Y. Wang, A. Gupta, G. Silberberg, and C. Wu

2004. Interneurons of the neocortical inhibitory system. *Nature reviews. Neuroscience*, 5(10):793–807.

McCormick, D. A., B. W. Connors, J. W. Lighthall, and D. A. Prince

1985. Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of neurophysiology*, 54(4):782–806.

Meyer, A. H., I. Katona, M. Blatow, A. Rozov, and H. Monyer

2002. In vivo labeling of parvalbumin-positive interneurons and analysis of electrical coupling in identified neurons. *The Journal of neuroscience*, 22(16):7055–7064.

Miller, K. D. and T. W. Troyer

2002. Neural noise can explain expansive, power-law nonlinearities in neural response functions. *Journal of neurophysiology*, 87(2):653–659.

Miller, L. M., M. A. Escabí, and C. E. Schreiner

2001. Feature selectivity and interneuronal cooperation in the thalamocortical system. *The Journal of neuroscience*, 21(20):8136–8144.

Mitchell, S. J. and R. A. Silver

2003. Shunting inhibition modulates neuronal gain during synaptic excitation. *Neuron*, 38(3):433–445.

Monier, C., F. Chavane, P. Baudot, L. J. Graham, and Y. Frégnac

2003. Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. *Neuron*, 37(4):663–680.

Moore, A. K. and M. Wehr

2013. Parvalbumin-Expressing Inhibitory Interneurons in Auditory Cortex Are Well-Tuned for Frequency. *The Journal of neuroscience*, 33(34):13713–13723.

Morgan, M. L., G. C. DeAngelis, and D. E. Angelaki

2008. Multisensory integration in macaque visual cortex depends on cue reliability. *Neuron*, 59(4):662–673.

Morrone, M. C., D. C. Burr, and H. D. Speed

1987. Cross-orientation inhibition in cat is GABA mediated. *Experimental Brain Research*, 67(3):635–644.

Moser, T. and D. Beutner

2000. Kinetics of exocytosis and endocytosis at the cochlear inner hair cell afferent synapse of the mouse. *Proceedings of the National Academy of Sciences of the United States of America*, 97(2):883–888.

Mountcastle, V.

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

Movshon, J. A. and P. Lennie

1979. Pattern-selective adaptation in visual cortical neurones. *Nature*, 278(5707):850–852.

Murphy, B. K. and K. D. Miller

2003. Multiplicative gain changes are induced by excitation or inhibition alone. *The Journal of neuroscience*, 23(31):10040–10051.

Nagel, K. I. and A. J. Doupe

2006. Temporal processing and adaptation in the songbird auditory forebrain. *Neuron*, 51(6):845–859.

Nanou, E., A. Kyriakatos, A. Bhattacharjee, L. K. Kaczmarek, G. Paratcha, and A. El Manira

2008. Na+-mediated coupling between AMPA receptors and KNa channels shapes synaptic transmission. *Proceedings of the National Academy of Sciences of the United States of America*, 105(52):20941–20946.

Nelson, P. C. and L. H. Carney

2007. Neural Rate and Timing Cues for Detection and Discrimination of Amplitude-Modulated Tones in the Awake Rabbit Inferior Colliculus. *Journal of neurophysiology*, 97(1):522–539.

Nelson, S. B., C. Hempel, and K. Sugino

2006. Probing the transcriptome of neuronal cell types. *Current Opinion in Neurobiology*, 16(5):571–576.

Nicholson, C. and J. A. Freeman

1975. Theory of current source-density analysis and determination of conductivity tensor for anuran cerebellum. *Journal of neurophysiology*, 38(2):356–368.

Nowak, L. G., M. V. Sanchez-Vives, and D. A. McCormick

2008. Lack of orientation and direction selectivity in a subgroup of fast-spiking inhibitory interneurons: cellular and synaptic mechanisms and comparison with other electrophysiological cell types. *Cerebral cortex (New York, N.Y. : 1991)*, 18(5):1058–1078.

Nuttall, A.

1974. Tympanic muscle effects on middle-ear transfer characteristic. *J. Acoust. Soc. Am*, 56:1239–1247.

Ohshiro, T., D. E. Angelaki, and G. C. DeAngelis

2011. A normalization model of multisensory integration. *Nature Neuroscience*, 14(6):775–782.

Okun, M. and I. Lampl

2008. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature neuroscience*, 11(5):535–537.

Oláh, S., G. Komlósi, J. Szabadics, C. Varga, E. Tóth, P. Barzó, and G. Tamás

2007. Output of neurogliaform cells to various neuron types in the human and rat cerebral cortex. *Frontiers in neural circuits*, 1:4.

Olsen, S. R., V. Bhandawat, and R. I. Wilson

2010. Divisive normalization in olfactory population codes. *Neuron*, 66(2):287–299.

Olsen, S. R., D. S. Bortone, H. Adesnik, and M. Scanziani

2012. Gain control by layer six in cortical circuits of vision. *Nature*, 483(7387):47–52.

Olsen, S. R. and R. I. Wilson

2008. Lateral presynaptic inhibition mediates gain control in an olfactory circuit. *Nature*, 452(7190):956–960.

Ozeki, H., I. M. Finn, E. S. Schaffer, K. D. Miller, and D. Ferster

2009. Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron*, 62(4):578–592.

Pecka, M., Y. Han, E. Sader, and T. D. Mrsic-Flogel

2014. Experience-dependent specialization of receptive field surround for selective coding of natural scenes. *Neuron*, 84(2):457–469.

Peters, A. and B. R. Payne

1993. Numerical Relationships between Geniculocortical Afferents and Pyramidal Cell Modules in Cat Primary Visual Cortex. *Cerebral Cortex*, 3(1):69–78.

Peters, A., B. R. Payne, and J. Budd

1994. A Numerical Analysis of the Geniculocortical Input to Striate Cortex in the Monkey. *Cerebral Cortex*, 4(3):215–229.

Pettersen, K. H., A. Devor, I. Ulbert, A. M. Dale, and G. T. Einevoll

2006. Current-source density estimation based on inversion of electrostatic forward solution: effects of finite extent of neuronal activity and conductivity discontinuities. *Journal of neuroscience methods*, 154(1-2):116–133.

Pfeffer, C. K., M. Xue, M. He, Z. J. Huang, and M. Scanziani

2013. Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nature neuroscience*, 16(8):1068–1076.

Pi, H.-J., B. Hangya, D. Kvitsiani, J. I. Sanders, Z. J. Huang, and A. Kepecs

2013. Cortical interneurons that specialize in disinhibitory control. *Nature*, 503(7477):521–524.

Pinto, D. J., J. C. Brumberg, and D. J. Simons

2000. Circuit dynamics and coding strategies in rodent somatosensory cortex. *Journal of neurophysiology*, 83(3):1158–1166.

Pinto, D. J., J. A. Hartings, J. C. Brumberg, and D. J. Simons

2003. Cortical damping: analysis of thalamocortical response transformations in rodent barrel cortex. *Cerebral Cortex*, 13(1):33–44.

Pitts, W.

1952. Investigations on synaptic transmission. In *Cybernetics, Trans. 9th Conf. Josiah Macy Foundation*, H. von Foerster, ed., Pp. 159–166.

Poggio, T.

2012. The Levels of Understanding framework, revised. *PERCEPTION*, 41(9):1017–1023.

Polack, P.-O., J. Friedman, and P. Golshani

2013. Cellular mechanisms of brain state-dependent gain modulation in visual cortex. *Nature neuroscience*, 16(9):1331–1339.

Polley, D. B., M. A. Heiser, D. T. Blake, C. E. Schreiner, and M. M. Merzenich

2004. Associative learning shapes the neural code for stimulus magnitude in primary auditory cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 101(46):16351–16356.

Pouille, F., A. Marin-Burgin, H. Adesnik, B. V. Atallah, and M. Scanziani

2009. Input normalization by global feedforward inhibition expands cortical dynamic range. *Nature Neuroscience*, 12(12):1577–1585.

Pouille, F. and M. Scanziani

2001. Enforcement of Temporal Fidelity in Pyramidal Cells by Somatic Feed-Forward Inhibition. *Science (New York, N.Y.)*, 293(5532):1159–1163.

Puig, M. V., M. Ushimaru, and Y. Kawaguchi

2008. Two distinct activity patterns of fast-spiking interneurons during neocortical UP states. *Proceedings of the National Academy of Sciences of the United States of America*, 105(24):8428–8433.

Rabinowitz, N. C., B. D. B. Willmore, A. J. King, and J. W. H. Schnupp

2013. Constructing noise-invariant representations of sound in the auditory pathway. *PLoS biology*, 11(11):e1001710.

Rabinowitz, N. C., B. D. B. Willmore, J. W. H. Schnupp, and A. J. King

2011a. Contrast gain control in auditory cortex. *Neuron*, 70(6):1178–1191.

Rabinowitz, N. C., B. D. B. Willmore, J. W. H. Schnupp, and A. J. King

2011b. Contrast Gain Control in Auditory Cortex. *Neuron*, 70(6):1178–1191.

Rabinowitz, N. C., B. D. B. Willmore, J. W. H. Schnupp, and A. J. King

2012. Spectrotemporal contrast kernels for neurons in primary auditory cortex. *The Journal of neuroscience*, 32(33):11271–11284.

Rees, A. and A. R. Møller

1983. Responses of neurons in the inferior colliculus of the rat to AM and FM tones. *Hearing Research*, 10(3):301–330.

Reichardt, W., T. Poggio, and K. Hausen

1983. Figure-ground discrimination by relative movement in the visual system of the fly. *Biological cybernetics*, 46(1):1–30.

Reynolds, J. H. and R. Desimone

2003. Interacting roles of attention and visual salience in V4. *Neuron*, 37(5):853–863.

Reynolds, J. H. and D. J. Heeger

2009. The normalization model of attention. *Neuron*, 61(2):168–185.

Rieke, F., D. A. Bodnar, and W. Bialek

1995. Naturalistic Stimuli Increase the Rate and Efficiency of Information Transmission by Primary Auditory Afferents. *Proceedings. Biological sciences / The Royal Society*, 262(1365):259–265.

Ringach, D. L.

2009. Spontaneous and driven cortical activity: implications for computation. *Current Opinion in Neurobiology*, 19(4):439–444.

Ringach, D. L.

2010. Population coding under normalization. *Vision research*, 50(22):2223–2232.

Ringach, D. L. and B. J. Malone

2007. The operating point of the cortex: neurons as large deviation detectors. *The Journal of neuroscience*, 27(29):7673–7683.

Rose, H. J. and R. Metherate

2005. Auditory Thalamocortical Transmission Is Reliable and Temporally Precise. *Journal of neurophysiology*, 94(3):2019–2030.

Rothman, J. S., L. Cathala, V. Steuber, and R. A. Silver

2009. Synaptic depression enables neuronal gain control. *Nature*, 457(7232):1015–1018.

Rothschild, G., I. Nelken, and A. Mizrahi

2010. Functional organization and population dynamics in the mouse primary auditory cortex. *Nature Neuroscience*, 13(3):353–360.

Rudy, B., G. Fishell, S. Lee, and J. Hjerling-Leffler

2011. Three groups of interneurons account for nearly 100% of neocortical GABAergic neurons. *Developmental neurobiology*, 71(1):45–61.

Ruggero, M. A. and N. C. Rich

1991. Furosemide alters organ of corti mechanics: evidence for feedback of outer hair cells upon the basilar membrane. *The Journal of Neuroscience*, 11(4):1057–1067.

Runyan, C. A., J. Schummers, A. Van Wart, S. J. Kuhlman, N. R. Wilson, Z. J. Huang, and M. Sur

2010. Response features of parvalbumin-expressing interneurons suggest precise roles for subtypes of inhibition in visual cortex. *Neuron*, 67(5):847–857.

Rust, N. C., O. Schwartz, J. A. Movshon, and E. P. Simoncelli

2005. Spatiotemporal elements of macaque v1 receptive fields. *Neuron*, 46(6):945–956.

Rutkowski, R. G., T. M. Shackleton, J. W. H. Schnupp, M. N. Wallace, and A. R. Palmer

2002. Spectrotemporal receptive field properties of single units in the primary, dorsocaudal and ventrorostral auditory cortex of the guinea pig. *Audiology & neuro-otology*, 7(4):214–227.

Ryugo DK, Willard FH, F. D.

1981. Differential afferent projections to the inferior colliculus from the cochlear nucleus in the albino mouse. *Brain Res.*, 210:342–9.

Sachs, M. and P. Abbas

1974. Rate versus level functions for auditory-nerve fibers in cats: tone-burst stimuli. *The Journal of the Acoustical Society of America*, 56:1835?47.

Sadagopan, S. and D. Ferster

2012. Feedforward origins of response variability underlying contrast invariant orientation tuning in cat visual cortex. *Neuron*, 74(5):911–923.

Saganich, M. J., E. Machado, and B. Rudy

2001. Differential expression of genes encoding subthreshold-operating voltage-gated K+ channels in brain. *The Journal of neuroscience*, 21(13):4609–4624.

Sahani, M. and J. F. Linden

2003. Evidence optimization techniques for estimating stimulus-response functions. *Advances in neural information processing*, 15:317–324.

Sakata, S. and K. D. Harris

2009. Laminar Structure of Spontaneous and Sensory-Evoked Population Activity in Auditory Cortex. *Neuron*, 64(3):404–418.

Sanchez-Vives, M. V., L. G. Nowak, and D. A. McCormick

2000a. Cellular mechanisms of long-lasting adaptation in visual cortical neurons in vitro. *The Journal of neuroscience*, 20(11):4286–4299.

Sanchez-Vives, M. V., L. G. Nowak, and D. A. McCormick

2000b. Membrane mechanisms underlying contrast adaptation in cat area 17 in vivo. *The Journal of neuroscience*, 20(11):4267–4285.

Saviane, C. and R. A. Silver

2006. Fast vesicle reloading and a large pool sustain high bandwidth transmission at a central synapse. *Nature*, 439(7079):983–987.

Schnupp, J. W. H., J. A. Garcia-Lazaro, and N. A. Lesica

2015. Periodotopy in the gerbil inferior colliculus: local clustering rather than a gradient map. *Frontiers in neural circuits*, 9:37.

Schnupp, J. W. H., T. D. Mrsic-Flogel, and A. J. King

2001. Linear processing of spatial cues in primary auditory cortex. *Nature*, 414(6860):200–204.

Schroeder, C. E., A. D. Mehta, and S. J. Givre

1998. A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. *Cerebral Cortex*, 8(7):575–592.

Schweizer, H.

1981. The connections of the inferior colliculus and the organization of the brainstem auditory system in the greater horseshoe bat (Rhinolophus ferrumequinum). *Journal of Comparative Neurology*, 201(1):25–49.

Schwindt, P. C., W. J. Spain, and W. E. Crill

1989. Long-lasting reduction of excitability by a sodium-dependent potassium current in cat neocortical neurons. *Journal of neurophysiology*, 61(2):233–244.

Sclar, G.

1987. Expression of "retinal" contrast gain control by neurons of the cat's lateral geniculate nucleus. *Experimental Brain Research*, 66(3):589–596.

Sclar, G., J. H. R. Maunsell, and P. Lennie

1990. Coding of image contrast in central visual pathways of the macaque monkey. *Vision research*, 30(1):1–10.

Sessolo, M., I. Marcon, S. Bovetti, G. Losi, M. Cammarota, G. M. Ratto, T. Fellin, and G. Carmignoto

2015. Parvalbumin-Positive Inhibitory Interneurons Oppose Propagation But Favor Generation of Focal Epileptiform Activity. *The Journal of neuroscience*, 35(26):9544–9557.

Seybold, B. A., E. A. K. Phillips, C. E. Schreiner, and A. R. Hasenstaub

2015. Inhibitory Actions Unified by Network Integration. *Neuron*, 87(6):1181–1192.

Shamma, S. A., J. W. Fleshman, P. R. Wiser, and H. Versnel

1993. Organization of response areas in ferret primary auditory cortex. *Journal of neurophysiology*, 69(2):367–383.

Shannon, C.

1948. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656.

Shapley, R. and C. Enroth-Cugell

1984. Chapter 9 Visual adaptation and retinal gain controls. *Progress in retinal research*, 3:263–346.

Shapley, R. M. and J. D. Victor

1978. The effect of contrast on the transfer properties of cat retinal ganglion cells. *The Journal of physiology*, 285(1):275–298.

Shapley, R. M. and J. D. Victor

1981. How the contrast gain control modifies the frequency responses of cat retinal ganglion cells. *The Journal of physiology*, 318:161–179.

Shapley, R. M. and D. Xing

2013. Local circuit inhibition in the cerebral cortex as the source of gain control and untuned suppression. *Neural networks : the official journal of the International Neural Network Society*, 37:172–181.

Sharpee, T. O., H. Sugihara, A. V. Kurgansky, S. P. Rebrik, M. P. Stryker, and K. D. Miller

2006. Adaptive filtering enhances information transmission in visual cortex. *Nature*, 439(7079):936–942.

Simmons, J. A., E. G. Wever, and J. M. Pylka

1971. Periodical cicada: sound production and hearing. *Science (New York, N.Y.)*, 171(3967):212–213.

Simon, J. Z., D. A. Depireux, D. J. Klein, J. B. Fritz, and S. A. Shamma

2007. Temporal symmetry in primary auditory cortex: implications for cortical connectivity. *Neural Computation*, 19(3):583–638.

Simoncelli, E. P. and D. J. Heeger

1998. A model of neuronal responses in visual area MT. *Vision research*, 38(5):743–761.

Simoncelli, E. P. and B. A. Olshausen

2001. Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24(1):1193–1216.

Smirnakis, S. M., M. J. Berry, D. K. Warland, W. Bialek, and M. Meister

1997. Adaptation of retinal processing to image contrast and spatial scale. *Nature*, 386(6620):69–73.

Sohal, V. S., F. Zhang, O. Yizhar, and K. Deisseroth

2009a. Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, 459(7247):698–702.

Sohal, V. S., F. Zhang, O. Yizhar, and K. Deisseroth

2009b. Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, 459(7247):698–702.

Solomon, S. G., J. W. Peirce, N. T. Dhruv, and P. Lennie

2004. Profound contrast adaptation early in the visual pathway. *Neuron*, 42(1):155–162.

Somogyi, P., G. Tamás, R. Lujan, and E. H. Buhl

1998. Salient features of synaptic organisation in the cerebral cortex. *Brain research. Brain research reviews*, 26(2-3):113–135.

Sousa, V. H., G. Miyoshi, J. Hjerling-Leffler, T. Karayannis, and G. Fishell

2009. Characterization of Nkx6-2-derived neocortical interneuron lineages. *Cerebral cortex (New York, N.Y. : 1991)*, 19 Suppl 1(suppl 1):i1–10.

Spassova, M. A., M. Avissar, A. C. Furman, M. A. Crumling, J. C. Saunders, and T. D. Parsons

2004. Evidence that rapid vesicle replenishment of the synaptic ribbon mediates recovery from short-term adaptation at the hair cell afferent synapse. *Journal of the Association for Research in Otolaryngology : JARO*, 5(4):376–390.

Sternson, S. M. and B. L. Roth

2014. Chemogenetic tools to interrogate brain functions. *Annual Review of Neuroscience*, 37:387–407.

Stratford, K. J., K. Tarczy-Hornoch, K. A. Martin, N. J. Bannister, and J. J. Jack

1996. Excitatory synaptic inputs to spiny stellate cells in cat visual cortex. *Nature*, 382(6588):258–261.

Sugino, K., C. M. Hempel, M. N. Miller, A. M. Hattox, P. Shapiro, C. Wu, Z. J. Huang, and S. B. Nelson

2006. Molecular taxonomy of major neuronal classes in the adult mouse forebrain. *Nature neuroscience*, 9(1):99–107.

Sun, Y. J., G. K. Wu, B.-h. Liu, P. Li, M. Zhou, Z. Xiao, H. W. Tao, and L. I. Zhang

2010. Fine-tuning of pre-balanced excitation and inhibition during auditory cortical development. *Nature*, 465(7300):927–931.

Szymanski, F. D., J. A. Garcia-Lazaro, and J. W. H. Schnupp

2009. Current source density profiles of stimulus-specific adaptation in rat auditory cortex. *Journal of neurophysiology*, 102(3):1483–1490.

Tamás, G., E. H. Buhl, A. Lörincz, and P. Somogyi

2000. Proximally targeted GABAergic synapses and gap junctions synchronize cortical interneurons. *Nature neuroscience*, 3(4):366–371.

Taniguchi, H., M. He, P. Wu, S. Kim, R. Paik, K. Sugino, D. Kvitsiani, D. Kvitsani, Y. Fu, J. Lu, Y. Lin, G. Miyoshi, Y. Shima, G. Fishell, S. B. Nelson, and Z. J. Huang

2011. A resource of Cre driver lines for genetic targeting of GABAergic neurons in cerebral cortex. *Neuron*, 71(6):995–1013.

Thomson, A. M. and C. Lamy

2007. Functional maps of neocortical local circuitry. *Frontiers in Neuroscience*, 1(1):19–42.

Traub, R. D., A. Bibbig, F. E. N. LeBeau, E. H. Buhl, and M. A. Whittington

2004. Cellular mechanisms of neuronal population oscillations in the hippocampus in vitro. *Annual Review of Neuroscience*, 27:247–278.

Troyer, T. W., A. E. Krukowski, N. J. Priebe, and K. D. Miller

1998. Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intracortical connectivity. *The Journal of Neuroscience*, 18(15):5908–5927.

Trussell, L. O. and G. D. Fischbach

1989. Glutamate receptor desensitization and its role in synaptic transmission. *Neuron*, 3(2):209–218.

Tsodyks, M. V. and H. Markram

1997. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences of the United States of America*, 94(2):719–723.

Uematsu, M., Y. Hirai, F. Karube, S. Ebihara, M. Kato, K. Abe, K. Obata, S. Yoshida, M. Hirabayashi, Y. Yanagawa, and Y. Kawaguchi

2008. Quantitative chemical composition of cortical GABAergic neurons revealed in transgenic venus-expressing rats. *Cerebral cortex (New York, N.Y. : 1991)*, 18(2):315–330.

Ulrich, D.

2003. Differential arithmetic of shunting inhibition for voltage and spike rate in neocortical pyramidal cells. *The European journal of neuroscience*, 18(8):2159–2165.

Varela, J. A., K. Sen, J. Gibson, J. Fost, L. F. Abbott, and S. B. Nelson

1997. A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *The Journal of Neuroscience*, 17(20):7926–7940.

Varga, V., A. Losonczy, B. V. Zemelman, Z. Borhegyi, G. Nyiri, A. Domonkos, B. Hangya, N. Holderith, J. C. Magee, and T. F. Freund

2009. Fast synaptic subcortical control of hippocampal circuits. *Science (New York, N.Y.)*, 326(5951):449–453.

Viemeister, N. F.

1974. Intensity discrimination of noise in the presence of band-reject noise. *The Journal of the Acoustical Society of America*, 56(5):1594–1600.

Viemeister, N. F.

1983. Auditory intensity discrimination at high frequencies in the presence of noise. *Science (New York, N.Y.)*, 221(4616):1206–1208.

Viemeister, N. F.

1988. Intensity coding and the dynamic range problem. *Hearing Research*, 34(3):267–274.

Vruwink, M., H. H. Schmidt, R. J. Weinberg, and A. Burette

2001. Substance P and nitric oxide signaling in cerebral cortex: anatomical evidence for reciprocal signaling between two classes of interneurons. *The Journal of comparative neurology*, 441(4):288–301.

Vucurovic, K., T. Gallopin, I. Férézou, A. Rancillac, P. Chameau, J. A. van Hooft, H. Geoffroy, H. Monyer, J. Rossier, and T. Vitalis

2010. Serotonin 3A receptor subtype as an early and protracted marker of cortical interneuron subpopulations. *Cerebral cortex (New York, N.Y. : 1991)*, 20(10):2333–2347.

Wang, Y., A. Gupta, M. Toledo-Rodriguez, C. Z. Wu, and H. Markram

2002. Anatomical, physiological, molecular and circuit properties of nest basket cells in the developing somatosensory cortex. *Cerebral Cortex*, 12(4):395–410.

Wang, Y., M. Toledo-Rodriguez, A. Gupta, C. Wu, G. Silberberg, J. Luo, and H. Markram

2004. Anatomical, physiological and molecular properties of Martinotti cells in the somatosensory cortex of the juvenile rat. *The Journal of physiology*, 561(Pt 1):65–90.

Wehr, M. and A. M. Zador

2003. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature*, 426(6965):442–446.

Wen, B., G. I. Wang, I. Dean, and B. Delgutte

2009. Dynamic range adaptation to sound level statistics in the auditory nerve. *The Journal of neuroscience*, 29(44):13797–13808.

Wen, B., G. I. Wang, I. Dean, and B. Delgutte

2012. Time course of dynamic range adaptation in the auditory nerve. *Journal of neurophysiology*, 108(1):69–82.

Westerman, L. A. and R. L. Smith

1984. Rapid and short-term adaptation in auditory nerve responses. *Hearing Research*, 15(3):249–260.

White, E. L. and A. Keller

1989. Cortical circuits: synaptic organization of the cerebral cortex: structure, function, and theory.

Williams, S. R. and S. J. Mitchell

2008. Direct measurement of somatic voltage clamp errors in central neurons. *Nature neuroscience*, 11(7):790–798.

Willmore, B. D. B., J. E. Cooke, and A. J. King

2014. Hearing in noisy environments: noise invariance and contrast gain control. *The Journal of physiology*, 592(Pt 16):3371–3381.

Wilson, N. R., C. A. Runyan, F. L. Wang, and M. Sur

2012. Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature*, 488(7411):343–348.

Winer, J. and C. Schreiner

2005. *The Inferior Colliculus*. Springer, New York.

Winer, J. A., L. M. Miller, C. C. Lee, and C. E. Schreiner

2005. Auditory thalamocortical transformation: structure and function. *Trends in neurosciences*, 28(5):255–263.

Winkowski, D. E. and P. O. Kanold

2013. Laminar transformation of frequency organization in auditory cortex. *The Journal of neuroscience*, 33(4):1498–1508.

Woodruff, A., Q. Xu, S. A. Anderson, and R. Yuste

2009. Depolarizing effect of neocortical chandelier neurons. *Frontiers in neural circuits*, 3:15.

Wu, G. K., R. Arbuckle, B.-h. Liu, H. W. Tao, and L. I. Zhang

2008. Lateral sharpening of cortical frequency tuning by approximately balanced inhibition. *Neuron*, 58(1):132–143.

Wu, G. K., H. W. Tao, and L. I. Zhang

2011. From elementary synaptic circuits to information processing in primary auditory cortex. *Neuroscience and biobehavioral reviews*, 35(10):2094–2104.

Xiao, F, D. J. C. P. W. B.

2002. High dynamic range imaging of natural scenes. In *Tenth Color Imaging Conference: Color Science, Systems, and Applications*, Pp. 337–342.

Xiao, F., J. M. DiCarlo, P. B. Catrysse, and B. A. Wandell

2002. High Dynamic Range Imaging of Natural Scenes. *Color and Imaging . . . .*

Xu, X. and E. M. Callaway

2009. Laminar specificity of functional input to distinct types of inhibitory cortical neurons. *The Journal of neuroscience*, 29(1):70–85.

Xu, X., K. D. Roby, and E. M. Callaway

2006. Mouse cortical inhibitory neuron type that coexpresses somatostatin and calretinin. *Journal of Comparative Neurology*, 499(1):144–160.

Xue, M., B. V. Atallah, and M. Scanziani

2014. Equalizing excitation-inhibition ratios across visual cortical neurons. *Nature*, 511(7511):596–600.

Yates, G.

1995. Cochlear structure and function. In *Hearing*, B. Moore, ed., Pp. 41–74. Academic Press, San Diego.

Yates, GK, R. D. J. B.

1985. Very rapid adaptation in the guinea pig auditory nerve. *Hear. Res.*, 17:1–12.

Yin, P., J. B. Fritz, and S. A. Shamma

2014. Rapid spectrotemporal plasticity in primary auditory cortex during behavior. *The Journal of neuroscience*, 34(12):4396–4408.

Yoshimura, Y. and E. M. Callaway

2005. Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nature neuroscience*, 8(11):1552–1559.

Zaghloul, K. A., K. Boahen, and J. B. Demb

2005. Contrast adaptation in subthreshold and spiking responses of mammalian Y-type retinal ganglion cells. *The Journal of neuroscience*, 25(4):860–868.

Zariwala, H. A., L. Madisen, K. F. Ahrens, A. Bernard, E. S. Lein, A. R. Jones, and H. Zeng

2011. Visual tuning properties of genetically identified layer 2/3 neuronal types in the primary visual cortex of cre-transgenic mice. *Frontiers in systems neuroscience*, 4:162.

Zhang, F., L.-P. Wang, M. Brauner, J. F. Liewald, K. Kay, N. Watzke, P. G. Wood,

E. Bamberg, G. Nagel, A. Gottschalk, and K. Deisseroth

2007. Multimodal fast optical interrogation of neural circuitry. *Nature*, 446(7136):633–639.

Zhou, M., F. Liang, X. R. Xiong, L. Li, H. Li, Z. Xiao, H. W. Tao, and L. I. Zhang

2014a. Scaling down of balanced excitation and inhibition by active behavioral states in auditory cortex. *Nature Neuroscience*, 17(6):841–850.

Zhou, M., F. Liang, X. R. Xiong, L. Li, H. Li, Z. Xiao, H. W. Tao, and L. I. Zhang

2014b. Scaling down of balanced excitation and inhibition by active behavioral states in auditory cortex. *Nature neuroscience*, 17(6):841–850.

Znamenskiy, P. and A. M. Zador

2013. Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature*, 497(7450):482–485.

Zoccolan, D., D. D. Cox, and J. J. DiCarlo

2005. Multiple object response normalization in monkey inferotemporal cortex. *The Journal of neuroscience*, 25(36):8150–8164.