NEURAL MECHANISMS AND FUNCTIONAL SIGNFICANCE OF PERI-SACCADIC RESPONSE MODULATION

by

JONATHAN SIMON GUEZ

A Dissertation submitted to the

Graduate School-Newark

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Behavioral and Neural Sciences

written under the direction of

Bart Krekelberg

and approved by

Dr. Jorge Golowasch _____

Dr. Bruce Katz _____

Dr. Denis Paré _____

Dr. Horacio Rotstein _____

Dr. Bart Krekelberg _____

Newark, New Jersey

October, 2015

ABSTRACT OF THE THESIS

NEURAL MECHANISMS AND FUNCTIONAL SIGNFICANCE OF PERI-SACCADIC RESPONSE MODULATION

by

JONATHAN SIMON GUEZ

Thesis Director:

Bart Krekelberg

Active vision involves fast eye movements (saccades) with brief inter-saccadic fixations. This presents two interesting problems. (1) During each saccade the moving eye creates motion signals on the retina, and yet we are unaware of this sweeping visual input. (2) Upon each brief fixation the visual system is input with a new scene, and tasked to quickly encode many stimulus features. Regarding the first problem, a decrease in contrast sensitivity during saccades (saccadic suppression) is thought to contribute to our lack of intra-saccadic perception. Chapters 2 & 3 of this thesis seek to further understand the neural mechanisms of saccadic suppression. Regarding the second problem, a post-saccadic change in neural activity is thought to specialize processing of newly fixated stimuli. Chapter 4 investigates the changes in visual response properties that occur after a saccade.

Chapter 2 used a signal-detection model that describes the psychophysical phenomenon of saccadic suppression in computational terms. The model is built up of visual detectors

in which gain, noise, and spatial uncertainty can be varied. Thus saccadic suppression is recast in terms that provide testable predictions of neural activities. We found that saccadic suppression is the result of reduced detector gain. Chapter 3 studied neural responses from permanently implanted multi-electrode arrays in V1 of macaques. Based on our Chapter 2 results, we looked specifically at how contrast responses change during saccades. We found saccadic gain reduction that begins before saccade onset, suggesting that V1 is a neural site of suppression, and that saccadic gain reduction is the result of a corollary discharge signal. Chapter 4 studied the effect of post-saccadic modulation on contrast response properties. Using model fitting and signal detection measures, we showed that post-saccadic modulations in V1 lead to an increased range of discriminable contrasts. We argue that this increased operating range gives a functional benefit when encountering newly fixated stimuli. Chapter 5 concludes by relating the saccadic gain reduction shown in Chapters 2 & 3 to the post-saccadic response changes shown in Chapter 4 – arguing that saccadic suppression can be viewed as part of a more general process to improve post-saccadic vision.

Table of Contents

List of Illustrations

Chapter 1. **Introduction**

Human vision is active: In order to maximize the amount of relevant information gleaned from the visual environment, the high-acuity central part of the retina (fovea) must frequently be redirected. To this end primates make fast, ballistic eye movements (saccades) about three times per second. During each saccade the eye is quickly moved across a visual scene and this generates spurious motion signals on the retina. This leads to the question: How are we unaware of the sweeping visual input that the moving retina induces during saccades? Chapters 2 and 3 of this dissertation address a related question – of how visual input during saccades is poorly detected (saccadic suppression) – by studying how the brain processes stimuli differently during saccades.

The neurons that constitute early visual processing have receptive fields that move with the retina, and so after each saccade (i.e., upon each new fixation) they are input with a different set of stimuli. When we consider the wide variation of stimuli that the visual system is tasked to encode during brief inter-saccadic fixations, and the perceptual biases imposed by stimuli at the previous fixation, we may ask: Does the visual system make use of the knowledge that upon a new fixation it will encounter new stimuli? In other words does the visual system process stimuli differently – for example by resetting any adaptation from the previous fixation – at the onset of a new fixation? Chapter 4 addresses this second question by studying how the brain processes stimuli differently after saccades.

This chapter will introduce the notion of visual stability as it is discussed in much of the scientific literature on saccadic suppression. It will follow with an overview of saccadic suppression in order to motivate the specific approach taken in Chapter 2, in which we apply a signal-processing framework to better understand the psychophysical phenomenon that is saccadic suppression. This introduction will also provide some background for Chapter 3, which will show saccadic modulation of contrast responses in the macaque primary visual cortex (V1). More background on studies of suppression in V1, and other visual areas, is given in the introduction of Chapter 3, and elaborated in detail in its discussion.

1.1 Visual Stability

Each saccade causes a blurred motion signal because the eye quickly moves across the visual scene. Also, after each saccade (i.e., upon a new fixation) the image on the retina is displaced from the previous fixation. Despite these two facts, we generally perceive a stable visual world – that is we do not perceive full-field motion during each eye movement, nor do we perceive that the visual world has been displaced after each eye movement. The first of these two perceptual phenomena (no awareness of intra-saccadic motion) is referred to as saccadic omission (Campbell & Wurtz, 1978); it is discussed in later sections along with the related psychophysical phenomenon of saccadic suppression. The second phenomenon – perception of a stable visual world despite the displacement of the image from one fixation to another – will be referred to as saccadic displacement (Wurtz, 2008), and discussed here briefly.

As we move our eyes around we perceive each fixation as a snapshot, integrated in a larger, persisting (i.e., stable) visual world. This suggests that at some level of neural

processing, information about the image on the retina is integrated with information about where the eye is (or will be) pointing. How is this achieved? One hypothesis is that information about the change in eye position (a saccadic displacement vector) provides a reference across saccades of how each fixational snapshot is spatially related within a more global, spatial map. Another hypothesis is that some neurons code in 'real' position coordinates: rather than having receptive fields located in retinal space that move with each saccade these neurons would have receptive fields located in a world centered reference frame. (For a good review of the origins and experiments addressing these hypotheses see Wurtz, 2008). Finally it has been proposed that eye position could modulate responses of neurons in the visual sensory pathway. These 'eye position gain fields' have been shown to carry enough information distributed across multiple neurons to allow eye position to be decoded from neurons with receptive fields located in retinal coordinates (Morris, Bremmer, & Krekelberg, 2013). This suggests that a neural population could simultaneously compute both the position of the eye and the position of the image (on the eye) – and so spatially stable vision could be the result of such a distributed neural computation.

Again, visual stability has been studied as two separate questions – saccadic displacement (described above) and saccadic omission. This thesis is more concerned with saccadic omission and the related psychophysical phenomena of saccadic suppression.

1.2 Saccadic Suppression

How is it that we do not see the visual scene move each time our eyes make a saccade across it? The question has been in the scientific literature for over a century, and led to decades of experiments investigating changes in visual perception during saccades (see

Matin, 1974 for a review of early work). An example is an experiment by Volkmann et al. (1968) that showed a decrease in the detection of light flashes when they are presented during saccades; this decrease in sensitivity during saccades is called saccadic suppression. A more recent experiment by Diamond et al. (2000) nicely illustrated the time course of saccadic suppression. Figure 1.1 (orange squares, black trace) shows that contrast sensitivity appears to decrease around 75 ms before the onset of a saccade and reaches maximal suppression at the time of saccade onset (0 ms). An important aspect of suppression studies is to disentangle the reduced visibility due to the motion-induced smearing of the stimulus (this will be referred to as a passive, retinal, or visual component) from a potential active suppressive component. Diamond et al. (2000) controlled for potential visual (or passive) components of saccadic suppression by comparing detection performance during saccades, with equivalent retinal stimulation during fixation (achieved by swiveling a mirror reflecting the stimuli at saccadic velocity). The control for visual factors, along with the pre-saccadic onset of suppression in their study, suggested a non-visual (or active) control signal involved in the perceptual changes during saccades. Subsequent electro-physiological work proved that a signal arrives in MT (and other extra-striate areas) before saccade onset, and has a modulatory effect on visual input (Bremmer et al., 2009). By flashing stimuli at different times relative to a saccade and recording neural responses in extra-striate areas (Figure 1.1, colored traces), Bremmer et al (2009) were able to show that there is a decline in neuronal excitability before the saccade begins. This is solid evidence that some kind of signal pre-empts the saccade and arrives in certain visual areas where it modulates responses to stimulation during saccades. The origin of this signal is discussed in detail

in Chapter 3 (section 3.4). I will refer to it throughout this dissertation as an extra-retinal (or non-visual) signal, an active mechanism, or a corollary discharge – depending on the context. Corollary discharge (CD) refers to a type of active signal that originates in areas mediating motor control and reaches sensory areas. It is sent in order to cancel, suppress, or somehow modify the sensory stimulation resulting from self-motion. In our case that self-motion is the retina moving as a result of the saccade. The possible regions supplying a CD are discussed in more detail in Chapter 3 (section 3.4)

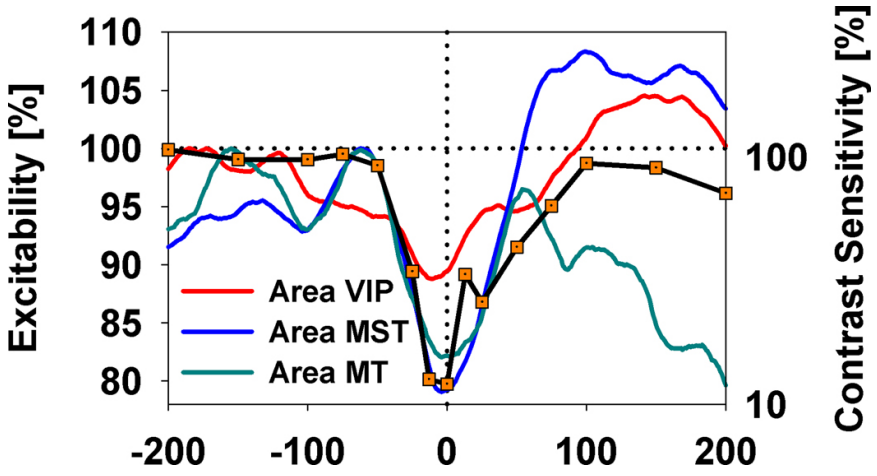

Figure 1.1 – Time course of contrast sensitivity and extra-striate activity during saccades. For the behavioral time course (orange squares, black trace; ordinate on right) we see contrast sensitivity in a detection task at different stimulus presentation times, aligned to saccade onset. Contrast sensitivity appears to decline ~75 ms before saccade onset, reaching peak suppression at onset (0 ms). (Behavioral data taken from Diamond, Ross, & Morrone, 2000.) Colored traces show excitability of extra-striate areas relative to saccade onset. Stimuli were flashed at different times relative to saccade, and population responses averaged from recordings made in the ventral intraparietal (VIP), medial

superior temporal (MST), and middle temporal (MT) areas. (Reprinted from Bremmer, Kubischik, Hoffman, & Krekelberg, 2009)

Figure 1.1 shows a match between the time course of peri-saccadic neural modulation and that of (psychophysical) saccadic suppression. This is a promising lead in terms of studying the neural origins of suppression. But there is an explanatory gap: how does this change in excitability in MT (or any other area) lead to a loss of contrast sensitivity behaviorally? One good approach to this question is to model the psychophysical phenomenon in terms more easily matched to neural measures (e.g., input gain, multiplicative noise) in order to help make predictions. Chapter 2 does this by describing saccadic suppression in a signal-processing framework. The result is that saccadic suppression is best described as a gain reduction. The electro-physiological experiments in Chapters 3 and 4 – in studying how neural responses changes during and after saccades – make use of this knowledge by focusing extensively on changes in the contrast gain functions of neurons.

Chapter 2. Saccadic suppression in a signal-processing framework

2.1 Introduction

Saccades are the fast and frequent eye movements that we make in order to direct the high-resolution fovea at regions of visual interest. We remain almost completely unaware of the sweeping visual input that the moving retina induces during saccades – a perceptual phenomenon termed saccadic omission (Campbell and Wurtz, 1978). This omission of self-induced retinal motion is crucial for perceptual stability (Wurtz, 2008). Saccadic suppression – defined as a decrease in visual sensitivity in the ~75 ms (Ross et al., 2001) leading up to and during saccades – is thought to contribute to this omission. This study focused on decreased sensitivity to visual stimuli presented during the saccadic eye movement (intra-saccadic suppression).

The computational and neural mechanisms that underlie intra-saccadic suppression remain debated. Some researchers emphasize a retinal origin (Richards, 1969; Castet et al., 2001), others an extra-retinal, or central, origin (Wurtz and Goldberg, 1972; Ibbotson et al., 2008), or the combination of central and retinal mechanisms (Ibbotson and Krekelberg, 2011). One long-standing theory states that brain areas mediating eye movements send a corollary discharge that modulates activity in visual brain areas during saccades (for review see Ibbotson and Krekelberg, 2011). An alternative view holds that spatial uncertainty induced by saccades can explain suppression (Matin, 1974; Greenhouse and Cohn, 1991). Which of these mechanisms is

dominant and how these mechanisms interact (if at all) during intra-saccadic suppression has not been established.

We expanded on previous work (Watson and Krekelberg, 2011) that focused on *pre*-saccadic suppression, by applying Lu and Dosher's (Lu and Dosher, 1998a) perceptual template model (PTM) to intra-saccadic suppression. The PTM describes a visual detection process in stages. A template stage attempts to filter out irrelevant sensory inputs, such as those arising from locations in space or time that do not contain signal; a gain stage scales all inputs (i.e., signal and noise); a multiplicative noise-injection stage adds noise in a stimulus-dependent manner; an additive noise-injection stage adds noise in a stimulus-independent manner. Specifically, the PTM allows us to establish whether intra-saccadic suppression arises as a consequence of a change in the template stage, a decrease in gain, or an injection of noise into the visual system.

Watson and Krekelberg (2011) have shown that saccadic suppression of stimuli presented just *before* a saccade is best explained by a gain reduction mechanism. Here we applied the PTM to intra-saccadic suppression. We expected that an uncertainty mechanism could play a larger role than it does in pre-saccadic suppression. This expectation came from known factors that could lead to spatial uncertainty during saccades, such as peri-saccadic changes in receptive fields (Duhamel et al., 1992; Tolias et al., 2001; Krekelberg et al., 2003), peri-saccadic mislocalization (Honda, 1989; Lappe et al., 2000; Binda et al., 2009), and a temporarily inaccurate internal representation of eye position (Honda, 1989; Dassonville et al., 1992; Morris et al., 2012). Contrary to our hypothesis, the results showed that gain reduction is the dominant mechanism in intra-saccadic suppression.

2.2 Methods

2.2.1 Participants

A total of eight subjects (ages ranging from 20 to 30) participated in three experiments ("wide-signal", "narrow-signal", and "high-noise"). All subjects had normal or corrected-to-normal vision. All subjects, except one author (S1), were naïve as to the experiment's purpose. Four subjects (two females) participated in the wide-signal experiment. Four subjects (one female) participated in the narrow-signal experiment. Four subjects (one female) participated in the high-noise experiment. The author (S1) participated in all three experiments. The experiments were in compliance with the protection of human subjects as outlined in the Declaration of Helsinki and were approved by Rutgers University's Institutional Review Board.

2.2.2 Stimuli

One challenge in intra-saccadic experiments is presenting a stimulus such that it will appear the same to a fixating and a moving retina. The effect of a (world-referenced) stationary grating moving across a translating retina is called smear. While smear certainly contributes to lowered visibility during saccades, we were interested in the internal mechanisms leading to suppression. To minimize the effects of retinal smear, our task involved a horizontal saccade across a horizontal grating (Volkmann et al., 1978; Burr et al., 1994). Another factor that contributes to lowered visibility during saccades in everyday vision is forward and backward masking of the intra-saccadic scene by the structured pre- and post-saccadic scenes. By presenting a uniform gray background before and after the stimulus we minimized this influence.

The stimulus was a horizontal grating (to be detected) embedded within a noise pattern. Stimuli were presented on a 30 x 40 cm Sony FD Trinitron (GDM-C520) CRT monitor with a resolution of 1024 x 768 pixels and a refresh rate of 120 Hz. The target grating was oriented horizontally, had a spatial frequency of 0.1 cycles per degree, and was vignetted by a Gaussian contrast envelope with a standard deviation of 2° in the vertical direction (Figure 2.1b). The peak of the Gaussian envelope was located 4° above or below fixation. The contrast of the sine-wave component was varied from trial-to-trial according to a Bayesian-adaptive method (Kontsevich and Tyler, 1999) to optimize threshold estimation. The stimulus was flashed for one frame (8ms).

The external noise pattern consisted of horizontal bars. The noise pattern was added to both the target grating and the background – extending across the entire monitor (Figure 2.1b). Each bar's luminance offset was chosen from a normal probability distribution with a mean of zero and a standard deviation (σ_e) expressed in terms of percent background luminance.

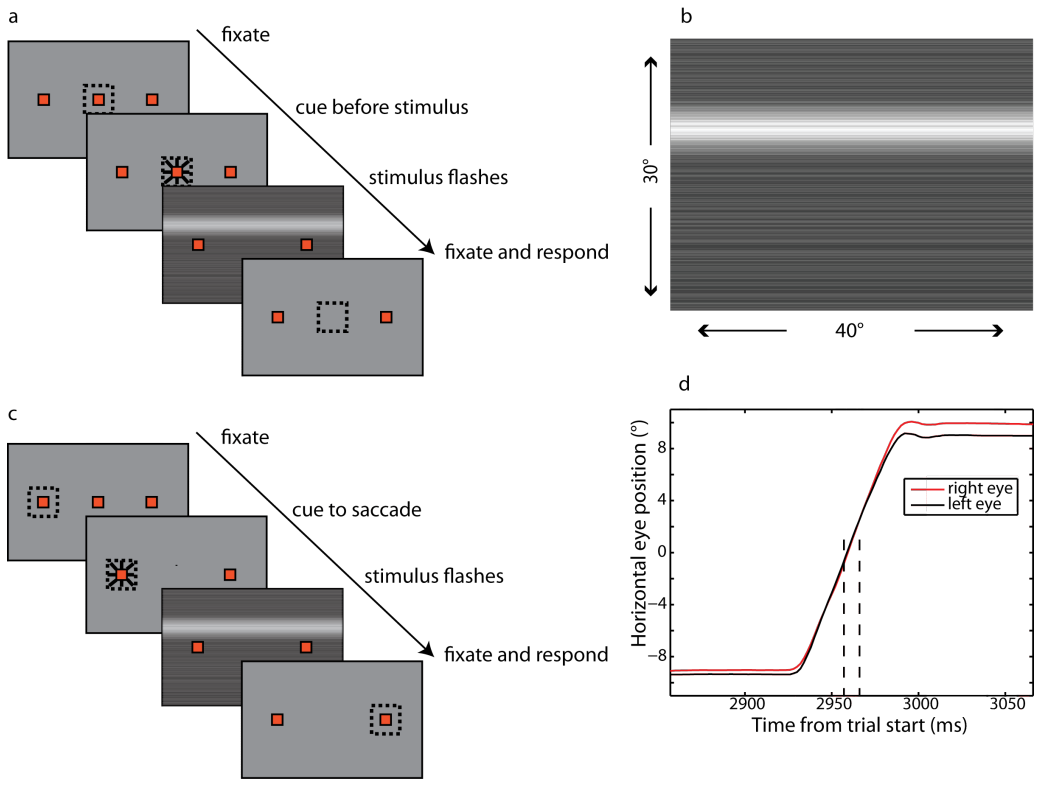

Figure 2.1 – Stimulus and Procedure. (a) The fixation condition began with subjects fixating a central dot, which disappeared before a horizontal grating flashed on the screen for one frame. (c) The saccade condition began with subjects fixating a left dot, which blinked, cuing the subject to make a rightward saccade. When the subject's gaze crossed an invisible trigger, the stimulus flashed on the screen for one frame. Subjects completed the saccade and fixated the rightmost dot before responding with a key press. (b) A horizontal grating was used. The vertical coarseness of the entire stimulus (grating and noise) was chosen according to the subject's ability to execute horizontal saccades (See Methods). (d) The eyes' horizontal position versus time, illustrating the cued, 18 degree saccade. The dashed, vertical lines denote the grating stimulus' on- and offset.

In the wide-signal experiment the target grating's horizontal extent covered the entire monitor. The vertical extent of each noise bar was 0.2°, and the external noise level was varied over the range 10, 15, 30, 60, and 90%.

In the narrow-signal experiment the target grating's horizontal extent was 12°. The external noise was the same as in the wide-signal experiment.

In the high-noise experiment the target grating's horizontal extent covered the entire monitor (as in the wide-signal experiment). The vertical extent of individual bars of the external noise was 0.4°, and external noise was set to 2% or 60%.

2.2.3 Procedure

Subjects were seated 57 cm from the monitor in a dark booth with a molded bite bar used to restrict head movement. An Eyelink II (SR Research) camera was used to monitor eye movements at a sample rate of 500 Hz and a nominal spatial resolution of 0.1° of visual angle. Trials in which eye position strayed beyond the windows specified below were discarded from analysis.

In all experiments, subjects performed a two-alternative-forced-choice (2AFC) task in which they decided whether a horizontal grating was presented in the lower or upper half of the screen. The grating was presented either during fixation (fixation condition) or mid-way through a saccade (saccade condition). Subjects responded using a keyboard press, and received auditory feedback (low/high pitched beep for wrong/right answer) after each response. For saccade conditions (Figure 2.1c), two red dots (0.3°) appeared on the screen – one 9° to the left, the other 9° to the right of the screen center on the horizontal meridian. Subjects fixated the left dot, which vanished between 640-960

ms (jittered) after fixation and reappeared 83 ms later, cuing the subject to make a saccade to the right dot. To ensure that the stimulus was presented mid-way through the saccade, the stimulus was presented when the subject's gaze crossed a screen positional threshold. The trigger's position was adjusted such that the stimulus flashed during the middle of the saccade (Figure 2.1d). Fixation of the rightmost dot had to be maintained within a 2 x 2 degree window for 100 ms after the saccade.

For fixation conditions (Figure 2.1a), the same dots as in the saccade condition appeared on the screen (to create a comparable stimulus in both conditions) along with a central fixation dot. Subjects fixated the central dot and after a jittered delay of 640-960 ms the dot disappeared and the grating flashed approximately 160 ms later for one frame (approximately 8 ms). To ensure that the fixation condition and saccade condition had comparable temporal uncertainty (Morris et al., 2010), the delay between the central dot disappearing and the stimulus appearing was adjusted to approximately match the delay between the cue and the stimulus onset in the saccade condition. Trials from the five external noise levels and two eye-movement conditions (saccade and fixation) were randomly interleaved within the same run.

2.2.4 Analysis

The aim of our analysis was to first estimate an observer's contrast threshold at each level of external noise, and then to estimate the parameters of the PTM (β, w, σ_m) that best explained the observed threshold-versus-noise (TvN) relationship. These steps were performed separately for the fixation and saccade conditions.

Using the adaptive threshold estimation algorithm (Kontsevich and Tyler, 1999), we determined the maximum a-posteriori psychometric function for each external noise level. We chose the 75% correct performance level on these curves to define threshold contrast. The three free parameters of the PTM (β, w, σ_m) were fit to the thresholds across the different external-noise levels using a (inverse-variance) weighted least-squares curve-fitting algorithm, (lsqnonlin in MATLAB). Once parameters for the fixation and saccade conditions were estimated, a Wilcoxon signed-rank test was performed to see if any of the parameters differed significantly between fixation and saccade conditions.

2.3 Results

2.3.1 Simulation

The PTM (Lu and Dosher, 1998b) has been used extensively to model the effects of attention on contrast sensitivity. We apply it here to understand the signal detection mechanisms underlying intra-saccadic suppression. The model includes an input stage (Figure 2.2a), a set of signal processing stages (Figure 2.2b, c, d), and an output (Figure 2.2e) resulting in a decision variable (DV). The elegance of the PTM lies in the qualitatively distinct changes in TvN curves predicted for variations of the parameters that represent each of the processing stages (Figure 2.2b, c, d). We applied the PTM to detection thresholds obtained during saccades and during fixation with the goal of ascertaining which of the processing stages were responsible for intra-saccadic suppression.

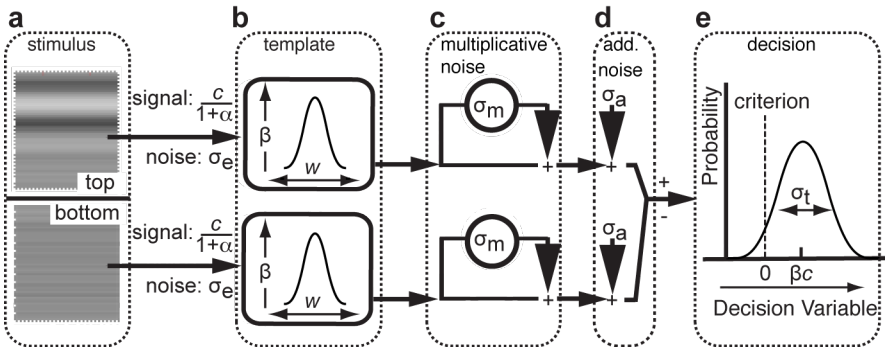

Figure 2.2 – The Perceptual Template Model (PTM). (a) Detectors for the top and bottom halves of the visual field are modeled as signal and noise processed through a series of stages. (b) The first template stage can be thought of as a signal gain and noise exclusion stage – increasing β increases the output of the template; narrowing w will exclude noise (non-relevant stimuli). (c) Multiplicative noise injects noise into the system in a stimulus-dependent manner. (d) Additive noise injects noise into the system in a stimulus-independent manner. (e) A decision variable is produced by summing the top and bottom detectors; it has a standard deviation, σ_t.

The PTM implemented here consists of two identical spatial channels: one for the screen location that receives signal-plus-noise and one for the screen location that receives only noise (Figure 2.2a). The variable *c* represents signal contrast (or strength). The α parameter is a novel addition to the PTM, which we discuss below. Each of the spatial channels proceeds independently until the final stage at which their outputs are subtracted (Figure 2.2e) to produce DV. If DV is greater than zero, then the model responds that the signal was in the upper screen position. If DV is less than zero then the model chooses the lower screen position. DV's distribution has a total variance (σ_t^2) that

can be expressed as a function of the external noise (noise added to the signal before it enters the system) and the model parameters. The detector's sensitivity (d') is then given by the ratio of the signal, $\beta c/(1+\alpha)$, and the standard deviation of the DV: σ_t. Analogous to the derivation in (Watson and Krekelberg)their equations 1-3), this allows us to express contrast threshold for a given level of performance (d') as a function of external noise (σ_e) and the model parameters:

$$c = (1+\alpha)\sqrt{2\frac{(w\sigma_e)^2(1+\sigma_m^2)+\frac{\sigma_a^2}{\beta^2}}{\left(\frac{1}{d'^2}-\sigma_m^2\right)}}. \qquad (1)$$

Figure 2.3 shows the shape of TvN curves that the PTM predicts if intra-saccadic suppression were determined by a change in only one of the parameters, or dominated by that parameter. In other words, these are the quantitative hypotheses that our experiments tested.

The first processing stage is the perceptual template stage (Figure 2.2b). This stage's output depends on how well the template matches the input signal – in effect, modeling selectivity. The template's two features are a gain parameter (β) and a tuning parameter (w) controlling the width of the template. The gain (β) amplifies the input, which contains either signal plus external noise (Figure 2.2a, top) or noise only (Figure 2.2a, bottom). The template can be thought of as an exclusion term: in general, a tighter template (a smaller w) implies a more focused perceptual analysis on the true timing, spatial position, or other characteristic (e.g., spatial frequency) of the stimulus (Lu and Dosher, 1998a). In our model, an increase in the template stage (w) refers to an increase in spatial uncertainty about the signal. A narrow template (less spatial uncertainty)

focuses the detector on only the relevant spatial location, and thereby excludes any external noise appearing outside of the signal's location. Therefore an increase in external noise has a minimal effect on thresholds for a narrow template. Now consider a wide template: external noise existing outside of the signal's location would be allowed into the detector. At low external noise levels, this would only have a minimal influence on thresholds. However at high external noise, the wider template's inclusion of noise would have a dramatic effect on thresholds (Figure 2.3a), allowing high levels of noise (outside of the signal's location) into the detector. Therefore if saccadic suppression occurs because of an increase in spatial uncertainty we expect saccade and fixation thresholds to diverge as external noise increases (Figure 2.3a).

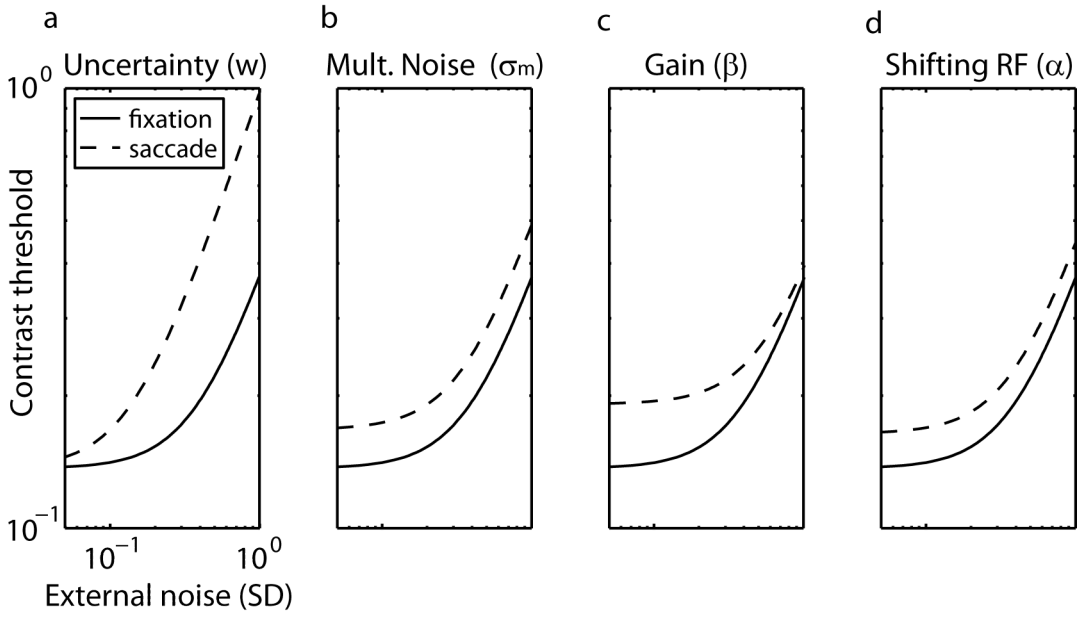

Figure 2.3 – The trends of the PTM under four different parametric accounts of saccadic suppression. (a) An increase in uncertainty (i.e., a wider template) will lead to an increase in detector noise as external noise increases. This is because the increasing external noise will be allowed into (not excluded by) the system. At low external noise

levels this lack of exclusion will not play a critical role in determining thresholds. (b) An increase in multiplicative noise will increase the detector noise in a stimulus-dependent manner, leading to a constant separation (in log units) of thresholds as external noise is increased. (c) Gain reduction will lead to converging thresholds at high external noise. This is because the gain factor amplifies both signal and external noise; hence it becomes ineffective when the input is dominated by external noise. (d) We extended the PTM with a term that reflects the influence of a shifting receptive field (α), which we modeled as a loss of signal. Our simulation reveals a trend that is qualitatively similar to multiplicative noise injection.

The second processing stage is a multiplicative noise injection (Figure 2.2c). The noise added to each channel at this stage is stimulus dependent because the output of the template stage is scaled by the parameter (σ_m). Because this noise injection is scaled by both the signal and noise, the PTM predicts that detection thresholds will be influenced equally at both low and high external noise levels (Figure 2.3b). Therefore, if saccadic suppression is caused by a stimulus-dependent noise injection, we expect a constant amount of suppression (in log units) across external noise conditions.

The third stage is a stimulus-independent additive noise injection (Figure 2.2d). Noise with standard deviation σ_a is added to the system, independent of the signal or external noise. It follows that as the external noise is increased, this term will have less of an influence on the threshold. Therefore the PTM predicts that varying this term will lead to different thresholds at low external noise, when the internal noise is dominated by

this additive injection. As external noise is increased (thereby increasingly influencing the system's response) while the additive noise remains constant, the thresholds should converge (Figure 2.3c). Because this term is mathematically equivalent to gain (β) reduction (See Equation 1) we will consider its effects under the gain (β) parameter.

To model the possible influence of RF shifts that are known to occur at the time of saccades (Duhamel et al., 1992; Tolias et al., 2001), we introduced a novel parameter into the PTM. A spatially shifted detector should reduce the signal, but not the external noise (assuming the shifted position is also somewhere on the display), hence we introduced the term at the earliest stage (Figure 2.2a). The parameter (α) represents the amount of shift between the detector and the stimulus (when there is no shift, α equals zero, and the input to the detector is equal to c). Our simulation of the influence of α (Figure 2.3d) shows a trend similar to that of multiplicative noise injection (Figure 2.3b). Because it is qualitatively similar to multiplicative noise injection (σ_m) we do not consider α in the fits.

2.3.2 Experiments

In the first – wide-signal – experiment, we measured contrast thresholds for horizontal gratings covering the entire width of the monitor in varying levels of contrast noise while subjects fixated, and during 18 degree cued saccades. The psychometric curves of one subject obtained for two noise conditions (10% and 60%) are seen in the two panels in Figure 2.4. The 75% performance level thresholds were extracted for each level of external noise. In each panel, the dotted curve (saccade condition) is shifted to

the right from the solid curve (fixation condition), indicating an increase in threshold during saccades (saccadic suppression) for both external noise conditions. In this study we were concerned with how this saccadic suppression changed as a function of external noise level. This example shows that the increase in threshold induced by the saccade was larger at low (10%) than at high (60%) external noise.

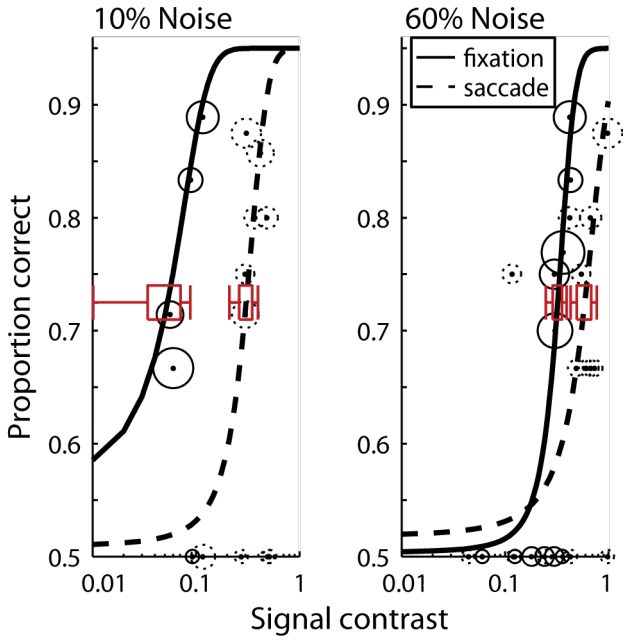

Figure 2.4 – Psychometric curves for the 10% and 60% external noise conditions, and for saccade (dashed line) and fixation (solid line) data for one subject (S5). Circle size around data points represents number of trials at the given contrast. The curves were obtained using the Bayesian-adaptive method for threshold estimation (see Methods). 95% Confidence intervals are plotted in red for the 75% performance threshold. A total of 1621 trials were run for this subject, with at least 100 trials for each of the 10 conditions (5 external noise conditions for both saccade and fixation).

Figure 2.5 shows thresholds across all external noise levels and for all subjects for the wide-signal stimulus. At low external noise levels there was a clear separation between saccade and fixation thresholds (Figure 2.5) but this gap decreased as noise increases. This is consistent with a gain reduction mechanism (Figure 2.3c). However the thresholds did not converge completely for all subjects.

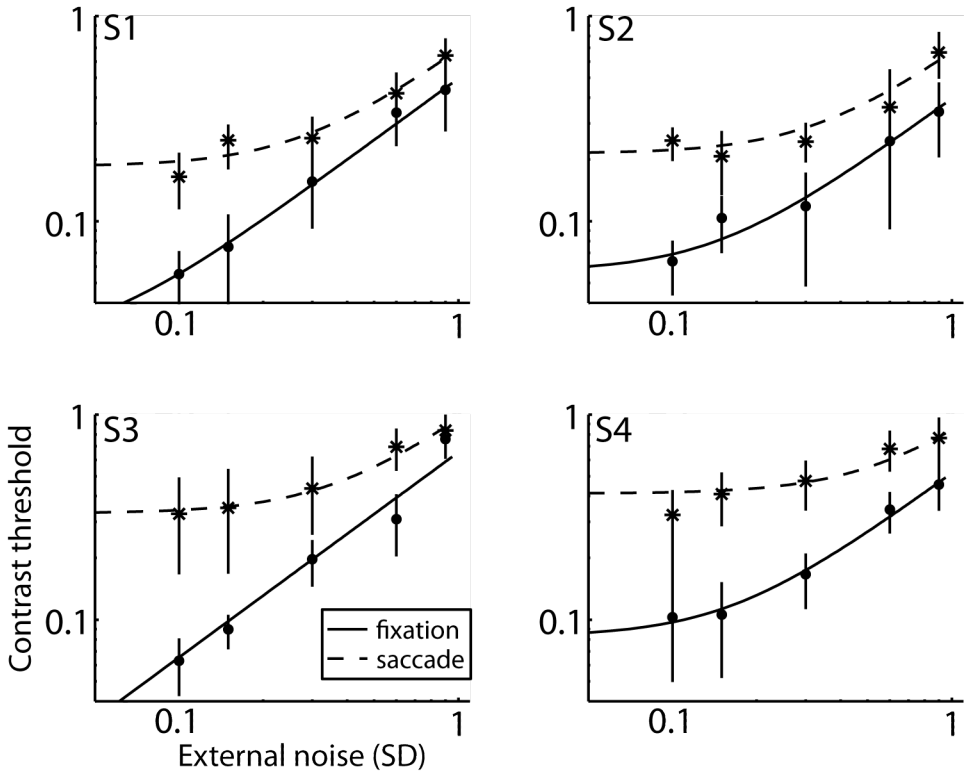

Figure 2.5 – TvN data for the wide-signal stimulus with PTM fits. A horizontal grating extending the width of the monitor (40°) was flashed for one frame (8 ms) during fixation and cued saccades. Fixation data are shown as dots, with the PTM fits drawn solid. Saccade data are shown as asterisks with PTM fits drawn dashed. Each data point is drawn with 95% confidence intervals (thresholds and confidence intervals obtained as in Figure 2.4). The external noise levels were 10, 15, 30, 60, and 90%.

The stimulus of the wide-signal experiment was chosen to minimize retinal smear, and therefore was uniform along the horizontal meridian and extended horizontally 40° (the entirety of the display). One could argue that such a wide target stimulus limits the possible contributions of an uncertainty mechanism. Essentially, any variation in horizontal (but not vertical) extent or location of the template could go unnoticed because the stimulus does not vary in that dimension. To address this we performed another experiment, now with a horizontally narrow grating (12°). The results of this "narrow-signal" experiment are shown in Figure 2.6. The trends were qualitatively similar to the results of the wide-signal experiment (Figure 2.5): suppression decreased as external noise increased. Again, this is consistent with a gain reduction mechanism (Figure 2.3c). To analyze this possibility, we fit each of the TvN curves from the wide-signal and narrow-signal experiments with the PTM, separately for the fixation and saccade conditions (Figures 2.5 and 2.6, solid and dashed lines).

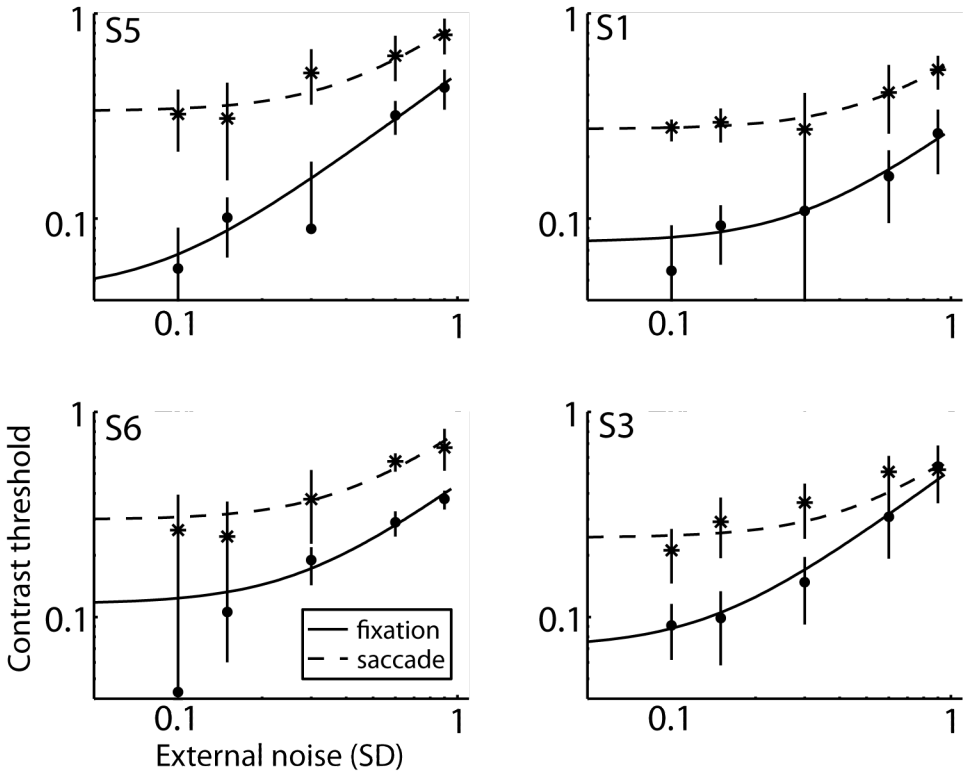

Figure 2.6 – TvN data for the narrow-signal stimulus with PTM fits. A horizontal grating extending 12° was flashed for one frame (8 ms) during fixation and cued saccades. All plotting conventions are the same as in Figure 2.5.

Figure 2.7 compares the best-fitting PTM parameters in the fixation and saccade conditions (wide-signal data points shown as squares, narrow-signal data points shown as triangles). A unity line is plotted for comparison purposes; points lying below the line represent a decrease in that parameter during saccades. The uncertainty parameter and the multiplicative noise parameter did not increase or decrease consistently during saccades, but the gain parameter decreased in seven of the eight subjects. One wide-signal data point (S3) is not displayed on the gain plot because it was two orders of magnitude greater during fixation (and therefore rendered the plot visually less informative); it still obeyed the trend of the other data, decreasing during saccade. A

statistical analysis of these parameter changes across all subjects first confirmed the qualitative impression based on Figures 2.5 and 2.6: there was a significant intrasaccadic reduction in the gain (β) parameter (Wilcoxon T=3.00, p=0.039). Second, even though the multiplicative noise (σ_m) parameter appears somewhat reduced during saccades (Figure 2.7), this effect was not significant (T=8.00, p=0.200). We note furthermore that such an intrasaccadic multiplicative noise reduction would not be consistent with saccadic suppression (i.e., it predicts lower thresholds during saccades). Third, the uncertainty parameter (w) was not significantly different during saccades (T=16.00, p=0.844). The same statistical conclusions (a significant effect of gain, but not of multiplicative noise, or uncertainty) were also reached after excluding outliers from the statistical analysis.

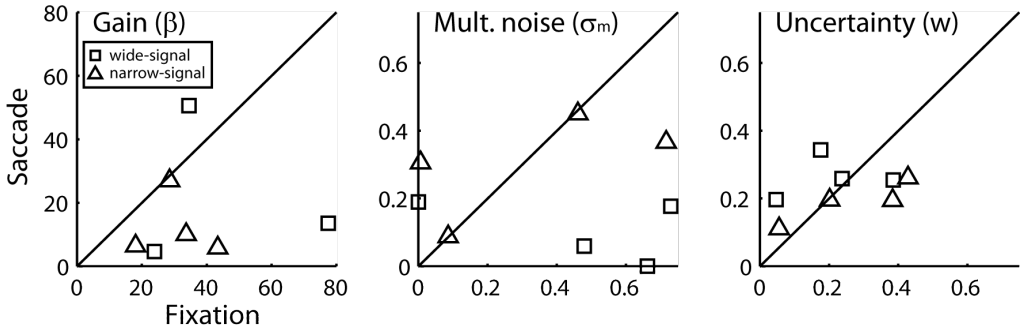

Figure 2.7 – Changes in PTM parameters between fixation and saccade conditions. Each plot shows a PTM parameter's saccade vs fixation value. Each data point represents one subject (wide-signal shown as squares; narrow-signal as triangles). Points below the line represent a decrease in that parameter during saccades. A Wilcoxon signed-rank test revealed a significant difference in gain during saccades (T=3.00, p=0.039), but not in multiplicative noise (T=8.00, p=0.200) or uncertainty (T=16.00, p=0.844).

Rather than relying solely on this population statistical analysis we used the PTM to generate a prediction that more clearly disentangles gain on the one hand, and multiplicative noise and uncertainty mechanisms on the other. Notably, if saccadic suppression were dominated by gain, then TvN curves should converge at higher levels of external noise. To test this prediction, we performed a third experiment – high-noise – in which the size of the external noise bars was increased to match the spatial frequency of the target stimulus more closely (it was not possible to add further contrast noise because it was beyond the capability of the display). Because the template stage (acting as a matched filter) is presumably matched to the signal's characteristics, the external noise is more effective (less of it is filtered out) as its spatial frequency approaches the signal's spatial frequency. Figure 2.8 shows the results of the high-noise experiment: at low external noise, the saccadic thresholds were significantly higher than at fixation, but at high levels of external noise the saccade had no significant influence. This is again consistent with a gain change (Figure 2.3c), but inconsistent with the other mechanisms: an active uncertainty mechanism (Figure 2.3a), multiplicative noise (Figure 2.3b), and a shift in the detector (Figure 2.3d) all predict suppression at high external noise. This confirms that the intra-saccadic suppression is dominated by a gain change.

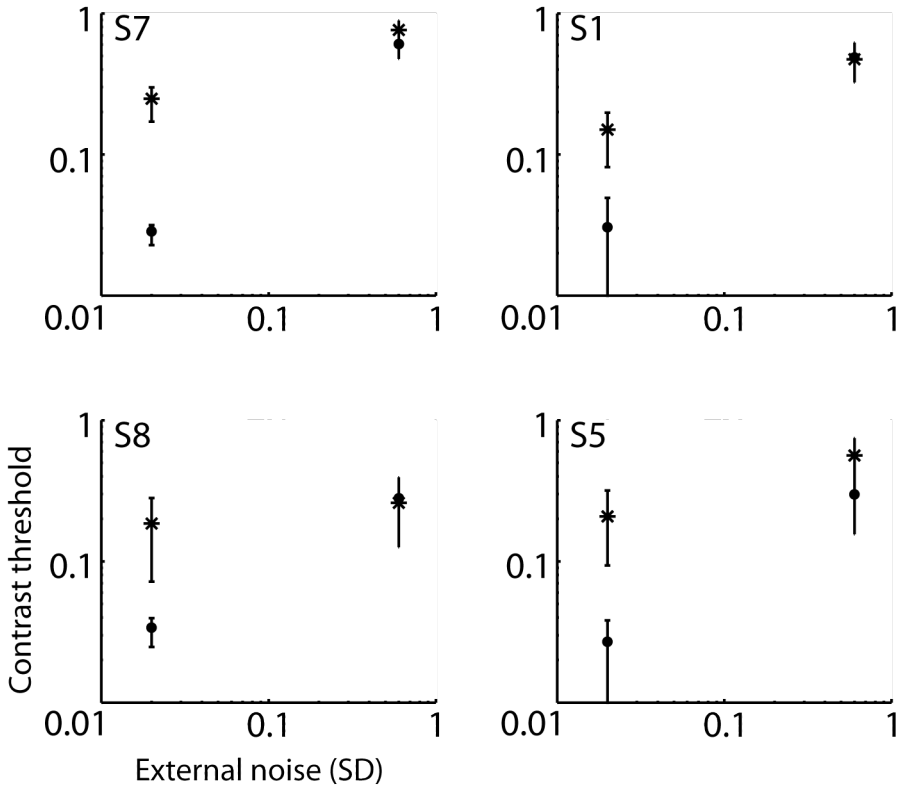

Figure 2.8 – High Noise Experiment. Fixation thresholds are shown as dots, saccade thresholds as asterisks. The external noise levels were 2 and 60%, but the coarseness of the external noise was increased to more closely match the grating's spatial frequency – increasing the effectiveness of the noise. As predicted by a gain mechanism, the thresholds fully converged at high noise, indicating a gain reduction mechanism.

2.4 Discussion

We set out to provide a description of intra-saccadic suppression in terms of basic signal detection mechanisms. The main finding of our experiment was that a gain reduction is the dominant mechanism in explaining an increase in contrast thresholds for stimuli presented during saccades, as it is for stimuli presented before saccade onset (Watson and Krekelberg, 2011). This conclusion is based on 1) consistency of gain

reduction across subjects and 2) the convergence of fixation and saccadic contrast thresholds at very high levels of external noise. The fact that uncertainty about the spatial location of the stimulus contributed little to detection thresholds is surprising as it implies that phenomena such as peri-saccadic receptive field shifts and mislocalization have at best a minor influence on intra-saccadic stimulus detection.

2.4.1 Signal-Detection Mechanisms

We found that gain reduction (β) was sufficient to explain an increase in contrast thresholds during saccades. Our conclusion is consistent with a previous study using the PTM that found that gain reduction (β) was sufficient to explain pre-saccadic suppression (Watson and Krekelberg, 2011) and earlier proposals that gain reduction (Burr et al., 1994) or (the formally equivalent) stimulus independent noise injection (Diamond et al., 2000) in early visual areas may be responsible for increased thresholds during saccades. In addition, these findings are compatible with electrophysiological studies showing pre and intra-saccadic reduction of neural responses. (For review see Ibbotson and Krekelberg, 2011)

We set out to test the hypothesis that uncertainty about a stimulus' location is another mechanism that could contribute to intra-saccadic suppression. We considered spatial uncertainty in two qualitatively different ways. One way is a reduction in signal resulting from peri-saccadic variability in stimulus localization – such as could arise from RF shifts (Duhamel et al., 1992; Tolias et al., 2001; Kusunoki and Goldberg, 2003), or from peri-saccadic error in eye-position signals (Morris et al., 2012). We refer to this as signal mislocalization because the increased detection thresholds result purely from lost signal. Simulations of such a detector show that signal mislocalization predicts

suppression effects both at high and at low external noise (Figure 2.3d). We also considered spatial uncertainty as a change in the detection *process* that aims to compensate for the expected intra-saccadic uncertainty about the stimulus' location. For instance, a good detection strategy would be to widen the population from which a detector receives input. This change in the detection strategy is captured in the PTM by a widening of the search template (w). The widening, however, also allows more noise into the system, leading to increasing thresholds as external noise increases (Figure 2.3a). Hence, although these two uncertainty mechanisms predict quantitatively different TvN curves (Figure 2.3a vs 2.3d), qualitatively they both predict that threshold differences should be found even at high external noise. Our data did not match this prediction, hence we reject the hypothesis that uncertainty is a dominant contributor to intra-saccadic suppression.

Our conclusion is opposite to that of Greenhouse & Cohn, who found uncertainty to be sufficient to explain a decline in detectability at the time of saccades (Greenhouse and Cohn, 1991). It is important here to distinguish between the different characterizations of uncertainty used in our respective studies. Greenhouse & Cohn found reduced ROC slopes for stimuli flashed during saccades – a result explained by an observer's uncertainty about stimulus parameters. While their methods show that uncertainty plays a role, they do not necessarily show that uncertainty is a dominant mechanism, or that it is sufficient to explain suppression. While our use of the PTM allows us to compare the contributions of different mechanisms, the Greenhouse & Cohn study investigates only one mechanism (uncertainty) and cannot eliminate the possibility that another mechanism is *more* responsible for saccadic suppression. In an experimental

variant designed to prove that uncertainty is sufficient to explain suppression, Greenhouse & Cohn did show a reduction of suppression by flashing a pedestal at the spatial location of the stimulus on both signal and noise (blank) trials. However in this experimental variant they also increased the signal intensity to achieve the same detectability. This implies that the pedestal effectively served as a source of external noise. Hence their finding that the pedestal reduced suppression can be interpreted in the equivalent noise framework as a convergence of TvN curves at high external noise – which is consistent with our findings and points to gain reduction. In a final variant, Greenhouse & Cohn attempted to reduce uncertainty without increasing stimulus intensity by using spatial markers at the site that the stimulus would appear. However, the results of this variant were inconclusive as one of the two subjects tested continued to show significant suppression even when the markers were in place. In summary, Greenhouse & Cohn's results are not inconsistent with ours, but our application of the PTM allows us to compare between mechanisms, and to conclude that gain is dominant.

Another main difference between our study and Greenhouse & Cohn's, is that their stimulus was only 1 degree. This raises the interesting possibility that an uncertainty mechanism could play a more dominant role at a small spatial scale (and therefore in a brain region with smaller receptive fields), but that its influence on the overall detection process is diminished by selecting a sufficiently large stimulus size. We note, however, that a small stimulus necessarily has sharper edges which move across the retina during the saccade. This provides a qualitatively different visual cue that complicates a fair comparison with the fixation condition. The larger stimulus with a

smooth Gaussian edge in our wide-signal experiment greatly reduces this potential confound.

Our conclusion that uncertainty is not a dominant factor in intra-saccadic suppression appears at odds with the physiological findings that motivated our study. However, it is quite easily conceivable that RF shifts (Duhamel et al., 1992), eye position uncertainty (Morris et al., 2012), or RF widening (Tolias et al., 2001), are too small to affect detection relative to the effects of gain reduction (Ibbotson and Krekelberg, 2011). It is also possible that, because not all neurons display peri-saccadic changes such as RF shifts, the effect of cells that do shift are less critical for detection at the population level.

2.4.2 Neural Mechanisms

Equivalent noise analysis can identify dominant mechanisms in terms of signal processing, but it cannot identify the underlying *neural* mechanisms. In the context of saccadic suppression it is important to point out that it cannot distinguish between the contributions of so-called active and passive backward masking mechanisms. An active gain reduction of visual neurons is certainly compatible with the data (Bremmer et al., 2009), but our data do not exclude a contribution from passive mechanisms as long as these result in a behavioral effect that is dominated by a gain change. Such neural mechanisms could include backward masking, or the Stiles-Crawford effect. The Stiles-Crawford effect reflects an intra-saccadic tilt of the photoreceptors resulting in less light absorption – this scales both signal and external noise and therefore results in a reduction in gain at the earliest stage of the detection process. We note, however, that gain reduction has also been shown to account for pre-saccadic suppression (Watson and Krekelberg, 2011). In that study, stimuli were flashed just before the saccade, while the

eyes were still stationary, so that the gain reduction could not be explained by the Stiles-Crawford effect. Considering the gradual time course of suppression – both behaviorally (Diamond et al., 2000) and in studies of its neural correlates – the most parsimonious explanation is that the same gain reduction mechanism that plays a role pre-saccadically also operates intra-saccadically. The importance of gain control is also supported by the work of (Burr and Morrone, 1996). They studied impulse response functions during saccades to stimuli processed predominantly by magno- or parvo-cellular pathways and showed that differences in contrast gain control in these pathways are compatible with behavioral differences in saccadic suppression.

Even though equivalent noise analysis cannot identify a neural mechanism, it can aid the search for such mechanisms by providing a quantitative target. In other words, our findings show that to qualify as an explanation of the behavioral effect of saccadic suppression, any proposed neural mechanism should result in a gain change at the behavioral level. A direct way to incorporate our findings into studies of neural mechanisms of suppression is to compare signal-to-noise ratios of neural responses to a target stimulus embedded in external noise with the predictions of the PTM. This could be a worthwhile extension of single cell studies in sensory areas already known to show intra-saccadic modulation of activity (reviewed in Ibbotson and Krekelberg, 2011), but also in functional imaging studies (Kleiser et al., 2004; Sylvester et al., 2005; Vallines and Greenlee, 2006) that have the ability to reveal changes in multiple brain regions at the same time. A quantitative investigation of the neural mechanisms of saccadic suppression within the equivalent noise framework may be able to decide which (if any)

among the plethora of proposed neural mechanisms is the dominant neural mechanism of saccadic suppression.

Chapter 3. Gain reduction in primate V1 during saccades

3.1 Introduction

Evidence that saccadic suppression is a result of reduced detector gain (Watson and Krekelberg, 2011; Guez et al., 2013) has guided our investigation of a neural correlate of suppression toward a specific characterization of the neural response: in this chapter we analyzed saccadic modulation of sustained neural responses to a range of contrasts. We recorded from chronically implanted multi-electrode arrays in the macaque primary visual cortex (V1). We specifically tested whether peri-saccadic contrast responses exhibited gain reduction compared to contrast responses during steady fixation.

We recorded in V1 because it is the first cortical processing stage, and shows other behaviorally relevant contrast response properties such as fast contrast adaptation (Levy et al., 2013) and attentional modulations (Thiele et al., 2009). V1 also shows some evidence of gaze-contingent gain modulation – possibly by input from non-cortical areas (Durand et al., 2010; Buisseret and Maffei, 1977) – as well as evidence of subcortical input (Doty, 1983). These are important given studies suggesting a corollary discharge (or more generally an active) mechanism in suppression (Sommer and Wurtz, 2008; Bremmer et al., 2009).

Most of the experiments on behavioral suppression involve detection of static luminance-modulated contrast gratings (for review see Wurtz, 2008) – and yet to our knowledge there has been no direct measurement of the peri-saccadic contrast response function in V1. These considerations along with our conclusions from Chapter 2 – that the

computational mechanism of suppression is gain reduction – suggested that we should explore peri-saccadic response properties in macaque V1 for neural correlates of saccadic suppression

3.2 Methods

3.2.1 Recording

Electrophysiology. 32-channel Floating Microelectrode Arrays (FMA) were implanted in parafoveal primary visual cortex (V1) of two adult male macaques (*Macaca mulatta*). Monkey *M* had one array in the right hemisphere with receptive fields at x:0.27°, y:-3.07° from the fovea; Monkey *Y* had two arrays in the left-hemisphere – one array with receptive fields at x:-3.2°, y:-3.2°, and the second with receptive fields at x:-1.6°, y:-1.6° from the fovea. The experimental and surgical protocols were approved by the Rutgers University Animal Care and Use Committee, and were in agreement with the National Institute of Health guidelines for the humane care and use of laboratory animals. The neural signals were amplified, filtered and sampled at 30 kHz using a Grapevine neural interface system and Trellis software (both produced by Ripple). Wavelet transforms were used for spike detection (Nenadic and Burdick, 2005).

Eye position. Eye position was recorded with an infrared tracker at a sampling frequency of 250 Hz (EyeLink2000; SR Research). Saccades were cued by presenting fixation targets at different locations (see Procedure); once a target was fixated, eye position had to remain within a 2°x2° square area centered at the target. A trial was discarded if any fixation breaks (such as blinks or straying eye position) occurred.

3.2.2 Visual stimuli

All stimuli were computed in Neurostim (http://neurostim.sourceforge.net), and displayed on a 20-inch CRT monitor (Sony GDM-520) 57 cm from the eyes. The display was 40° x 30°, with a resolution of 1024 x 768 pixels; it had a refresh rate of 150 Hz (~6.67 ms per frame) and a mean luminance of 30 cd/m^2.

Grating Stream. Our purpose was to compare contrast responses before saccades with those during steady fixation. Therefore the stimulus was a stream of luminance-modulated contrast gratings with intermittent cued saccades. Sine-wave gratings had a duration of 50 frames (~333.33 ms), a spatial frequency of 1.5 cycles/°, and extended 20° x 20° positioned at the center of the monitor. The stimulus was a stream of seven of these sine-wave gratings (~2333.33 ms duration), presented consecutively (no blank frames), and randomly chosen with replacement from a set of 7 Michelson contrasts (0, 8, 16, 32, 48, 64, 96%). Because different units recorded by the array could have different orientation preferences we chose one of four possible orientations (0, 45, 90, 135°) for two consecutive grating streams – that is, 14 gratings of different contrast but the same orientation were presented before changing the gratings' orientations.

3.2.3 Procedure

Subjects sat in a primate chair with their heads stabilized and received a reward for each block of seven trials. Subjects maintained fixation on a red target located at the center of the fovea for ~ 267 ms, initiating the grating stream. Saccades were cued before the 3rd, 5th, and 7th grating – so that saccade and fixation conditions were interleaved. Saccades

were cued by presenting a new fixation point 6° away from the initial point, along the axis of the grating's orientation (Figure 3.1, top) – this was to minimize retinal smear. Saccades were cued such that saccade onset began at least 200 ms after grating onset – so that transient effects of onset responses would have desisted and we could observe the effect of saccades on sustained neural activity. After a block of 7 contrasts subjects were given a reward and a 500 ms break before a new block began.

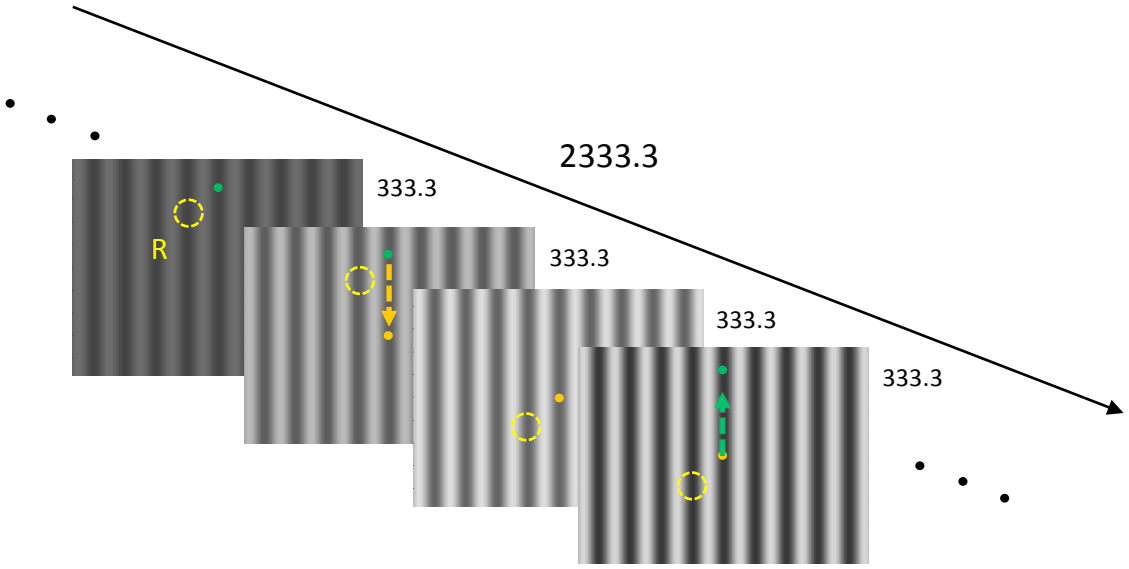

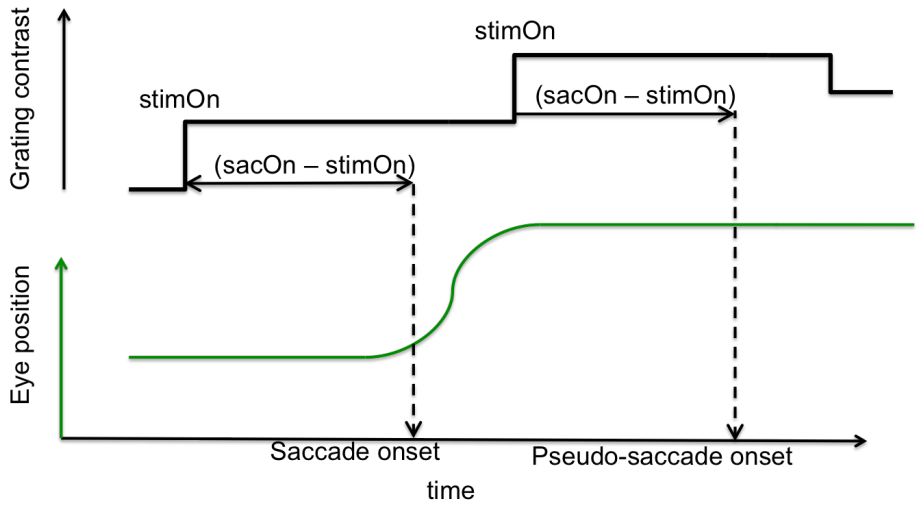

Figure 3.1 – Procedure. (top) Different contrast gratings in a block of trials. Each contrast is presented for 333.3 ms. Saccades are cued during every other grating, at least 200 ms after the onset of that grating so that the peri-saccadic responses to that grating are during sustained (i.e., not transient) stimulus activity. (bottom) Time intervals between grating onset and saccade onset (sacOn – stimOn) were calculated. These were added to grating onset in fixation trials thereby creating 'pseudo-saccade' onset times – allowing us to comparably average responses in the two conditions.

3.2.4 Data analysis

24 sessions were recorded from Monkey M and 11 sessions from Monkey Y. Each session consisted of at least 5000 trials. Units from the same electrode but different recording dates that had similar waveforms (and therefore were suspected of being the same unit across different days) were compared functionally on the basis of orientation tuning. If the functional properties were similar then we pooled these into the same unit – thereby giving us a conservative estimate on sample size. Of the pooled units we

included only those exhibiting a reliable stimulus-evoked response (signal), as compared to spontaneous firing (noise). We calculated d' as a measure of signal-to-noise ratio (SNR); we excluded any units with a d' < 1.4. After the aforementioned sorting, pooling, and SNR criteria we had 18 units in Monkey M, and 35 units in Monkey Y. For each unit we determined a preferred orientation and then only analyzed trials in which that orientation was presented.

For saccade trials we aligned neural responses to saccade onset; for fixation trials we aligned responses to a 'pseudo-saccade' time (see Figure 3.1, bottom), calculated in the following way: we found the difference between saccade onset and stimulus onset for all saccade trials in a session and then randomly redistributed these intervals to all fixation trials in that session, adding them to fixation trial stimulus onsets. Thus, stimulus responses were comparably averaged across trials in each of the two eye-movement conditions.

To minimize the effects of stimulus onset transients and brief timescale adaptation, our analysis window began at least 175 ms after stimulus onset. We confirmed that responses remained steady from this time point with a two-factor (contrast and time) ANOVA on the time course of the fixation-condition contrast responses, beginning at 80 ms before the pseudo-saccade and ending when the saccade was complete (~40 ms from onset). This revealed an effect of contrast (d.f.=6, $F=346.11$, $p<0.0001$) but no effect of time (d.f.=6, $F=0.56$, $p=0.76$).

3.2.4.1 Contrast response functions

We fit Naka-Rushton contrast response functions (CRF) to the average firing rate at each contrast using a least-squares error method in MATLAB (http://www.mathworks.com). The equation was,

$$r(c) = r_{max}\frac{c^n}{c_{50}^n + c^n} + b$$

In the case where the analysis was purely descriptive (section 3.3.1, Suppression dynamics) the free parameters were the maximum firing rate (r_{max}), the semi-saturation contrast (c_{50}), the exponent (n), and the baseline response (b). In the case where the CRF was fit to address a specific type of gain reduction (3.3.2, Contrast gain or response gain), the free parameters were r_{max} and c_{50}.

3.3 Results

We studied the influence of saccades on the contrast responses of V1 units by presenting a stream of different contrast gratings while the subject made cued saccades (saccade condition) or maintained steady fixation (fixation condition). Figure 3.2 (top panel) shows raster plots of a V1 unit's activity aligned to pseudo-saccade (blue) and saccade (red) onset. The unit is being stimulated with a 96% contrast grating shown at least 200 ms before saccade onset (exact time between grating and saccade onsets vary). The bottom panel shows the average spike rate across trials for the two conditions: this unit shows a reduction in response in the saccade condition. In order to study the effect of saccades on population activity we ran a three factor ANOVA (saccade or fixation, 7 test contrasts, and 7 time points at which spikes were binned) on all unit responses from both subjects (53 units total). We found that making a saccade had a significant effect on

neural response (F=71.82, p<0.0001). We found an overall effect of contrast (F= 573.81, p<0.0001), and of time (F=6.39, p<0.0001). We also found significant interactions between eye condition and contrast (p=0.0003), as well as eye condition and time (p<0.0001) – prompting us to look more closely at how suppression affects the contrast response function as well as the dynamics of suppression at each contrast.

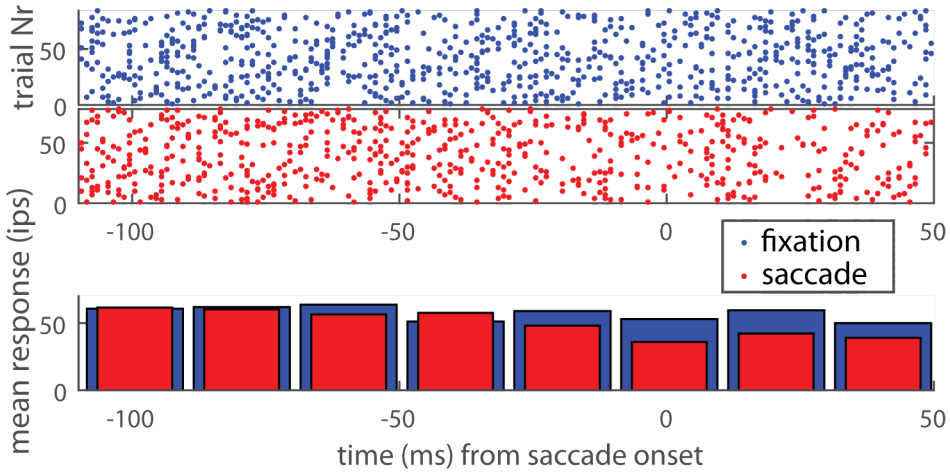

Figure 3.2 – Raster plot showing saccadic response decrease for a V1 unit. Top panel shows spike responses to a 96% contrast grating in all fixation (blue) and saccade (red) trials aligned to pseudo-saccade and saccade onset, respectively. Bottom panel shows spike rates in 20 ms bins. This unit shows response reduction in the saccade condition; minimum response is centered at saccade onset (0 ms).

3.3.1 Suppression dynamics

We studied suppression dynamics in each subject. We normalized each unit's responses between 0 and 1 and then averaged over all units – to study the time course of fixation and saccadic responses at each level of contrast stimulation. Figure 3.3 shows the

difference between peri-saccadic and fixation neural responses – a negative value indicates a smaller neural response in the saccadic condition and therefore suppression. At low contrasts we see that there is no effect: the difference between saccade and fixation responses (solid blue line; 99% confidence intervals plotted as dashed red line) is around zero (dashed black line). At the highest contrasts we see a clear increase in suppression as saccade onset approaches. We see there is no suppression at 80 ms before the saccade and that suppression gradually increases, reaching a maximum around the time of saccade onset. Viewing these panels from lowest to highest contrast there appears to be a trend of increasing suppression with increasing stimulus contrast.

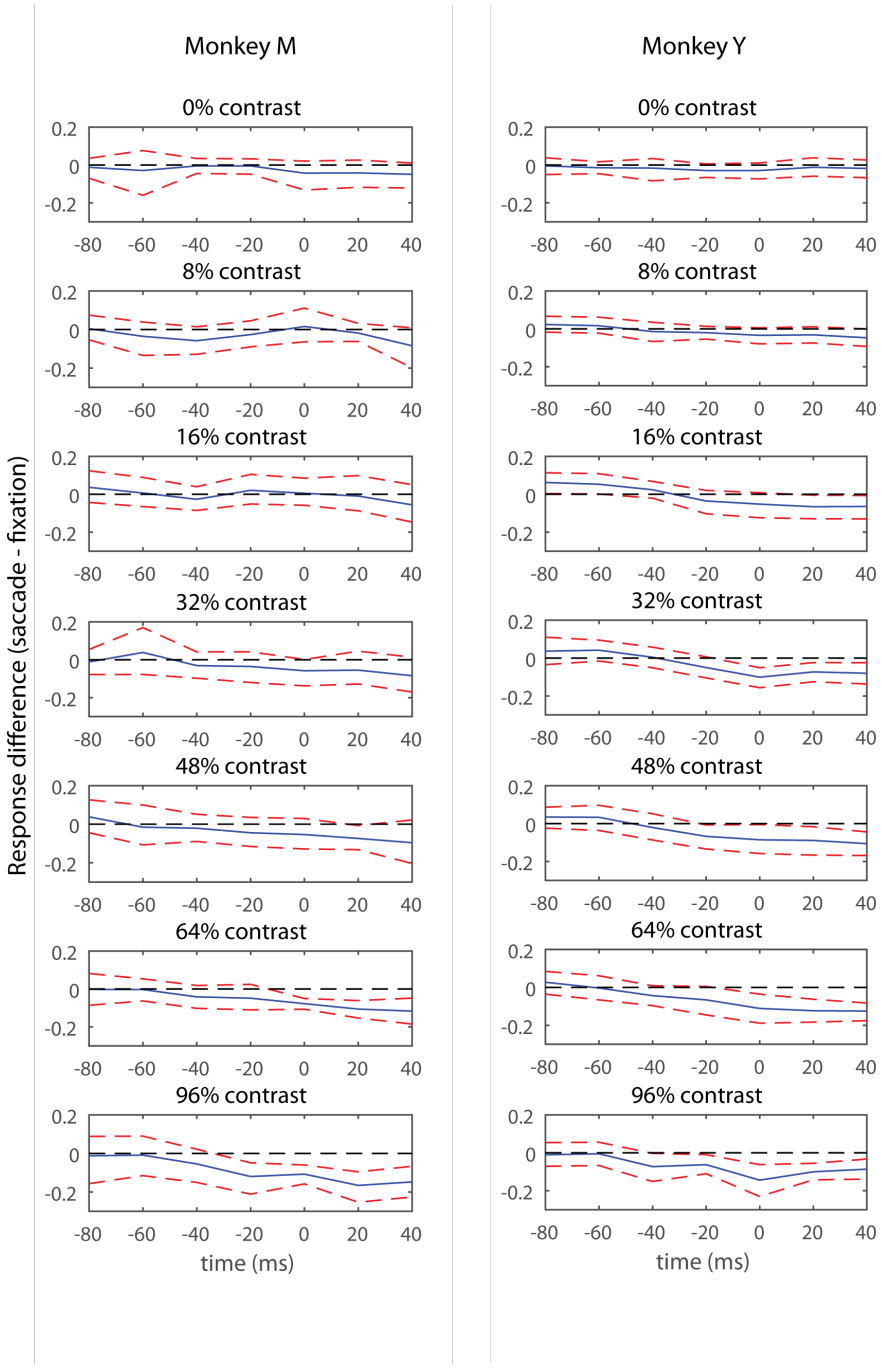

Figure 3.3 – Time course of suppression at different contrasts. We aligned time courses to saccade onsets (saccade condition) or pseudo-saccade onset (fixation condition, see Methods) and analyzed differences in neural response at each of the presented contrast levels. Response difference is plotted in blue; 99% confidence intervals (bootstrap)

shown in dashed red; zero value (no response difference) extended in dashed black. Spikes were binned in 20 ms time windows. Low contrasts (top panels) show no suppression. High contrasts (bottom panels) show suppression starting 60 ms before saccade onset.

In Figure 3.4 we express these data as the time course of a contrast response function (CRF). The results are qualitatively similar in both subjects – the CRF dynamics confirm a gradual reduction in the response to intermediate and higher contrasts.

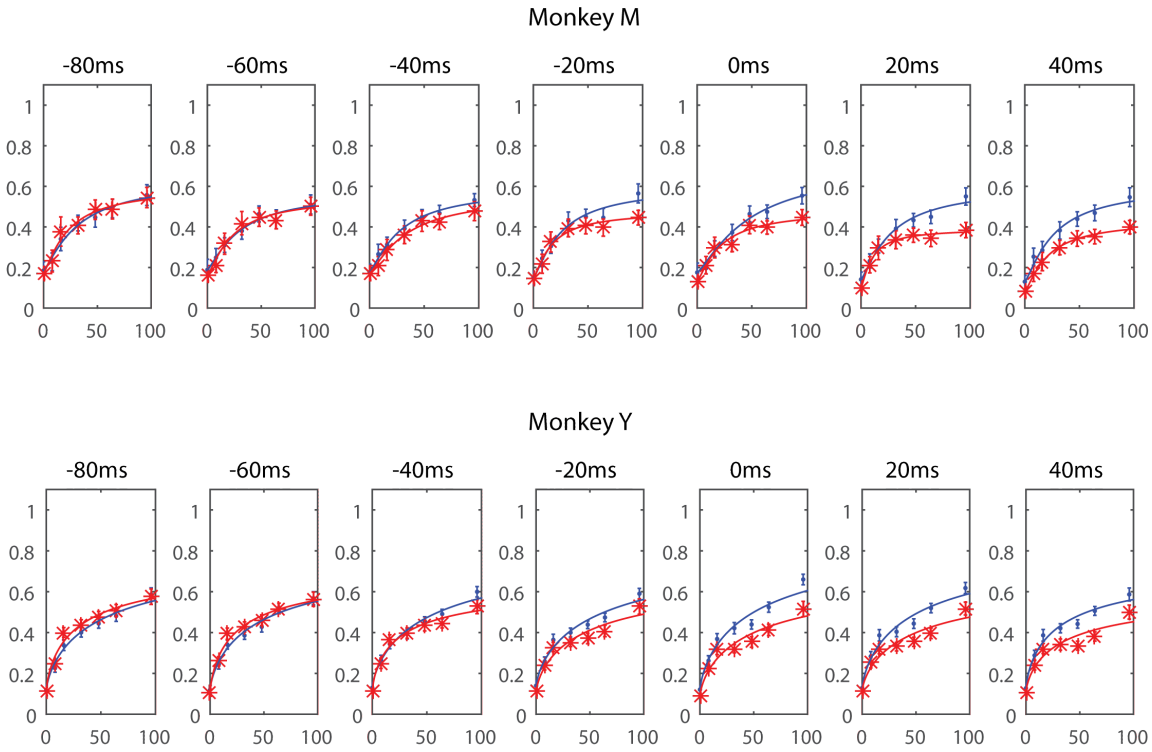

Figure 3.4 – Time course of population averaged CRF during saccades, aligned to saccade onset (0 ms). The top row shows contrast response functions for Monkey M. Saccade data (red asterisk) and fixation data (blue dots) are plotted with Naka-Rushton

functions (for visual comparison) – fit independently to each 20 ms time bin. The bottom row shows the same for Monkey Y. Both CRF time courses show a gradual gain reduction, especially at higher contrasts.

3.3.2 Contrast gain or response gain

Up to this point we have shown evidence that a V1 population exhibits peri-saccadic gain reduction. Gain reduction can be characterized as a shift in one or more Naka-Rushton parameters. More specifically, an increase in c_{50} (a horizontal shift of the CRF) is equivalent to a reduction in the input contrast, and therefore called a reduction in contrast gain. A decrease in r_{max} is a reduction of the response function by a scaling factor, and therefore called a reduction in response gain. We wanted to know whether the population gain reduction could be explained in terms of these more specific modes of gain reduction in the individual units. For analysis we chose the time point of maximal suppression, which according to our data (Figure 3.3) was saccade onset (0 ms). We excluded any unit that had a very noisy (i.e., confidence intervals included values outside the bounds of parameter search space) parameter estimate in either condition (saccade or fixation). Figure 3.5 shows a scatter plot of r_{max} (left panel) and c_{50} (right panel) shifts for the remaining 14 units. The left panel shows that most units have a lower r_{max} value during saccade than they do during fixation. A paired t-test confirmed a significant reduction in saccadic r_{max} values ($p<0.025$, Bonferroni-corrected). The right panel shows most units to be near the unity line indicating no significant change in c_{50} ($p>0.025$, Bonferroni-corrected).

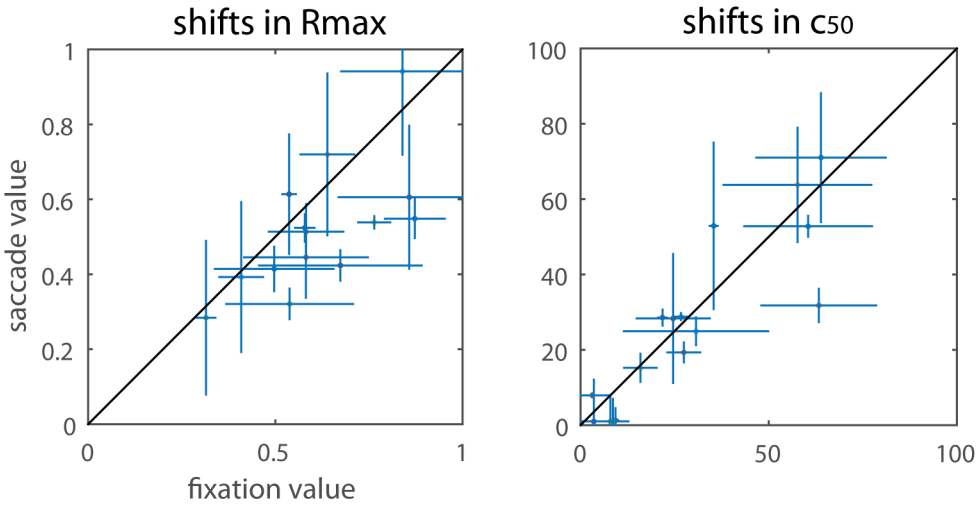

Figure 3.5 – CRF parameter shifts between fixation and saccadic conditions. Data were binned in a 20 ms window centered around saccade onset (saccade condition) or pseudo-saccade onset (fixation condition). Standard errors are shown for saccade and fixation parameter fits (vertical and horizontal lines). The left panel shows shifts in the r_{max} parameter; most units show a decrease in r_{max} during saccades – indicating a reduction in response gain. Right panel shows shifts in the c_{50} (half-saturation). A paired t-test revealed saccades have a significant effect on r_{max} ($p=0.0176$) but not on c_{50} ($p=0.4920$); $n=14$.

3.3.3 Attentional effects

Saccades and changes in spatial attention are coupled under natural circumstances; our study was not specifically designed to disentangle them. However the pronounced attentional effects on V1 contrast response functions depend on differences in attended vs. non-attended RF stimulation: attending to a stimulus within a neuron's RF increases responses to that stimulus, especially at higher contrasts (Thiele et al., 2009). We wanted to make sure that pre-saccadic gain reduction was not the result of attention being drawn away from the RF as the subject prepared to make a saccade. In our study we included saccades from the center of the monitor to a new fixation target 6° away, and then back toward the original fixation target. Assuming that attention is drawn to the saccade target, any RF along the saccade axis would be differentially attended when the subject made a saccade away vs toward the center of the monitor. We were able to find 16 units that had RFs along a given saccade axis, and that preferred the grating orientation presented with that saccade direction. (Remember that we presented one of four different grating orientations, with saccades cued along that grating's orientation; see Methods.) We created two conditions: trials in which the saccade target was closer to the RF than the fixation target (meaning attention would be to the RF when the saccade target appeared) and those in which the saccade target was farther (attention drawn away from RF). We reasoned that if our effect was mainly due to attentional modulation then these two conditions would show opposite effects because as the saccade was cued, attention would be drawn to the RF leading to response enhancement in one condition, and drawn away leading to response reduction in the other. We compared responses in each saccade condition ('to-RF' and 'away-RF') to the fixation condition – choosing the time point (0

ms) and contrast (96%) that showed maximal effect in our previous analyses. The result was that both 'to-RF' (22% average reduction; $p<0.017$, Bonferroni-corrected) and 'away-RF' (10% average reduction; $p<0.017$) showed significant saccadic gain reduction compared to fixation. We therefore concluded that spatial attention response modulation was not responsible for our saccadic gain modulation. Interestingly, 'away-RF' showed less average suppression (10% reduction) than 'to-RF' (22% reduction). If spatial attention was a significant factor interacting with suppression in our paradigm then we would have expected a slight response increase in the 'to-RF' condition compared to 'away-RF' – meaning 'away-RF' would have shown more suppression. However the differences between the two suppression conditions was not significant ($p>0.017$).

3.4 Discussion

We showed that peri-saccadic contrast responses in V1 exhibit gain reduction. The effect was especially pronounced at higher contrasts – consistent with a reduction in response gain. Patterns across individual units also suggested that V1 gain reduction is the result of decreasing response gain.

Super et al. (2004) also studied peri-saccadic, sustained activity in V1 – but in the context of how neural activity influences the selection of saccade targets. Contrary to our result they found a gradual increase beginning on average 176 ms before saccade onset. This pre-saccadic response increase precedes the decrease that we report by at least 100 ms. They used a delayed saccade paradigm – where the subject is cued toward a saccade target, then cued again to execute the saccade. Spatial attention is critically altered in

such a paradigm, so it is possible that what they interpreted as pre-saccadic activity is (or is combined with) attentional modulation during the delay period as the subject waited to make the saccade. They raise this possibility in their discussion but unfortunately do not disentangle the two effects in their results (the closest they get is their Figure 5c, however the two conditions in which the RF is stimulated are not directly compared).

Most other previous studies of suppression in V1 were designed to study 'saccade-like' visual stimulation (Wurtz, 1969; Fischer et al., 1981; Kagan et al., 2008) – for example comparing the neural response when a saccade moves an RF across a bar of light to the response when a bar of light is swept across the RF, producing the same retinal stimulation but during fixation. In such a design, any extra-retinal (i.e., non-visual) purely modulatory signal is only observed as an interaction with the visual response – but earlier dynamics of such a signal would be missed. Our design is better suited to searching for extra-retinal gain modulations because of the ongoing activity. Our grating onset occurred at least 200 ms before saccade onset – so that we analyzed the influence of saccades once activity had reached stable rates.

Our results show that gain reduction begins ~60 ms before saccade onset. This is strong evidence of a corollary discharge signal. The origin of this signal could be direct sub-cortical input – perhaps from the inferior pulvinar of the thalamus (Doty, 1983) or from the LGN, or both. The LGN does show some pre-saccadic modulation (Reppas et al., 2002) and is the main relay of visual information to V1. It could be that some of the pre-saccadic modulations observed in this study are the result of feed-forward input from the LGN – particularly the pre-saccadic decrease in firing. However, layer 6 of V1 has been shown to modulate visual response gain through both intra-cortical and cortico-thalamic

circuts (Olsen et al., 2012). This raises the interesting possibility that a top-down signal (perhaps arriving at more superficial layers) is influencing visual responses through layer 6 modulation – and that therefore it is V1 causing the pre-saccadic modulations in LGN. Rajkai et al. (2008) studied the effect of saccades on V1 activity in total darkness. The study focused on post-saccadic effects and so the time courses were aligned to fixation onset (and not saccade onset). The different alignment, along with the fact that saccades in their study were of widely varying durations, makes a direct comparison to onset-locked responses difficult. However they do show a suppression of multi-unit activity (MUA) during saccades, in agreement with our general findings. Close inspection of their laminar recordings reveals a current sink (indicating net inward flowing current) in supra-granular layers during saccades. This could be a top-down intra-cortical corollary discharge. While they show some suppression in total darkness, our results show no suppression when V1 is minimally driven (Figure 3.3, top panel). However their MUA activity is averaged over different cortical layers including the superficial layers that show the current sink. Again, close inspection of their laminar recordings reveals the strongest intra-saccadic MUA suppression at the most superficial layers, with almost no suppression at layer 4 and slight suppression in layer 6. This suggests that an intra-cortical signal could be influencing feed-forward thalamic visual input (layer 4) through layer 6 gain modulation. Rajkai et al.'s (2008) study cannot detect said gain modulation because subjects were in total darkness (i.e., no layer 4 visual input).

There has been no established connection between saccadic suppression and contrast adaptation, but it has been suggested that similar neural mechanisms could be responsible (Burr et al., 1994). Slower contrast adaptation (Ohzawa et al., 1985) has been shown to

shift V1 units' CRFs horizontally along the contrast axis (varying c_{50}) – to match the dynamic range of the CRF to the local contrast distribution. An increase in c_{50} also has the effect of reducing contrast response as if the input contrast was lowered (i.e., shifted in the opposite direction as c_{50}). It could explain the elevated contrast thresholds during saccades (observed psychophysically) as the result of neuronal CRFs shifting higher along the contrast axis, effectively lowering input contrasts during saccades. Again, such a shift would be seen as a horizontal shift of the CRF, and would affect responses at intermediate responses most drastically (Herrmann et al., 2010). Both our analysis of individual unit CRFs as well as population responses do not have these characteristics, but rather show strongest effects at higher contrasts at the population level (Figure 3.4) and reduction of r_{max} at the individual unit level (Figure 3.5). This reduction in response gain is more characteristic of fast contrast adaptation.

Finally our result is important when considering area MT as one of the better established targets of suppression. Neurons in MT are highly tuned to speed and direction of motion, and their activity has been tied more directly to motion perception (Krekelberg et al., 2006) – therefore any saccade-induced motion should have dramatic effects on visual stability at that site. Thiele et al. (2002) found that many MT neurons were suppressed by saccades but fired robustly when the same image motion was played back during fixation – demonstrating that MT processed stimuli differently during saccades. Bremmer et al. (2009) showed that this change in processing began before the saccade – providing strong evidence of an active (corollary discharge) mechanism that drives this saccadic modulation. If MT received this modulatory signal in a feed-forward manner, then the LGN and V1 would probably show similar strength of suppression, but they do

not (Reppas et al., 2002; Wurtz, 1969; reviewed in Wurtz, 2008). Perhaps the modulatory signal arrives in MT directly from sub-cortical inputs (Wurtz, 2008) or through top-down signaling from areas mediating eye movements such as the FEF (Bremmer et al., 2009). Berman and Wurtz (2011) provided evidence for a signal from the superficial layers of the superior colliculus (SC) – which show strong suppression of saccade-induced visual responses – relayed through the pulvinar to area MT. This is strong evidence that at least part of the suppression shown in MT is the result of a pathway that bypasses V1. MT and V1 are densely recurrently connected, so the gain reduction we found may reflect top-down modulation from MT (as discussed earlier). To summarize this view: (1) Saccades produce spurious motion signals that should be eliminated. (2) Motion sensitive cortical areas effectively eliminate this motion because they are strongly suppressed through (3) a sub-cortical corollary discharge signal that bypasses V1. However (4) this signal influences V1 activity through top-down modulation.

The function of peri-saccadic V1 gain reduction can therefore be explained within the above framework of dampening motion signals. Indeed if MT is densely recurrently connected with V1 and relies on its feed-forward signals to process motion, then a suppressive corollary discharge signal reaching MT will (from a connectivity standpoint) and should (from a processing standpoint) reduce visually driven activity in V1. This framework treats visual stability as its functional end (Wurtz, 2008). It treats intra-saccadic motion as a threat to that stability – as a perceptually jarring visual experience – and explains the interesting psychophysical and neural changes during saccades as working to omit intra-saccadic stimuli from perception. However in the next chapter we

explore changes in post-saccadic contrast responses, and show how some of these changes may be beneficial when new visual input is encountered (as happens after a saccade, at a newly fixated scene). Based on the results of the following chapter, we will try to argue that saccadic suppression (i.e., reduced contrast sensitivity during saccades) and its potential neural correlates (e.g., gain reduction in V1) may be worth investigating outside the context of saccadic omission, and in the larger context of active vision.

Chapter 4. Post-saccadic contrast responses in primate V1

4.1 Introduction

Human vision is active: the retina is directed at different regions of the environment with fast, ballistic eye movements (saccades) about three times per second. This ever-shifting visual input is relayed to neurons in the primary visual cortex (V1). Recent work has shown that V1 neurons have a different response to the same visual stimulation after a saccade (i.e., at fixation onset) than during steady fixation. This post-saccadic modulation has been found in general firing rate (Kagan et al., 2008) and spike synchronization (Maldonado et al., 2008), as well as changes in non-visually evoked activity (Rajkai et al., 2008). Is there some functional significance to these differences? Asked another way, does post-saccadic modulation have an effect on V1 tuning properties? This is an especially important question to address if we are to understand V1 processing in the context of active vision. While there has been speculation about this in the above studies none have directly shown any functional significance to the different kinds of neural responses observed in the post-saccadic time window.

In the present study we set out to address this question of functional significance. To approach it we considered the following. First, there is evidence of a saccadic extra-retinal input arriving in V1 (Rajkai et al., 2008) – with speculation about its role in 'priming' the visual system for the new input arriving at the onset of fixation. There is also psychophysical evidence that eye movements reduce the perceptual bias imposed by stimuli at the preceding fixation (Paradiso et al., 2012); here the speculation was that a

corollary discharge signal is critical to this phenomenon. We also considered that neurons in V1 exhibit fast contrast adaptation – reducing their gain in a stimulus-dependent manner (Levy, Fournier, & Fregnac, 2013). Taken together, (1) a post-saccadic modulation arrives in V1 to somehow improve its processing of new stimuli, and (2) contrast responses in V1 are influenced by previous stimuli (adaptation). This led us to investigate whether post-saccadic contrast responses exhibit some form of reset, or release from adaptation induced at the previous fixation.

To test the above hypothesis we recorded from V1 units while stimulating them with a range of contrasts, presented at durations matching those of natural viewing. Our design included a random stream of different contrast gratings, with intermittent cued saccades. This allowed us to study how adaptation at varying contrasts (i.e., adaptation condition) affected contrast response when an adapter was presented and followed by a test contrast during steady fixation (fixation condition) and when an adapter was presented and then a saccade was executed just before the test contrast (post-saccadic condition). We analyzed responses using parametric model fitting (CRF analysis) as well as more direct contrast discrimination measures (ROC analysis). Overall we found no evidence of a post-saccadic reset. Our analysis did however reveal that post-saccadic responses gave neurons a wider operating range: trends in the CRF analysis showed increased firing over a wider range of contrasts after saccades. And post-saccadic post-stimulus response distributions were better at discriminating a wider range of contrasts than were those during steady fixation.

4.2 Methods

The visual stimuli, participants, procedure, and electrophysiological methods are the same as those in Chapter 3. In the present study, the analysis window was 50 ms long and began 50 ms after stimulus onset. Figure 4.1 (top trace) shows the time course of randomly changing grating contrasts (axis on right, % contrast). The bottom trace shows eye traces (axis on left, degrees of visual angle). The middle line shows spikes and the fixation (blue) and post-saccadic (red) analysis windows.

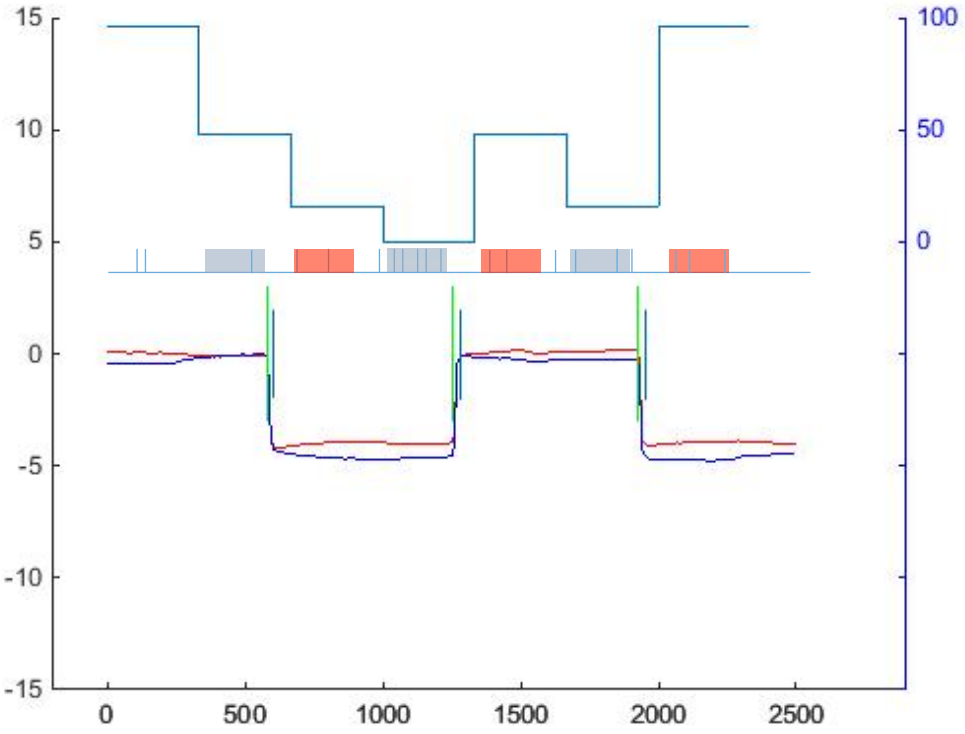

Figure 4.1 – Stimulus and eye position. Our design included a random stream of different contrast gratings, with intermittent cued saccades. Each contrast (top trace, axis on right) acted as both test contrast as well as adapter for the next test contrast. Bottom trace shows eye position for left and right eyes. Following a saccade, a post-saccadic

(red) analysis window can be seen on the middle trace showing recorded spikes. After the contrast changes during steady fixation a fixation (blue) analysis window can be seen over the spike trace.

4.2.1 Data analysis

4.2.1.1 Contrast response functions

We fit Naka-Rushton contrast response functions (CRF) to the average firing rate at each contrast using a least-squares error method in MATLAB (http://www.mathworks.com). The equation was,

$$r(c) = r_{max}\frac{c^n}{c_{50}^n + c^n} + b$$

The free parameters were the maximum firing rate (r_{max}), the semi-saturation contrast (c_{50}), the exponent (n), and the baseline response (b). We categorized a given contrast as low- or high-contrast adapted for a given unit by one of two methods. In the parametric method, if a contrast was less than the unit's semi-saturation contrast it was low-adapting (and high-adapting if it was greater). In the non-parametric method, we found the two neighboring contrasts that led to the largest difference in response; all contrasts less than or equal to the lower of these two contrasts were low-adapting, and all other contrasts were high-adapting. The parametric method of dividing between low- and high-adapting contrasts was used unless the semi-saturation contrast was above the highest test contrast; in this case the non-parametric method was used.

Figure 4.2 shows the effect of changing a single CRF parameter on a reference (black) response curve. In each panel the red curve is the result of increasing the parameter, and

the blue is the result of decreasing it. The r_{max} parameter (top left) scales the entire gain function, and is especially important in determining the maximum response. Changing c_{50} (top right) shifts the curve along the contrast axis. This term is also referred to as the semi-saturation term because it is the contrast value at which the function yields half of its maximum (or saturation) response. The baseline term (b; not shown) acts as an offset along the response axis, shifting the curve up for increasing values, and down for decreasing values. The non-linearity term (n) affects the steepness of the curve about the semi-saturation value: a high 'n' (bottom left, red trace) leads to a steeper curve than a low 'n' (blue trace). One of the implications of this increased steepness can be seen in the left panel of Figure 4.3 – notice how a contrast increment in the range where the high 'n' (red) curve is steepest (c1) leads to a large increment in response (R1) when compared to the response increment produced by the shallower low 'n' (blue) curve. Notice also how a high 'n' leads to wider range of contrasts at which the red curve is relatively flat. These two flat portions of the red curve – the lower contrast portion yielding the sub-threshold, baseline response and the higher contrast portion yielding the saturation, maximum response – show no change in response (i.e., have little to no slope). This is shown in Figure 4.3 (left panel): the red curve produces almost no change in response (R2) to a contrast increment near its saturation range (c2). Notice how the blue curve with its shallower curve still produces some change in response at c2. The relationship between response increment and discrimination is clarified later.

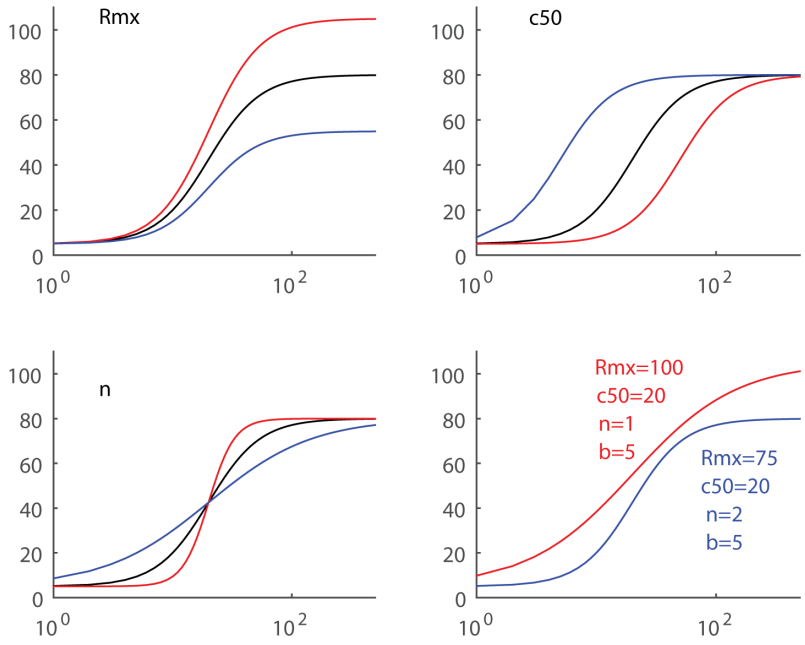

Figure 4.2 – The effect of parameter shifts on the contrast response function (CRF). In each panel the same reference CRF is plotted in black (r_{max} =75, c_{50}=20, n=2, b=5). Changing the r_{max} parameter while keeping the others the same results in a multiplicative scaling of the entire function (top left). Changing the c_{50} parameter results in a horizontal shift of the curve (top right). The 'n' term (bottom left) is the non-linearity and steepens the curve with increasing value. The bottom right panel shows the effect of increasing r_{max} while decreasing 'n' (red curve) compared to a reference curve (in this case blue).

4.2.1.2 Receiver operator characteristic

Our CRF analysis is a parametric approach to describing the contrast encoding of neurons (a transfer function between input contrast and average response across many trials). One of its potential drawbacks (besides being parametric) is that it does not take the variability

of stimulus-evoked responses into account – but this is a crucial aspect of how neurons encode stimuli, as will be demonstrated below.

The left panel in Figure 4.3 shows a low 'n' (blue) and a high 'n' (red) CRF. The right panel shows hypothetical distributions of responses that could have produced the model CRFs in the left panel (because we fit CRFs to the average responses). Note that on both the middle axis and the rightmost axis the blue distributions have the same means and therefore correspond to the same points along the blue CRF (i.e., would have produced the same model fit). However, the amount of overlap between the two pairs of blue distributions is very different in the two cases. In the case where the response across trials is less variable (rightmost axis) there is much less overlap between the blue distributions than in the case where the response across trials is more variable (middle axis; blue distributions). Remember that the two distributions are separated by the same means, and so this change in overlap is strictly due to response variability.

But how is the overlap between a neuron's response distributions to two different stimuli related to discrimination performance? To demonstrate this let us consider the R2 (top) red distributions along the middle axis of Figure 4.3, and try to quantify how reliably an observer (with access only to the neuron's stimulus responses, R21 and R22) can discriminate between the two stimuli, c21 or c22. To do this we establish a simple response criterion for deciding which contrast was shown: if the response is above this criterion value then c22 (the higher contrast) is chosen. Let us place the criterion as shown (dashed black line between the red distributions). How often will we choose the correct contrast? This will be the proportion of trials in which c22 is shown that evoke responses above our criterion value, represented by the area of R22 lying above the

dashed line – also called the true positive (TP) rate. How often will we choose the wrong contrast? This will be the proportion of trials in which c21 is shown that evoke responses above our criterion, represented by the area of R21 lying above the dashed line – also called the false positive (FP) rate. If these two areas (the TP and FP rates) were exactly the same – as they would be with perfectly overlapping distributions – then we would choose the correct contrast as often as we would the wrong one. In other words we would discriminate at chance performance. In our example the response distributions (R21 and R22) are highly overlapping and so discrimination, using our chosen response criterion, would be very close to chance. In ROC analysis we vary the response criterion, plotting the TP against FP at each value – thereby quantifying discrimination performance across all criteria. The area under this curve gives us a more general measure of discrimination performance. As demonstrated above, the more overlapping two distributions are, the worse discrimination will be. Therefore to capture how changes in a unit's mean firing and variability may affect the contrast discrimination capabilities of V1 neurons we used ROC analysis. We computed ROC curves from the raw stimulus-dependent spike distributions, and reported the area under the ROC curve as a measure of discrimination performance (DP).

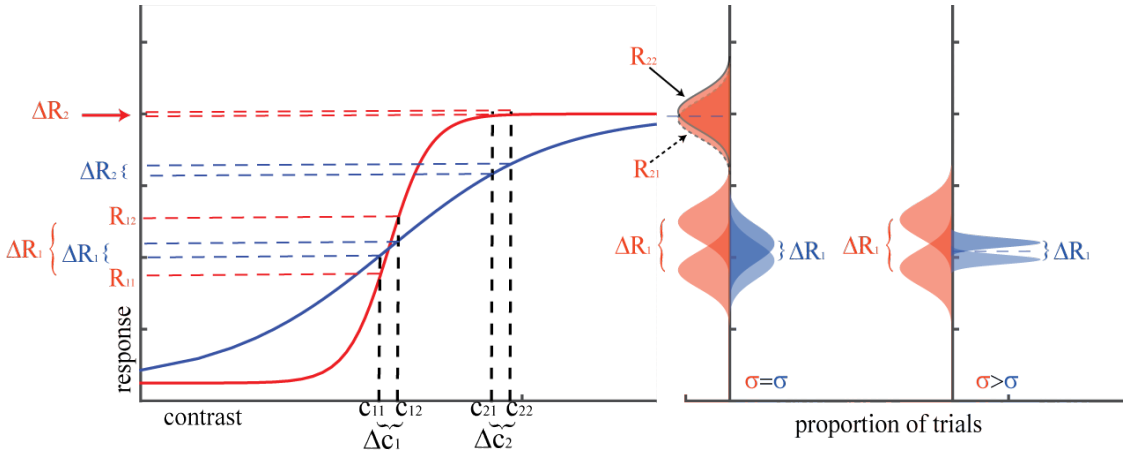

Figure 4.3 – CRF to ROC. The panel on the left shows a high 'n' (red) CRF and a low 'n' (blue) CRF. Response increments (R1, R2) produced for two different contrast increments (c1, c2) by the two curves are shown along the response axis (ordinate). Note the high 'n' (red) curve produces a relatively large response increment (R1) for a contrast increment near its steepest portion (c1), but almost no response increment (R2) for a contrast increment near its saturation range (c2). The low 'n' (blue) curve produces a smaller response (R1) compared to the high 'n' curve at c1, but a greater response (compared to high 'n') at c2. The right panel shows hypothetical response distributions with means separated by intervals as indicated from the curves in the left panel. The middle axis shows the case where the red and blue distributions have equal variability. In this case the blue distributions overlap more than the red distributions. The rightmost axis shows the case where the blue distributions are much narrower (i.e., the stimulus-evoked responses are much less variable) than the red distributions. In this case the blue distributions overlap less than the red distributions. Importantly, the distance between the red distributions (red R1) is the same in both middle and rightmost cases, and so is the distance between the blue distributions (blue R1) – the only factor affecting the change in

overlap is the change in response variability. The middle axis also shows the major overlap of distributions in the case of the small red R2 increment.

4.3 Results

We analyzed a total of 76 units – 18 units in Monkey M, and 58 units in Monkey Y (array and receptive field location as described in Chapter 3 Methods). We set out to investigate whether V1 shows a post-saccadic resetting, or release, from adaptation to stimuli from the previous fixation. We found no 'hard reset' – but we did find some significant changes in post-saccadic contrast tuning. Our main result is that V1 units show post-saccadic improvement at discriminating between contrasts further from an adapting stimulus. During fixation, discrimination between two contrasts is best after adaptation to the lower of the two contrasts – and performance drops as the adapter moves further from the test pair (i.e., discrimination performance is adapter-specific). After saccades, the same general pattern of adapter-specificity is true, but the performance drop is less dramatic as test contrasts and adapting contrast are moved apart (less adapter-specific). Direct comparison between the two conditions shows that this post-saccadic decrease in specificity does not come at the expense of discrimination performance at the adapter. In other words units still performed equally well post-saccadically in the conditions that fixation performance was specialized at, but also improved under other conditions – indicating an overall increase in operating range.

4.3.1 CRF analysis

For the parametric analysis we fit CRFs to V1 responses during steady fixation (fixation condition) and during fixation onset (post-saccadic condition), under low- and high-

contrast adaptation conditions. First we confirmed that adaptation had an effect on CRFs during fixation. For each unit we divided the adaptation conditions into two categories: low-adapted and high-adapted. We reasoned that any contrast that evokes half the cell's maximum firing rate is driving it to a highly-adapted state and any contrast that cannot, keeps it in a low-adapted stated (see Methods). We looked at effects on fixation CRFs (Figure 4.4, left) after low-contrast adaptation (light blue) and high-contrast adaptation (dark blue). The example units show a decrease in overall firing (decrease in r_{max}), a rightward shift along the contrast axis (increase in c_{50}), a rise in baseline firing (b), and a decrease of the non-linearity term (n).

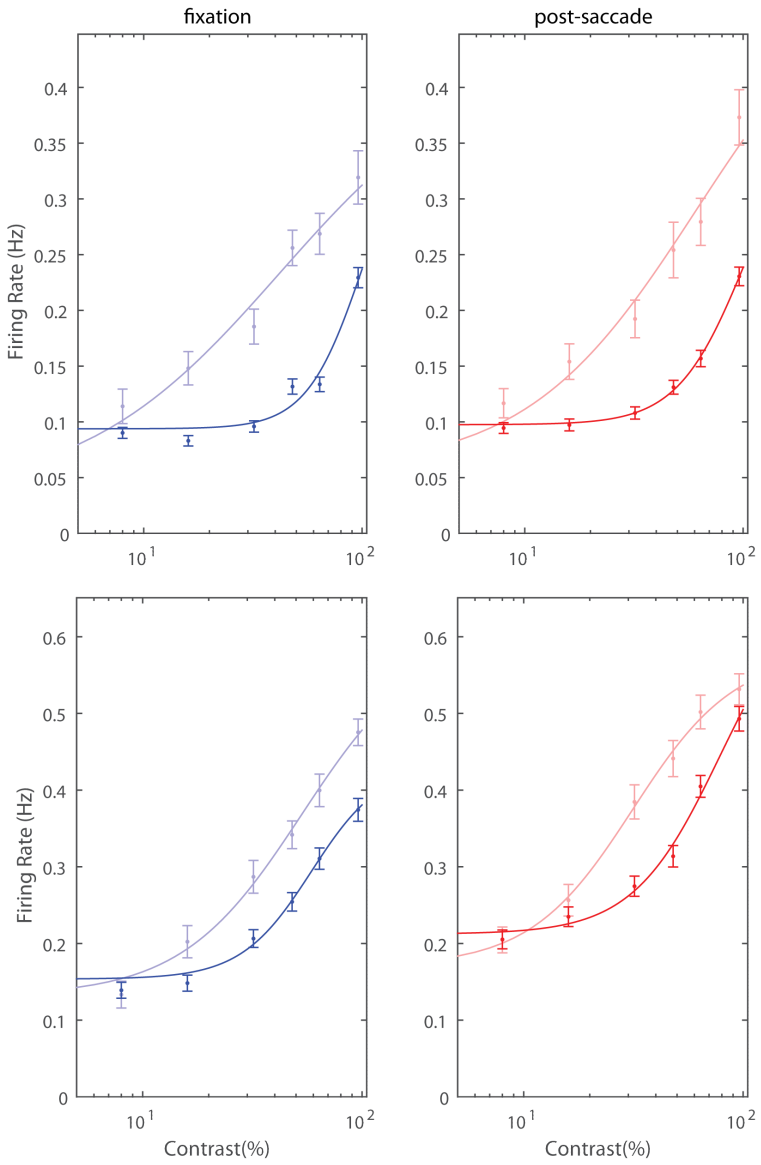

Figure 4.4 – High and low adaptation CRFs of two representative units during fixation and post-saccade. In both fixation (left) and post-saccade (right) conditions, high-adaptation (dark) shows a reduction in r_{max}, an increase in baseline, and a rightward shift in c_{50} compared to low-adaptation (light).

The pattern across the population (Figure 4.5, left) is consistent with the example units in Figure 4.4: during fixation high-contrast adaptation leads to a general decrease in max

firing, a rightward shift in the half-saturation contrast, and an increase in non-linearity ($p<0.001$ for all shifts, paired t-test). These results show that high-contrast adaptation has a pronounced effect on contrast tuning during fixation when compared to low-contrast adaptation. If post-saccadic modulation imposed a 'hard reset' on V1 contrast tuning then adaptation (i.e., tuning biases imposed by stimulus history) from the previous fixation would have no effect on contrast responses at the post-saccadic location. To check for this we looked for post-saccadic changes in how low and high-contrast adaption affected CRF parameters. The same example unit described for fixation is shown in the post-saccadic condition (Figure 4.4, right). A hard reset effect would cause the two curves to be completely overlapping (no effect of adaptation); even a substantial reset effect would bring the two curves closer together than they are in the fixation condition (left panel). However the two CRF curves appear very different, with qualitatively similar changes between them as during fixation. The results across the population (Figures 4.5, right panels) show that adaptation to a high contrast adaptation has a significant effect on all CRF parameters when the adapter is presented at one fixation and the test contrast is presented upon a new fixation ($p<0.001$ for all shifts, paired t-test). This shows that there is no post-saccadic hard reset – adaptation at a previous fixation affects contrast tuning at a new fixation.

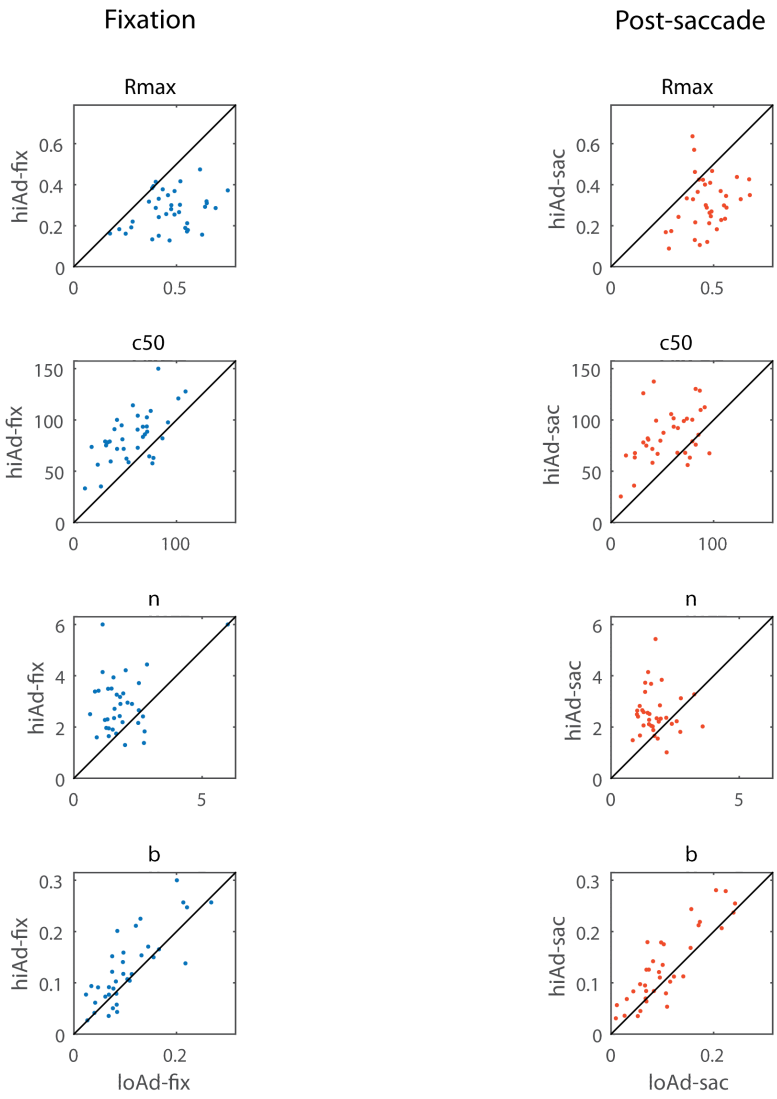

Figure 4.5 – Post-saccadic and fixation CRF shifts (across population) between low and high-contrast adaptation for fixation (left) and post-saccadic (right) data. The effect of adaptation appears to be qualitatively similar in both conditions: higher contrast adaptation decreases max response (r_{max}), while increasing baseline firing (b), non-linearity (n), and the semi-saturation point (c_{50}).

The above analysis demonstrated that post-saccadic modulation does not impose a hard reset on contrast adaptation. However close inspection between fixation (Figure 4.5, left

panels) and post-saccadic (right) conditions does suggest some other interesting patterns. For example it appears that for the panels showing changes in 'n' there is an overall decrease between fixation and post-saccade (a downward shift of the cluster along the diagonal). In order to further explore the effect that post-saccadic modulation has on contrast tuning we compared CRF parameters between the two conditions. In this analysis we were interested in predictions: what could shifts in the CRF parameters tell us about post-saccadic contrast responses? Therefore we restricted our analysis to units that were well-fit by the CRF (Rsq>0.85). Figure 4.6 shows differences in the CRFs for four example units: each shows a post-saccadic increase in r_{max} and decrease in non-linearity.

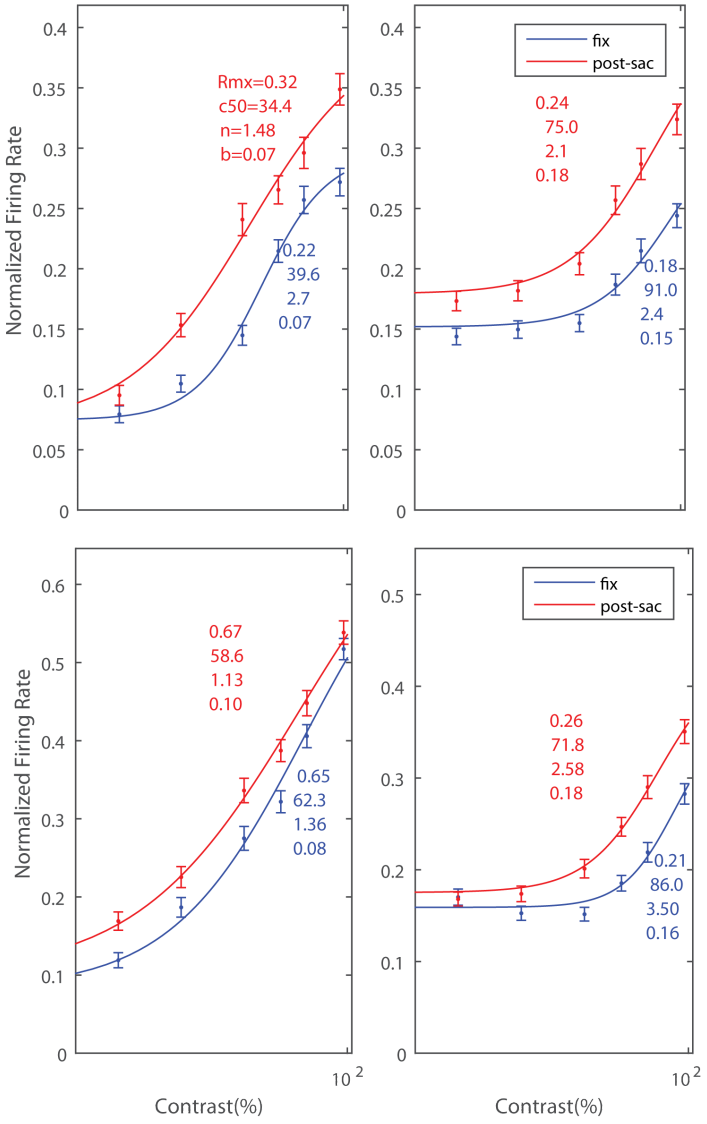

Figure 4.6 – CRFs for post-saccadic (red) and fixation (blue) conditions. CRF parameters are shown on each plot (key printed in upper left plot). Each unit shows a post-saccadic increase in r_{max} as well as a decrease in 'n'.

Figure 4.7 shows the changes between the fixation and post-saccadic CRF parameters for all units included (n=14). There appears to be an overall increase in maximum firing and a decrease in the non-linearity term. Because of the relatively low sample, we used a

Wilcoxon signed rank test to check for significant parameter shifts (r_{max}, p=0.1161; c_{50}, p=0.8022 ; n, p=0.0190; b, p=0.8734).

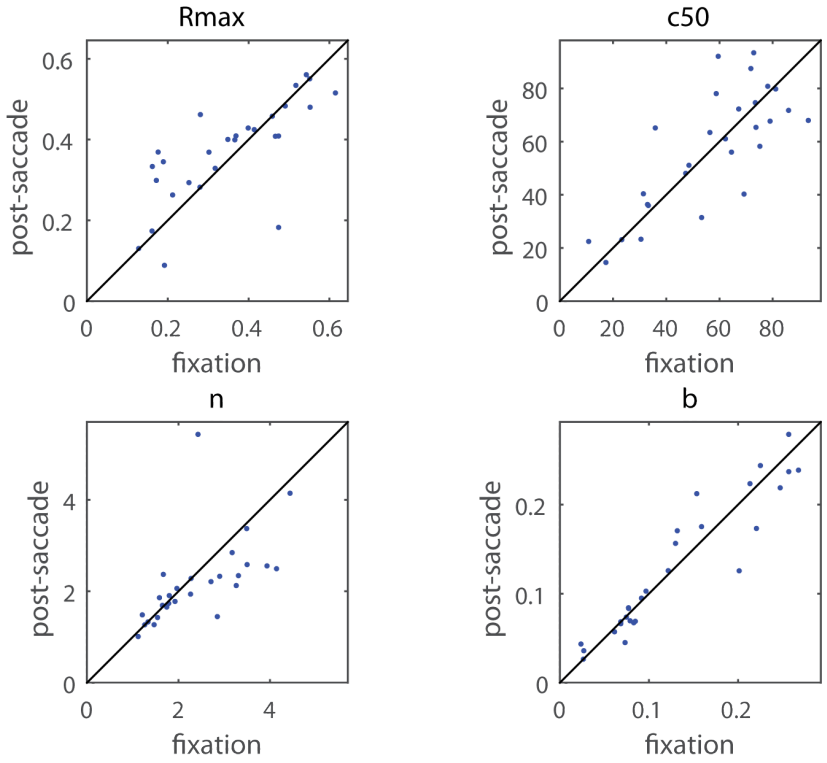

Figure 4.7 – Fixation vs post-saccade CRF parameter shifts (population). Parameter values are plotted for units well-described by Naka-Rushton function (Rsq>0.85, param CV>0.75). There is a tendency for r_{max} to increase post-saccadically, and for *n* to decrease.

Using the CRF parameter shifts shown in Figure 4.7 we constructed an illustration of post-saccadic CRF modulation. This is just a characterization, but it is constructed from trends observed over the population and we use it here as a predictive tool. The model post-saccadic CRF – with increased r_{max} and decreased n – has a wider operating range.

That is, there is a larger range of contrasts over which it generates an incremental response. This makes strong predictions of how well it will perform discriminations over a range of test contrasts when compared to the fixation CRF. In order to extend our investigation of a reset mechanism as well as to directly test some of the predictions made about increased operating range we used ROC analysis.

4.3.2 ROC analysis

CRF analysis is a parametric approach to studying tuning properties. As stated in the Methods section, it does not model changes in response variability. ROC analysis allows us to assess how reliable a unit's responses are for discriminating between stimuli (see Methods) – and this takes into account the entire distribution (including variability) of responses to those stimuli. Again, as with the CRF analysis, we wanted to first confirm that adaptation had an observable effect on an ROC characterization of contrast responses. We did this in the following way. We calculated discrimination performances (DP) between each pair of neighboring contrasts (test pair), for all adapters; this yielded a DP matrix of 6 (the number of nearest-neighbor test pairs) by 7 (the number of adapters). For a given test pair (i.e., column) we then divided by the average DP across all adapters (column average). The resulting matrix (Figure 4.8, left) shows that during fixation, a given pair of contrasts are best discriminated after adaptation to the lower of the pair. This can be observed by choosing a test pair (e.g., 8:16% column) and checking which adapting contrast (i.e., row) leads to the highest DP value (8%, in this case). Another way of describing this effect is that detection of an incremental (i.e., smallest interval increase) change at a given contrast is best after adaptation to that contrast. The effect is consistent across all test pairs as is apparent from the pattern across the DP matrix's

diagonal: for each test pair the best discrimination is achieved at the adapter. We can therefore say that discrimination performance during fixation has the quality of being adapter-specific. We then asked, how does an adapting contrast (at a previous fixation) affect discrimination at a new fixation? To address this we repeated the above analysis, this time using post-saccadic firing-rate distributions. Again, a resetting of visual detectors would eliminate such a relationship between adapter and test pair as was demonstrated during fixation. In fact a hard reset would eliminate all effect of adapter on test pair and result in an arbitrary pattern (because of noise). The results (Figure 4.8, right panel) however appear to retain the basic diagonal structure as during fixation. The same type of test pair (column) inspection shows that for low contrasts, post-saccadic discrimination has the same qualities as during fixation – with the best discrimination between a test pair being after adaptation to the lower of the two. One difference with fixation is that the relationship between improved discrimination and adapter appears less specific: in the fixation (left) panel there is a more pronounced improvement in discrimination specifically at one adapter, whereas the post-saccadic (right) panel shows a shallower drop-off of performance when adapted to neighboring (above and below) values. It therefore appears that post-saccadic discrimination is less adapter-specific.

This is also consistent with the CRF analysis showing a trend of reduced ‘n’ in post-saccadic tuning (Figure 4.7, bottom left panel). Figure 4.3 (left panel) shows qualitatively how varying n leads to changes in the steepness of the tuning curve; a lower ‘n’ (blue curve) leads to a shallower contrast tuning curve. Lowering ‘n’ alone (keeping all other parameters the same) would have the effect of increasing the range of contrasts that evoke a minimum criterion response (Figure 4.3, left panel – compare low-n blue R2

to high-n red R2) – but at the expense of increased response criterion at some other contrasts (compare low-n blue R1 to high-n red R1). In other words lowering 'n' increases the operating range (assuming equal variability) but at the cost of 'expertise'. Expertise in this case is a smaller range of contrasts (Figure 4.3, left panel – c1) over which incremental responses are very large (high-n red R1) and therefore discrimination very reliable (again, assuming no changes in variability). However our CRF data also indicate an increase in r_{max} (Figure 4.7, top left panel). Changing both of these parameters together makes the apparent tradeoff between operating range and expertise more complex. In any case a proper investigation of the above CRF-based predictions should take response variability into account (explained in Methods) so we proceeded with further ROC analysis.

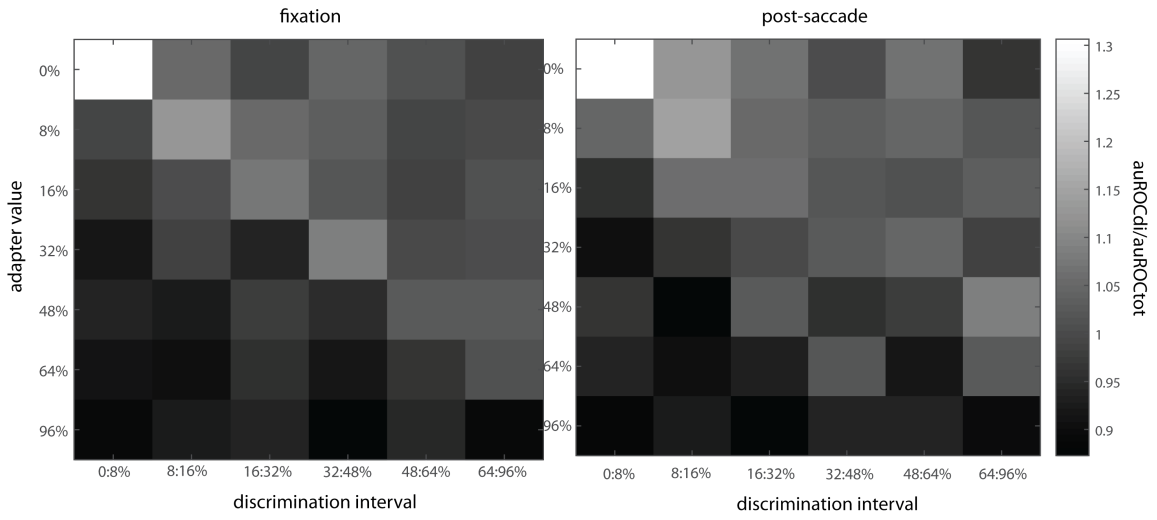

Figure 4.8 – Adapter specificity of contrast discrimination. DP was calculated during fixation (left), and during saccade (right) between each test pair (columns) after each adapter (row), then divided by the total adapter-averaged performance for a given test

pair. The result illustrates which adapting value makes discrimination best for a given test pair. Discrimination of a given test pair during fixation (left) is best after adaptation to the lower of the two contrasts being discriminated – meaning discrimination is adapter-specific. The same is true for low contrast discrimination after saccades (right) – but the adapter-specificity appears weaker. When discriminating between high contrasts there seems to be no pattern of specificity.

It is important to note that the above is a qualitative depiction of adaptation's effect on discrimination during fixation and post-saccade. It does not allow for a direct comparison of discrimination performance between the two conditions. Perhaps the post-saccadic increase in operating range comes at the cost of expertise (i.e., higher performance over a limited range of contrasts). Another way to think of this is that during fixation a unit is a specialist at detecting contrast increments near the adapter, but after a saccade it loses this expertise, becoming a 'jack-of-all-trades'. In order to test this we directly compared between post-saccadic and fixation contrast discrimination as the test pair was moved further from the adapter. For each unit we calculated the difference in DP between post-saccadic and fixation conditions – a positive value indicated a higher post-saccadic DP. Figure 4.9 shows histograms of these differences between post-saccadic and fixation DP for the entire population. The groups (top to bottom panels) are defined by the distance between the test pair and the adapter – with distance starting at a minimum (adapter is one of the test pair), and increasing to a maximum (see Figure 4.9 inset). The results show that at contrast pairs one value away from and adapter, the population is on average better after a saccade than during fixation (paired t-test, $p<0.008$). This is in agreement with the decrease in adapter-specificity shown previously

(Figure 4.8). Indeed the shallower drop off in performance in the post-saccadic condition (Figure 4.8) leads to better performance in discriminating contrasts further from the adapter (Figure 4.9, 2nd panel). Interestingly, performance in the post-saccadic condition is not worse than fixation at the adapter (Figure 4.9, top panel); this shows that there is no loss in 'expertise' when discriminating around the adapting contrast. The conclusion is that post-saccadic contrast responses show an increased operating range compared to responses during fixation.

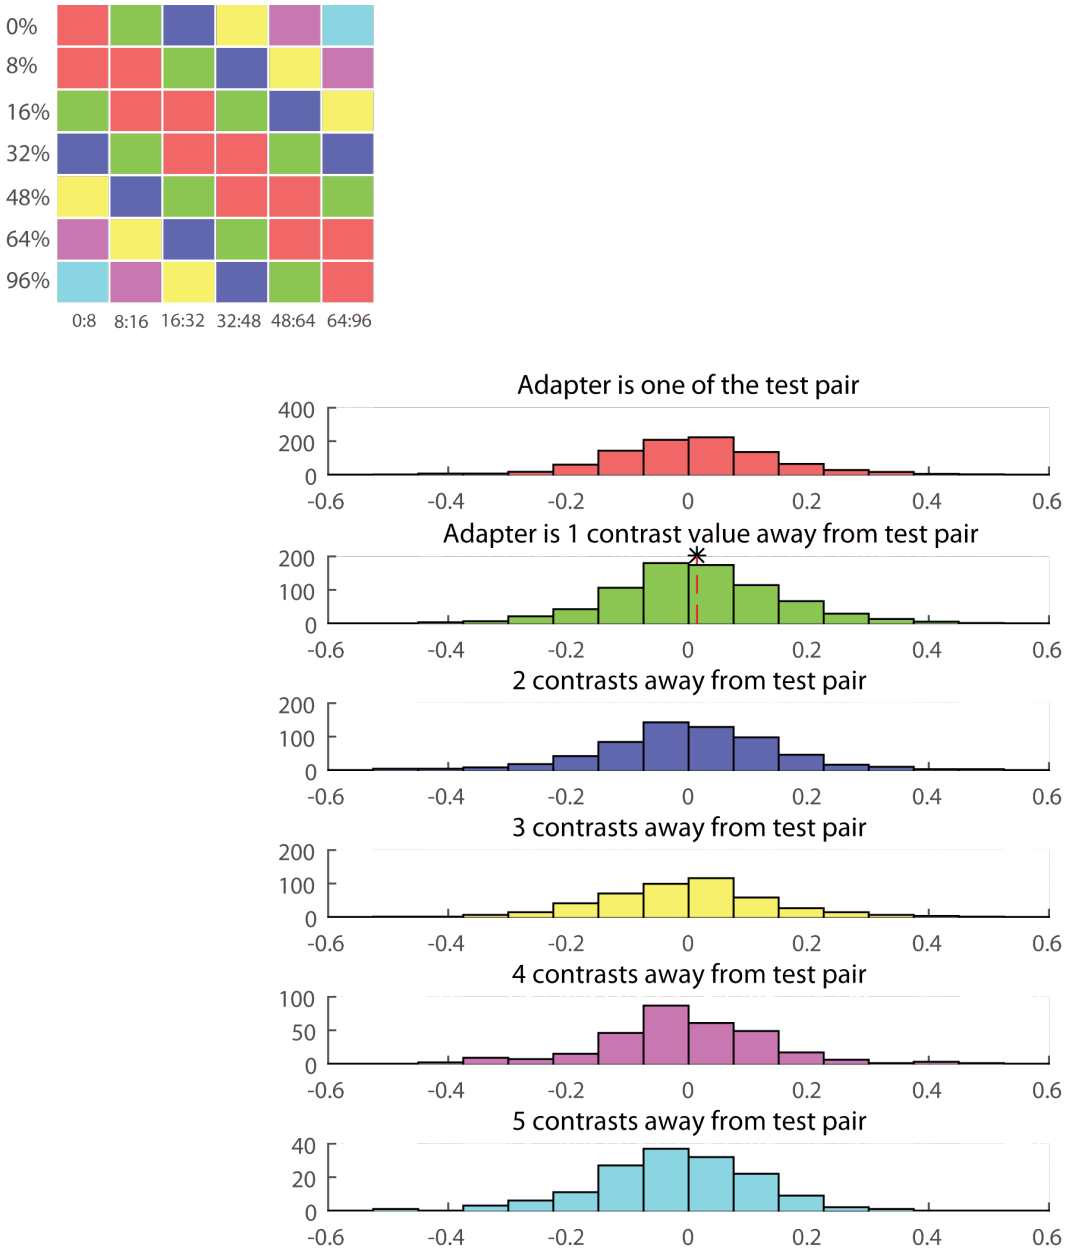

Figure 4.9 – Distribution of difference between post-saccadic and fixation discrimination (population). The data are split into panels as distance between test pair and adapter is increased. The top panel shows the case where the adapter is one of the contrasts being discriminated; the 2nd panel shows the case where the adapter is one value away from one of the contrasts being discriminated; etc. Significance (*) indicated when $p<0.008$ (Bonferroni corrected $p<0.05$).

4.4 Discussion

Our results did not show a complete reset of contrast adaptation from previous fixations. Parametric analysis showed qualitatively similar adaptive shifts in both fixation and post-saccadic CRFs. A direct comparison of fixation to post-saccadic CRFs did reveal a post-saccadic increase in maximum firing and a decrease in non-linearity. This suggested an increase in operating range, which we confirmed with ROC analysis. Therefore our results directly show a novel functional significance to the post-saccadic modulations found in V1: they expand the range of discriminable contrasts. We will discuss how our results relate to the notion of a post-saccadic reset. We will also discuss the possibility that the post-saccadic change in contrast processing that we observed is itself an optimization strategy – a shift better-suited to cope with the uncertainty about stimulus parameters at a newly fixated scene.

Paradiso et al (2012) have suggested a 'reset' or release from perceptual bias induced at a pre-saccadic location. They measured the extent to which orientation perception at fixation onset was biased by an orientation presented at the previous fixation. They did not find a complete reset – subjects' judgments were still systematically affected by stimuli at the previous fixation. But they did find a significant reduction in bias (compared to a simulated saccade control). We also show that the influence of stimuli from a previous fixation are not completely reset at a new fixation – but that there is a post-saccadic reduction in the previous fixation's influence. In our case this is shown as a reduction in post-saccadic adapter specificity. Paradiso et al. (2012) make it a point that their effect is due to influence – but not adaptation – from previously

fixated stimuli. In their study the stimulus at the previous fixation is presented for 250 ms, and they state that this is too short a duration to induce a significant tilt after-effect. This is probably to differentiate between their effect and the previously shown remapping of orientation adaptation that relies on the tilt after-effect (Melcher, 2007). Our stimulus duration was 333 ms; while this is too brief to induce slow contrast adaptation (Ozhawa, 1985), it is certainly long enough to induce fast contrast adaptation (Levy, Fournier, & Fregnac, 2013).

Previous work on post-saccadic modulation has speculated on its functional significance. Paradiso et al. (2012) concluded that post-saccadic modulation functions as a parsing mechanism to organize discrete percepts from the ongoing stream of neural activity during active vision – in other words the role is scene segmentation. Maldonado et al. (2008) studied V1 responses to natural scenes and found an increase in firing as well as spike synchronization at fixation onset. Similarly to Paradiso et al. (2012), they suggested that this increased synchrony played a role in scene segmentation during natural viewing. Rajkai et al. (2008) studied post-saccadic effects on background (non-visually stimulated) activity in V1 and found that neuronal oscillation phase is more organized at fixation onset. They speculated that this is due to a saccade-related nonvisual input whose role is to prime the visual system for new stimuli encountered upon fixation. We take a similar view to Rajkai et al. (2008). A more specific view on how this priming mechanism could improve contrast processing is as follows.

Adaptation to the local distribution of contrasts in a scene improves visual processing: Dannemiller & Stephens (2000) showed that contrasts near that of a short adapting stimulus (on the time scale of fixations) are easiest to discriminate. But because

of the statistics of natural scenes, adaptation to contrast at one fixation is likely to be suboptimal at the next (Frazor and Geisler, 2006). Perhaps post-saccadic reduction in adapter-specificity serves to pre-emptively relax these suboptimal settings. A more elaborate network level explanation can be given in terms of contrast gain normalization – a phenomenon by which the contrast response of a single neuron is normalized by the pooled activity of other neurons in its locale (Carandini and Heeger, 2012). Contrast normalization is thought to explain important response properties of single units in V1, but it has also been shown to play a critical role in increasing the efficiency with which neurons code natural scenes (Schwartz and Simoncelli, 2001). Furthermore fast contrast adaptation (on the time scales studied in our experiments) in V1 is thought to underlie this process (Levy, Fournier, & Fregnac, 2013). These points together show that in V1 fast contrast adaptation works on the time scale of typical fixation durations to increase coding efficiency through network level computations (i.e., normalization by contrast responses of a surrounding population). Again, if in natural scenes there is very little correlation of contrast in a receptive field from one fixation to the next (Frazor & Geisler, 2006) then divisive normalization – which hinges on the contrast distribution falling on a group of receptive fields – will have to be recomputed. Then the post-saccadic increased operating range (i.e., range of contrasts to which a neuron will discriminably respond) could help facilitate contrast normalization at new fixations as follows. Consider a pattern of different contrasts falling on a group of neighboring receptive fields. Contrast normalization tells us they will each respond to both the contrast within their receptive field along with the normalization factor – a function of the whole group's responses. Our results also tell us that during fixation they will adapt to the contrast within their

receptive field and subsequently be specific in discriminably responding to contrasts near to that value. If they are suddenly presented with a new scene (but without any of the post-saccadic changes we showed) they will each likely be presented with a very different contrast – and because of their limited operating range will be less likely to respond to that contrast in any discriminable way. This is especially disastrous because the response properties of each neuron rely on the renegotiation of the entire group – but each is now saying very little about the contrast within its receptive field. So the slightly increased operating range that we show post-saccadically could pay out in a big way at the network level, especially considering the uncertainty about contrast distributions upon a new fixation.

Paradiso et al. (2012) suggested that a corollary discharge signal is likely responsible for the post-saccadic perceptual un-biasing effect that they observed. Rajkai et al. (2008) also concluded that the post-saccadic changes in V1 were due to a CD. Our results here combined with those of Chapter 3 – showing a pre-saccadic gain modulation in V1 – support the notion as well. This raises the possibility that the pre-saccadic reduction in contrast responses that we observed in Chapter 3 play a role in the post-saccadic enhancement of contrast processing demonstrated here. This has novel implications for theories on saccadic suppression. Perhaps what has been referred to as saccadic suppression – and been considered a loss of contrast sensitivity subservient to 'visual stability' – could be viewed within a larger framework of active vision. As part of an optimal strategy to increase the operating range at the new, post-saccadic scene. This is discussed in more detail in Chapter 5 (General Discussion).

Chapter 5. General Discussion

Visual perception and neural responses change around the time of saccades. This dissertation included investigations of the neural mechanisms of saccadic suppression, and of post-saccadic neural response modulations.

Chapter 2 of this dissertation described the kinds of neural response changes that could lead to saccadic suppression at the psychophysical level. It showed that the lowered contrast sensitivity exhibited during saccades is the result of reduced detector gain. Chapter 3 studied sustained neural activity in V1 at different contrasts during saccades and found that the population showed a reduction in response gain. In the introductory chapter I posed the question of how visual input during saccades is poorly detected. In other words, what are the neural mechanisms of saccadic suppression? The conclusions of Chapters 2 & 3 together tell us that during saccades contrast sensitivity decreases in part because of reduced response gain in V1. Furthermore this reduced response is mediated by a corollary discharge signal, probably arriving in a top-down intra-cortical manner and affecting gain modulation of feed-forward visual input.

The experiment in Chapter 4 was partially motivated by questions about functional significance: post-saccadic modulation has been previously demonstrated in V1 – but does it have any value in the context of active vision? Our original hypothesis was that post-saccadic modulation serves to reset visual detectors. We did not expect a complete reset – but based on our reasoning (that any adaption from a previous fixation would be less informative at a new fixation), and recent psychophysical reports that

perceptual biases from previous fixations are reduced (Paradiso et al., 2012), we did expect some reduction in the effect of past stimulation on current responses. We showed as much in the form of a decrease in the adapter-specificity of contrast discrimination. We also showed an increase in operating range, and we argued that this property is particularly advantageous (in terms of stimulus processing) when fixating a new scene. Therefore the conclusions of Chapter 4 have a direct functional significance within the framework of active vision. Chapters 2 and 3 however require more explanation.

5.1 The role of suppression

Saccadic suppression (an established psychophysical phenomenon) has been studied extensively, especially in the context of visual stability (for a review see Wurtz, 2008) – which is a perceptual framework. In other words the role of saccadic suppression is explained in terms of what it buys us on a perceptual level: without suppression we would perceive jarring (i.e., de-stabilizing) motion signals with each saccade. Even within this framework, saccadic suppression takes a supplementary role as follows. Everyday vision is experienced as an uninterrupted series of inter-saccadic fixations – as if the saccadic intervals were omitted from visual experience. Campbell & Wurtz (1978) referred to this phenomenon as saccadic omission. In an attempt to understand saccadic omission (and whatever role saccadic suppression may have in it) they were able to show that the dominant factor in omission was an effect known as temporal masking: the fixation periods before and after a saccade have the effect of eliminating (or masking) from perception the smeared or 'greyed-out' image generated during saccades. This masking effect depended on the relative contrasts of the masking and the masked stimuli, as well as their durations. However there may be circumstances in which the relative

contrasts of the (masking) fixation and (masked) intra-saccadic stimuli are not favorable for temporal masking – for example fixating a cloudless sky (low contrast) and then making a saccade across the hard line of the horizon (high contrast edge) to fixate a largely flat prairie (low contrast). In this case saccadic suppression – by reducing contrast sensitivity to intra-saccadic input – may further decrease the effective contrast of stimuli between fixations, further guaranteeing that they will be masked and therefore omitted from awareness (Ibbotson and Cloherty, 2009).

Here we can expand on the role of suppression by taking a computational view. Instead of assuming that visual stability would deteriorate without saccadic suppression we may ask, what specifically (i.e., in computational terms) does saccadic suppression improve about active visual processing? Perhaps the gain reduction that we found during saccades (Chapters 2 & 3) works to increase the post-saccadic operating range (Chapter 4). In this view, saccadic suppression (the psychophysical phenomenon) is the behavioral fallout (i.e., an epiphenomenon) of a slew of detectors pre-emptively changing their gain functions to better cope with the uncertain contrast distributions at the post-saccadic scene. The above theory has not been directly proven in this thesis. My post-hoc attempt based on the data gathered was as follows. I calculated a suppression index for each V1 unit (Chapter 3) based on how much saccadic gain reduction it showed at saccade onset (the time at which the population showed maximal suppression). I calculated a discrimination index for each V1 unit (Chapter 4) based on how much post-saccadic discrimination improvement it showed. I reasoned that a correlation between neurons' saccadic gain reduction and post-saccadic discrimination improvement would lend support to the above theory – however I found no such correlation. It is possible that the

corollary discharge (reaching extra-striate areas through direct sub-cortical input in a more temporally precise manner) arrives in V1 in a top-down manner over a relatively longer time window. In that case different units may show maximal suppression at slightly different times, and my selection criterion based on the time bin that showed slightly more suppression on the population level (0 ms) would lead to an inaccurate ranking of suppression indices (and therefore no significant correlation with the discrimination indices).

My working hypothesis is that during saccades an intra-cortical top-down signal arrives in V1 (Rajkai et al., 2008) from various extra-striate areas shown to exhibit strong saccadic response modulation (Bremmer et al., 2009). This signal begins pre-saccadically (Chapter 3) but may be temporally spread (because of multiple cortical sources) over the course of the saccade. It reduces the response gain on neurons, therefore most dramatically decreasing the responses of neurons with the highest contrast responses (Chapter 3). These high contrast response neurons had (before gain reduction) the largest effect on contrast normalization at the local network level: contrast normalization works by intra-cortical inhibition (Heeger, 1992) so that active neurons are more suppressive on the group. It is worth noting that it is the specific kind of gain reduction (i.e., response gain) shown in Chapter 3 that would most effectively lead to a network level re-normalization. A slight change in c_{50} could leave many neurons that were firing at saturation before gain reduction to still be firing maximally afterwards (Figure 4.2, top right panel – note that highest contrast responses would remain near maximum firing on any of the three CRFs). These units would still exert a large inhibition on surrounding ones. On the other hand a reduction of r_{max} (shown at the unit

level, Figure 3.5) can lead to re-normalization by reducing highest contrast responses the most (Figure 4.2, top left panel). This also suggests that saccadic gain reduction works to re-normalize by relaxing lateral inhibition. It is interesting to note that response saturation in V1 is attributed to lateral inhibition (Heeger, 1992); removing GABA mediated inhibition causes cells to fire as much as twice their normal saturated rates (Debruyn and Bonds, 1986). Going back to the saccade-aligned population CRFs (Figure 3.4) we carried out the time courses until the approximate time when the suppression signal begins to rebound in extra-striate areas (Bremmer et al., 2009; Figure 1.1). We reasoned that if this signal was the controlling source of top-down gain reduction (discussed in Chapter 3) then as it let up we would see high contrast responses increase once again, and that furthermore if it had indeed relaxed lateral inhibition then these high contrast responses would be (temporarily) at their highest – briefly unchecked by the saturating effects of lateral inhibition (Heeger, 1992; DeBruyn & Bonds, 1986). This time course is shown for both subjects in Figure 5.1. The CRFs are still aligned to saccade onset (0 ms), but we leave out the pre-saccadic time windows discussed in Chapter 3 (Figure 3.4). Instead we looked at the time course from 20 ms after saccade onset until 80 ms after saccade onset – when extra-striate responses have almost completely rebounded (Figure 1.1). At ~80 ms we expected CRFs to show the post-saccadic increases described above. Monkey M (Figure 5.1, top) nicely shows a post-saccadic increase at 80 ms, and the CRF is notably unsaturated. Monkey Y (Figure 5.1, bottom) also shows an increase but in the previous time bin (60 ms), with highest contrast responses already saturating by 80 ms. Perhaps our choice of time bins did not match the window of maximal post-saccadic response in Monkey Y, and by 80 ms the inhibitory

processes had already led to a more saturated CRF at the highest contrasts (Figure 5.1, bottom right panel). But overall this appears to support the hypothesis that the gradual course of saccadic gain reduction leads to a disabling of surround suppression and culminates (once top-down modulation releases) in a temporary, high-firing, and relatively unsaturated CRF. The time bin at which this CRF springs up in our data (60-80 ms) corresponds roughly to when the suppressed activity recovers in extra-striate areas (see Figure 1.1). Therefore it seems that while V1 is receiving this top-down modulation there is a gain reduction of high contrast responses, leading to a relaxing of lateral inhibition as described earlier. Once this top-down modulation ceases, the feed-forward responses are free of both top-down gain reduction as well as lateral inhibitory normalization, leading to a post-saccadic, mostly feed-forward driven (i.e., uninhibited by pooled surround) increase in activity. And this short window of primarily feed-forward driven response – free of the pooled, contrast-dependent influence from previously fixated stimuli – is what allows neurons to temporarily respond to contrasts more discriminably (Chapter 4) and therefore supports an efficient re-normalization of responses to newly fixated stimuli.

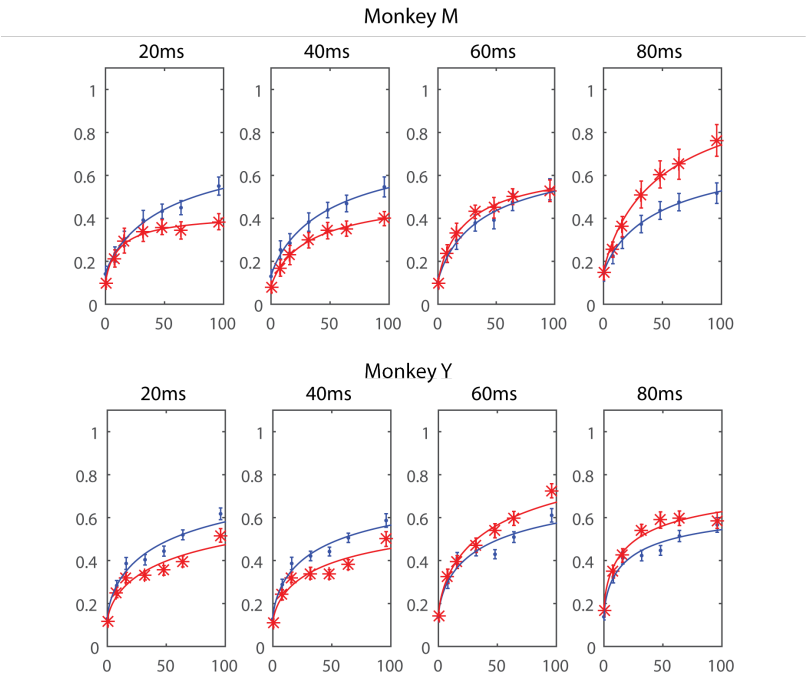

Figure 5.1 – Time course of population averaged CRF (post-saccade) aligned to saccade onset (0 ms, not shown), from 20 ms to 80 ms. The top row shows contrast response functions for Monkey M. Saccade data (red asterisk) and fixation data (blue dots) are plotted with Naka-Rushton functions (for visual comparison) – fit independently to each 20 ms time bin. The bottom row shows the same for Monkey Y.

5.2 Future Experiments

Using the discussions in the above sections and a more detailed description of a V1 contrast normalization model, I will lay out some future experimental predictions in this section. Figure 5.2 shows a simple schematic of a contrast normalization model. The first stage acts as a linear filter and this output is then divided by the summed activity of a population of neurons. Importantly, this (normalizing) population can have different tunings for orientation and spatial frequency than the (normalized) neuron (Heeger, 1992).

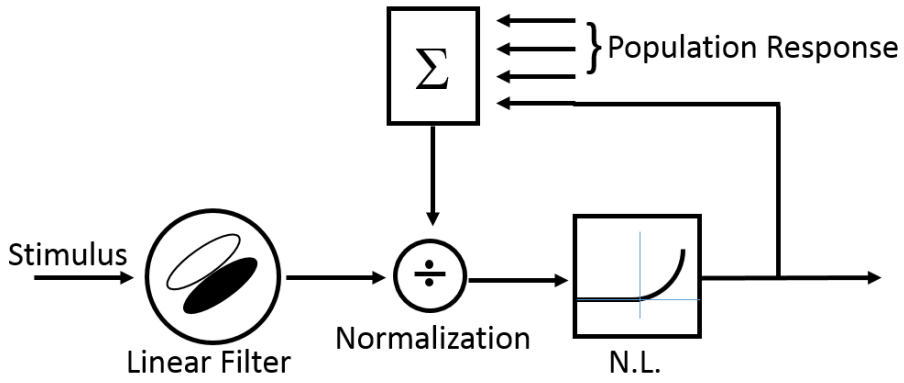

Figure 5.2 – Contrast normalization model. Response for a neuron in V1 can be modeled as a linear response, divisively normalized by population activity, and then passed through a non-linearity (N.L.) stage (based on Heeger, 1992).

If (as explained in section 5.1) there is a brief post-saccadic re-normalization window, during which a neuron's response is dominated by feed-forward input, with minimal divisive inhibition, then some of the response properties which have been previously explained by contrast normalization should be temporarily altered. For example, the different saturated responses (Rmax) of a neuron when tested at different spatial frequency gratings should not be possible. This is a demonstrated property of striate neurons that has been explained by Heeger's normalization model (1992). If for example Neuron A is not selective for a particular spatial frequency then its linear filter stage has little output even as the contrast is increased, however other neurons in the population that are selective for this spatial frequency will respond increasingly with contrast and inhibit Neuron A. The result is that Neuron A will have a relatively low Rmax when tested at this spatial frequency. When Neuron A is tested at a spatial frequency for which it is selective relative to the population then it will have a relatively high Rmax. Again, because this property results from an interaction between the neuron's feed-forward

response and the normalizing populations', I predict that just after a saccade striate neurons will show the same Rmax regardless of which spatial frequency they are tested at: the saturated response will be a function of the linear filter (feed-forward) stage and the (static) non-linearity stage (Figure 5.2) and therefore not be able to shift as a result of the relative spatial frequency selectivities present in the population. Again, this would support the idea (explained in section 5.1) that saccadic gain reduction works most on the highest responding neurons (i.e., the most influential normalizers), and that when it is released there is a brief time window during which contrast responses are primarily feed-forward. In order to further demonstrate the functional benefits of the renormalization hypothesis – that saccadic gain reduction facilitates renormalization – I would also analyze onset responses to various combinations of center and surround stimulation during fixation and after saccade. I predict that responses reach a given information content criterion faster after a drastic change in the relative center-surround contrast when that change is brought into the receptive field by a saccade vs. when it is changed in the receptive field during steady fixation. This is because a drastic change in center-surround contrasts should lead to some renormalization. In the saccadic case the gain reduction should pre-emptively reduce responses to some of the previous contrasts, therefore getting a head start on renormalization to the (new) uncertain contrasts.

Finally, in Chapter 1 I described the two issues that are frequently related to visual stability: saccadic omission and saccadic displacement. If we understand neural changes during saccades (usually explained as serving saccadic omission) as serving to improve visual processing during fixation, then we can try to relate them to how the saccadic displacement problem is solved. For example if neurons in the early visual pathway do

indeed carry eye position signals – and these are crucial for a stable visual representation – then perhaps saccadic gain reduction facilitates the accurate representation of eye position upon a new fixation. Or perhaps saccadic gain reduction should be considered in the context of receptive field remapping: If the visual system solves the problem of displacement by having certain neurons briefly but simultaneously respond to the present as well as future (retino-centric) receptive fields then perhaps it is necessary to reduce these neurons' gain during this time window – or else intermediate contrasts at both locations would lead to maximal (non-discriminable) responses in the neuron. These are just rough examples. The point is to try to relate these different phenomena – saccadic gain reduction and the range of proposed mechanisms underlying spatially stable vision – as part of a unified dynamic process.

5.3 Conclusion

In conclusion we have demonstrated saccadic gain reduction in V1 (Chapter 3), predicted by a computational description of saccadic suppression (Chapter 2) – and we have linked this effect on a mechanistic level with the post-saccadic improvements in processing demonstrated in Chapter 4. The more general conclusion is that the expanded framework of active vision can provide a useful perspective. In this thesis the phenomenon of saccadic suppression, until now defined as a type of sensory information loss, was explained in terms of a more general strategy to improve stimulus processing and thereby gain information.

References

Berman RA, Wurtz RH (2011) Signals Conveyed in the Pulvinar Pathway from Superior Colliculus to Cortical Area MT. J Neurosci 31:373-384.

Binda P, Cicchini GM, Burr DC, Morrone MC (2009) Spatiotemporal distortions of visual perception at the time of saccades. J Neurosci 29:13147-13157.

Bremmer F, Kubischik M, Hoffmann KP, Krekelberg B (2009) Neural Dynamics of Saccadic Suppression. J Neurosci 29:12374-12383.

Buisseret P, Maffei L (1977) Extraocular proprioceptive projections to the visual cortex. Exp Brain Res 28:421-425.

Burr DC, Morrone MC (1996) Temporal impulse response functions for luminance and colour during saccades. Vision Res 36:2069-2078.

Burr DC, Morrone MC, Ross J (1994) Selective suppression of the magnocellular visual pathway during saccadic eye movements. Nature 371:511-513.

Campbell FW, Wurtz RH (1978) Saccadic omission: why we do not see a grey-out during a saccadic eye movement. Vision Res 18:1297-1303.

Carandini M, Heeger DJ (2012) Normalization as a canonical neural computation. Nat Rev Neurosci 13:51-62.

Castet E, Jeanjean S, Masson GS (2001) 'Saccadic suppression'- no need for an active extra-retinal mechanism. Trends Neurosci 24:316-318.

Dannemiller JL, Stephens BR (2000) Contrast discrimination under temporally varying contrast conditions. Percept Psychophys 62:1039-1050.

Dassonville P, Schlag J, Schlag-Rey M (1992) Oculomotor localization relies on a damped representation of saccadic eye displacement in human and nonhuman primates. Vis Neurosci 9:261-269.

Debruyn EJ, Bonds AB (1986) Contrast Adaptation in Cat Visual-Cortex Is Not Mediated by Gaba. Brain Res 383:339-342.

Diamond MR, Ross J, Morrone MC (2000) Extraretinal control of saccadic suppression. J Neurosci 20:3449-3455.

Doty RW (1983) Nongeniculate afferents to striate cortex in macaques. J Comp Neurol 218:159-173.

Duhamel JR, Colby CL, Goldberg ME (1992) The updating of the representation of visual space in parietal cortex by intended eye movements. Science 255:90-92.

Durand JB, Trotter Y, Celebrini S (2010) Privileged processing of the straight-ahead direction in primate area V1. Neuron 66:126-137.

Fischer B, Boch R, Bach M (1981) Stimulus versus eye movements: comparison of neural activity in the striate and prelunate visual cortex (A17 and A19) of trained rhesus monkey. Exp Brain Res 43:69-77.

Frazor RA, Geisler WS (2006) Local luminance and contrast in natural images. Vision Res 46:1585-1598.

Greenhouse DS, Cohn TE (1991) Saccadic suppression and stimulus uncertainty. J Opt Soc Am A 8:587-595.

Guez J, Morris AP, Krekelberg B (2013) Intrasaccadic suppression is dominated by reduced detector gain. J Vis 13.

Heeger DJ (1992) Normalization of cell responses in cat striate cortex. Visual Neurosci 9:181-197.

Herrmann K, Montaser-Kouhsari L, Carrasco M, Heeger DJ (2010) When size matters: attention affects performance by contrast or response gain. Nat Neurosci 13:1554-1559.

Honda H (1989) Perceptual localization of visual stimuli flashed during saccades. Percept Psychophys 45:162-174.

Ibbotson M, Krekelberg B (2011) Visual perception and saccadic eye movements. Curr Opin Neurobiol 21:553-558.

Ibbotson MR, Cloherty SL (2009) Visual perception: saccadic omission--suppression or temporal masking? Current biology : CB 19:R493-496.

Ibbotson MR, Crowder NA, Cloherty SL, Price NS, Mustari MJ (2008) Saccadic modulation of neural responses: possible roles in saccadic suppression, enhancement, and time compression. The Journal of neuroscience : the official journal of the Society for Neuroscience 28:10952-10960.

Kagan I, Gur M, Snodderly DM (2008) Saccades and drifts differentially modulate neuronal activity in V1: effects of retinal image motion, position, and extraretinal influences. J Vis 8:19 11-25.

Kleiser R, Seitz RJ, Krekelberg B (2004) Neural Correlates of Saccadic Suppression in Humans. Curr Biol 14:386-390.

Kontsevich LL, Tyler CW (1999) Bayesian adaptive estimation of psychometric slope and threshold. Vision Res 39:2729-2737.

Krekelberg B, van Wezel RJ, Albright TD (2006) Adaptation in macaque MT reduces perceived speed and improves speed discrimination. J Neurophysiol 95:255-270.

Krekelberg B, Kubischik M, Hoffmann KP, Bremmer F (2003) Neural correlates of visual localization and perisaccadic mislocalization. Neuron 37:537-545.

Kusunoki M, Goldberg ME (2003) The time course of perisaccadic receptive field shifts in the lateral intraparietal area of the monkey. J Neurophysiol 89:1519-1527.

Lappe M, Awater H, Krekelberg B (2000) Postsaccadic visual references generate presaccadic compression of space. Nature 403:892-895.

Levy M, Fournier J, Fregnac Y (2013) The role of delayed suppression in slow and fast contrast adaptation in V1 simple cells. J Neurosci 33:6388-6400.

Lu Z-L, Dosher BA (1998a) External Noise Distinguishes Attention Mechanisms. Pergamon 38:1183-1198.

Lu ZL, Dosher BA (1998b) External noise distinguishes attention mechanisms. Vision Res 38:1183-1198.

Maldonado P, Babul C, Singer W, Rodriguez E, Berger D, Grun S (2008) Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. J Neurophysiol 100:1523-1532.

Matin E (1974) Saccadic suppression: a review and an analysis. Psychol Bull 81:899-917.

Melcher D (2007) Predictive remapping of visual features precedes saccadic eye movements. Nature Neurosci 10:903-907.

Morris AP, Kubischik M, Hoffmann KP, Krekelberg B, Bremmer F (2012) Dynamics of eye-position signals in the dorsal visual system. Curr Biol 22:173-179.

Morris AP, Liu CC, Cropper SJ, Forte JD, Krekelberg B, Mattingley JB (2010) Summation of visual motion across eye movements reflects a nonspatial decision mechanism. J Neurosci 30:9821-9830.

Nenadic Z, Burdick JW (2005) Spike detection using the continuous wavelet transform. IEEE Trans Biomed Eng 52:74-87.

Ohzawa I, Sclar G, Freeman RD (1985) Contrast gain control in the cat's visual system. J Neurophysiol 54:651-667.

Olsen SR, Bortone DS, Adesnik H, Scanziani M (2012) Gain control by layer six in cortical circuits of vision. Nature 483:47-52.

Paradiso MA, Meshi D, Pisarcik J, Levine S (2012) Eye movements reset visual perception. J Vis 12:11.

Rajkai C, Lakatos P, Chen CM, Pincze Z, Karmos G, Schroeder CE (2008) Transient cortical excitation at the onset of visual fixation. Cereb Cortex 18:200-209.

Reppas JB, Usrey WM, Reid RC (2002) Saccadic eye movements modulate visual responses in the lateral geniculate nucleus. Neuron 35:961-974.

Richards W (1969) Saccadic suppression. J Opt Soc Am 59:617-623.

Ross J, Morrone MC, Goldberg ME, Burr DC (2001) Changes in visual perception at the time of saccades. Trends Neurosci 24:113-121.

Schwartz O, Simoncelli EP (2001) Natural signal statistics and sensory gain control. Nat Neurosci 4:819-825.

Sommer MA, Wurtz RH (2008) Visual perception and corollary discharge. Perception 37:408-418.

Super H, van der Togt C, Spekreijse H, Lamme VA (2004) Correspondence of presaccadic activity in the monkey primary visual cortex with saccadic eye movements. Proc Natl Acad Sci U S A 101:3230-3235.

Sylvester R, Haynes J-D, Rees G (2005) Saccades Differentially Modulate Human LGN and V1 Responses in the Presence and Absence of Visual Stimulation. Curr Biol 15:37-41.

Thiele A, Henning P, Kubischik M, Hoffmann KP (2002) Neural Mechanisms of Saccadic Suppression. Science 295:2460-2462.

Thiele A, Pooresmaeili A, Delicato LS, Herrero JL, Roelfsema PR (2009) Additive effects of attention and stimulus contrast in primary visual cortex. Cereb Cortex 19:2970-2981.

Tolias AS, Moore T, Smirnakis SM, Tehovnik EJ, Siapas AG, Schiller PH (2001) Eye movements modulate visual receptive fields of V4 neurons. Neuron 29:757-767.

Vallines I, Greenlee MW (2006) Saccadic Suppression of Retinotopically Localized Blood Oxygen Level-Dependent Responses in Human Primary Visual Area V1. J Neurosci 26:5965-5969.

Volkmann FC, Schick AM, Riggs LA (1968) Time course of visual inhibition during voluntary saccades. J Opt Soc Am 58:562-569.

Volkmann FC, Riggs LA, White KD, Moore RK (1978) Contrast sensitivity during saccadic eye movements. Vision Res 18:1193-1199.

Watson T, Krekelberg B (2011) An equivalent noise investigation of saccadic suppression. J Neurosci 31:6535-6541.

Wurtz R (2008) Neuronal mechanisms of visual stability. Vision Res 48:2070-2089.

Wurtz RH (1969) Comparison of effects of eye movements and stimulus movements on striate cortex neurons of the monkey. J Neurophysiol 32:987-994.
Wurtz RH, Goldberg ME (1972) Activity of superior colliculus in behaving monkey. 3. Cells discharging before eye movements. J Neurophysiol 35:575-586.