UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

The role of spatial context in rat vision

Permalink

https://escholarship.org/uc/item/77t4k6n2

Author

Meier, Philip Martin

Publication Date

2011

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA, SAN DIEGO

The Role of Spatial Context in Rat Vision

A dissertation submitted in partial satisfaction of the requirements
for the degree Doctor of Philosophy

in

Neurosciences
with Specialization in Computational Neuroscience

by

Philip Martin Meier

Committee in charge:

Professor Pamela Reinagel, Chair
Professor Thomas Albright
Professor Timothy Gentner
Professor Donald MacLeod
Professor Terrence Sejnowski

2011

Copyright
Philip Martin Meier, 2011
All rights reserved.

The Dissertation of Philip Martin Meier is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____

Chair

University of California, San Diego

2011

DEDICATION

To my parents, for relentlessly supporting my dreams.

TABLE OF CONTENTS

Signature Page iii

Dedication iv

Table of Contents v

List of Abbreviations viii

List of Figures ix

List of Tables xii

Acknowledgements xiii

Vita and Publications xv

Abstract of the Dissertation xvi

Chapter 1: Introduction 1

1.1 What is vision? 1
1.2 Early vision 1
1.3 Spatial context: motion, color & orientation 2
1.4 Feed forward and recurrent vision 4
1.5 Surround processing: influences on contrast 7
1.6 What the field needs 12
1.7 First steps 14
1.8 In this thesis 17

Chapter 2: High throughput behavioral training of rats on visual tasks 20

2.1 Motivation: a scalable behavioral system 20
2.2 Previous methods 22
2.3 Description of use 24
2.4 Considerations for setting up rodent behavior 26
2.4.1 Response modality 27
2.4.2 Explicit stimulus requests 29
2.4.3 Feedback informs the subject of the consequence of an action. 30
2.4.4 Structure of a trial 31
2.4.5 Motivating rodents 35

2.4.6 Managing bias 36
2.4.7 Choosing a display device 38
2.5 Observations of behavioral data: quality and volume 40

Chapter 3: Collinear features impair visual detection by rats 57

Abstract 57
3.1 Introduction 58
3.2 Results 61
3.2.1 Rats can report the presence of a small oriented grating in the presence of flankers 61
3.2.2 Collinear flankers impair detection more than other arrangements 63
3.2.3 Collinear flankers cause rats to miss the target 66
3.3 Discussion 67
3.3.1 Potential confounds 69
3.3.2 Comparison to humans 73
3.3.3 Relating our findings to natural scene statistics 75
3.3.4 Rats as a model system for vision research 78
3.4 Experimental procedures 79
3.4.1 Animal subjects 79
3.4.2 Training apparatus 79
3.4.3 Training protocol 80
3.4.4 Performance stability, data filtering, and performance measures 86
3.4.5 Statistical tests 87
3.5 Acknowledgements 89
3.6 Supplementary 90
3.6.1 Flanking stimuli impair target detection performance 90
3.6.2 Choosing stimulus parameters 92
3.6.3 Statistics for all pair wise comparisons of flanker configurations 94
3.6.4 No influence of phase 97
3.6.5 Water restriction, training schedule and environment 98
3.6.6 Training system 99
3.6.7 Stimulus Display 102
3.6.8 Considering signal detection theory 103

Chapter 4: Rat performance on visual detection task modeled with divisive normalization and adaptive decision thresholds 122

Abstract 122
4.1 Introduction 123

4.2 Methods 125
4.2.1 Animal training 125
4.2.2 Task 126
4.2.3 Stimuli 128
4.2.4 Model fitting 129
4.3 Results 132
4.3.1 Target contrast improves detection 132
4.3.2 Flanker contrast impairs detection 134
4.3.3 Fitting a model to the data 136
4.4 Discussion 144
4.4.1 What is the meaning of the parameter k_F? 145
4.4.2 What is the meaning of the parameter λ? 147
4.4.3 Theory 149
4.4.4 Future recommendations 151
4.4.5 Applications 153

Chapter 5: Neural Recordings from the Lateral Geniculate Nucleus of the Rat 167

5.1 Choosing a preparation 168
5.2 Eye tracking 171
5.3 Isolating single cells 175
5.4 Localizing individual receptive fields 177
5.5 Influence of the surround under anesthesia 179
5.6 The influence of isoflurane anesthesia 187
5.7 Future 193

Chapter 6: Discussion 208

6.1 Summary of findings 208
6.2 Significance of collinear effect 210
6.3 Low level visual explanations 211
6.4 Higher level visual explanations 214
6.5 Flankers and divisive normalization 216
6.6 Why flankers might facilitate 220
6.7 Bursts in detection 222

Appendix A: Natural scene statistics 225

A1.1 A theoretical interpretation of collinear impairment 225
A1.2 Methods used for natural scene statistics 227

Appendix B: Human psychophysics consistent with literature 231

References 238

LIST OF ABBREVIATIONS

BIC	Bayesian Information Criterion
CRT	Cathode Ray Tube (display monitor)
d'	d-prime, detection sensitivity
ISI	Inter-spike Interval
LGN	Lateral Geniculate Nucleus
OKR	Optokinetic Response
OLED	Organic Light Emitting Diode (display monitor)
ROC	Receiver Operating Characteristics
SDT	Signal Detection Theory
STA	Spike Triggered Average
V1	Primary Visual Cortex
YNFC	Yes/No Forced Choice
2AFC	Two Alternative Forced Choice

LIST OF FIGURES

Figure 2.1: Schematic diagram of the training system's architecture 46

Figure 2.2: Diagram of a single chamber in which rats are trained to perform perceptual tasks 47

Figure 2.3: Components of a single lickometer port 48

Figure 2.4: Trial volume, performance, bias and response time 49

Figure 2.5: Trial rate decreases over each two hour session 50

Figure 2.6: The percent correct performance of a trained rat does not change across two hour sessions 51

Figure 2.7: The performance of a trained rat is stable over days 52

Figure 2.8: The learning curves of eight rats progressing through four shaping steps 53

Figure 2.9: Overview of rat performance from data collected in May 2008 54

Figure 2.10: The volume of data collected by a single researcher over three years 55

Figure 2.11: Examples of stimuli used to train rats 56

Figure 3.1: Training rats to detect an oriented target in the presence of flanking distractors 105

Figure 3.2: A schematic model indicating how the presence of flankers might influence rat's decisions 106

Figure 3.3: Grouping flanker stimuli into conditions 107

Figure 3.4: Collinear flankers impair detection more than other arrangements .. 108

Figure 3.5: Collinear flankers cause rats to miss the target 109

Figure 3.6: Schematic model of pattern sensitive contrast normalization 111

Figure 3.S1: Photograph of a rat viewing a target with collinear flankers 112

Figure 3.S2: Stimuli used for target detection, with or without flankers 113

Figure 3.S3: Rats are worse at detecting the target when the flankers are present 114

Figure 3.S4: Detection performance depends on target spatial frequency, target contrast, flanker contrast and flanker proximity 115

Figure 3.S5: Additional pair-wise comparisons between the flanker conditions. 116

Figure 3.S6: Multiple comparison adjustments 117

Figure 3.S7: Phase of grating has no influence on detection with flankers..... 118

Figure 3.S8: Percent correct and d′ yield the same conclusions..... 119

Figure 4.1: Experimental design for yes/no 2AFC target detection 155

Figure 4.2: Higher target contrasts increase detection performance and bias rats to choose yes..... 156

Figure 4.3: Higher flanker contrasts decrease detection performance and bias rats to choose yes..... 157

Figure 4.4: The impairment and bias caused by flankers is present for all target contrasts..... 158

Figure 4.5: Various decision criteria 159

Figure 4.6: Spatial contrast normalization improves the model fit 161

Figure 4.7: Bias allows model to generalize to different subjects..... 163

Figure 4.8: Confidence intervals for the model parameters 164

Figure 5.1: Eye movements in anesthetized and awake rats 194

Figure 5.2: Detecting and clustering spikes 195

Figure 5.3: A localized LGN receptive field from an awake rat 196

Figure 5.4: Response of an LGN cell while flashing gratings with varying target and flanker contrast 197

Figure 5.5: Properties of putative bursts in an LGN cell in response to flashed gratings with varying target and flanker contrast 198

Figure 5.6: Interspike intervals of cell with putative Calcium bursts..... 199

Figure 5.7: Sensitivity to the contrast of a target centered over the receptive field 200

Figure 5.8: Average evoked firing rate is lower under isoflurane anesthesia compared to no anesthesia..... 201

Figure 5.9: LGN cells recorded under isoflurane have less spatial selectivity ... 202

Figure 5.10: The peak response of the spike-triggered average has a longer delay for cells that are recorded under isoflurane anesthesia 203

Figure 5.11: Flashed gratings that approximate the stimulus shown to other rats performing a behavioral detection task 204

Figure 5.12: The firing rate increases when the rat exits anesthesia, increasing sensitivity to low contrasts 205

Figure 5.13: The response of four different cells in the LGN of an anesthetized rat to flashed gratings of varying contrast 206

Figure A1.1: Co-occurrence of flanker and target features in natural images. ... 230

Figure A2.1: A comparison of the contrast sensitivity in humans and rats..... 234

Figure A2.2: The influence of collinear stimuli compared to pop-out stimuli in humans and rats..... 235

Figure A2.3: The influence of collinear stimuli compared to pop-out stimuli in humans and rats have the same interpretation with signal detection theory metrics..... 236

Figure A2.4: Behavioral performance of humans and rats on a target detection task in the presence of flankers 237

LIST OF TABLES

Table 3.1: Shaping sequence and training steps 120

Table 4.1: Shaping sequence and training steps 165

ACKNOWLEDGEMENTS

I would like to thank my committee members, especially my advisor Pam Reinagel for her support, flexibility, and wisdom. Over the years she has consistently provided valuable advice on both strategic and experimental levels, despite the fact that my interests carried me out of her domain of expertise in neurophysiology. I extend my gratitude to all of the undergraduates and technicians for their assistance training rats and assisting with projects in the lab. Thank you Liz Murphy, Duc Nguyen, Holly Vo, Aria Jafari, Yuli Wang, Alee Bowden, Marv Zech, Ginger Beriones, Dan Parks, Fan Li, Nareg Kalajian, Danielle Dickson and Sarah Meder. I would also like to thank all the graduate students in the lab for their camaraderie and valuable feedback. Erik, Balaji and Claire - lunch at the price center or the coop would never have been the same without you. I would specifically like to thank Erik Flister for his passion, vision and hard work to establish the high throughput behavioral training system, and Balaji Sriram for his enthusiasm in the long hours shared working on physiological recordings. Finally, I would like to thank all my peers in the neuroscience program, especially Corinne Teeter, Sam Nummela and David Matthews for their support while I was working on my dissertation.

Chapter 3, in full, appears as it was published in the *Journal of Vision*, 2011, Meier, P. M., Flister E.D., Reinagel, P. Philip Meier and Erik Flister contributed equally to this work. Each made essential contributions to the conception, design, and implementation of the general method and technology for high-throughput automated training and testing of rats in 2AFC visual tasks, with help and critical discussion from

Pamela Reniagel. Philip Meier was responsible for the conception of the flanker experiment, design and implementation of the task, training and data collection, data analysis, interpretation of the results, and writing of the manuscript, with help and critical discussion from Erik Flister and Pamela Reinagel.

Chapter 4, in full, appears as it was submitted for publication in the *Journal of Vision*, Meier, P. M., Reinagel, P. The dissertation author was the primary investigator and author of this paper and the co-author listed in this publication directed and supervised the research which forms the basis for this chapter.

Chapter 5, in part, contains data presented at the conference Computational Systems Neuroscience (Cosyne), 2011, Sriram B.S., Meier, P. M., Reinagel, P. The dissertation author was a co-first author of this poster, the other author contributed equally, and the last co-author listed on the poster directed and supervised the research which forms the basis for part of this chapter.

VITA

1999	Harvard Westlake School Los Angeles
2003	B.F. in Printmaking & Drawing, 2nd Major in Philosophy-Neuroscience-Psychology Washington University in St. Louis
2005-2007	IGERT Fellowship: Computational Neurobiology National Science Foundation
2007-2009	IGERT Fellowship: Vision & Learning in Humans & Machines National Science Foundation
2009-2011	Graduate Student Research Fellowship National Science Foundation
2011	Ph.D. in Computational Neuroscience, University of California, San Diego

PUBLICATIONS

Cox D., Meier P., Oertelt N., DiCarlo J. 'Breaking' position-invariant object recognition. *Nature Neuroscience.* 2005 Sep;8(9):1145-7.

Meier P., Flister E. & Reinagel P. Collinear features impair visual detection by rats. *Journal of Vision.* 2011 Mar 31;11(3). pii: 22. doi: 10.1167/11.3.22.

Meier P., & Reinagel P. Rat Performance on Visual Detection Task Modeled with Divisive Normalization and Adaptive Decision Thresholds. *Journal of Vision.* Submit March 18, 2011.

ABSTRACT OF THE DISSERTATION

The Role of Spatial Context in Rat Vision

by

Philip Martin Meier

Doctor of Philosophy in Neurosciences

with Specialization in Computational Neuroscience

University of California, San Diego, 2011

Professor Pamela Reinagel, Chair

I ask whether rats have pattern vision. Is the rat's visual system influenced by the spatial context surrounding a feature? Specifically, is the processing of an oriented target feature influenced by the relative orientation and relative position of nearby visual features?

In chapter 1, I provide an introduction to the early visual system and the role of contextual processing. I review the influence of spatial context at the level of neural responses, anatomy, theory, and psychophysics.

In chapter 2, I describe the experimental paradigm we developed in the lab. The methods were sufficient to train rats on a range of visual tasks, including training

them to detect small faint oriented visual targets in the presence of nearby flanking stimuli ('flankers').

In chapter 3, I ask whether the geometric relationship between the flankers and the target influences the rats' ability to detect an oriented target. Indeed, of all the spatial configurations I tested, one condition was harder than the rest: when all the stimuli where collinear. This finding shows that rats are influenced by spatial context when detecting a faint target.

In chapter 4, I ask whether a difference in the contrast of target and flanker will reverse the impairment observed in collinear stimuli. Evidence from human behavior and neural responses suggests that performance might improve when targets are faint and the surrounding contrast is strong. In no case did the presence of flankers improve the rat's ability to detect a target. A model is fit to the behavioral data to explain the animal's impairment and bias.

In chapter 5, I include a pilot study to validate the feasibility of physiological recordings that isolate surround processing in the lateral geniculate nucleus (LGN) of the visual thalamus of rats. In one experiment, changes in surround contrast and luminance did not reliably influence the neural response to the target, suggesting that contrast normalization in the LGN is not substantial for briefly flashed targets with flankers. I also characterize basic response properties of LGN neurons in the awake and anesthetized rat. The differences in the response properties validate the importance of recording from neurons in awake animals.

Chapter 1:

Introduction

1.1 What is vision?

Vision is determining *what* is *where* in a scene. Vision lets me identify my keys resting on the table two feet to the right of me. But it does more: it indicates properties of objects. The table is fake wood. They keys are a dull metal with blue tape wrapped around them. Vision confirms the nature of my body's movement through space. It helps me keep my balance. It detects the presence of objects and the category of their movement. Vision warns of a lunging assailant, and captures concern of my friend's frown. The evolutionary utility of vision spans many tasks.

1.2 Early vision

Early vision refers to the processing present in the retina, the visual thalamus and sometimes the primary visual cortex. In a temporal chain of processing these are the first regions to receive and transform visual signals. The function of early vision is to select the visual features used for subsequent processing, and to control the sensitivity of these features. More abstractly we can say that the goal of early vision is to encode the visual world by means of patterns of spiking neurons, such that the maximum amount of relevant information is transmitted at each moment in time. Intuitively, if a neuron is silent, it is not communicating much information. Additionally, if it is saturated at its maximum firing rate it also is not providing much

information. A sensory channel is most informative when in fluctuates across its entire dynamic range. More formally, if the number of responses in a channel is discrete, the maximum amount of information is transmitted when there is an equal probability of each of the different responses. Interestingly, even without addressing the content of neurons, we can identify if there are inefficiencies in the neural code that reduce information transfer.

The optic nerve is a bottle neck of information from the eye to the brain. Not all of the information about the visual world can be sent. To send the most information, or even simply to send the most information for a given metabolic cost, the encoding representation of signals should be sensitive to the statistical frequency of stimuli in natural scenes (Barlow, 2001).

1.3 Spatial context: motion, color & orientation

Neurons in the early visual system are sensitive to localized patches of the visual world. This is the neuron's "receptive field." Neurons in the primary visual cortex are tuned to respond to different kinds of stimuli in their receptive field. For example, a given neuron may be most responsive to luminance patterns of a particular spatial scale and orientation at localized region with respect to a subject's gaze (Hubel & Wiesel, 1968). The primary driving force of spikes comes from the localized classical receptive, but individual neurons are modulated by stimuli outside this region (Fitzpatrick, 2000).

Consider a diagonal edge of an object translating horizontally. From a limited aperture, this motion appears diagonal and not horizontal. Given only the evidence from a single receptive field, diagonal motion is the best statistical estimate assuming the world it is dominated by slow motion, which it is (Weiss, Simoncelli, & Adelson, 2002). Determining the global motion of an object requires integration of a larger field of view (Huang, Albright, & Stoner, 2007).

Consider a local patch of color. An accurate estimate of an object's reflectance is influenced by the local regions in the scene, both the average chroma (Helmholtz, 1866) and the variance (Brown & MacLeod, 1997). The influence of the surround may have to do with discounting the illumination in the surface (Delahunt & Brainard, 2004), though it may also have to do with integrating evidence that spans a homogenous region of space. For example, the blue cones play an influential role in color processing and may be noisy at any given location do to the sparse spatial sampling of the blue cone mosaic (Roorda & Williams, 1999). Biased sampling from the cone mosaic could be minimized by integrating over larger regions. Additionally color evidence is collected along the edges of objects and influences the perception of the color in the center of the object (Pinna, Brelstaff, & Spillmann, 2001). Thus estimating a color is likely to involve information outside of the region being estimated whether surrounding regions belong to the same surface or different ones.

Consider a local patch containing the edge of an object. The oriented power in the image patch is a useful cue, but sometimes the orientation of an object edge remains ambiguous. The surrounding context of a scene can help identify if the

luminance edge is truly an object edge, and if so, what orientation it is. The edge of a shadow may only be recognized as such given the context of a scene. The edge of an object may be occluded, and yet humans can still group separated edges together quite well even when the context is restricted to two local patches (Geisler & Perry, 2009). The perception of an oriented edges is influence be nearby orientations as evidenced by the tilt illusion (Wenderoth & Johnstone, 1988). The influence of perceived tilt can be explained by a statistical interpretation of natural scenes (Schwartz, Sejnowski, & Dayan, 2009).

Local patches remain ambiguous. In all cases above, in order to estimate a fact about a local region in an image, it helps to have access to information at adjacent regions in the image. How does the brain integrate this information? How does the brain combine information at distant regions in a scene?

1.4 Feed forward and recurrent vision

A prevalent framework to understand visual processing is as a feed forward hierarchy (Fukushima, 1980). In early stages, individual neurons (or nodes) have high spatial precision and limited invariance. At later stages the individual processing units combine information. This is reflected in the size of the receptive fields in the cortical processing stream of macaques (Zeki, 1978) though not always (DiCarlo & Maunsell, 2003). Additionally, the notion that only feed forward circuitry is necessary for scene perception and object detection is defended on the speed of visual processing (Thorpe, Fize, & Marlot, 1996). Notably, the rapid presentation of images in a sequence does

not substantially impair our perception of individual images, which one might expect if visual processing was highly dependent on feedback loops. In a purely feed forward system, the integration of spatial context would be performed by the large receptive fields in the later stages of a visual hierarchy.

While feed forward processing may be sufficient for some visual tasks, normal visual processing is likely to involve other computations as well. These computations might benefit from recurrent connections that allow neural activity to propagate up and down hierarchies and well as spread laterally within a cortical area. From anatomy we know that there are substantial projections of neurons from higher cortical areas to lower ones (Rumberger, Tyler, & Lund, 2001; Salin, Girard, Kennedy, & Bullier, 1992). Also, there are reciprocal connections from each cortical area to its corresponding thalamic area (Sherman & Guillery, 2002; Van Horn, Erisir, & Sherman, 2000). Finally, there are rich lateral connection with a given cortical area, especially within, but not limited to, layer 2/3 (Bosking, Zhang, Schofield, & Fitzpatrick, 1997; Hellwig, 2000). Rather than a single pass through hierarchy of layers, the entire network may iterate to allow higher, lower and lateral connection to constrain activity (Bullier, 2001). During a visual detection task, missed trials contained reduced activity in the sustained response of V1 neurons more than the initial response (Super, Spekreijse, & Lamme, 2001). Thus sustained activity is a stronger correlate of perceptual report that the initial feed forward volley.

Lateral and feedback processing are active during visual stimulation, and are likely to have an influence, especially on a steady state response. It has been argued

that these lateral and feedback circuits act more as modulators of neural activity, and less as drivers (Sherman, 2005). That is to say, individual lateral connections might not reliably induce responses, but the correlated activity may influence the response of neurons through the mean and variance of background activity (Shu, Hasenstaub, Badoual, Bal, & McCormick, 2003). Elegant work in slices of tissue have included stochastic conductance changes that mimic lateral processing in cortical networks (McCormick et al., 2003) and feedback processing from cortex the thalamus that change the transfer efficiency of spikes in from the retina (Wolfart, Debay, Le Masson, Destexhe, & Bal, 2005).

The activity of nearby neurons, tens to hundreds of milliseconds before a potential spike, will determine the state of the network, and how it passes information from one area to the next. What determines this activity? Let us combine three assumptions: that the activity in the network is dominated by transient events, and that neurons in the visual system maintain retinotopic maps, and that neurons make more local than distal connections. This suggests that the state of the network is primarily determined by activity that is nearby in visual space and time. For example, the barrage of background synaptic activity in a V1 cell may be the result of processing of the spatially adjacent regions of the scene immediately preceding. In other words, the perceptual challenges posed by the aperture problem could be solved by computations that appear to be the result of many weak lateral modulators combing to shift the gain of a neurons transfer function. This suggests that individual neurons would be

influenced by the spatial context of a stimulus, ultimately influencing the behavioral performance of a subject.

1.5 Surround processing: influences on contrast

Perceptual effects of spatial processing may be mediated by "extra classical processing" in visual neurons. Extra classical processing encompasses influences that come from beyond the classical receptive field. An extensive literature documents the effects of surround contrast, surround orientation, and surround configuration on local visual processing. Contrast normalization is a basic and ubiquitous form of contextual processing in which the contrast of a local signal is normalized to the contrast in a surrounding temporal or spatial window. Visual processing at each image location is also sensitive to the orientation of the contrast in the surround relative to that of the center. Even more specifically, surround effects can depend on both the location and orientation of visual features in the surround, thus reflecting configuration of visual elements. Here I review key evidence at the level of neurophysiology, neuroanatomy, computational theory, and visual perception.

The amplitude of neural responses in the retina, LGN, and V1 are normalized to spatially nearby contrast (Bonin, Mante, & Carandini, 2006; Carandini, Heeger, & Movshon, 1997; DeAngelis, Robson, Ohzawa, & Freeman, 1992; Finn, Priebe, & Ferster, 2007; Heeger, 1992; Shapley & Victor, 1979). Several neural mechanisms of spatial contrast normalization have been proposed, such as spatially pooled feed-back inhibition, or lateral feed-forward inhibition, or changes in membrane conductance

(Angelucci, Levitt, & Lund, 2002; Angelucci, Levitt, Walton et al., 2002; Bonds, 1989; Carandini, 2004; Chisum & Fitzpatrick, 2004). These mechanisms can be sensitive to contrast both within and beyond the classical receptive field.

Thus, while the inputs that primarily drive a neuron are presumed to come from its classical receptive field, the neurons sensitivity (the gain of its transfer function) may be influenced by a larger region. The literature occasionally refers to the "extra-classical receptive field," which I will simply refer to as "the surround," Please note that it is often larger than and distinct from the center-surround structure of the classical receptive fields in the retina and LGN. In many cases, the presence of stimuli in the surround suppresses a neuron, and the amount of the suppression is tuned to the relative orientation of the surround stimuli with respect to the driving stimulus (Bonds, 1989; Cavanaugh, Bair, & Movshon, 2002; Chen, Kasamatsu, Polat, & Norcia, 2001; Polat, Mizobe, Pettet, Kasamatsu, & Norcia, 1998; Sillito, Cudeiro, & Murphy, 1993). However, it is also reported that some cells in the appropriate contrast conditions can increase their spiking activity when the orientation of stimuli in the surround is matched to the driving stimulus (Bakin, Nakayama, & Gilbert, 2000; Chen et al., 2001; Polat et al., 1998; Sillito et al., 1993). Often oriented surround stimuli are presented as an annulus. However, surround effects can depend on the angular position of the flanker with respect to the target orientation, such that specific configuration like collinearity are differentially suppressed or facilitated (Cavanaugh et al., 2002; Polat et al., 1998).

At several levels in the mammalian visual system, anatomic studies reveal

circuitry that could support normalization to contrast including regions outside a cell's classical receptive field (Shepherd, 2004). In the retina, lateral connectivity through horizontal cells and amacrine cells (Rodieck, 1998) provides an anatomical substrate for normalization to surround contrast. Additional surround normalization in the lateral geniculate nucleus could arise from wide-field inhibitory neurons projecting from the perigeniculate nucleus (Cruikshank, Landisman, Mancilla, & Connors, 2005; Dubin & Cleland, 1977; Uhlrich, Cucchiaro, Humphrey, & Sherman, 1991). While these circuits are capable of explaining the normalization to surround contrast, cortical circuits are likely to be involved in orientation-specific normalization. There is evidence that connectivity among V1 cells depends on their relative orientation tuning (Das & Gilbert, 1995, 1999; Gilbert & Wiesel, 1989; Matsubara, Cynader, Swindale, & Stryker, 1985). Configural effects depend on position and cannot be explained by contrast and orientation alone. These might arise from lateral connections in V1 that are sensitive to a cells orientation preference and lack radial symmetry in a retinotopic map (Chisum, Mooser, & Fitzpatrick, 2003). Alternately configural effects could be explained by comparably structured feedback from higher cortical areas to V1 (Angelucci, Levitt, Walton et al., 2002), or from cortex to the thalamus (Murphy, Duckett, & Sillito, 1999). It remains unknown if these orientation-selective circuits are found in rats. Rodents have orientation tuned cells in V1 but lack orientation columns (Ohki, Chung, Ch'ng, Kara, & Reid, 2005; Van Hooser, Heimel, Chung, & Nelson, 2006) and rats differ from primates in V1 micro-circuitry (Zarrinpar & Callaway, 2006).

From a theoretical perspective there are several reasons it would be advantageous for representation of local features to be sensitive to nearby image context (Series, Lorenceau, & Fregnac, 2003). In natural scenes, local image features such as luminance, contrast, and orientation are correlated at nearby locations (Field, 1987; Ruderman & Bialek, 1994; Tolhurst, Tadmor, & Chao, 1992; van der Schaaf & van Hateren, 1996). When features are spatially correlated, surround processing can optimize the fidelity or efficiency of image estimation (Geisler, 2008). For example divisive normalization from a nearby population of cells (Heeger, 1992) can allow a neuron to better adapt its sensitivity, reduce redundancy with its neighbors, and thus maximize information transfer (Schwartz & Simoncelli, 2001). Surround processing could also enhance salience of relevant features such as continuous contours (Field, Hayes, & Hess, 1993; Geisler, Perry, Super, & Gallogly, 2001) or statistically surprising features (Itti & Koch, 2000). These theories and others predict that different patterns in the surround should have distinct influences on a visual target's neuronal representation, even if lower order statistics like luminance and contrast are matched.

Human perception of oriented targets is influenced by the contrast, spatial frequency and orientation of nearby stimuli, both in contrast discrimination tasks (Cannon & Fullenkamp, 1996; Ejima & Takahashi, 1985; Xing & Heeger, 2001) and target detection tasks (Chen & Tyler, 2008; Polat & Sagi, 1993, 2007; Solomon & Morgan, 2000; Williams & Hess, 1998; Zenger-Landolt & Koch, 2001; Zenger & Sagi, 1996). This influence is typically largest when the surrounding stimuli match

the orientation and spatial frequency of the target, both for annuli that completely surround a target and for discrete flankers (Cannon & Fullenkamp, 1996; Chubb, Sperling, & Solomon, 1989; Polat & Sagi, 1993). The sign of the influence from the surround (contrast enhancement vs. contrast reduction, detection facilitation vs. detection suppression) is dependent on parameters like the size of the surround and the contrast of the target (Williams & Hess, 1998; Xing & Heeger, 2001). Notably, results from contrast discrimination tasks may differ from detection tasks because the target contrast is lower in detection tasks. Human psychophysical studies have studied detection tasks and shown that surround stimuli may either facilitate target detection or suppress it, depending on many factors. These factors include: relative orientation, relative contrast, relative phase, collinearity, surround size, distance between target and surround, central vs. peripheral vision, experimental paradigm (spatial 2-alternative forced-choice, temporal 2-interval forced-choice, single presentation yes/no), blocked vs. interleaved trials, experience–dependant expertise, and individual differences between subjects (Kurki, Hyvarinen, & Laurinen, 2006; Polat & Sagi, 1993, 2007; Solomon & Morgan, 2000; Williams & Hess, 1998; Zenger-Landolt & Koch, 2001). Though the phenomena are complex, these studies consistently find that performance is sensitive to the relative orientation between a target and its surroundings (Polat & Sagi, 2007; Williams & Hess, 1998). The influence of stimulus arrangement and phase are more variable and subject to experimental paradigms.

1.6 What the field needs

Visual research remains fragmented in disciplines. The field needs experimental paradigms that unify awake behavior with tools to selectively and temporarily change neural circuits. A new paradigm is likely to benefit by shifting to a species that supports high throughput visual behavior. However, in order to address questions of spatial contextual processing, extra experiments are required to confirm the presence of the surround processing in the new species.

Human psychophysics excels in providing behavioral evidence, but knowledge of mechanisms is generally implicit, inferred, and not directly observed. Systems neuroscience is rarely coupled with awake behaving animals. The exceptions are noteworthy: challenging experiments which most researchers would be hesitant to add to their complexity. It would be desirable to combine these experiments with new emerging techniques for manipulating neural circuits, but the investment of training animals can be prohibitive when coupled with new risky techniques. Exploration is limited. Part of the bottle neck is the sheer amount of time of training animals, part is the choice of species, and another aspect is the method of training.

Much of the research in system neuroscience correlates stimuli with neural responses, and does very little to modify the nervous system directly. Our knowledge of the underlying mechanisms is constrained by our ability to modify neural circuits: for example, by removing, stimulating or silencing known parts. Previous methods for modifying cortical activity include electrical stimulation, cooling, and pharmacological manipulations. Ablation and removal are quite drastic manipulations

that are not selective or reversible. Electrical stimulation has precise temporal control, and moderate spatial localization, but indiscriminately activates nearby neurons. Cooling is also quite useful because it is reversible, but is quite spatially diffuse. Pharmacological manipulations have been used to great benefit in slice physiology where it is easier to control the external cellular medium. In vivo, injections of GABA agonists have been used to temporally inactivate brain regions (Nummela & Krauzlis).

Molecular tools and optogenetics promise to modify neural activity with unprecedented specificity. The specificity comes in four forms: location, time, cell type and connectivity. The spatial and temporal specificity comes from the placement of an activating laser (Papagiakoumou et al.). Additional spatial selectivity may come from the intersection of selective receptor expression and the laser. Alternately, a viral injection in one brain area can be used to infect a subset of neurons in an adjacent brain area that it receives input form (Lechner, Lein, & Callaway, 2002). Ultimately location of the virus, and the resulting receptor under optical control (J. Wang, Borton, Zhang, Burwell, & Nurmikko), can select neurons that make a particular connection in the network, neurons that were created at a certain time during development (Saito & Nakatsuji, 2001) or neurons of a particular genetic type (Gong et al., 2003).

The application of these new tools will enable for direct "circuit breaking." However, new tools always bring risk to experiments. A likely consequence is that most experiments will avoid animal behavior that requires extended training. To address this we decided to develop a behavioral paradigm that scaled well for high

throughput training, thus complimenting other techniques in neuroscience. To scale well it would require small animals, inexpensive technology, and software automation. We considered many rodent species, which we narrowed down to rats, mice, squirrels and degus. Squirrels and degus are more visual creatures; they have higher cone densities in their retinas (Jacobs, Calderone, Fenwick, Krogh, & Williams, 2003; Kryger, Galli-Resta, Jacobs, & Reese, 1998). Rats and mice are "less visual" and we questioned the ease that we would be able to train them on visual tasks. Pilot studies revealed that degus and mice were more challenging to train, while both squirrels and rats were easily trained by our methods. Unfortunately squirrels rarely breed in captivity, and thus it is difficult to reliably obtain experimental subjects. Rats emerged as the most pragmatic choice to further experimentation. Extending the labs training methods to mice would be valuable for future projects, especially because of the genetic resources available with mice.

1.7 First steps

I have argued that rodent vision research has the potential to add substantially to the field of neuroscience. Rodent visual systems are less similar to human vision than primates; there are many questions that rats and mice will not be able to address. For example, rodents are a poor model system to study the exploration of visual scenes with eye movements. However, there are many questions about the brain, and about vision, that can be addressed with rodents. For example, as a model species, rats are amenable to answering question of neural coding, adaptation, reinforcement learning,

cortical processing and the role of the thalamus. What about spatial contextual processing? Are rats able to perform the tasks that are required to observe spatial contextual effects behaviorally?

A minimal paradigm for contextual processing requires a "target" and a "context." The subject performing the behavioral task must report an attribute of the target. The relevant property could be anything: orientation, color, contrast or an object's identity. While I am personally interested in color, it seemed more pragmatic to study orientations first. Rats only have two cone types and they are very sparsely distributed. Orientation is a significant, well studied visual property that is selected by the early visual system. Thus, I considered two routes: training rats to report differences in a target's orientation, and training rats to report the presence or absence of a target. In both cases, the target was an oriented grating.

The target would have to be spatially limited so that a "context" could be presented around it. There are many interesting experiments that study the contextual effects provided by spatially overlapping stimuli (orthogonal masks, drifting plaids). However, for overlapping stimuli, interactions between target and context could result from saturating a common pool early in the stream of visual processing. Spatially separated stimuli invoke separate neural populations at transduction. Thus, when target and context do not overlap, behavioral observations about the influence of context are more informative about the interactions of neural responses at subsequent stages.

If rats can report small orientation differences easily, this opens the door to studying the tilt illusion. I successfully trained rats to discriminate orientated stimuli that differed +/- 15 degrees from vertical. Three of four rats failed to learn this task, and it was difficult to get the rats to generalize to small stimuli. One rat finally learned this task after two years of training. Thus it is possible for rat to perform reliable on tasks that require fine orientation discriminations, but I did not find a pragmatic training protocol. I never presented contextual stimuli surrounding a target of which a rat was discriminating the orientation. I also never generalized the task to a range of orientations in which all positively tilted orientations were rewarded equally. The later is important to observe a shift in the perceived orientation of stimuli.

An additionally reason convinced me not to focus on training rats to report orientations: the problem of interpretation. Consider a trained rat that reports the target's tilt, but is biased to report the tilt of the surround. How do we know if the rat is correctly reporting the target property (and revealing its illusory percept) or if the rat is misguidedly reporting the property of the context itself (which has nothing to do with its percept of the target)? In short, even if rats displayed evidence for a tilt illusion, it would be difficult to rule out degenerate scenarios that resulted in the same data. To make it worse, I am convinced that the "degenerate" scenarios are actually quite likely. I suspect that the key to understanding tilt illusions in rats (or any animal subject) will be to find evidence of targets at that are both attracted to the surround as well as repulsed from the surround, consistent with the phenomena characterized in humans (Wenderoth & Johnstone, 1988).

Instead, I would vary orientations but train rats to report the presence or absence of the target. No orientation would ever be reinforced by the task. In order to perform the task, the rats only had to perceive the presence of the target in the center. Using this approach it was possible to learn about pattern processing in rats (the interactions of oriented stimuli at offset spatial locations) without requiring the rats to report about the properties being studied.

1.8 In this thesis

In the next chapter I present general considerations about training rats on visual tasks, including training methods and their aptitude. The chapter serves to validate that rats are a valuable species for a range of different visual studies. I review alternative methods that we considered, and why I chose experimental paradigms. Hopefully the information is useful for other experimenters that wish to set up behavioral training systems, even though their object of study might not necessarily be focused on contextual spatial processing in vision. Additional training information including part number and specific shaping steps for the task I studied is available in the methods section of the third chapter.

In the third chapter I present a study on the influence of spatial patterns on rats' ability to detect an oriented target. This chapter contains the key observation that there is a selectivity to the influence of particular spatial patterns on rats detection performance. Not all surrounding patterns impair performance equally; collinear stimuli impair the most.

In the fourth chapter I modify three experimental parameters that might be relevant to the impairment of collinear stimuli. Specifically, I ask if, under favorable experimental conditions, collinear stimuli would ever improve rats' detection performance compared to the no target condition. One possibility is that the fixed contrast condition I had tested in the previous chapter was not in a regime that would display facilitation of detection. Maybe the target contrast was too high. Additionally, various experimenters had suggested that the relative contrast of the flanker to the target might also be important. Thus I tested all possible combinations for a range of target contrast and flanker contrast. Additionally, it is known that subject's performance is different if all conditions are randomly leaved on a trial to trial basis, as opposed to blocked (Polat & Sagi, 2007). The data comparing spatial patterns in Chapter 3 used stimuli with properties that were randomly chosen on each trial. In Chapter 4, stimuli were blocked so that the same contrast conditions were applied for a hundred and fifty trials in a row. Finally, spatial conditions were limited to a single category: collinear stimuli with a constant orientation on every trial. In no case did the collinear flankers improve the rats' detection performance. The focus of the chapter is to develop a parsimonious mathematical model capable of predicting the rats' responses.

In the fifth chapter I present preliminary data from physiological recordings from the visual thalamus of awake and anesthetized rats. I characterize responses to drifting gratings, spatiotemporal noise patterns and flashed gratings. I show that cells can have sufficient single cell isolation, eyes can be sufficiently stable and receptive

fields can be sufficiently quantified and localized. However, the yield of high quality data capable of addressing surround processing was low in the anesthetized case and not achieved in the awake preparation. The observed differences in the awake and anesthetized condition highlight the importance of working in the awake preparation, despite the experimental difficulties.

Chapter 2:

High throughput behavioral training of rats on visual tasks

In this chapter I provide the motivation for developing an automated behavioral training system to train rats on visual tasks, a history of training methods, a description of the use of the system, and a resource for other researchers that are contemplating setting up a behavioral training system. I review the major decisions in setting up a training system including choosing a response modality, determining an experimental paradigm and choosing a display device. I also discuss strategies for managing a subject's response bias and what works for motivating rodents to perform behavioral trials. I provide a basic overview of behavioral responses, focusing on the volume and quality of the data.

2.1 Motivation: a scalable behavioral system

In 1912 Karl Lashley failed to train rats to discriminate horizontal from vertical lines printed on pieces of paper (Lashley, 1912). Twenty years later, using different training methods, he got it (Lashley, 1931). The difference between success and failure can hide in the interaction between biology of motivation and the design of technology. Animal training chambers have been around for many decades. However, most training systems require a fair amount of human oversight, thus limiting automation.

The automation of software and hardware for high throughput rodent behavior is aptly described as a "mere" engineering challenge. However, when applied to a particular scientific question, the quantitative shift enables a qualitative change in the way that science is performed. Animal training ceases to be determined by the idiosyncratic intuitions of individual trainers. Instead, the training process is digitally documented and reproducible across subjects. And most importantly, the number of observations greatly increases. All of these are attributes of digital information systems that can scale to large numbers with small additional costs.

Training the first rat may have cost about ten thousand dollars, mostly in human labor. Five years later, the cost of training mostly involves technicians and veterinary costs, resulting in one trained animal for one hundred times cheaper. A fair estimate of cost is the price per trial averaged over the entire system, including all parts, labor and employees. This is approximately 3 cents/ trial when 100 rats are trained for 90 minutes each day. The sustained costs after development are even less. Thus, the win for science is not merely the supply of a large number of homogenously trained animals, but the opportunity to use statistical models that are only possible with large data sets.

Low cost behavioral training allows basic research to tackle more challenging questions. It also makes it possible to combine multiple techniques, without the cost of animal training becoming overwhelming for the individual researcher. The goal of developing this system was twofold: 1) to provide trained animals for subsequent visual physiology experiments, and 2) to investigate the capacity of rats to perform

visual tasks suitable to address questions of pattern vision, attention and decision making. A priority for the system was to be able to train many animals in parallel with limited oversight, and reliable results.

2.2 Previous methods

Here I review other systems in which rats are trained on a goal-oriented task in a two alternative forced choice (2AFC) paradigm.

Historically, progress in animal training was made by examining the swim path of a rat searching for a hidden pedestal (Morris, 1984; Robinson, Bridge, & Riedel, 2001). The swimming approach was adapted to a 2AFC experiment by letting a rat swim down a Y-shaped box to find a panel at one of the two locations (Douglas, Neve, Quittenbaum, Alam, & Prusky, 2006). A visual display placed at the end of each arm of the Y-maze enables the rat to choose which arm to swim to based on a visual cue. While water systems have a history of success, it requires constant human oversight to pick up the rat and set up each trial. It is hard to automate, and scales poorly. The use of water is also difficult to combine with physiology.

An alternative touch-screen technology has been applied to automated behavioral training so that a rat may choose one of the stimuli presented on a computer screen by touching a region near the screen (Bussey, Muir, & Robbins, 1994; Bussey et al., 2008; Cook, Geller, Zhang, & Gowda, 2004). Their training system provides food pellet rewards for successful choices, located in a hopper at the other end of the

chamber. A downside of this reward strategy is that the location of the reward involves a time delay which may be more difficult to reinforce with than rewards immediately following an action (within tens of milliseconds). Additionally, the act of chewing food causes artifacts in electrical recordings of brain activity, which are often desirable to collect in parallel to the animal's behavior. Finally, the touch screen responses and food pellet rewards are difficult to combine with neural recording that require the animal's head to remain motionless.

Rats have also been trained on visual tasks by using lever presses in exchange for water rewards (Munoz Tedo, Herreros de Tejada, & Green, 1994).

The training system we developed was inspired by the work of researchers at the Cold Spring Harbor Research Labs. The system of animal training they had developed had been applied to train rats on olfactory (Uchida & Mainen, 2003) and auditory tasks (Otazu, Tai, Yang, & Zador, 2009). Their system allowed for the independent request of a trial at a central port and the possibility of responding and receiving a reward at one of two symmetric locations. Our training system was different because focused on the visual system of rats. Additionally we made an effort to keep the cost per system particularly low (less than a thousand per unit). At a later date, another group of researchers (Zoccolan, Oertelt, DiCarlo, & Cox, 2009) also developed system that successfully trained rats on visual tasks. While they used contact sensors instead of infrared sensors, they also coupled rewards at the same spatial location as the response. In their training system, the left and right response ports were about a centimeter apart.

Proximal response sensors provide the possibility that rats can transfer their learned responses to a context in which they are not free to move their head. About 10% of rats that I trained used an operant chamber that required the rats to lick one of three sensors, all within one centimeter of one another. However, the rate of learning was faster with spatially separated sensors. Also, on average rats performed better when the sensors were farther apart. More precisely, when the task became harder, the rats using proximal sensors were more likely to resort to random responses. Thus, all of the data reported here are from operant chambers in which the animals move their upper body to make a response. Proximal sensors will probably be useful in the future for combing animal behavior with head fixed neural recording methods.

2.3 Description of use

I describe the process of using the training system. A schematic diagram of the architecture depicts the flow of experimental control and data collection (Fig. 2.1). The software architecture allows each researcher to use the same core behavioral training code, but define their own experiment. Researchers completely define each training protocol in software, using object oriented Matlab code. Protocols specify everything about the training process: the stimulus manager which renders images, the trial manager that determines the trial structure and experimental paradigm, the reinforcement manager that determines the rewards and penalties, and the graduation criteria that specifies when to advance to the next training step. After debugging and

testing code on a local computer, the software protocols are created and stored on a training server, and rats are assigned to the protocol in a database.

Training sessions were performed by research technicians. Once a day the water level was checked and the water pressure set to 350mmHg. Rats were weighed to confirm they were maintaining a healthy body weight. Before each session, the rats' home cage was transferred so that it was adjacent to the training chamber (Fig. 2.2). A door was opened in the side of the home cage, and connected to a tube that allowed free passage between home cage and training chamber.

Each training station had a training chamber with a dedicated computer and monitor. Training stations were organized into racks of 6-9 stations. Four to six heats were run every day. Technicians began a session for all the stations in a rack by using a software graphical user interface to specify what heat was ready, and starting the program. The technician's training server polled a database to look up which rats were assigned to each training station. The server transferred the training protocol to the client computer, along with the relevant training history for that subject. At this point the software began a training session and the ports were responsive. Three lickometer ports were arrayed along the wall of the training chamber. Each port detected a rats licking response with an infrared beam break (Fig. 2.3). When a rat licked the central request port, the software presented the appropriate stimulus.

During the penalty or reward period, the stimulus for the next trial was rendered and cached. Using this organization it is possible for stimuli to be

dynamically rendered appropriately depending on the subject's previous response. Yet the caching prevents the computational time required to render stimuli from interfering with the reliable presentation of the stimulus.

At the end of a session the behavioral data was sent to a data server. Additionally, an automated process extracted the compressed description of the raw data. The raw data was a good repository of information because it was exhaustive and immutable; however, its size and structure were unwieldy for analysis spanning multiple days or many subjects. The compressed format only documented the most relevant details, was optimized for subsequent use, and could be regenerated from the raw data in a consistent form (to accommodate past data formats and future analysis needs).

2.4 Considerations for setting up rodent behavior

Setting up a new training system involves many choices. Some choices are constrained by experimental needs. But often researchers seek a method that simply works reliably. Here I provide the choices we made and our reasons for making them. We occasionally explored alternatives, but never exhaustively investigated alternative paradigms once we discovered something that worked. This section is organized by response modality, experimental paradigm, reinforcement method, display technology, and bias prevention method.

2.4.1 Response modality

How do you measure your subject's behavior? What constitutes a choice? To study the properties of a visual system requires the subject to learn a relationship between some sensory evidence and a particular motor response. In some cases it is possible to exploit reflexive behaviors such as the dilation of an eye in response to a flash of light (Trejo & Cicerone, 1982), or the movement of the head in response to optical translations (Prusky, Alam, Beekman, & Douglas, 2004). However, a reflexive behavior need not involve an active choice, is not learned, and does not generalize to different stimulus patterns.

Choosing which motor response to engage is pragmatic and specific to each species. The best behaviors to train are the ones that are natural, frequent, active, symmetric and specific. Natural behaviors are good not only because their bodies can generate them with ease, but also because it is likely that subjects represent them internally. It is harder to learn what past action resulted in a reward if you don't have an internal representation of it. Frequent behaviors are beneficial because they provide more opportunities for learning associations. To examine choices requires at least two possible responses. If one response is to "do nothing" - this is less desirable because the experimenter can't disambiguate if the subject was distracted from the task or intentionally chose not to act. This problem is avoided if both measured responses are active. Additionally, it is desirable if both actions are symmetric. For example, the actions would not be symmetric if one response was to navigate to a nearby location and another response was to navigate to a location that was farther

away. It's likely that subjects have an innate preference for actions that require less effort or less waiting. It would be better if the two locations were equally far away from the subject. While the example of distance is bland, the principle extends to any other asymmetry of two possible responses. Finally, it helps if the action is specific because it enables the measurement to be specific. Ideally there should be no doubt that the action, once performed, was exactly what the subject intended to do. It is bad if other similar actions are treated as false positives by the experimentalist. The later increases the perceived variability of environmental contingencies and invites idiosyncratic strategies between subjects. It also reduces the experimenter knowledge of what the subject's choice actually was.

One advantage of training rats, or other scavenging creatures, is that they explore their environments actively: they move around, smell things, touch surfaces, poke their noses into holes. Good candidate behaviors to train are the position of the body, the position of the head, the actions of the mouth, or the activity of the paws. While any of these behaviors could be assessed by a camera, it is certainly more convenient to design an environment that reliably detects a behavior with a local sensor. We applied this strategy to the rat's action of licking the reward spigot. Each lick is detected by an infrared sensor; the sensor responds to the shadow cast by the tongue. The reinforced behavior is the action of licking the spigot on either the left side or the right side of the training chamber.

Exploring the ports is active, specific, and natural for rats. The spatial location of the ports allows for symmetry of the response between the left and right. However,

after initial exploration, rats will not visit the port frequently. In order to increase the frequency of licking, we initially provide rewards at the port location, without any visual contingency. The goal is to establish a high frequency, such that other contingencies can be learned by the subject later on. Thus, the rat learns first that licking ports will provide rewards, and next that moving between ports will refresh the capacity of ports to provide rewards. At this point, the rat is producing the motor activity of a trained rat: moving between ports, licking the spigots and collecting rewards. The next step is for the animal to understand the structure of a trial: that the stimulus is contingent on licking the centrally located request port, and that the stimulus contains information about which side will have a reward or not.

2.4.2 Explicit stimulus requests

In our experiments, stimuli only appeared when subjects requested them. The alternative to having a defined request behavior is to present to continuous stimulus, or one that appear at random moments in time. This has some advantageous, such as reducing the total number of possible actions from three down to two. However, it probably comes at a cost. The subject is less likely to be "ready to process" the stimulus if he is not in control of when it appears. If the subject request the stimulus, learning might be quicker because the subject is more likely to have encoded the visual context which will determine if his subsequent action is correct or not.

2.4.3 Feedback informs the subject of the consequence of an action

It is usefully for humans to receive auditory feedback while dialing a phone number. Similarly we inform the rat of completing every intended action. This should allow the rat to separate failed licks (intended actions that were not registered by the sensors) from mistakes (actions that were wrong and produce no reward). From visual observation, I know that failed licks are quite rare: much less than one in a hundred. However, their incidence might be higher if the rats did not have feedback to learn about the efficacy of their motor actions.

Rich auditory feedback may also help the rats learn. The particular sound indicates the state of the training system. The request port makes a sound when the stimulus appears. This sound indicates both the precise moment in time that contains relevant stimuli, and that the next action should be to go to one of the side ports. The rat should know that it is on the path to receiving a reward soon. If the rat continues to request another stimulus he will get a different sound. The sound indicates a completed lick, but that continuing that particular action is a dead end: it alone will never result in a reward. A third and fourth sound indicate correct and incorrect trials. The meaning of all four sounds could be described as ‘keep going’, ‘wrong way’, ‘correct choice’ and ‘wrong choice.’

I suspect that the auditory feedback increases a rat’s sense of control over his environment: auditory events occur if and when he wants. Additionally, sounds might help rats identify salient moments in time, thus honing their attention on the visual stimulus before they have learned that it is predictive.

2.4.4 Structure of a trial

The structure of a single trial provides the evidence for a subject to make a choice. What is the nature of the evidence? What is the nature of the choice? Is there one stimulus or multiple stimuli? Are they spatially or temporally separated? I review experimental paradigms beginning with ones that we have validated, and concluding with other options that are worth considering in future experiments.

The first task that we train animals on is typically to go to the location of a stimulus. The rat indicates its choice by licking a response port directly under the stimulus. The task is easier for the rats to learn: they only have to notice the location of the stimulus, and not its properties. After subjects learn that the images on the computer monitor are relevant for obtaining water, we then switch to a different visual task. For example, we add a second stimulus with different properties. Thus, while performing the easy task of going to the stimulus, the rat had the opportunity to build associations with the stimulus' properties (e.g. the orientation of a target or the pattern of an image). The addition of a second stimulus forces the rat to make use of the differences between the stimuli to solve the task. When both a target and a distractor stimulus are presented at the same time, the localization task is referred to as a spatial two alternative forced choice (spatial 2AFC).

Discrimination tasks require subjects to report a property of a single stimulus. For example, the property could be the orientation of a target, the direction of motion of a pattern, or the proximity of an image to an exemplar. The subject must learn the

association of a response with a particular orientation, motion or image. Discrimination tasks take more time for rats to learn. For example, training a rat to discriminate between clockwise and counter clockwise oriented stimuli took more than six months of gradual learning until performance was greater than 85% correct. Only two of four rats learned the task: two other rats were removed from the study for never performing above chance. The exception to this is the motion task. In one case, a rat learned the motion task in three days. However, individual variability in learning can be high: other rats have taken weeks to learn the motion task.

In a detection task the subject reports the presence or absence of a stimulus, such as a grating. The task uses the "forced choice" paradigm if the subject is required to respond on every trial: one response indicates the present of the target ("yes") and another indicates the absence of the target ("no"). We refer to detection task with a forced choice as a yes-no force choice paradigm (YNFC). An advantage of a detection task is that the response can be independent of the target's property or location. For example, for the task use in Chapter 3, the targets orientation was different on each trial, but this fact did not influence what the correct choice was. For detection tasks it is particularly important to have an auditory tone paired with the stimulus interval, indicating that the stimulus interval did occur even when there was no change in the visual stimulus.

The flanker task is a special case of YNFC in which additional task-irrelevant stimuli are presented alongside the location of the target, whether or not the target is there. The advantage of using a flanker task is that it enables the researcher to infer

interactions between spatially distinct stimuli and also the neurons that represent them. A down side of the flanker task is that is adds 4-8 weeks of training time.

The go-no-go paradigm has the advantage of only requiring a single action. Thus it is most compatible with a head-fixed context in which animals only perform a reduced set of actions. Creating a device that allows a rat to lick three different spatial locations without moving their head has proven to be challenging. Prototypes of three port devices only allow the rats to explore two of the ports. From direct observations this limitation appears to be due to geometric constraints and not motivational ones. The disadvantage of go-no-go is that the responses are asymmetric and one of the choices ('no go') is passive. Failure to respond from inattention can be confused with intentionally withholding a response. It is easy to accumulate many trials in which the animal produces no response. It does not seem appropriate to treat every "trial" as a "choice." This go-no-go data is challenging to incorporate into a clean statistical description of performance.

I believe it would be possible to train rodents on other paradigms including a two interval force choice, odd man out, or delayed match to sample. However, I have not personally attempted to train any of these paradigms.

A two interval force choice allows the stimuli to be placed in a single spatial location, which might belong to a single receptive field, and it can provide a referent on every trial. For a spatial forced choice the target stimulus and the distracter stimulus are in different locations, and it is unlikely that one could get a neural response for both stimuli on a single trial. There are two downsides to using temporal

intervals: the influence of temporal adaptation, and the challenge of training rats to wait. Temporal adaptation can be reduced by increasing the delay between the first and the second interval. However, increasing this delay will require rats to learn to wait longer. If rats are able to respond immediately after the first interval, they will never have the opportunity to make a comparison to the second interval.

The odd man out paradigm requires the subject to select the one stimulus that is different from the other two. A given image would consist of three stimulus patches, each associated with a response port. The odd man out paradigm offers the possibility of a powerful generalization. No property is special: there is not a target to be followed, nor is there an association between the specific property and a particular response. If a rat has understood the abstract rule, then the experimenter can explore many different stimuli properties. A negative aspect of the odd man out paradigm is that it is asymmetric. All three ports are used for responses, and thus a rat does not start a trial from a fixed position. For example, sometimes the rat views the monitor from the right port and must choose whether to stay put, move a little to the center, or move a lot to the left port. The opportunity to stay at one port (correct 1/3 of the time) may be too tempting, and thus it might be necessary to only allow the unique stimulus to appear at one of the other ports. At this point, the paradigm resembles non-match to sample: the rat must choose the stimulus that is different from the one located at the port it last responded. Other labs have trained mice to perform an odd man out task for color discrimination (Jacobs, Williams, Cahill, & Nathans, 2007).

2.4.5 Motivating rodents

How do you convince a rat to perform your task? Applying reinforcement learning requires rewarding the subject after performing a task correctly, or punishing the subject after performing a task incorrectly.

We avoided using direct forms of punishment. Instead we initiated a "time out" penalty period after the animal made an incorrect response. Presumably this was undesirable to the subject because they had to wait before the next opportunity to request a trial. I typically used penalty periods of four seconds, though sometime I reduced them to as low as one second or as high as ten seconds.

Interestingly, having a penalty period is not necessary for animals to initially learn an easy task. Rats learned to go to the side of the chamber that contained a stimulus, even when the penalty period was zero. However, there is always a risk that an animal's performance reduces to chance, especially if the task is difficult. The presence of a penalty period increases the consequence of an error because it decreases the frequency that a subject can perform trials. For example, an animal performing at chance might only be able to perform ten trials in a minute: five correct trials about two seconds each and five incorrect trials lasting about ten seconds each. Thus, penalty periods are useful for sustaining motivation when the task is difficult, and the subject may be tempted to just guess. After some experimenting I ran all my rats with 4 second penalties. I increased this duration to 10 seconds only if a rat was struggling to learn a hard task, or if their performance dropped on a task for an unknown reason.

What is a good strategy to deliver positive reinforcement? There are three main options to choose from: food rewards, water rewards, or other tasty liquids. It is possible that food rewards have the strongest reinforcing capacity. However it is more challenging to rapidly deliver small volumes of solid food, especially if there is not a constant reward size for each correct trial. I do not have experience training animals with food, but I would be concerned that it might be hard to motivate subjects to do as many trials a day with food rewards. With water rewards, rats reliably perform hundreds of trials a day. With a little effort on the experimenter's part, using small reward volumes and short penalty periods, it is possible to motivate rats to perform almost a thousand trials within a two hour window. The number of trials per session, averaged over each rats compete training history, ranged from 200 to 800 trials a day.

I have found water restriction and water rewards to be ample for motivating rats to do trials. Other researchers who focus on mice suggest that limiting their food intake and provide sucrose sweetened water is helpful to motivate mice (Johnson, Crombag, Smith, & Ramanan). Interestingly, they claim that even though mice will still perform trials for water, they won't learn as well. Another researcher with experience training a range of rodent species, insisted that fluids with protein are better than fluids with just sugar, and recommended soy milk (Jacobs et al., 2007).

2.4.6 Managing bias

Ideally each trial is truly a random draw. That is, the correct response has no correlation with a property of the past. Unfortunately, during initial learning rats

commonly get stuck by perseverating on a single side. This is an extreme form of bias in which the rat always chooses one response over the other. The dangerous thing about this bias behavior is that it is self reinforcing: the rat continues to benefit from a 50% variable reinforcement schedule. The risk is that the rat will never explore, and thus never learn.

I used correction trials to prevent these side biases from forming. If the animal makes a mistake, a fraction of the time initiates a correction trial that requires the animal to switch sides. Correction trials persist until the rat finally responds differently. Thus, the correction trial will prevent a sustained side bias from being reinforcing. For most of my rats, the probability of a correction trial after an error was fifty percent during training. Correction trials were either removed during testing, or excluded from list of trials that are valid for analysis. Unfortunately, the use of correction trials can induce compensatory behavior, whereby the rats learn they are more like to get a reward by switching sides after an error. The rat's incentive to exploit correction trials is lower if they occur less frequently. For this reason I ultimately preferred to set the probability of a correction trial at only ten or twenty percent of the trials after an error.

Additionally, during training I modified the random sequence of trials to exclude runs of more than four in a row on the same side. Since the trials were not random, the forced switches were also removed from all analysis of performance. The reason is that I wanted to guarantee that there was no way the subject could generate above chance behavior without actually learning the visual task.

2.4.7 Choosing a display device

All stimulus display technologies have imperfections. Depending on the experimental requirements, some of these imperfections can have severe consequences on interpreting one's results. I chose to use a CRT because it provided better control over the amplitude of contrast, and produced a consistent temporal response for all contrasts.

The biggest concern about many flat screen displays is that they have a strong angle dependency of their output. From the point of view of energy efficiency, the angular dependence is a good thing. Most of the energy is oriented directly in front of the screen, where most users tend to be. However, this can cause problems interpreting experimental results because the effective contrast produced by the monitor is not known. If the subject is close to the screen, as rodents often are in our experiments, then the pixel in front of the rat produces more light than the pixel to the side of the screen. As a consequence, the contrast of stimulus will vary with its position on the screen. CRTs are much less influenced by the viewing angle.

Second, some flat screens like LCDs have a strong temporal dependence of their output. For example, a gray value of 127 will produce a difference amount of candelas if the previous value was 0 as opposed to 255. Engineers have attempted to fix this problem by creating a custom set of rules that minimizes the errors for all the transitions. To make matters worse, the specifications of gaming world used a metric that incentives manufacturers to transition as fast as possible to the final state,

producing overshoots and undershoots that might influence the neural responses of the early visual system. And all this might be occurring unbeknownst to the researcher! An example where this could be a concern is during the reverse correlations of neural response in response to a flickering stimulus that is presumed to be independent. Also, it could influence what is salient in a detection task.

Finally, pixels are not spatially independent (Pelli, 1997). This is true for CRT monitors as well as other display devices I tested. It is noticeable when comparing a high spatial frequency stimulus that is horizontal as opposed to vertical: vertical stimuli are lower contrast by about 5% (Garcia-Perez & Peli, 2001). Thus, orientation discrimination task should avoid using horizontal and vertical stimuli, as this task may reduce to contrast discrimination and have nothing to do with orientations. A simple solution is to keep stimuli symmetric about the vertical axis. This matches the horizontal spatial frequency of the different orientations, and prevents contrast from being a useful learning cue.

CRT monitors are becoming more difficult to obtain, especially in the large numbers of identical models which are preferred for a high throughput system. OLED monitors may be promising display technology in the future, though early prototypes available for heads up displays seems to have substantial pixel specific noise. Given the current state of technology in 2011, I would consider using projectors in order to provide stimuli with good control of timing and angular independence. Experiments that do not require precise control of contrast with time varying stimuli can probably work fine with most display devices.

2.5 Observations of behavioral data: quality and volume

Good psychophysical data permits measuring changes in the subject's performance with respect to an experimental manipulation. Measuring changes is easier when the variability in performance is less from day to day and moment to moment. The ideal psychophysical subject has a constant performance over time and the primary influence on performance is the experimental manipulation. Additionally it is useful if a subject consistently produce large volumes of data because this provides more statistical power. Are rats good experimental subjects? Will they perform consistently across days? Will they perform consistently from the beginning to the end of a session?

Rats are good psychophysical subjects for visual tasks because they perform many trials, can attain high levels of performance, have manageable bias, and respond quickly to stimuli. Here I take a more detailed view of a single subject's behavior in the first two months of data collected. If rats perform enough trials, it is easier to estimate their performance accurately in a limited time. During initial learning the number of trials per day was more variable (20-400 trials per day) but stabilized to 200 trials per day after the rat learned the task (Fig. 2.4a). Average daily rates were between 200-800 trials for trained rats. The range in trial volume was mostly influenced by the subject and protocol, and not by daily fluctuations. After learning, most rats maintained stable trial volumes for months at a time (Fig. 2.4b). Rats can obtain performance levels greater than 95% correct (Fig. 2.7b). Often however, rats'

performance will saturate around 80-90% correct, even for easy tasks. Rats commonly have a bias to choose one side more than another, especially early in training (Fig. 2.4c). This bias is automatically improved with correction trials. Finally, it would be bad if subjects took a long time to respond to a stimulus. Subjects that wait long periods of time to respond may be distracted and less likely to learn about the stimulus. Rats however, respond fairly rapidly - in less than a second (Fig. 2.4d). Most of this time is spent moving to the response port, so the decision process is likely to take a few hundred milliseconds.

At the beginning of each session rats perform trials faster than at the end. Specifically, they spend less time waiting between trials, resulting in higher rates per minute (Fig. 2.5). This suggests that their motivation might be changing over time. Indeed, this is likely because the rats are thirsty from water restriction, and this thirst may abate with each correct trial. Does this change in trial rate have implications for performance over time? Are rats less motivated to perform well on hard tasks at the end of a session? A well trained rat performing a difficult discrimination task maintains a performance of 80% correct; this does not change from the beginning to tend of the session (Fig. 2.6).

To quantify performance stability within a session for all rats, I compare the performance of the first and last hundred trials of a session. There were 9296 sessions in which rats performed at least 200 trials and were significantly above chance (lower bound of the confidence interval for binomial proportions > 0.5). Of these, 60% of the time, the performance at the end of the session was higher than at the beginning. The

increase in performance was significant (Kruskalwallis, $p<10^{-10}$), though modest (2.2%). A small increase may be due to learning on a short time scale, though it mostly disappears the next day. Alternatively, it could be that the subjects are very thirsty, and perform sub-optimally on the initial trials. This is consistent with the slow rise in performance over the first 50-100 trials (Fig. 2.6).

My initial concern was that subject's performance might degrade on later trials, possibly even dropping to chance performance. In fact the opposite is true: subjects perform better at the end of a session than at the beginning. While a small difference exists, it does not interfere with data collection: it still makes sense to pool that data from the entire session.

Rats reliably learned detection tasks, but their learning rates were variable, even among brothers that trained in the same station. One rat may learn the task in two days, while another may take two weeks (Fig. 2.8). Variability in learning was influenced both by rats that learned late and rats that learned slowly. Some rats spent many days at chance behavior without demonstrating any learning at all; then, all of a sudden they would begin to improve on the task. Other rats began to learn from the very first day but displayed a gradual rise in performance each day, obtaining 85% correct after more than a week. Two of ten rats that began in the comparative study of learning were removed because they failed to engage with the system, and were placed on free access to water. While the natural variability of individual differences presents challenges, the use of a high throughput behavioral system makes it possible to obtain

results despite this variability. Additionally, it opens the door to the possibility of studying the variability itself.

How much intervention was required to train animals? Was the training process *entirely* automated? Intervention was most useful in the first training step in which the rats began to engage with the system, before the stimulus was included. Without close monitoring as many as half of the rats could end up failing to learn in a reasonable amount of time (for example, they might fail to perform above chance for a month, and then be removed from the study). Interventions involved minor software changes like returning a rat to a previous step, or advancing him manually. For example, initially every port produced a reward on every visit. On the next step the rules would change, creating some risk that the animal will give up, or cease to explore each port. The researcher has to estimate when the animal is sufficiently engaged that it will continue to explore even when some ports appear to be inactive. In principle, the automated graduation could estimate this. The software confirmed engagement by checking that the rat had performed a sufficient trial rate (5 trials per minute) for a sustained period (5 continuous minutes). Sometimes the software held a rat on the same step for too long. Humans were better than the software at estimating the rat's engagement. In summary, constant oversight in the first two weeks of training substantially increases the number of rats that learn. However, after the first visual learning was achieved, constant oversight was not necessary. Only two variables were actively managed: the reward quantity and the penalty duration. Performance was monitored by inspecting the data (and not the rats) daily to weekly.

Changes were made to reinforcement parameters about once a month, usually affecting less than a quarter of the subjects. Some rats never required any personalized changes to their reinforcement schedule.

In a given month it was possible for a researcher to oversee the training of twenty subjects, where each subject performed a session every day. This was achieved by running four or five different subjects per training chamber each day. A high level overview shows the performance of about 10,000 trials for each subject collected over one month (Fig. 2.9). During this month some subjects began to learn advanced tasks, some subjects performed stably during data collection, one rat reduced its performance, and others had performance that fluctuated as I explored different stimulus parameters. Data is the average of performance across all trial types which include randomly interleaved difficulty levels for some rats. Over the course of three years, the rats that I trained collectively performed a little over eight million trials (Fig. 2.10).

Rats learned to perform a range of visual tasks, attesting to the generality of their visual system and the methods for training (Fig.2.11). Rats learned to respond to the location of stimuli, the angle of a grating, or the motion of drifting dots. Rats can learn to detect the present of a grating in isolation, as well as when it was flanked by other gratings. Rats can discriminate between natural images by selecting the location of a target image over a distractor image.

Comforted that the training could be generalized to different tasks, I focused on the detection task in the presence of flankers. This task provided independent

modification of target stimulus and the surrounding stimuli that established a spatial context. The next chapter documents how the geometric properties of the surrounding stimuli influence the detection task.

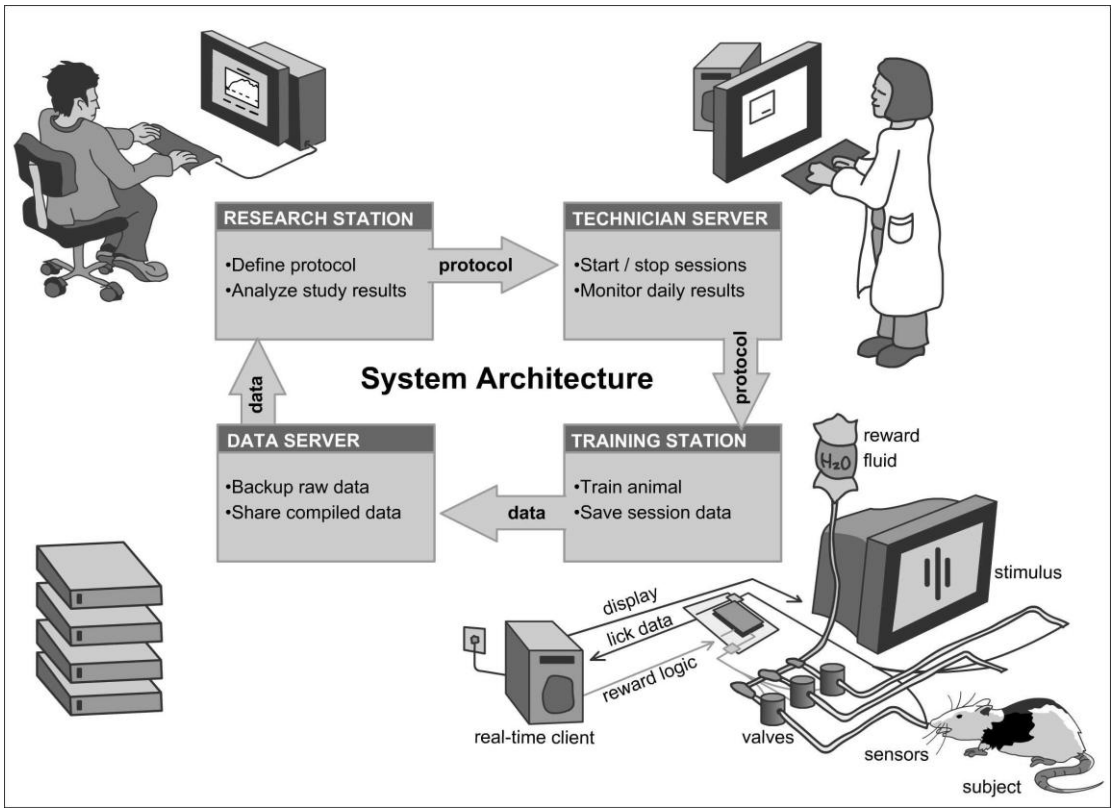

Figure 2.1. Schematic diagram of the training system's architecture. Experimental control is issued from the researcher to the technician's server, and relayed to each client training station which renders stimuli and maintains the appropriate reward contingencies. The data of interest are the states of the animal's nervous system, which can be inferred from its behavior. The animal's choices are measured by detecting the shadow cast by its tongue as it licks each response port. Data is temporarily stored on the client training station. Data is relayed to the data server after each animal finishes a session, which is typically two hours long. An automated process compiles the relevant details of the behavioral data to facilitate subsequent access by researchers and technicians. A graphical user interface allows the technician to monitor and the researcher to analyze the behavioral data.

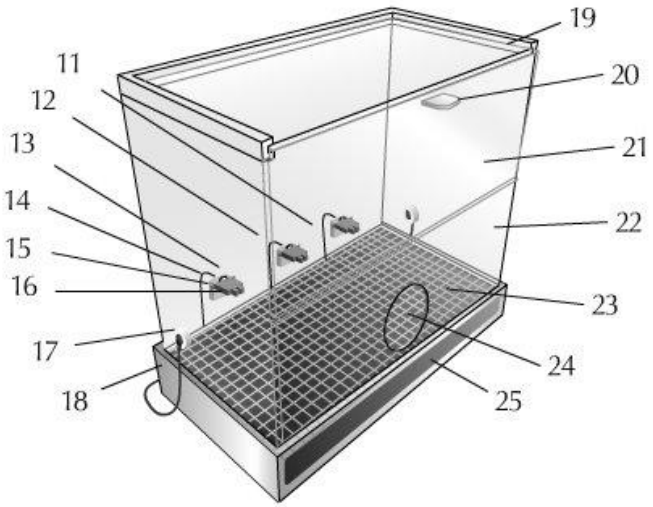

Figure 2.2 Diagram of a single chamber in which rats are trained to perform perceptual tasks. In front of the visual display is a left port (11), central port (12), and right port (13) which rats lick to receive rewards. Each port is composted of a fluid reward tube (16a), an infrared lick sensor (16b), a protective covering (16c), and a quick release knob (15). Two earbud audio speakers (17) held by an adjustable housing (17b) are mounted on the lateral walls of the training chamber. A collection tray (18) is below the floor (24) and chamber walls (22) which have a built in L-shaped lid (21) with a pull tab (20) that slides the lid out of the rails (19). An access slot (25) facilitates cleaning of the collection tray (18). Subjects enter via the portal (24). The dimensions of the chamber are flexible but symmetry should be maintained between ports for most training purposes, and the height of the sensor should be ergonomic for the species.

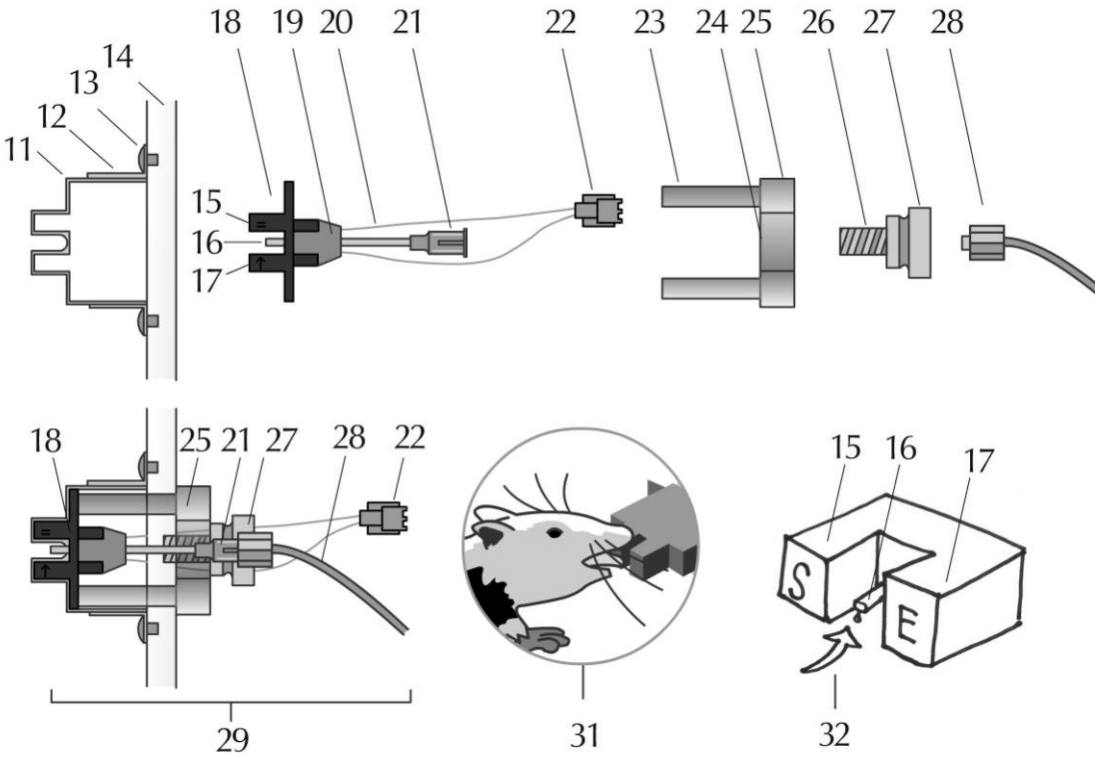

Figure 2.3. Components of a single lickometer port. In the context of the training chamber, a single assembled lickometer is referred to as a "port" (29) which senses animal responses and delivers fluid rewards. A protective housing (11) is fastened to the chamber wall (14) by means of a small L-bracket (12) and screw (13). The sensor (18) may be inserted into the housing and held in place by the positioning device (25) which is held in place by a screw (26) with a broad head (27). A luer lock (28) connects the water line to the reward spigot (21), which is a modified syringe tip (20) that protrudes through the sensor (18) so that the tip (16) is adjacent to, but not blocking the beam from the emitter (17) to the sensor (15). A rigid fixative (19) prevents the tip (16) from moving. Power and data are transferred by wires to the circuit board via a telephone jack (22). An animal uses the port by licking the tip (16) and triggering the sensor. The sensor (15) detects licks because the path of the tongue (32) interrupts a beam from emitter (17).

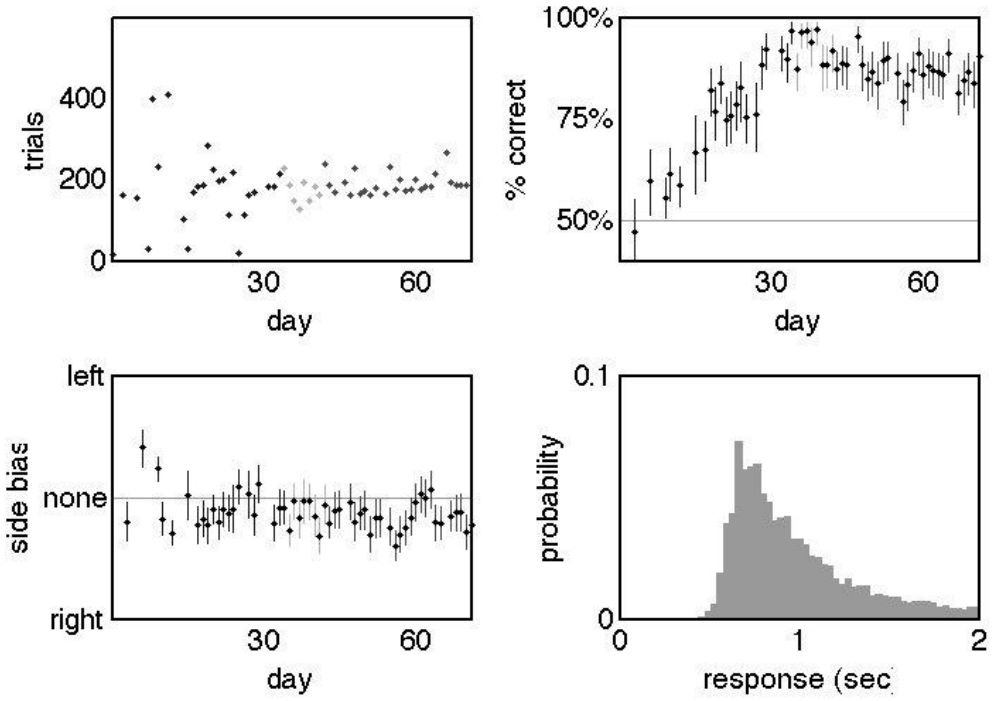

Figure 2.4. Trial volume, performance, bias and response time. Data is collected from a single rat during training and testing of a motion discrimination task for about two months. a) The number of trials that a single animal performs is plotted for each session. About 200 trials each day were valid for further analysis. Sessions were once per day and lasted 90 minutes. b) Performance is plotted as the fraction of correct trials each day. Over 54 sessions, the animal performed 14,074 trials and improved from chance performance up to 95% correct. The animal's performance decreased when the stimulus was made harder (day 33), but remained steady across days. c) The fraction of trials in which the animal chose the left port. The running average indicates that the animal initially favored the left side, but that this bias disappeared during learning. d) The duration between the request and response of each trial is summarized with a histogram over all trails. The modal response time is about three quarters of a second. Responses are typically faster than a second, but are occasionally longer when the animal does something else before responding. Responses longer than 2 seconds are not included in this graph.

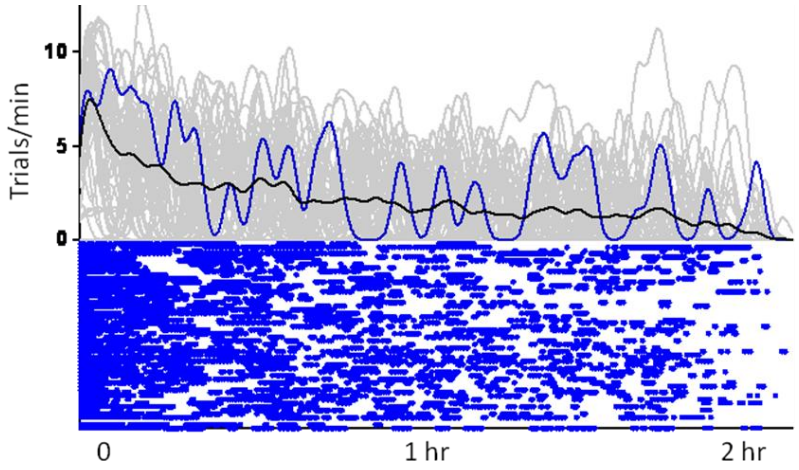

Figure 2.5. Trial rate decreases over each two hour session. The blue trace shows the trial rate during a single session, calculated by convolving the trial times with a Gaussian kernel (σ=5min). The X axis indicated the times since the beginning of a session. Each session ends after 2 hours. The gray traces display the trial rate data from a single rat for each session over two months. The average of these rates is shown in black. Below, a raster of the same data is provided. Each dot corresponds to a trial. Each row along the Y axis corresponds to a single session, with the first session on top and the most recent on the bottom. The reduction in trials in the latter half of the session is caused both by a reduction in the peak sustained rates, and by a higher probability that the subject is taking a break.

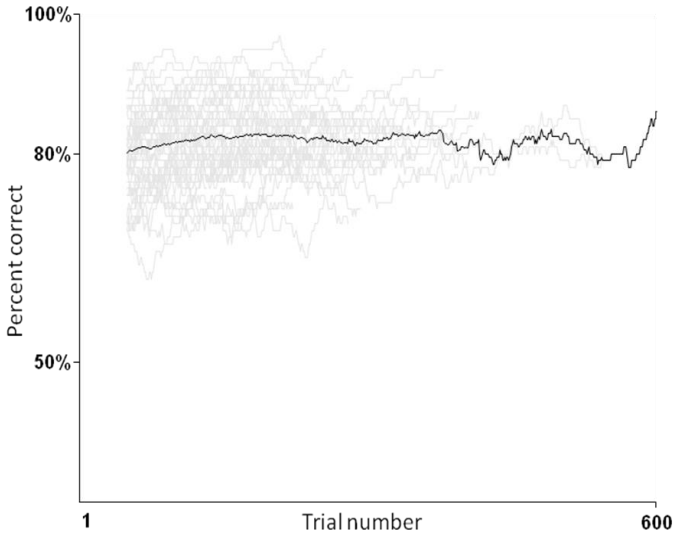

Figure 2.6. The percent correct performance of a trained rat does not change across two hour sessions. The percent correct performance is shown as a 50 trial moving average. Each gray trace corresponds to the performance from a single session. All data are from a single rat performing an orientation discrimination task over a two month period. Gray lines vary in length because the number of trials performed each day varies. The average performance is plotted in black. The average performance remains constant throughout a session. The data are the same trials as in Figure 2.5.

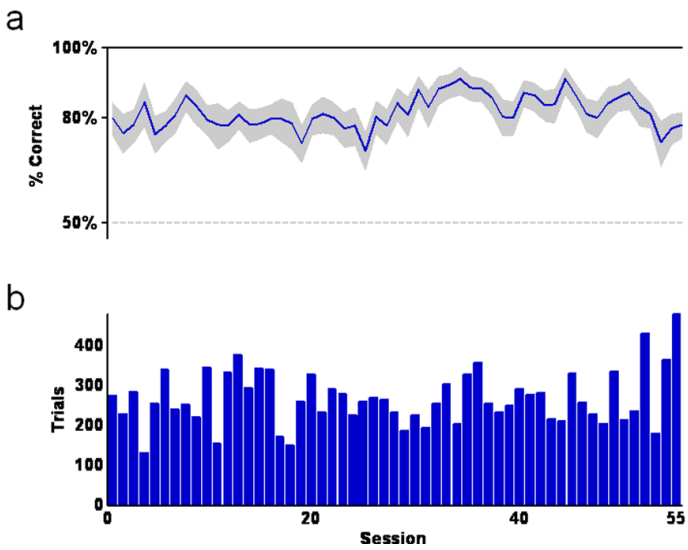

Figure 2.7. The performance of a trained rat is stable over days. a) The percent correct performance is plotted for 55 sessions within a 2 month period of time. During this time the stimulus difficulty was held constant and the protocol was not changing. The rat was performing an orientation discrimination task. The data are the same trials as in Figure 2.5. The blue line shows the average performance per session. The grey error bars indicate the 95% confidence interval for a binomial proportion. The changes in the rat's performance between days are modest fluctuations of +/- 5% correct. The subject never falls to chance behavior. b) This subject performs about 250 trials per session, and always more than a hundred trials per session.

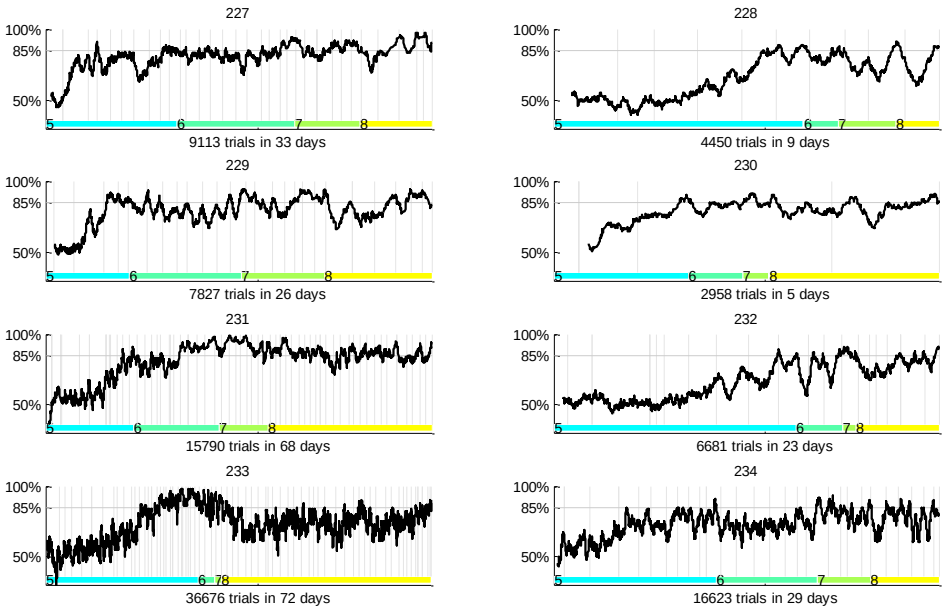

Figure 2.8. The learning curves of eight rats progressing through four shaping steps. Prior to the data presented here, the animals progressed through four shaping steps to become familiar with using the response ports and performing trials (See Chapter 3.4.3). Step 5 in cyan documents the first visual learning in which the rats detect a large high contrast grating. Step 6 corresponds to a reduction in spatial frequency. On step 7 the monitor was linearized, reducing the effective contrast. On step 8 the size of the target was reduced to a third of its size by decreasing the standard deviation of the Gaussian window. Training time ranged from 5 to 72 sessions. Subjects had a single two hour session each day. All rats that learned the initial learning proceeded to perform well on the harder detection task with smaller, thinner, fainter stimuli. Two rats began training at the same time as these eight, and were not included, because they never learned the initial visual task. Rats graduated when they were more than 85% correct for 200 trials in a row. Correction trials have been removed in the plots shown here.

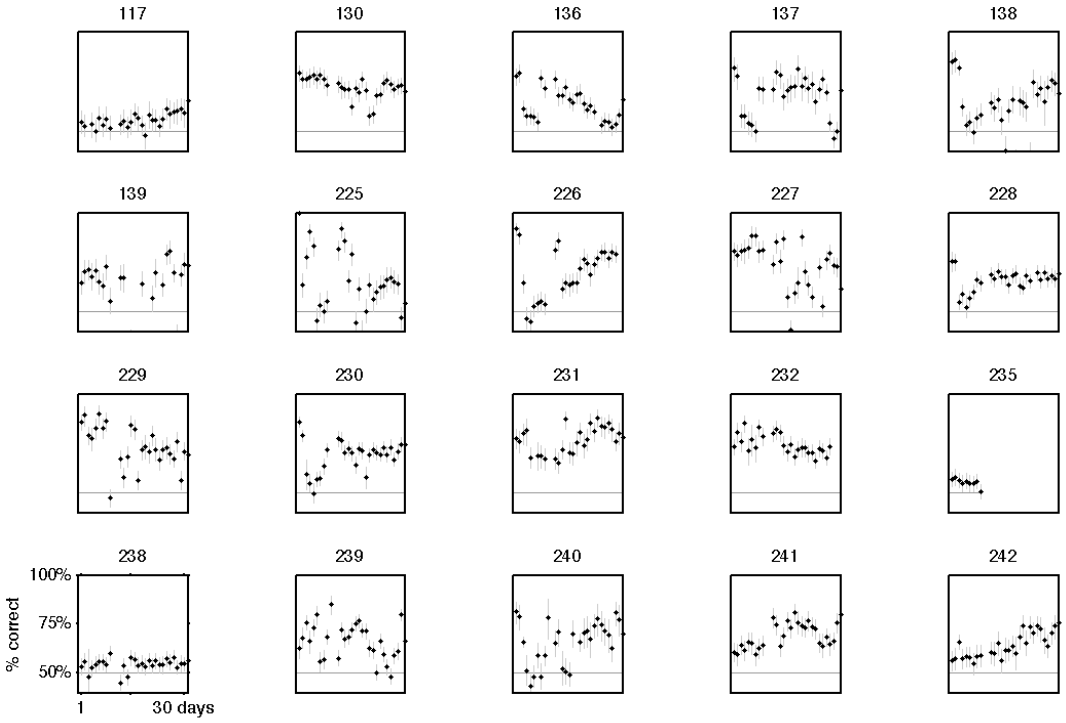

Figure 2.9. Overview of rat performance from data collected in May 2008. During this month, 20 rats performed about 10,000 trials each for a total of 213,884 trials. Each subplot shows data from a single animal, with the title indicating its ID. Each dot indicates the percent correct performance of the animal for a single day. The top of the graph is perfect performance and chance is the grey line. Error bars show the confidence interval for a binomial proportion estimated each day. The horizontal axis is time, spanning one month. The rats are not all on the same protocol: some are learning an easy task, others are being shaped by an automated process, and others are collecting data on a testing phase.

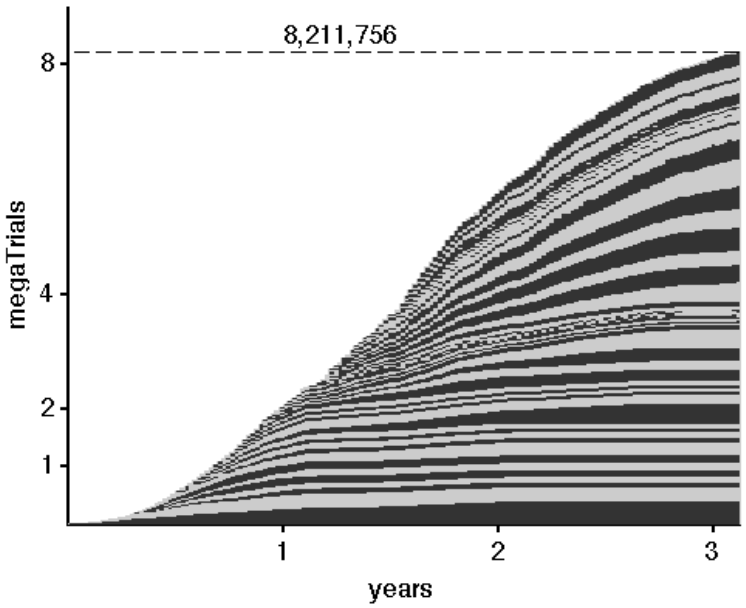

Figure 2.10 The volume of data collected by a single researcher over 3 years. Each stripe is the cumulative data produced by one rat. Twenty rats produced the majority of the 8 million trials that I collected during my graduate work.

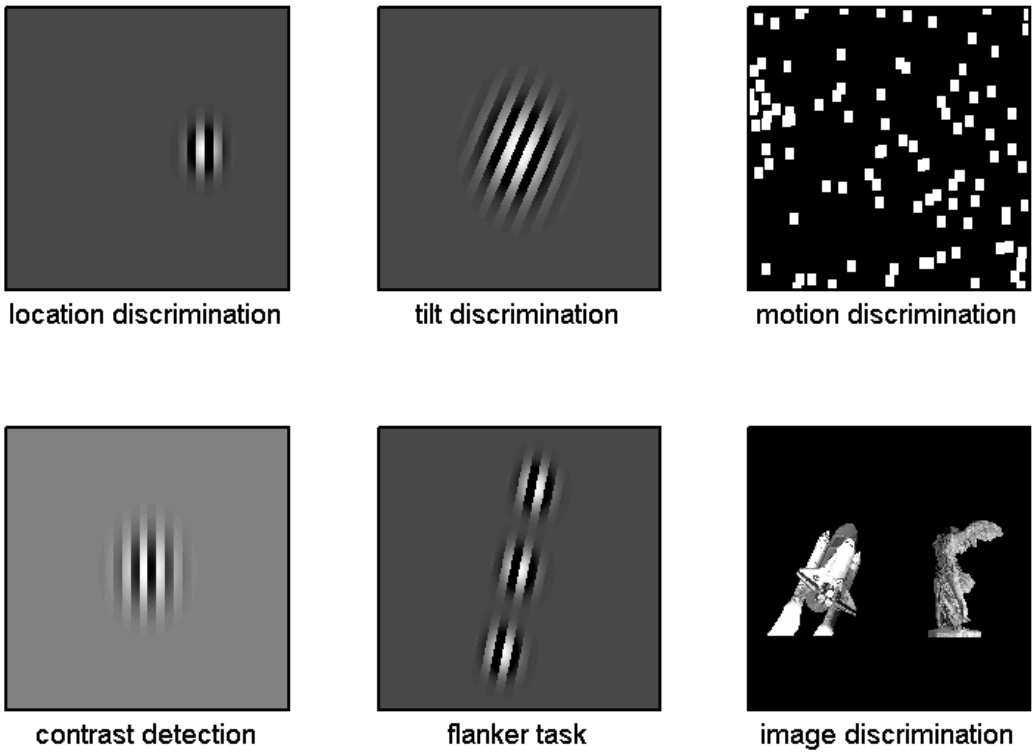

Figure 2.11. Examples of stimuli used to train rats. Each stimulus class would require more than on image. The example stimuli displayed here are the ones that indicate to the subject to choose the port on the right hand side. The correct answer would be the left port if the stimuli were different in the appropriate aspect. For example, the subject must learn to respond to the position of the stimulus patch in location discrimination, the orientation of the grating in tilt discrimination, direction of the dots in motion discrimination, the presence or absence of the stimulus in contrast detection, the presence or absence of the central stimulus patch in the flanker task, and the location of the target image (statue vs. space shuttle) in image discrimination.

Chapter 3:

Collinear features impair visual detection by rats

Abstract

We measure rats' ability to detect an oriented visual target grating located between two flanking stimuli (“flankers”). Flankers varied in contrast, orientation, angular position, and sign. Rats are impaired at detecting visual targets with collinear flankers, compared to configurations where flankers differ from the target in orientation or angular position. In particular, rats are more likely to miss the target when flankers are collinear. The same impairment is found even when the flanker luminance was sign-reversed relative to the target. These findings suggest that contour alignment alters visual processing in rats, despite their lack of orientation columns in visual cortex. This is the first report that the arrangement of visual features relative to each other affects visual behavior in rats. To provide a conceptual framework for our findings, we relate our stimuli to a contrast normalization model of early visual processing. We suggest a pattern-sensitive generalization of the model which could account for a collinear deficit. These experiments were performed using a novel method for automated high-throughput training and testing of visual behavior in rodents.

3.1 Introduction

In this study, we consider a classic task of visual psychophysics, the discrimination between the presence and absence of a visual target at a known location. Human perception of oriented targets is influenced by the contrast, spatial frequency and orientation of nearby stimuli, both in contrast discrimination tasks (Ejima & Takahashi, 1985; Xing & Heeger, 2001) and target detection tasks (Chen & Tyler, 2008; Polat & Sagi, 1993, 2007; Solomon & Morgan, 2000; Williams & Hess, 1998; Zenger & Sagi, 1996). This influence is typically largest when the surrounding stimuli match the orientation and spatial frequency of the target, both for annuli that completely surround a target and for discrete flankers (Cannon & Fullenkamp, 1996; Chubb et al., 1989; Polat & Sagi, 1993). The influence of stimulus arrangement and phase are more variable and subject to experimental paradigms. The behavioral influence of flanking stimuli has occasionally been studied in non-human primates (Li, Piech, & Gilbert, 2006), but never in rodents.

The psychophysical effect of flankers may be caused by surround processing in visual neurons, whereby features outside of the classical receptive field modulate neural responses (Angelucci, Levitt, Walton et al., 2002; Chisum & Fitzpatrick, 2004). The amplitude of neural responses in the retina, thalamus, and visual cortex are normalized to spatially nearby contrast (Carandini et al., 1997; Heeger, 1992; Shapley & Victor, 1979), likely due to lateral connectivity at each level. In many cases oriented stimuli in the surround suppress activity, and suppress most when surround

orientation matches the driving stimulus (Bonds, 1989; Cavanaugh et al., 2002; Polat et al., 1998). However, some cells in the appropriate contrast conditions increase their spiking activity when the orientation of stimuli in the surround matches the driving stimulus (Li et al., 2006; Polat et al., 1998; Sillito et al., 1993). In many physiology experiments, oriented surround stimuli are presented in an annulus. However, physiological surround effects can depend on the angular position of the flanker with respect to the target orientation, specifically influencing geometric arrangements like collinearity (Cavanaugh et al., 2002; Polat et al., 1998). Cortical circuits are likely to be involved in orientation-specific surround processing (Chisum et al., 2003; Das & Gilbert, 1999; Gilbert & Wiesel, 1989). It remains unknown if the orientation-selective circuits described in cats and primates are also found in rats. Rodents have orientation tuned cells in V1 but lack orientation columns (Ohki et al., 2005; Van Hooser et al., 2006).

From a theoretical perspective there are several reasons it would be advantageous for representations of local features to be sensitive to nearby image context (Series et al., 2003). In natural scenes, local image features such as luminance, contrast, and orientation are correlated at nearby locations (Field, 1987; Geisler, 2008; Ruderman & Bialek, 1994). When features are spatially correlated, surround processing can optimize the fidelity or efficiency of image estimation (Barlow, 2001; Geisler, 2008). For example divisive normalization from a nearby population of cells (Heeger, 1992) can allow a neuron to better adapt its sensitivity, reduce redundancy with its neighbors, and thus maximize information transfer (Schwartz & Simoncelli,

2001). Surround processing could also enhance salience of relevant features such as continuous contours (Field et al., 1993; Geisler et al., 2001; Sigman, Cecchi, Gilbert, & Magnasco, 2001) or statistically surprising features (Itti & Koch, 2000). These theories and others predict that different patterns in the surround should have distinct influences on a visual target's neuronal representation, even if lower order statistics like luminance and contrast are matched.

In the interest of developing a rodent model for the study of surround processing, we trained rats to report the presence or absence of an oriented target when sandwiched between two flanking stimuli. Rats have previously been trained on visual tasks including grating detection (Birch & Jacobs, 1979; Keller, Strasburger, Cerutti, & Sabel, 2000), motion discrimination (Douglas et al., 2006), orientation discrimination (Cowey & Franzini, 1979), and object recognition (Bussey et al., 2008; Minini & Jeffery, 2006; Zoccolan et al., 2009), but never on tasks with flanking stimuli. The presence of flankers made the task difficult for rats, presumably for both cognitive and perceptual reasons. In this study we are interested in the differential effects of the arrangement of flankers when they are present. The flankers' contrast, size and separation were held constant, while we varied their orientation, angular position and sign in randomly interleaved trials. We ask if rats' detection performance is sensitive to the relative orientation, position, and sign with respect to the target. We report an effect specific to collinear arrangements irrespective of sign.

3.2 Results

3.2.1 Rats can report the presence of a small oriented grating in the presence of flankers

We developed an automated method to train rats to perform two-alternative forced-choice (2AFC) visual tasks (see Methods). We trained 7 male Long-Evans rats to detect an oriented grating target. The target was presented in the middle of a CRT display, and subjects were required to select one of two response ports to indicate that the target was either present or absent (Fig. 3.1a,b; photograph in Supplementary Fig. 3.S1).

The orientation of the target was tilted either clockwise or counter clockwise from vertical by a fixed angle; orientation was randomly chosen on each trial. Rats advanced automatically through a series of training steps that decreased target contrast, reduced its size, and increased its spatial frequency (Fig. 3.1c and d, steps 5-8). After rats learned the basic detection task in the absence of flankers, a brief testing period assessed the influence of the target's contrast and spatial frequency was tested on a subset of the trained rats (Supplementary Fig. 3.S4a,c). This identified suitable parameters (see Chapter 3.6.2) for the subsequent more difficult task involving distracting flankers.

Next we added two "flanking" gratings on either side of the target location (Fig. 3.1b, Fig. 3.1d steps 9-10). The target was absent on 50% of trials. Flankers were absent on 5% of the trials. During testing (step 10), flanker contrast was fixed at

40% and spatial parameters of the stimuli were independently varied (θ_T, θ_F, ω, S_T, and S_F, as described in Fig. 3.3). Rats learned to perform target detection even in the presence of distracting flankers. We collected >20,000 trials per rat over 2-5 months of the testing step, which are further analyzed and summarized in Figures 4 and 5, and Supplementary Figures 3.S3, 3.S5-S7. Throughout the testing step, performance on trials with flankers remained well above chance (e.g., step 10 in Fig. 3.1c). Performance was stable over the period of data collection used for analysis (see Experimental Procedures, Chapter 3.4.4).

The presence of flankers made the task substantially more difficult for rats (Supplementary Fig. 3.S3). This effect is significant for 7 of 7 rats *individually (Agresti-Caffo 95% confidence interval) and* the population as a whole ($p<0.01$ on 2-way ANOVA; $p<0.01$ on Friedman's). Presumably this is due to both cognitive and perceptual factors, which we have not disambiguated here, but are considered elsewhere (Meier P, 2010). A perceptual effect on detection could arise from spatial contrast normalization, a form of surround processing long-observed in other mammals (see Introduction). Flankers add contrast to the target's surround, which could lower the target's effective contrast through contrast normalization. We verified that lower target contrast impairs detection as expected (Supplementary Fig. 3.S4c), so detection should be sensitive to reduction in effective contrast. Flankers of higher contrast or closer proximity to the target should exert stronger contrast normalization, further reduce effective target contrast, and lead to larger impairments. Additional tests on a subset of rats confirmed both predictions (Supplementary Fig. 3.S4b,d).

Thus our task is a promising candidate for revealing effects that might depend on contrast normalization.

As indicated in the schematic of Fig. 3.1b, the rat's decision and response in our task obviously depends on contrast at the target location. A more complete schematic (Fig. 3.2) indicates that rats' decisions are also influenced by the presence of flankers. A contrast normalization component is indicated on the basis of past literature and in consistency with the data summarized above. The presence of flankers also affects performance for uncharacterized cognitive reasons, such as task confusion and compensating strategies. If performance is insensitive to the position and orientation of the flankers, a simple model like this would be sufficient.

3.2.2 Collinear flankers impair detection more than other arrangements

We next considered how performance depended on the arrangement of the flankers with respect to the target. There were two possible target orientations, as during training. Flankers always shared the same orientation and sign as one another, and were located symmetrically on either side of the target location on an imaginary line tilted either clockwise or counter clockwise (Fig. 3.3a). The target orientation ($\pm\theta_T$), flanker orientation ($\pm\theta_F$), and angular position of the flankers ($\pm\ \omega$) were chosen independently each trial for a total of 8 randomly interleaved stimulus configurations (Fig. 3.3b). The luminance signs of both target and flanker gratings ($\pm S_T$, $\pm S_F$) were also randomized for each trial.

We use the term "collinear" to refer to stimulus configurations in which the target and flanker orientations are both aligned with the flanker angular position ($\theta_T = \theta_F = \omega$), irrespective of the relative sign. This configuration could engage visual processing that relates line segments that fall along a common contour. We label non-collinear conditions as follows: "pop-out$_1$" ($\theta_T = -\theta_F = \omega$), "pop-out$_2$" ($-\theta_T = \theta_F = \omega$) and "parallel" ($\theta_T = \theta_F = -\omega$). For examples, see Figure 3.

The main finding of our study is that the collinear condition is consistently harder for rats than any of the other three configurations (Fig. 3.4a,b). This difference was true for each rat and significant at the population level even when adjusting for multiple comparisons ($p<0.01$ by Tukey-Kramer on 2-way ANOVA for all three comparisons; two of these comparisons were also significant by the more conservative Tukey-Kramer on Friedman $p<0.01$; see Supplementary Data in Chapter 3.6.3, Supplementary Fig. 3.S6). The other three conditions were not significantly different from one another (Supplementary Fig. 3.S5, 3.S6). In short, of all the flankers conditions tested, only the collinear condition was consistently most difficult.

All stimulus arrangements had the same flanker contrast and distance, so this difference cannot be explained by simple contrast normalization as illustrated in figure 2. We consider next how the contrast normalization framework could be extended to account for a collinear deficit at the level of early visual processing. The collinear flanker condition differs from the pop-out$_1$ condition only by flanker orientation, so the difference in performance (Fig. 3.4c) suggests orientation-sensitive surround processing. This could be explained by a simple modification of the model in Figure

2, such that the strength of contrast normalization (E) is stronger when the contrast in the surround shares the target's orientation. Such a model could account for our result that the collinear flanker condition was harder than either pop-out condition, but could not explain why performance in the parallel condition was significantly better than collinear (Fig. 3.4d), and indistinguishable from either pop-out (Supplementary Fig. 3.S5). Our data cannot be explained by simple orientation-dependent effect, nor by an angular position effect alone. Feature arrangement is important: flanker orientation and angular position interact. To capture differences between flanker conditions within the perceptual component of the model, it would be necessary to add a pattern-sensitive term (see Discussion, Fig. 3.6).

Flankers that are collinear to the target either had the same luminance sign ($S_F = S_T$), such that white bars of the target align with white bars of the flankers, or opposite signs, such that white bars align with black bars ($-S_F = S_T$). A reversal in sign is equivalent to a π shift in spatial phase. We could find no effect on performance of the relative or absolute luminance sign of the flanker and target gratings. In particular, whether the dark bars in the target aligned with the dark or light bars in the flankers, the specific impairment for the collinear configuration relative to other arrangements remained, and the amplitude of the effect was indistinguishable ($p>0.05$, 2-way ANOVA; $p>0.05$, Friedman's test; Supplementary Fig. 3.S7). In a pilot study, intermediate phase shifts also had no effect in a related task (Supplementary Fig. 3.S7). Therefore we do not include phase as a parameter in Figure 3.6.

The arrangement of nearby features has been shown to affect behavior and early visual processing in other species (see Introduction). This is the first demonstration of pattern sensitivity in a rodent, showing that such effects can occur even in species that lack orientation columns. This is also the first flanker study in any species in which both orientation and position were randomized in a single interleaved testing period, confirming that neither a position effect nor an orientation effect are sufficient to explain the collinearity effect.

3.2.3 Collinear flankers cause rats to miss the target

Performance reflects both the ability to say yes when the target is present (hits) and the ability to say no when the target is absent (correct rejections). We find the collinear condition decreases accuracy of both kinds. For each rat, the hit rate was lower for the collinear condition than the parallel condition (Fig. 3.5a; vertical axis in Fig. 3.5c,d). The false alarm rate was also higher for the collinear than the parallel condition (Fig. 3.5b; red horizontal axis in Fig. 3.5c,d). The average decrease in hit rate (3.7%, Fig. 3.5a) was about three times larger than the average increase in false alarms (1.1%, Fig. 3.5b). The decrease in hit rate is significant for 6 of 7 rats (Fig. 3.5a) whereas the increase in false alarms is significant for only 1 of 7 rats (Fig. 3.5b). The same trends are found when comparing collinear flankers to either pop-out condition (not shown). The net effect is that collinear flankers cause rats to report “no” more often than other flankers. They cause rats to miss the target.

The hit rates and false alarms for collinear and parallel conditions are also shown in an ROC space (Fig. 3.5c). Data from the one outlying rat (r2) is included next to a rat displaying a typical effect (r5) in the expanded view (Fig. 3.5d). Although subjects differ in overall performance and bias, the effect of collinearity is similar for all subjects, indicated by the consistent direction of the arrows. The increase in misses is the dominant effect in the population.

Signal detection theory interprets these raw data in terms of sensitivity (d') and bias (criterion). Applying this framework, sensitivity is consistently lower when flankers are collinear (Fig. 3.5e). This effect is significant for 6 of 7 rats (Agresti-Caffo 95% confidence interval) and for the population as a whole ($p<0.01$ by Tukey-Kramer on 2-way ANOVA). Rats also show a consistent shift in criterion, reflecting the greater bias to say "no" for collinear stimuli (Fig. 3.5f). The criterion shift is small compared to the change in sensitivity, and is not significant for any rat. We cannot confirm the assumptions of Signal Detection Theory hold in our study, but our conclusion (that collinear flankers cause misses) is observable in the raw data (Fig.3.5a-d) independent of these assumptions.

3.3 Discussion

These data show that detection of visual stimuli by rats is sensitive to the configuration of the flanking elements. In particular, flankers collinear to the target impair performance compared with other configurations. Agreement in sign between target and flanker gratings was not required for this effect. This result suggests

specialized processing of oriented image features that can be connected to form a continuous contour. It is noteworthy that this processing must occur in the absence of orientation columns, which are absent in rats (Ohki et al., 2005).

Contrast normalization is a powerful conceptual framework for explaining many surround effects in early visual processing. A pattern-sensitive generalization of contrast normalization could account for a collinear effect (Fig. 3.6). In this model the normalization strength (E) includes additional dependencies on the parameters of spatial configuration (θ_T,θ_F, ω). This extension of the model in Figure 2 allows the normalization strength (E) to be specific to orientation in the target location, and sensitive to the specific arrangement of flanking features. The dominant effect of pattern in our data could be explained by a single factor that selects for the alignment of all three experimental parameters: $\theta_T = \theta_F = \omega$. In this model, collinearity increases the normalization strength leading to greater performance impairment. We offer this as one plausible and parsimonious model which makes direct predictions that can be tested physiologically.

Other perceptual and/or cognitive models could also account for our behavioral data, if they incorporate a "collinearity" term sensitive to the interaction of position and orientation of flankers. Collinear effects could be ascribed to higher visual processing areas. For example mechanisms for binding features, processing gaps, or interpreting occlusions could play a role. In principle collinearity could differentially influence cognitive factors such as arousal, attention, motivation, or task strategy. In

order to account for the collinearity effect, these factors would have to switch on the timescale of seconds because trial types were randomly interleaved.

Collinear flankers cause rats to miss targets more than other flanker configurations (Fig. 3.5a). This is consistent with perceptual masking (rats not seeing the target), but we are reluctant to attribute the bias change to a perceptual process alone or a decision process alone without a measurement of an internal signal that represents the target. For example, if rats suppress perception of a target due to a lateral mask, or perceive a false target due to an illusory contour, they might learn to shift a downstream decision criterion to a new boundary that maximizes their reward rather than reporting their percept. Therefore we do not think the change in bias observed in our data supports any strong conclusions about perception. Nevertheless the data may constrain future models, so we report the raw values we observed for all conditions and rats in supplementary analysis (Supplementary Fig. 3.S8 a,d,g).

3.3.1 Potential Confounds

Our data show that rats' behavior is sensitive to the arrangement of oriented visual features above and beyond the effects of nearby contrast. In any flanker psychophysics study, one should address confounds that might arise from slow variation in performance, familiarity with stimuli, response biases, stimulus artifacts, cognitive confusion, or the influence of attention. Here we address each of these potential confounds.

During the test period, each rat's performance was approximately stable. Of course, performance does fluctuate, probably due to slow variations in motivation. Also, we cannot exclude a small effect of expertise learning. These correlations over time could influence blocked performance such that temporal variation would appear as differences across experimental conditions. For this reason we randomly interleaved all condition types using the method of constant stimuli. This method also balances subjects' exposure to long runs of the same flanker configuration (Fig. 3.3b), which may be difficult to accomplish with adaptive psychometric methods.

Second, it is possible that the exposure to certain orientations, in the recent past, or throughout a subject's lifespan, could influence their perceptual processing of that orientation (Kurki et al., 2006). If we had only used a single target orientation in training or testing, we could not rule out effects of orientation-specific familiarity or adaptation. Therefore subjects were exposed to the same distribution of target orientations and signs in the training steps as in the testing step. Moreover, whenever flankers were present they had the same distribution of properties as the target.

Third, if the rats' 'yes' and 'no' behavioral responses are inherently asymmetric, this would complicate interpretation. We avoided a go-no-go trial structure because it is likely different circuitry is required to initiate versus inhibit a response. Instead we used a 2-alternative forced-choice trial structure where both 'yes' and 'no' required initiation of a symmetric motor output. Reinforcement was also symmetric: correct trials were always rewarded and incorrect ones always initiated a time out, regardless of the target's presence. As a further control, two

subjects (r6 and r7) used the same experimental equipment as their brothers, but were trained with an inverted rule, such that "yes" and "no" were mapped onto the opposite sides. The results for these rats were not different (Fig. 3.4 and 3.5).

Fourth, target orientation or flanker properties might affect target visibility through artifacts of the monitor, rather than processing in the brain. Specifically, is it known that vertical gratings presented on any CRT monitor have lower effective contrast than horizontal stimuli, because RGB guns follow rasterized horizontal scans and lack perfect temporal resolution (Garcia-Perez & Peli, 2001). Had we used gratings that weren't symmetrically tilted about vertical, these artifactual differences in contrast could have been responsible for performance differences across conditions. We also designed the stimuli so that flankers and targets never shared horizontal scan lines to minimize their impact on each other's contrast.

Fifth, it is possible subjects' errors are not due to perceptual difficulty, but rather a failure to understand the intended task. We confirmed in 2/2 rats that their detection performance in all flanker conditions was sensitive to the target's contrast (Supplementary Fig. 3.S4c). Because their performance did not saturate with the contrast we used in our study, at least some incorrect responses were due to perceptual difficulty. We cannot rule out that cognitive difficulty may also have contributed to errors. For example, the decrease in performance when flankers are added (Supplementary Fig. 3.S3) could be explained if rats failed to understand that the target location contained the relevant feature, and also responded to gratings in non-

target locations. However, this confusion would not explain the consistent collinearity impairment observed in all individuals (Fig. 3.4).

Finally, spatial or feature-specific attention may play a role in some flanker tasks (Freeman, Sagi, & Driver, 2004). In our task, to ensure that feature-specific attention would not give the rats a differential advantage between stimulus conditions, target orientation was randomly chosen each trial. We did not employ a positional cue for target location because forward masking could affect target detection. In our task, flankers could improve allocation of spatial attention by reducing uncertainty about the target's location, ultimately improving performance compared to trials without flankers (Petrov, Verghese, & McKee, 2006). This could occur in our task, but if rats did benefit from spatial uncertainty reduction, other effects of the flankers overwhelmed this benefit, yielding net decreases in performance. Alternatively, rats' attentional allocation or visual representation might lack the spatial resolution to isolate flankers from the target location. These factors could explain the flanker-induced impairment (Fig.3.S3), but not the collinear specificity (Fig. 3.4).

In summary, we have controlled for slow variations in behavior, balanced stimulus familiarity, used symmetric responses, avoided CRT artifacts, confirmed that targets are perceptually challenging to detect, and avoided confounds due to orientation-specific attention. We conclude that the rat visual system is sensitive to pattern above and beyond the effects of nearby contrast. We attribute this sensitivity to the rats' visual system, as opposed to other sources of variability in the environment or the rats' cognition.

3.3.2 Comparison to humans

In some perceptual tasks, the presence of collinear flankers improves human performance (Chen & Tyler, 2008; Polat & Sagi, 2007). Why did the collinear flankers impair behavior in rats as opposed to improve it? The term "facilitation" and "suppression" refer to either increases or decreases in performance at a fixed contrast, as in this study, or the ability to match a constant performance in a new condition using a lower or higher target contrast. In this study collinear flankers suppressed detection in rats. On the other hand, collinear flankers facilitate detection for humans performing a two interval forced choice task in the lateral masking paradigm (Polat & Sagi, 1993; Solomon & Morgan, 2000; Williams & Hess, 1998) and dual masking paradigm (Chen & Tyler, 2008). We note that in different task paradigms, this facilitation is not found. The human study most similar to ours also used randomly interleaved trial conditions, fixed contrasts, a single stimulus yes/no paradigm, and oblique vs. collinear flankers, but found no collinearity effect at the target-flanker proximity we used (their Fig. 6, 3λ) (Polat & Sagi, 2007). It remains to be determined whether this difference is attributable to a difference in species, training experience, or stimulus parameters: they used sinusoidal gratings and a single target orientation.

Human studies show that flanker effects change in magnitude or even sign, depending on the contrast regime - which is low contrast for detection tasks and higher for contrast discrimination. The pattern of results in both can be cast in a contrast normalization framework, though the different contrast regimes may involve disparate

cellular and circuit mechanisms. The results above were obtained using a target contrast of 1.0 (where contrast is reported as the fraction of the linearized range of the display spanning 4-42cd/m^2). This contrast is near detection threshold for rats at the spatial frequency we used (Supplementary Fig. 3.S4a). Collinear facilitation is reported to be strongest at lower target contrasts. Thus we also analyzed the data collected from three lower contrasts (25%, 50%, 75%). In no case did any target contrast or any configuration of flankers improve detection compared to the target alone condition (data of Fig. 3.S4c, analysis not shown).

In our data we did not observe sensitivity to the relative sign of target and flankers. The two signs we used are equivalent to having one of two spatial phases. In human psychophysics, both phase-sensitive (Chen & Tyler, 1999; Ejima & Takahashi, 1985; Williams & Hess, 1998; Zenger-Landolt & Koch, 2001; Zenger & Sagi, 1996) and phase-insensitive (Chen & Tyler, 1999; Field et al., 1993; Xing & Heeger, 2001; Zenger & Sagi, 1996) collinearity effects have been described, perhaps reflecting the phase-sensitive (simple cell) and phase-invariant (complex cell) processing channels for orientation in V1. Differences among these studies that appear to be relevant include the distance of flankers from target (gap, no gap, or overlap), and whether the stimuli are presented in the fovea or periphery (Chen & Tyler, 1999). In rats, we only tested one distance with no overlap (3λ), one spatial frequency (0.22 cyc/deg), and two phases (aligned and reversed). We do not know what part of the retina rats used in the task, nor whether rats' central vision (Heffner & Heffner, 1992) is more similar to foveal or peripheral vision in primates. While it may

seem that phase sensitivity is useful for pattern processing, human psychophysics suggest invariance to phase is a hallmark feature of contour integration (Williams & Hess, 1998).

Though flanker effects reported in the human literature depend on stimulus and task details, our results agree with the consistent key finding across human studies: performance under collinear conditions is special.

3.3.3 Relating our findings to natural scene statistics

There is a substantial literature theorizing that early visual processing is optimized for the statistics of natural scenes (Barlow, 2001; Field, 1987; Geisler, 2008). These optimizations can impair performance in tasks that violate natural scene statistics (Howe & Purves, 2005; Schwartz et al., 2009; Weiss et al., 2002). In light of this theory, it is noteworthy that the condition that most affects rats' target detection, collinear flankers, corresponds to the feature conjunction that is statistically most frequent in natural scenes. Combining this perspective with a contrast normalization model leads to a speculation about how a pattern-specific normalization pool could be acquired by learning, without requiring anatomic segregation of orientation channels.

Suppose the visual system computes a prediction of target probability based on some function of the image at other locations, and this prediction is used to adjust the local representation of target. The theory implies that those surrounds that make the strongest predictions about the target in natural images should influence the representation most, and thus impair performance most in our task. This theory is

neutral about whether predicted targets should be suppressed (reducing redundancy) or enhanced (propagating beliefs) at the level of early vision. The direction of this influence cannot be predicted and may be species specific. The framework of normalization developed above implies suppression of predicted features.

It is biologically plausible that the visual system could perform this computation. For example, sensitivity to the separate pair-wise correlations of nearby local oriented features might be learned from the correlated firing of V1 neurons by activity-dependent mechanisms, without requiring orientation columns. Suppose that each local oriented feature's representation is normalized by the activity of all nearby local oriented features, in proportion to their statistical co-activation in natural images. In the statistics of natural images, contrast at one location is correlated with high contrast nearby (Ruderman & Bialek, 1994). The co-occurrence of oriented features depends on the relative orientation and position, and collinear features co-occur most often (Geisler et al., 2001; Itti & Koch, 2000; Sigman et al., 2001). Thus all flankers should normalize, and collinear features should normalize the most (Schwartz & Simoncelli, 2001). In our task, we find that all flankers impair detection and collinear flankers impair the most.

In this framework, one can think of the normalization strength in our model (E) as representing a predicted contrast at the target location ($\hat{C}_T$) estimated on the basis of the contrast in the flanking region (C_F). This suggests that the function of divisive normalization is to reduce the effective contrast for expected features, and amplify

unexpected features, thereby maximizing information transfer for natural scenes (Ruderman & Bialek, 1994).

Future studies could further test this ecological interpretation by correlating behavioral impairment with natural co-occurrence statistics of flanking features at other positions and orientations. In particular, it will be of interest to explore parallel flankers at other positions, and pop-out flankers with greater orientation differences. The more specific hypothesis of normalization makes the direct prediction that flanker features that are correlated with the target in natural images should reduce the firing rate of neurons that respond to the target early in the visual system (such as thalamus or primary visual cortex). Surround processing has not been studied in these neurons in rodents. In cat and primate V1, surround stimuli generally reduce firing (Bonds, 1989; Carandini et al., 1997; Cavanaugh et al., 2002; Heeger, 1992; Polat et al., 1998; Shapley & Victor, 1979), but some cells fire more when flanked by collinear features (Li et al., 2006; Polat et al., 1998; Sillito et al., 1993).

Contrast normalization would reduce redundancy, leading to more efficient codes, but this is not the only goal of vision. Pure contrast normalization may even be at odds with other visual goals. Statistical inference from surround stimuli could contextually de-noise the signal, leading to better parameter estimation through the combination of weak signals. This would also exploit the correlated signals of the natural world (Barlow, 2001) but with opposing effects. We presume that both processes occur and interact in natural vision; different tasks may emphasize one or

the other. We focus on the role of contrast normalization because it requires the least complexity to explain all of our data.

3.3.4 Rats as a model system for vision research

The impact of flankers on behavior has only been studied in humans and other primates, the physiology of the early visual system primarily in non-human primates and cats. Rats offer several advantages as a vision model. Rat husbandry is inexpensive and their behavior, neuroanatomy, and neurophysiology are extensively studied. We have demonstrated that they are easily trained to perform visual tasks that involve distracters and that their vision is sensitive to the spatial arrangement of features. This adds to the growing list of visual tasks demonstrated in rats (Birch & Jacobs, 1979; Cowey & Franzini, 1979; Douglas et al., 2006; Keller et al., 2000; Zoccolan et al., 2009). In this study, individual rats performed around 500 trials every day with stable performance over months and require little human supervision. Many powerful techniques are more feasible in rats than primates or cats, such as genetic, transgenic, viral, histological, optical, intracellular, and pharmacological methods. We conclude that rats provide a valuable complementary model system for studying contextual visual processing.

3.4 Experimental procedures

3.4.1 Animal subjects

Data were collected from seven male Long Evans rats (Harlan Laboratories). All experiments were conducted under the supervision and with the approval of the *Institutional Animal Care and Use Committee* at the University of California San Diego.

The rats included in this study were the 7 median performers from an initial cohort of 14 animals. Four animals were excluded from this study because they either remained at chance on the initial learning task or their performance never exceeded our automatic graduation criterion. They never saw flankers. Three of the remaining ten rats were high performers and were moved to another study before collecting a sufficient amount of data on the testing step. They performed 2, 7 and 16 sessions while other rats performed 60-150 sessions.

3.4.2 Training apparatus

We designed custom hardware and software for automation and parallelization of training. A broad overview of its design and architecture can be found in Supplemental Experimental Procedures 2. Each station consists of a CRT display adjacent to a transparent cage that interfaces easily with slightly modified standard vivarium rat cages (Fig. 3.1a). Rewards were spatiotemporally co-localized with response. Our initial training methods were adapted from previous studies of olfactory

and auditory tasks in rats (Otazu et al., 2009; Uchida & Mainen, 2003). Also see other similar methods (Zoccolan et al., 2009). In our study. behavior was detected via three infrared beam break detectors (Optek OPB980T11) in stainless steel housings for protection against chew damage. Water was delivered to a rounded 16 gauge stainless steel tube positioned just behind each beam by computer-timed solenoid valve opening (80 ms, Neptune Research, 161PO21-11, 161T01, Cooldrive) from a pressurized source (~300 mmHg, Infu-surg 4010, Ethox Corp.) through rigid tubing (CO_2 lines 8044, SurgiVet). Our reward volume was roughly 16 μL, and was typically delivered within 10 ms of a correct lick response. Occasionally rewards were larger, described in training protocol. Auditory feedback was provided with ear-bud headphones mounted on the side walls of the box. Different sounds indicated when a detector beam was broken, and differentiated responses as request, correct, incorrect, and inappropriate (left- or right- licks with no preceding center request lick). To present visual stimuli, we used PsychToolbox (Kleiner, 2007) to control standard OpenGL capable graphics cards (Nvidia GeForce 7600) via Matlab (Mathworks, Natick, MA). We did not track head position or gaze. The head position, head orientation, and eye level and is fairly consistent from trial to trial within subject, as determined from direct observation.

3.4.3 Training Protocol

We designed a series of shaping steps that gradually taught rats to perform the detection task with flankers. An overview of the steps is provided in Table 3.1. Details of the general training procedure such as water restriction, schedule, and environment

can be found in Supplemental Methods (Chapter 3.6.5). The specific shaping sequence used for the subjects in this study was as follows:

Learning to lick.

The goal of the first three steps was to teach the rats to use the detector/reward ports. We presented no visual stimuli, and rats obtain a reward by licking any port at any time. To encourage rats to move among the ports, only one consecutive reward was allowed per port. During step 1, the system occasionally stochastically generated a drop of water without any action from the rat, in order to generate interest in the ports. On step 2, these automatic rewards were turned off, so that rats must actively lick the ports. Rats graduated by sustaining 5 rewards per minute for 2 minutes. Step 3 was identical to step 2 except it required a stricter graduation; it need not be included in future studies. All rats in this study passed through step 3 in 1-2 sessions, with the exception of one rat that got stuck on step 3 for 6 sessions.

Learning two-alternative forced choice (2AFC) task structure.

From step 4 onward, trials had a 2AFC structure. A center lick initiated a trial (but was not rewarded), and the first subsequent lick at either the right ("target present") or left ("target absent") port determined the trial outcome. For 2 of the 7 rats, the present/absent port identities were reversed. Targets appeared on 50% of trials; for these trials, a response at the "present" port was rewarded. Otherwise a response at the "absent" port was rewarded. The tone that accompanied each trial

request provided the subject with confirmation that he had successfully initiated a trial and should proceed to ascertain and report target presence; a different tone and a flickering screen accompanied errors for the duration of the time-out penalty indicating the system was non-responsive. The gratings appeared on the same grey field background that was displayed between trials (after a response, during a reward, and before the next request).

We used two techniques to discourage guessing, side-biases, and other undesirable behavior patterns. First, 50% of incorrect responses were followed by "correction trials." During correction trials, trials with new stimuli but the same correct answer are repeated until the rat responds correctly. Correction trials can induce a rational strategy of switching response after an error (if 50% of errors are followed by correction trials, always switching responses after an error gives 75% performance for those trials – above chance performance without reference to the stimulus). The rats did, in fact, bias their responses in this way, so trials immediately following errors, including correction trials, were omitted from analysis. Strategies that ignore the visual stimulus can only impair performance in any trial that did not follow an error trial.

Second, we gave increasing rewards for consecutive correct answers. The first correct response after an error yielded an 80 ms valve opening (approximately 16 μL). The 2nd to 4th consecutive correct responses earned 100, 150, and 250 ms rewards. Consecutive responses thereafter earned 250 ms rewards; the first incorrect response reset this schedule to the beginning value of 80 ms.

For step 4, the first step with visual stimuli, the target was a large full-contrast square-wave grating masked by a circular Gaussian (21.6° standard deviation from the reference viewing location, with spatial frequency 0.11 cyc/deg, Fig. 3.1d, same stimulus as step 5). We used a square-gratings instead of a sine-gratings because we speculate that contour integration mechanism might engage most strongly for sharp edges. The Gaussian was truncated at 4 standard deviations (about the limit of 8-bit discretization). The grating bars were slanted clockwise or counterclockwise from vertical with equal probability (±15° for r1-r5, ±22.5° for r6 and r7). Results from r1-r5 and r6-r7 were combined because the trends and effect magnitudes were not significantly different ($p>0.05$ t-test, $p>0.05$ Kruskalwallis). For this and subsequent steps, monitor resolution was set to 1024x768 pixels at 100 Hz.

Because this step was designed only to establish the 2AFC trial structure (center request followed by side response), the graduation criterion required only high sustained trial rates, *not* above-chance performance.

Visual detection.

Step 5 introduced a penalty timeout period for incorrect responses (1-6 seconds, hand-tuned for each rat), during which the "incorrect response" sound played and rats could not initiate a trial. Graduation from steps 5-8 required 85% correct performance on the previous 200 trials or 80% on the previous 500 trials. Graduation from step 5 represents the first evidence that rats can perform visual detection in our apparatus (Fig. 3.1c).

Step 6 introduces a gamma correction table (previous steps leave the CRT's native gamma uncorrected) that linearizes the monitor's luminance output, so that the 2^8 grey levels correspond to equally spaced increments in cd/m^2 (See Supplemental Methods 3). After linearization, stimuli used a smaller luminance range and had a higher mean luminance than the earlier training steps. This reduced the effective contrast of the stimulus. All subsequent steps used linearized stimuli.

On step 7, grating spatial frequency increased to 0.22 cyc/deg (equivalent to λ=4.5° where λ is degrees per cycle, a standard unit for indicating flanker distance(Polat & Sagi, 1993). Pilot studies (Supplementary Fig. 3.S4c,d) identified this value as in the range where rats' detection performance was not saturated, but strongly sensitive to target contrast (ranging from 60-75% as target contrast ranges from 50%-100% of the linearized range). This was necessary in order to ensure that we could observe either performance improvements or impairments caused by flankers. This value is consistent with previous reports of rat acuity (Birch & Jacobs, 1979; Heffner & Heffner, 1992; Jacobs, Fenwick, & Williams, 2001; Keller et al., 2000; Prusky, West, & Douglas, 2000), which starts to roll off at 0.1 cyc/deg and shows no sensitivity above about 1.0 cyc/deg for high contrast displays. Our Gaussian mask reduces contrast throughout the grating except at the centermost pixel.

Step 8 reduces the grating size by about 70%, so that the standard deviation of the Gaussian mask was 6.8°. Subjectively, this left about 3 visible periods in the grating. Flankers were not yet included, but at this grating size and spatial frequency,

the monitor had room to display complete flankers (size and spatial frequency identical to the target) at a distance of up to 5λ.

Flanking targets were introduced at a distance of 3λ during step 9. Their contrast was slowly increased from 10% to 40% with training; each 10% was added after a performance criterion of 80% was reached. Performance remained above chance for higher flanker contrasts, but these were not chosen for testing because it is difficult to resolve differences in performance between conditions less than 60% correct. The increasing impairment of performance with flanker contrast could reflect either suppression of the target's apparent contrast below a detectable level, or confusion of high contrast flankers with rewarded targets. For this flanker distance and Gaussian mask size, flankers were non-overlapping, and subjectively appeared separated from the target (Supplementary Fig. 3.S1). For each trial, the target orientation (θ_T), flanker orientation (θ_F), and angular position (ω) were randomly and independently chosen to be either ±15° (rats r1-r5) or -22.5° (rats r6 and r7).

Testing Stimuli. Test stimuli were the same as those in step 9 with 40% contrast flankers (Fig. 3.3). For the purposes of analysis, we grouped all the flanker conditions into four categories depending on the relationship of the flankers to the target (Fig. 3.3b). During testing, error feedback and reward contingencies were not changed. Correction trials continued to occur after 50% of errors.

3.4.4 Performance stability, data filtering, and performance measures

Subjects performed 25,000-90,000 trials during the final testing phase, over the course of 67-136 training days. To assess the stability of performance over time, we calculated performance for each rat and condition in consecutive non-overlapping 500 trial windows. Considering the entire testing period, each subject's long-term performance trended slightly up or down, but this drift always amounted to <7% total change. The most unstable performance was exhibited by r5, whose behavior fell for unknown reasons from ~65% to ~55% for ~6000 trials (~12% of the data in his testing phase) and then recovered. Despite unstable performance over time, the influence of collinearity for this rat was typical of the population. The performance in 500-trial windows ranged between 52-72% for every rat and flanker condition, averaging 63% overall.

We excluded from analysis all trials following an error trial, because rats showed evidence of an alternation strategy after errors, perhaps due to correction trials (see above). Including non-correction trials that immediately followed errors reduces the effect of collinear flankers, but not below significance; this did not alter any trends or influence our interpretation. Considering only trials in the central 80% of reaction times can remove many aberrant trials where rats were either not on-task or very rapidly performing trials while apparently ignoring the visual stimulus. This improves performance and makes the collinearity effect appear stronger, but we do not filter the data in this way for the purposes of this publication.

We report performance in terms of percent correct (Fig. 3.4), but we confirmed that the metric d′ yields the same conclusions (Fig. 3.5 e, Supplementary Fig. 3.S8), and that the effects on performance are not an artifact of bias (Chapter 3.6.8).

3.4.5 Statistical tests

The performance of individual rats in each flanker condition is assumed to be the parameter of a stationary binomial distribution, so finding flanker-caused performance changes amounts to detecting differences in binomial parameters. When computing confidence intervals for differences of binomial parameters, we use the Agresti-Caffo method, in which one modifies a Wald interval by adding two successes and two failures to each estimated proportion(Agresti & Caffo, 2000). Like the Wald interval, the Agresti-Caffo confidence interval uses the Gaussian approximation for binomials, which is not valid if p is near 0 or 1, or n is too small. In our data, $0.55<p<0.85$, and $n>5{,}000$ for all conditions. We verified that alternative statistics agreed with the significance conclusions of Agresti-Caffo for representative test cases (specifically a permutation test and a Bayesian MCMC method, not shown). The Agresti-Caffo intervals graphed for each subject may exclude the point of zero difference; this can be used to test a hypothesis at $p>0.05$.

To determine if performance differences between conditions were significant at the population level, we used a 2-wayANOVA (anovan in Matlab, linear model, type 3 sum squared error). We tested for reliable changes across conditions accounting for the expected variability for each subject. For each subject and condition, we obtained

multiple estimates of the performance based on a non-overlapping sample of N trials selected randomly from the testing phase. We used N=200 trials per estimate. For interleaved trials without flankers (F- mix in Supplementary Fig. 3.S3) the number of estimates ranged from 3-13, compared with 78-253 estimates of F+. For comparing flanker conditions (Fig. 3.4 and 3.5) the number of independent estimates per flanker condition and subject ranged from 19-62. The resulting distributions were approximately Gaussian (e.g., when comparing flanker conditions typically 26/28 passed Lilliefors' test, close to the 5% failure rate expected on chance). Nevertheless, because distributions cannot be guaranteed to be Gaussian, we also we also report Friedman's test (friedman in Matlab), which has lower power but does not assume normality.

Whenever we made multiple comparisons, we used Tukey's honestly significant difference criterion (multcompare in Matlab) with a criterion of $p<0.01$. For example we tested all pairwise differences between the four stimulus conditions. Tukey's criterion for significance is more stringent, to adjust for the fact that making more comparisons increases the probability that one of them will cross significance by chance. A graph showing the critical values for hypothesis testing are shown in Supplementary Figure S5 where all six pairwise comparisons are reported. Results were not different with $p<0.05$ criterion. All tests that were significant with the 2-way ANOVA were also significant for Friedman's test except for one comparison (Supplementary Fig. 3.S5e), discussed there.

The 95% confidence intervals for differences in d′ or criterion within individual subjects (Fig. 3.5e,f) were determined using a Monte Carlo Markov Chain. For each subject and stimulus condition we sampled the posterior distributions of d′ constrained by the number of observed hits, misses, correct rejects and false alarms. The model assumes the hit and false alarm counts are independent binomial distributions and uses a uniform prior over hit and false alarm rates. The d′ or criterion posterior was estimated using WinBUGS and software written by Michael Lee(Lee, 2008). The distribution of d′ difference or criterion difference is sampled by taking the difference between independent random draws of two performance conditions. The confidence interval is the range of this data after removing the 2.5% highest and 2.5% lowest samples.

3.5 Acknowledgements

Philip Meier and Erik Flister contributed equally to this work. Each made essential contributions to the conception, design, and implementation of the general method and technology for high-throughput automated training and testing of rats in 2AFC visual tasks, with help and critical discussion from Pamela Reinagel. Philip Meier was responsible for the conception of the flanker experiment, design and implementation of the task, training and data collection, data analysis, interpretation of the results, and writing of the manuscript, with help and critical discussion from Erik Flister and Pamela Reinagel.

We thank Lung-hao Tai, Tony Zador, Carlos Brody, Zach Mainen, and Tim Gentner for generously sharing their behavioral training methods and hardware, on which our methods and apparatus are based. We thank Fan Li and Dan Parks for assistance with software engineering and Ginger Beriones and Nareg Kalajian for assistance in constructing training equipment. We thank Sarah Meder, Danielle Dickson, Alee Bowden, Marv Zech, Liz Murphy, Yuli Wang, Holly Vo, Aria Jafari, and Duc Nguyen for running daily behavioral sessions.

This work was supported by a Hellman Foundation Fellowship, an Innovative Research Grant from the Kavli Institute for Brain and Mind, a Scholar Award from the J. S. McDonnell Foundation, and NIH R01 Grant EY016856. Philip Meier and Erik Flister were both supported by NSF IGERT Grant DGE9987614; Philip Meier was also supported by NSF IGERT Grant DGE-0333451 to GW Cottrell/VR de Sa and by an NSF Graduate Research Fellowship.

3.6 Supplementary

3.6.1 Flanking stimuli impair target detection performance

The final step in our shaping sequence is the addition of flankers to the detection task (Fig. 3.S2b). This made the task more difficult for rats, presumably for both cognitive and perceptual reasons. This study was not designed to explain this effect, which will be a constant across our compared conditions.

To quantify the overall influence of flankers on rats' ability to detect a localized target, we compare detection with and without flankers, pooling over all other stimulus parameters (θ_T, θ_F, ω, S_T,S_F, see Fig. 3.4). During the continuous block of trials without flankers (Fig. 3.S2a), rats performed about 75% of trials correctly. A single rat's performance is shown in Figure S3a, and for all seven rats in Figure 3.S3b. Flankers impaired performance relative to both the previous block without flankers and the 5% of randomly interleaved trials in the testing step that omitted flankers. We note that performance on trials without flankers was lower when trials with flankers were interleaved as opposed to blocked, although the stimuli were identical. This implies that rats' visual processing or decision strategy depends on the distribution of stimuli over recent trials. For this and other reasons, all other results we show (except Supplementary Fig. 3.S4c) compare only conditions randomly interleaved in the same block.

The effect of flankers is summarized as the difference in performance on interleaved trials with and without flankers (Fig. 3.S3c). If detection were independent of stimuli outside the target location, the difference would be zero. These data, however, show the difference to be less than zero, indicating worse performance with flankers.

We note that the addition of flankers also increases the probability that rats respond yes (Fig. 3.S8a). This reflects a large increase in false alarms, and a small increase in the hit rate. We have not attempted to study or interpret this effect. These data are consistent with a modest reduction in "yes" responses due to contrast

normalization, together with a large increase in “yes” responses due to task confusion or strategy, but our data by no means establish this interpretation.

We suspect that flankers confuse rats and also exert contrast normalization. Additional experiments would be required to isolate these components**.** Therefore we draw no conclusions about the underlying cause of the net effect of adding flankers, but see (Meier P, 2010). For this study, the effect is a constant across the compared conditions, and performance with flankers is sufficiently high that differences in performance between flanker configuration can be detected.

3.6.2. Choosing stimulus parameters

In order to test for influence of the surround, a detection task must be difficult enough to resolve a difference between conditions, but not so difficult that rats fall to chance performance. Difficulty is influenced by the target’s size, contrast and spatial frequency as well as the proximity and contrast of the flanker. We trained rats on easy stimuli, and then varied one or two parameters at a time in separate experiments. These pilot experiments were done on different subsets of the subjects, as well as some subjects that were not in this study.

We tested spatial frequencies spanning 5 octaves (0.05-0.86 cyc/deg) and five contrasts (12.5-100%), for a total of 25 interleaved conditions. This test was performed using a large target with no flankers present (see Fig. 3.1d, step 7 for example stimulus). At 100% contrast, rats performed sufficiently well on 0.22 cyc/deg (green, Fig. 3.S3a), but not well enough for 0.43 or any higher spatial

frequencies. This is consistent with the visual contrast sensitivity previously reported for Long Evans rats; 0.22 cyc/deg is well within their capacity, but a little higher than their peak sensitivity. Gratings were 0.22 cyc/deg for all stimuli presented during the testing step in the main results, as well as all of the other sub-panels in Figure 3.S4.

When flankers were present, they induced a deficit in pe*rf*ormance that increased with their contrast (Fig. 3.S4c). This sub panel is the only data in this paper (other than Fig. 3.S3a,b) that was not randomly interleaved between conditions, but was derived from training steps in which the rats were automatically shaped to higher contrast flankers (see Fig.3.1d, step 9 for example stimulus). We know that the rat's behavior during this learning stage was not constant over time. In principle, learning effects could have increased a rat's performance at higher contrasts which were tested at a later date. Yet flanker impairment overwhelms any such learning improvements. We include this panel because it is informative about the rats learning and about the difficulty of the detection task in the presence of flankers. The trend is the same for other rats in which a range of flanker contrasts were revisited after learning (data not shown).

After rats were well trained on detecting targets in the presence of flankers, we examined the influence of target contrast. Interestingly, rats were sensitive to the contrast of the target across the entire range of our linearized monitor. If stimuli are supra-threshold on a detection task, one would expect performance to plateau. However, if increasing contrast improves performance, this is consistent with the target stimulus being at or near threshold for detection. Since contrast limited

performance, even at the upper end of our display, we chose a high contrast target, so that we could better resolved differences in performance (near 75% correct). We believe that the threshold for detection is shifted to a higher contrast by virtue of the flankers in the surround.

Flankers have a greater influence on target detection when they are closer. We note the possibility that subjects would be best when flankers are at a distance of 3λ, which they had prior training with. However, we observer rats perform worse at 2.5λ and better at 3.5λ. Despite the stronger effect at 2.5λ, our main results use a flanker distance of 3λ because previous work in human psychophysics and cat V1 neurophysiology find effects at this distance. Additionally, the tails of the Gaussian masks for target and flanker begin to overlap and 2.5λ, even though we truncate tails at 4 std. There is no overlap for our flankers and targets at a distance of 3λ.

Taken together these data suggest that the flanker stimuli used in our study are perceptually difficult for rats, because the spatial frequency is high for a rat, the grating patch small, and the flanker contrast high. The sensitivity to changes in contrast of the target or the contrast or distance of the flankers is all consistent with expectations from contrast normalization. Due to confounding cognitive factors we do not purport to measure the strength of contrast normalization from these data.

3.6.3 Statistics for all pairwise comparisons of flanker configurations

In the main paper we show statistics for comparisons between the collinear condition and two others: the condition that differs only in flanker orientation (Fig.

3.4c) and the condition that differs only in the flankers' angular position (Fig.3.4d). We consider all pairwise comparisons here.

For the four stimulus conditions tested, there are a total of six pair-wise comparisons. Rats performed worse on collinear stimuli than on any of the other three stimulus conditions (Fig. 3.S5a,c,e, left column). Each of these comparisons is the result of disrupting collinearity in a different way: by changing the flanker orientation (a, same as Fig. 3.4c), by changing the flankers' angular position (c, same as Fig. 3.4d), or by changing the target orientation (e). Taken alone, it seems possible that the visual system may be sensitive to each of these changes in isolation. However, the other three conditions provide a control (Fig. 3.S5b,d,f; right column); in each the geometry of the stimulus was held constant except for the flanker orientation (b), flanker location (d), or target orientation (f). No individual rats are significant in these three control comparisons (Agresti-Caffo 95% confidence interval), and they do not differ at the population level (Tukey's on ANOVA, $p>0.05$; Tukey's on Friedman's test, $p>0.05$).

To communicate significant differences at the population level, we present the marginal mean performance for all rats (Fig. 3.S6). Each error bar represents the half with of the critical value from Tukey's test of Honestly Significant Difference, which accounts for multiple comparisons. If the error bars from two conditions do not overlap vertically, then the difference between conditions is significant. The figure shows significance for $p<0.05$; we also assessed significance at $p<0.01$. All multiple comparisons in this paper use this same method; here we provide a graphical view to

facilitate understanding. Figure S6a summarizes the results for the ANOVA, and Figure S6b summarizes Freidman's test.

In one instance (collinear vs. pop-out$_2$, Fig. 3.S6), the difference is significant at the population level by 2-way ANOVA but not by Friedman's test after adjusting for multiple comparisons (Both tests on this condition are significant at $p<0.01$ before adjusting for multiple comparisons). Friedman's test is non-parametric and only acts on the rank-ordered average performance. With only seven subjects, it is quite conservative. It is actually surprising that two of the six comparisons are significant by this test even after adjusting for multiple comparisons. We interpret the collinear condition to be harder for rats than the pop-out$_2$ condition, based on the results of the 2-way ANOVA and the facts that this was true for 7 of 7 rats, and significant for 5 of 7 rats.

We find it somewhat surprising that there is no discernable effect of pop-out with respect to the parallel condition (Fig. 3.S5b,f). One might have expected pop-out conditions to be easier. Maybe the presence of only two flankers does not indicate the uniqueness of the center target as much as a field of distracters would (Nothdurft, 1991; Smith, Kelly, & Lee, 2007). Alternatively, the orientation difference between target and surround may be important. We used 30° (open symbols in Fig. 3.S3) and 45° (filled symbols), and did not observe a difference. A pop-out effect might be found at more orthogonal angles (Nothdurft, 1991; Schwartz et al., 2009).

In summary, our only significant differences involved comparisons to the collinear condition. Thus, impairment on the collinear stimulus is the primary effect of pattern-specific processing that we observed.

3.6.4 No influence of phase

In our experiment, for each flanker orientation and position we also varied sign of the grating ($S_F = \pm 1$) corresponding to inverting the luminance of the grating such that dark bars are switched for light bars. Changing the sign is equivalent to a π shift in the spatial phase of the grating. Neither the absolute sign of the target or flanker nor the relative sign between them had an effect on performance. Here we show there is no difference between the phase aligned and phase reversed collinear stimuli using the same data from the seven rats reported in the main paper (Fig. 3.S7a,c). This experiment used one spatial frequency (0.22 cyc/deg), one flanker distance (3λ, 13.6°), and two relative phases (0, π). Our data do not exclude the possibility that the rat visual system is sensitive to phase differences with other choices of these parameters.

In a pilot study (N=2) we tested four different absolute grating phases for target and flanker, using the same class of flanker stimuli as this study, but with some differences: the target contrast was weaker (60% instead of 100%) and rat performed a two alternative force choice between two simultaneously presented stimuli, instead of a single stimulus yes-no task. The relative phase between target and flanker did not influence the detection of the target (Fig. 3.S7b,d).

Since there is no evidence in our data of sensitivity to sign or phase, we have no reason to include it as a relevant parameter in the schematic model presented in the main paper (Fig. 3.6). The model is sufficient to explain all of the results in this paper, without adding a term in the surround for the relative phase.

The influence of oriented contrast in the surround is not phase sensitive in many previous psychophysical (Xing & Heeger, 2001) and physiological experiments (Webb, Dhruv, Solomon, Tailby, & Lennie, 2005). However, sometimes the relative phase of the surround does play a role (Ejima & Takahashi, 1985; Williams & Hess, 1998). For primates performing contour detection tasks, spatial pattern summation is phase insensitive in the fovea, but in the periphery it is not (Chen & Tyler, 1999). While phase sensitivity attests to the specificity of surround processing, it has been argued that invariance to phase is a useful aspect of surround processing, and such invariance continues to be applied in computational models of contour integration (Hansen & Neumann, 2008).

3.6.5 Water restriction, training schedule and environment

Beginning around post-natal day 30, rats were restricted from free access to water, instead working for water reward in the training environment. Water was earned solely by correctly performing tasks ("closed economy" training) as long as animals maintained adequate health and hydration. This was assessed by daily weight and health inspections. Supplemental water (or hydrating snacks, e.g. carrot slices) were provided as needed, for less than 25% of rats during initial shaping, and rarely

after 2 weeks. On most days, including weekends, animals were transferred to the training chamber where they could freely perform trials for around 90 minutes. On days where sessions were skipped, rats were provided with about an hour of free access to water, but we avoided this because it reduced their motivation to perform trials on the following day. Continuous access to rat chow was provided both in their home cages and the training chamber to stimulate desire for hydration. We phase-reversed the light cycle in the room so that rats trained during the day in dark rooms illuminated almost exclusively by the glow of 7 computer displays. When overhead lighting was required for working in the room during daytime, we typically filtered out the high wavelengths visible to rats (Jacobs et al., 2001) and worked in red light (Encapsulite Intl., red filter, 625nm cutoff, 48SOR20T12).

3.6.6 Training system

The system consists of an array of stations, each controlled by a computer running Windows XP-Pro and Matlab. An additional computer is used as a single point of control for all the stations (via custom TCP/IP code in Java/Matlab). This allows sessions to be started/stopped en masse, management of subject information and trial records in coordination with a database (Oracle 10g Express Edition), and centralized management of individualized training sequences and parameters for each subject.

The system design facilitates either live-in training or easy swap-in of groups of animals. The animals in this study were members of one of 6 groups that were

swapped in daily. Since rats can earn hydration adequate for health in roughly one hour of trials, we found it efficient to limit each animal's daily access to trials to sessions of about this duration. This focuses their motivation for performing trials correctly throughout the session and multiplies the number of subjects that can be trained using a given amount of available hardware.

Each training box is 35 cm wide x 18 cm deep x 30 cm high, with detector/reward ports positioned along the front wall spaced 9 cm apart 6 cm above the floor grating. The detector housings protrude 2.5 cm into the box, and the CRT is positioned 5 cm behind their wall. At the center port, a rat's eye is roughly 10 cm from the monitor, 10 cm below its center. At this position, the monitor display subtends 104° of visual angle, and the target grating at the CRT center is roughly 14 cm away and spans 6.8° per std of the Gaussian mask. This is closer than the distance established for maximum behavioral visual acuity in rats (20-30 cm; Wiesenfeld, Branchek 1976).

Accurate reward delivery requires a system that does not store pressure in unexpected ways. At such small volumes and flow times, tiny differences in port geometry likely affect the amount of reward and its accessibility due to being wicked away or the force with which it is ejected, and reward volumes/utility may not be linear with respect to open valve duration. This can cause strong side biases. We monitored behavioral trends for side bias and found we could control it using timeouts and correction trials alone; we only sporadically and imprecisely verified that left and right reward drops looked roughly similar in size. We have since added a syringe-

pump based reward system that is more accurate and requires less maintenance (New Era Pump Systems, Inc. NE-500).

We used the parallel port for electronic interface of the valves and sensors with the computer, using PortTalk (Craig Peacock, beyondlogic.org/porttalk/) via an open source Matlab wrapper (psychtoolbox.org/wikka.php?wakka=FaqTTLTrigger).

The sound for requests is a tone with energy at each octave spanning the frequency range of the speakers (some range within 20-20,000 Hz); for correct responses, a tone with the same harmonic structure is played a perfect fourth above the request tone, creating a harmonic resolution. The sound for incorrect responses is a chord made from two tones of the same harmonic structure separated by a tritone, the maximally dissonant interval. The sound for inappropriate responses is broadband noise. Sometimes a drop of water lodged in a port is sufficient to break the infrared detector beam, in which case, the corresponding sound plays continuously until the water is cleared by the rat (after some experience in the box, rats are observed to clear ports in this way even during a period of no trial activity, indicating a possible preference for silence to the continuous sounds). No provisions were made for sound isolation. Sounds from adjacent boxes are quite audible in any given box, but quite attenuated relative to that box's local sounds. The CRTs for adjacent boxes are not visible when operating the ports because stations are separated by opaque dividers.

Our software architecture facilitates the design of arbitrary task structures and manages each subjects' progress through their own training protocol. A *protocol* specifies a sequence of *training steps*, each with a *trial manager* (defining task

structure – e.g. two-alternative-forced-choice), *stimulus manager* (defining audiovisual stimuli and their parameters), *reinforcement manager* (specifying reward/timeout reinforcement rules), and *graduation criteria* (regulating progression through steps).

3.6.7 Stimulus display

We used CRT displays (NEC FE992, 19") because of their fast refresh and phosphor decay times; CRTs are generally better suited for visual psychophysics than LCD flat panels because their timing artifacts and brightness artifacts are more consistent and better understood. The CRT linearization table was created by fitting a power law with gain and offset ($y=b*x^{g}+m$) to photodiode measurements (Thorlabs, PDA55), and then computing the inverse function ($x=[(y-m)/b]^{1/g}$) for each RGB channel independently. Each value was measured as the average height of the smoothed peak of the phosphor decay curve recorded with that value presented in a rectangle occupying the central 60% of screen for 9 frames at 100 Hz and 1024x768 resolution. Gun values were measured in increasing order rather than randomized. The resulting tables were then verified to have linearized grayscale output to within 0.5% using the same method.

Since the background grey level was set equal to the mean luminance of the grating, linearizing has the consequence that a blurry optics system with spatial resolution lower than the spatial frequency of the gratings will not perceive a luminance difference at any point in the Gaussian patch of the grating, but only a

smooth grey at the same luminance as the background. We confirmed that this was the case for sufficiently high spatial frequencies for human observers. The linearization range was chosen such that the gratings were effectively lower contrast than in previous steps and the mean grating luminance/background grey was brighter than in previous steps. The minimum, mean, and maxiumum luminance were set to 4, 42 and 80 cd/m^2, respectively (Colorvision, spyder2express).

3.6.8 Considering signal detection theory

In the main paper we summarize performance as percent correct, but using d′ does not alter the character of the results (Fig. 3.S8). Signal detection theory measures attempt to unconfound bias and sensitivity, both of which affect percent correct. The metric of percent correct can be misleading when comparing results of subjects with very different overall biases. In this study we do not directly compare any rat's absolute performance to another rat's. Rather we test within subject for significant differences between stimulus conditions that are randomly interleaved. When the hit or false alarm rate is near 0 or 1, it is possible that the metrics of d′ and percent correct could give different trends or statistical significance when comparing conditions. In our data, the hit and false alarm rates are bounded by [0.15 0.85] for all rats and conditions. Thus our data do not occupy the extreme region where one would expect significant differences between the metrics.

One advantage of using percent correct, when the data permit, is that it avoids the need to make assumptions which we cannot verify. Specifically, signal detection

theory separates bias from sensitivity by assuming normality and homoscedacity of the stimuli distributions being distinguished – d′ summarizes their discriminabilty as the separation of their means in units of their common standard deviation. Another reason we prefer percent correct is that, unlike d′, it monotonically increases with reward; thus, tracking what rats naturally prioritize. We report the assumption-free data on hit rate and false alarm rate from which d′ and other metrics of signal detection theory are easily calculated.

Nevertheless comparisons based on d′ are provided for comparison (Fig. 3.S8). Across every major comparison in this study, the two metrics agree in individual significance, population trends, and relative magnitude. Specifically, we find that both measures indicate all flanker types impoverish performance (Fig. 3.S6b,c), collinear flankers impair performance more than parallel flankers (Fig. 3.S6 e,f), and collinear flankers maintain this impairment regardless of the relative contrast sign between target and flanker (Fig. 3.S6 h,i).

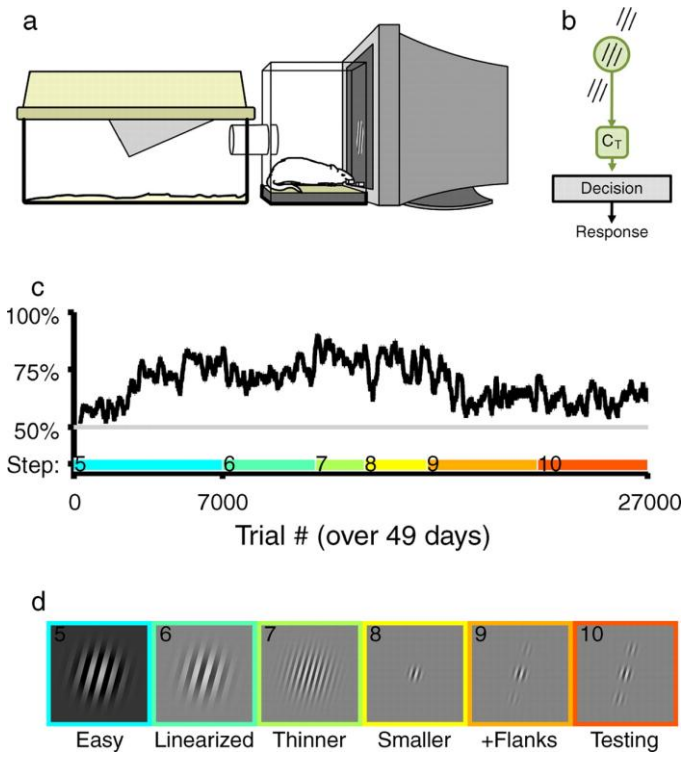

Figure 3.1. Training rats to detect an oriented target in the presence of flanking distractors. a) A diagram of the training environment. Rats could initiate trials by licking a sensor centered in front of the display monitor. This immediately rendered a stimulus. Rats were rewarded with water for correctly reporting the presence or absence of a target by licking one of two response spigots located on the left or right side of the chamber. b) A simple schematic indicating the final task the rat is being shaped to perform. The target location is denoted by a green circle. The target is present on 50% of the trials; the target contrast, C_T, is either 0 or 1. The target's presence, irrespective of the flanker configuration, informs the rat which decision will result in a water reward. c) To achieve the final task which has small, low contrast targets and distracting flankers, rats are shaped through a sequence of training steps that increase in difficulty. Performance for a single rat is plotted as a 200 trial running average, starting from the first easy visual trials (step 5) to the testing phase (step 10). The first four steps involved associating the response ports with rewards, and did not involve any visual stimuli (see Table 3.1 for details). When the rat's performance exceeded a preset criterion (>85% correct), he was automatically graduated to a new training step to ensure rapid progression and avoid over learning. d) Sample stimuli from each step are color-coded to match the performance plot, and named to emphasize the change from the previous training step. The addition of dim flankers (step 9) is displayed for a linearized contrast of 20% in b) but was increased from 10% up to 30% in step 9. All testing was performed with 40% linearized flanker contrast. For simplicity only one of the two orientations is shown, but all rats had equal training exposure to both. For more testing stimuli, see Figure Fig. 3.3. For a photograph of a rat performing the final task, see Figure 3.S1.

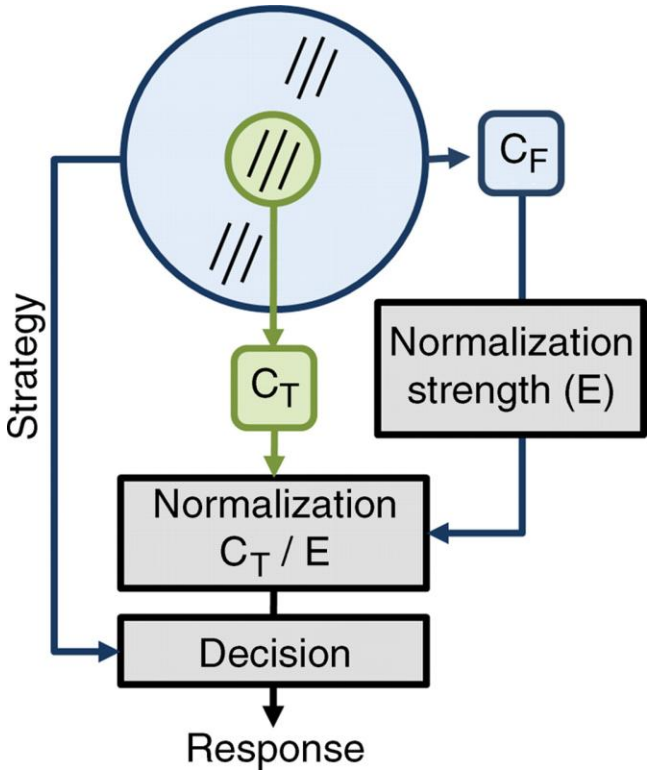

Figure 3.2. A schematic model indicating how the presence of flankers might influence rat's decisions. On this model, the contrast of the flankers (C_F) contributes to the normalization strength (E) such that the effective contrast of the target is reduced. The presence of spatial contrast normalization probably contributes to the deficits in performance associated with the presence of flankers. The presence of flankers also influences performance for cognitive reasons ("strategy"). This model is blind to orientation and position of flankers.

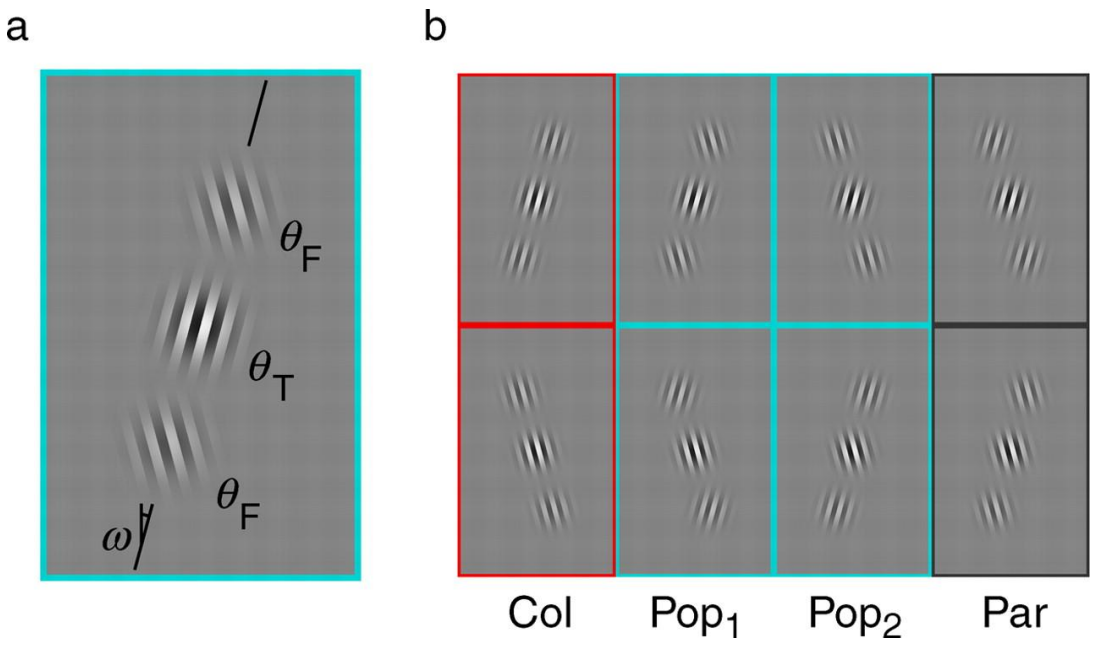

Figure 3.3. Grouping flanker stimuli into conditions. a) An example flanker stimulus, with labeled spatial parameters (θ_T, θ_F, ω). Flanker stimuli where generated by independently varying target orientation (θ_T), flanker orientation (θ_F), and the angular position of the flankers (ω). The two flankers always had the same orientation. b) The different stimuli were grouped in four conditions that preserved geometric relationships. The top and bottom row are mirror images of one another. In the collinear condition, the target and flanker orientation align with the angular position ($\theta_T = \theta_F = \omega$). Collinearity was disrupted by changing one of the parameters in each of the remaining flanker conditions: "pop-out$_1$" ($\theta_T = -\theta_F = \omega$), "pop-out$_2$" ($-\theta_T = \theta_F = \omega$) and "parallel" ($\theta_T = \theta_F = -\omega$). Each condition was equally likely during the testing. Every stimulus has a matching case in which the target was absent (not shown here, but see Supplementary Fig. 3.S2b). During all training and testing, the luminance sign of the target and flankers were randomized; the case of $S_T=S_F=+1$ is shown.

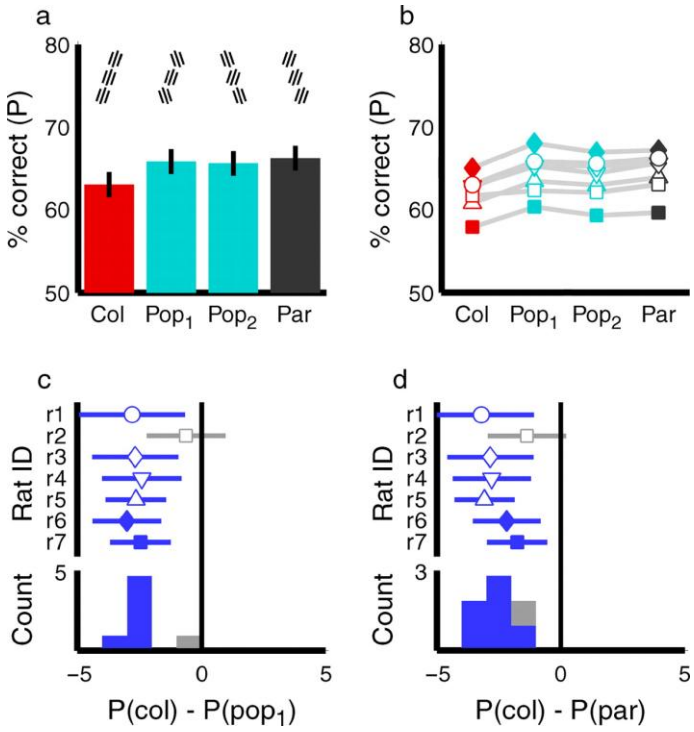

Figure 3.4. Collinear flankers impair detection more than other arrangements. a) A single rat's performance (r1) on the four conditions: collinear, and three patterns that disrupted collinearity. In **pop-out$_1$** collinearity was disrupted by changing the flanker orientation (θ_F) to be different from the target (θ_T). **Pop-out$_2$** maintained the same difference in flanker and target orientations ($|\theta_F-\theta_T|$) as pop-out$_1$, but the angular position of the flanker was different (ω). In **parallel** the collinearity was disrupted only by changing the angular position of the flanker. Error bars indicate 95% binomial confidence intervals. b) Performance of all seven rats on all four conditions. c) The difference in percent correct between collinear and pop-out$_1$ for all rats (r1-r7). Collinear is more difficult. Error bars indicate 95% confidence interval using a modified Wald interval described in methods (Agresti & Caffo, 2000).One rat's data are rendered gray to indicate that the difference in performance is not significant. The other six rats are rendered blue because they are each significant (Agresti-Caffo 95% confidence interval). Filled symbols indicate subjects in which the mapping between yes-no and left-right was inverted (See training protocol in methods). Subjects with inverted training rules had no different effects. d) The difference in percent correct between collinear and parallel. Again collinear is harder. For both comparisons (panels c and d), the difference between conditions is significant for six of seven rats individually (Agresti-Caffo 95% confidence interval), and for the population as a whole (2-way ANOVA with Tukey's, $p<0.01$). For all possible pairwise comparisons, see Figure S5.

Figure 3.5. Collinear flankers cause rats to miss the target. a) The change in hit rate between collinear and parallel configurations, h(col)-h(par). Symbols left of the zero line indicate that the hit rate is lower when flankers are collinear. The effect is shown for each subject: the symbol location indicates mean difference in hit rate, horizontal lines indicate 95% confidence interval (Agresti & Caffo, 2000). For all but one rat, the deficit with flankers is statistically significant. b) The change in false alarm rate, f(col)-f(par), using the same conventions as panel a. Rats display more false alarms on the collinear stimuli. However, the change in false alarms was smaller than the change in hit rate show in panel a. c) A geometric representation of the raw data in a and b is shown by plotting the Receiver Operator Characteristics (ROC). The individual ellipses show the boundary of the 95% confidence intervals of a binomial distribution for hit rate and false alarm rate for one subject and condition. For each subject, an arrow in ROC space summarizes the difference between the collinear condition (red) and randomly interleaved trials of the non-collinear reference condition (parallel, gray). The arrow points to the change in responses induced by the collinear feature with respect to the reference condition. Most arrows consistently point down and to the right indicating a decrease in hits, and a small increase in false alarms. d) Detail of subpanel c to provide better resolution for a typical rat (r5) and an outlier rat (r2). Interestingly, histology (not shown) revealed that the rat displaying the outlier effect (r2) had a naturally occurring tumor which compressed and displaced about a quarter of the ventral thalamus. We do not know if this played a role in the animal's behavioral differences. e) The difference in detection sensitivity between collinear and parallel conditions. Errorbars are the 95% confidence intervals of samples drawn from a Monte Carlo Markov Chain. These d' measurements are consistent with the report in Fig. 3.4d: collinear stimuli have targets that are harder to detect. f) The change in criterion between collinear and parallel conditions. Negative values indicate that rats are more likely to report the absence of the target in the collinear condition. The change in bias is consistent for all rats, but not significant for any individual rat (MCMC 95% confidence interval). The size of the bias difference is small compared to the change in d' in panel e.

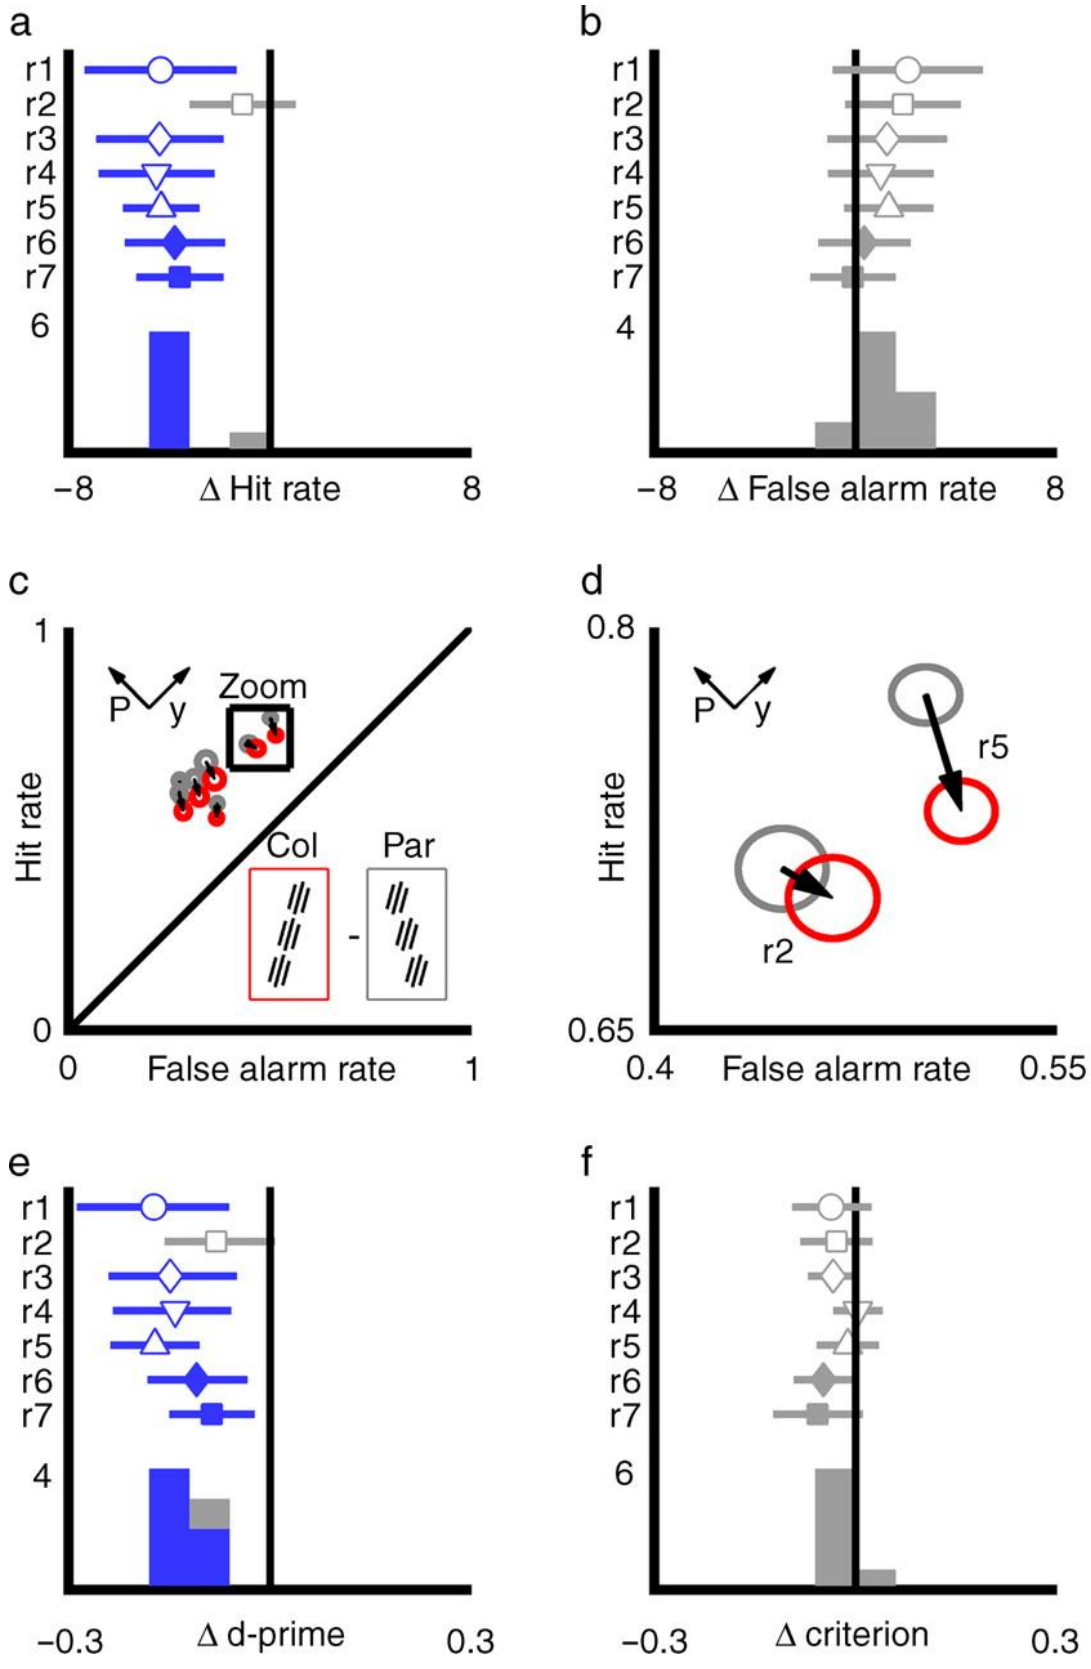
a
r1
r2
r3
r4
r5
r6
r7
6
−8
Δ Hit rate
8
b
r1
r2
r3
r4
r5
r6
r7
4
−8
Δ False alarm rate
8
c
1
Hit rate
0
P
y
Zoom
Col
Par
-
0
False alarm rate
1
d
0.8
Hit rate
0.65
P
y
r5
r2
0.4
False alarm rate
0.55
e
r1
r2
r3
r4
r5
r6
r7
4
−0.3
Δ d-prime
0.3
f
r1
r2
r3
r4
r5
r6
r7
6
−0.3
Δ criterion
0.3

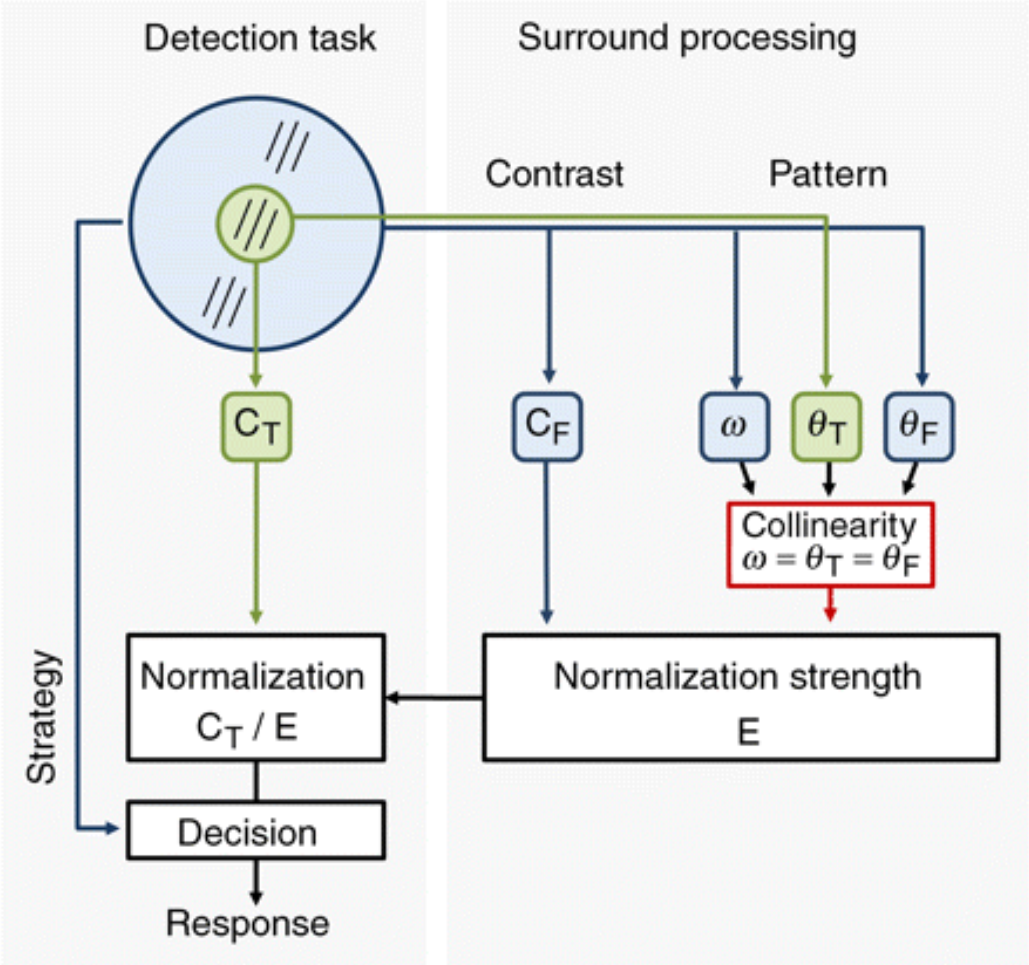

Figure 3.6. Schematic model of pattern sensitive contrast normalization. The detection task is summarized in the left hand region, unchanged from Figure 2. The contrast in the target region (C_T) is represented by a neural signal which is normalized by surround processing in early vision before the decision is made. Then the rat responds left or right to indicate the target's presence or absence. In this model, cognitive effects of the surround ("strategy") depend only on the presence of flankers and insensitive to their configuration. Surround processing contains two aspects: sensitivity to contrast, and sensitivity to pattern. In this study, surround contrast is determined by the experimental parameter for flanker contrast (C_F, same as Fig. 3.2), and was held constant during testing. The pattern-sensitive component must include at least three terms to account for our data: the angular position of the flanker (ω) the orientation of the target (θ_T), and the orientation of the flanker (θ_F). While there are many ways that these terms could interact, we only require a dependence of collinearity to explain the rat's behavior. The processing of contrast and pattern in the surround is combined into a single normalization term (E). This determines the gain of the neural signal that is used for detection. An argument that the normalization term E could be interpreted as an expected contrast ($\hat{C}_T$) is considered in the discussion.

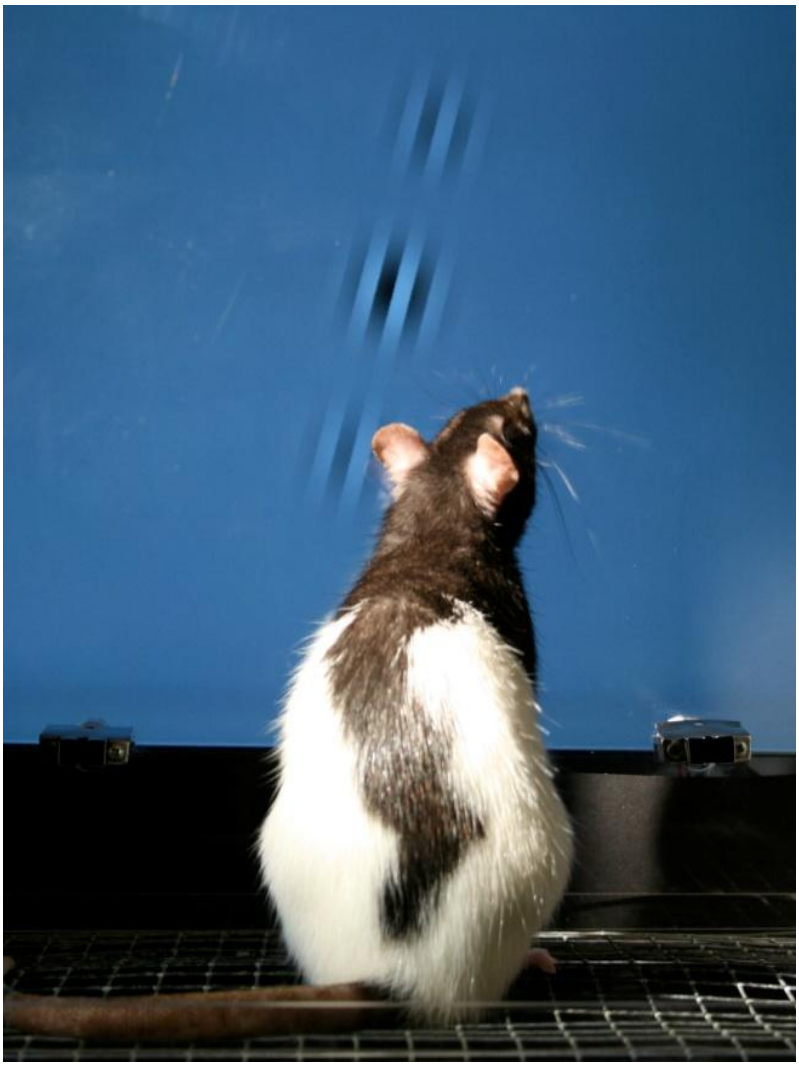

Figure 3.S1. Photograph of a rat viewing a target with collinear flankers. The rat is trained to receive water from the port on the right when the central target is present, and the port on the left if it's absent.

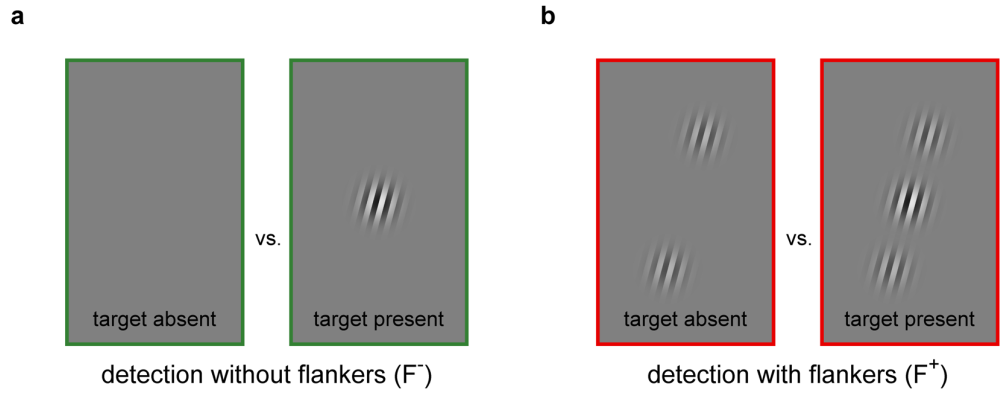

Figure 3.S2. Stimuli used for target detection, with or without flankers. a) During training (e.g. step 8), the rats perform a detection task when no flankers are present (F^-). When the rat initiates a trial, a target appears 50% of the time. When it is present, it is always located in the center of the screen. A rat would receive a water reward by correctly choosing to lick a sensor on the left side if the target were absent, or on the right side if the target were present. b) During the testing phase flankers are present (F^+) on 95% of the trials. The flanker's presence, location and orientation carry no information about the correct response. Only one flanker configuration is shown here, though all types were presented during testing (for more types see Fig. 3.4). In the first analysis all flanker conditions are grouped together.

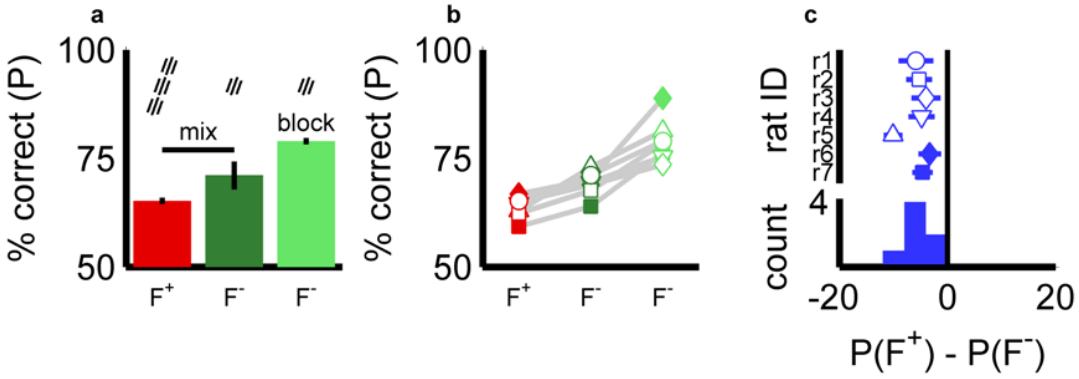

Figure 3.S3. Rats are worse at detecting the target when the flankers are present. a) The percent correct performance (P) of a single rat on three conditions with identical target properties. Error bars indicate 95% binomial confidence intervals. Performance from a block of randomly interleaved trials with flankers (F^+) and without flankers (F^-) is shown in red and dark green respectively. Light green indicates performance from a continuous block of trials without flankers (F^-) before flanking stimuli were introduced. b) Performance of seven rats in the same task, symbols colored as in (b), lines connect symbols representing a single subject (r1-r7). c) The effect size, measured as the difference in percent correct for interleaved trials with and without flankers: $P(F^+) - P(F^-)$. Symbols left of the zero line indicate that percent correct is lower when flankers are present. Effect is shown for each subject (symbol location indicates mean difference, horizontal lines indicate Agresti-Caffo 95% confidence interval, see Methods). For all individual subjects, the deficit with flankers is statistically significant. The effect is also significant if the difference between conditions is assessed at the population level ($p<0.01$, Tukey-Kramer on 2-way ANOVA; $p<0.01$, Tukey-Kramer on Friedman's test, see Methods).

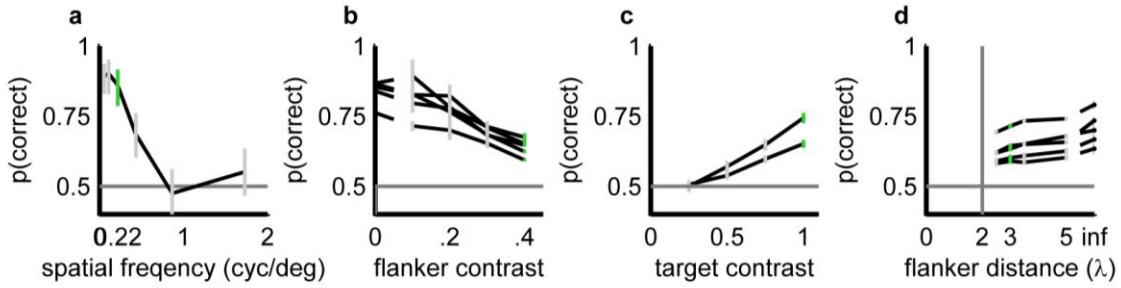

Figure 3.S4. Detection performance depends on target spatial frequency, target contrast, flanker contrast and flanker proximity. a) Spatial frequency dependence of a single rat's ability to detect a single large grating patch (Gaussian window of width 21.6 deg/std). Data were collected while randomly interleaving 5 contrasts and 5 spatial frequencies, but only data from 100% contrast (the target contrast used in testing) are shown. Performance was near chance for 0.86 or 1.73 cycles per degree (cyc/deg) at all contrasts, and was near maximum at 0.22cyc/deg at 100% contrast (green), which was used for all other tests in this paper for all rats. Error bars are the 95% binomial confidence interval. b) Detection performance for five rats while learning to detect targets in the presence of flankers of increasing contrast,using 0.22 cyc/deg gratings. Each rat is indicated by a separate curve. As rats demonstrated aptitude on lower flanker contrasts, they automatically graduated to the next higher contrast. All data points involve many hundreds of trials, with the exception of a single data point from one rat at 10% contrast involving only 75 trials. Performance for that rat may appear to be greater than performance with zero contrast flankers, but this difference is not significant (Agresti-Caffo 95% confidence interval) All other tests in this paper used 40% flanker contrast (green). c) Detection performance from two rats with five randomly interleaved target contrasts, while flanker contrast remained at 40%. All other tests used a target contrast of 100% (green). d) Detection performance from five rats on stimuli with four randomly interleaved flanker distances. The line at $\lambda=2$ represents the distance that full contrast target and flankers begin to perceptually overlap for humans. At $\lambda=5$ both flankers are still on the screen, but are close to the edges. The data at flanker distance of infinity indicates trials where there was no flanker on the screen. All other data in this paper uses $\lambda=3$ (green).

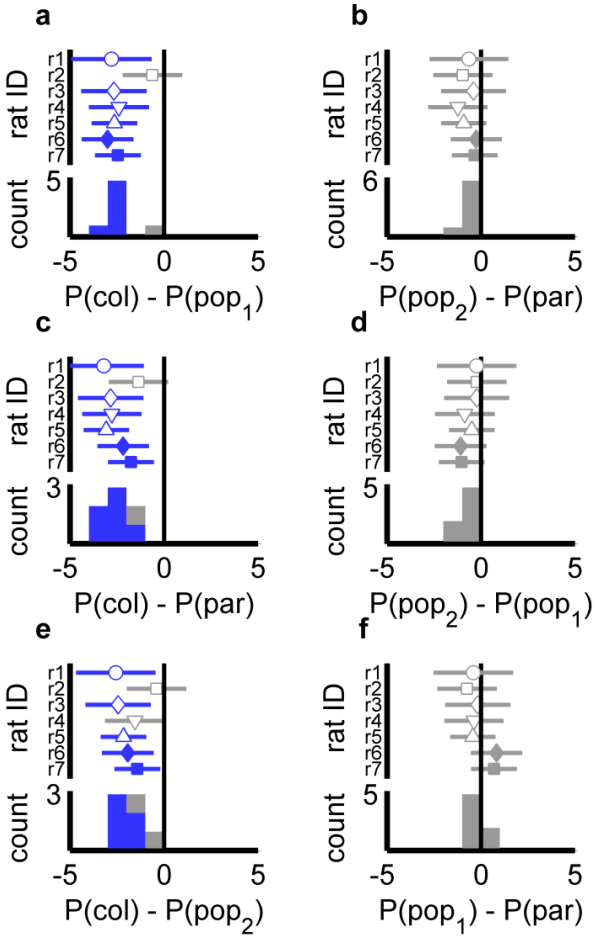

Figure 3.S5. Additional pair-wise comparisons between the flanker conditions. We present all six comparisons between the four stimulus conditions measured by the difference in percent correct. Panels a and c are the same as Figure 4c,d; they are reproduced here to facilitate comparisons. The vertical midline indicates that the performance on the two conditions is not different. Each horizontal line represents the 95% confidence interval for a single subject's data on a given comparison (Agresti-Caffo). The top row of sub panels isolates comparisons between pairs of conditions in which the location of the flanker differs, but other parameters were held constant: a) collinear – parallel and b) pop-out$_2$ – pop-out$_1$. The middle row isolates comparisons in which only flanker orientation differs: c) collinear – pop-out$_1$ and d) pop-out$_2$-parallel. The bottom row captures comparisons in which only target orientation differs: e) collinear –pop-out$_2$ and f) parallel – pop-out$_1$. The three comparisons that contain subjects with significant effects (Agresti-Caffo 95% confidence interval) are best explained by the reduction of performance for collinear stimuli which affects all the comparisons in the left column. The comparisons in the right column are not significant for the rats individually, nor at the population level (see Fig. 3.S6).

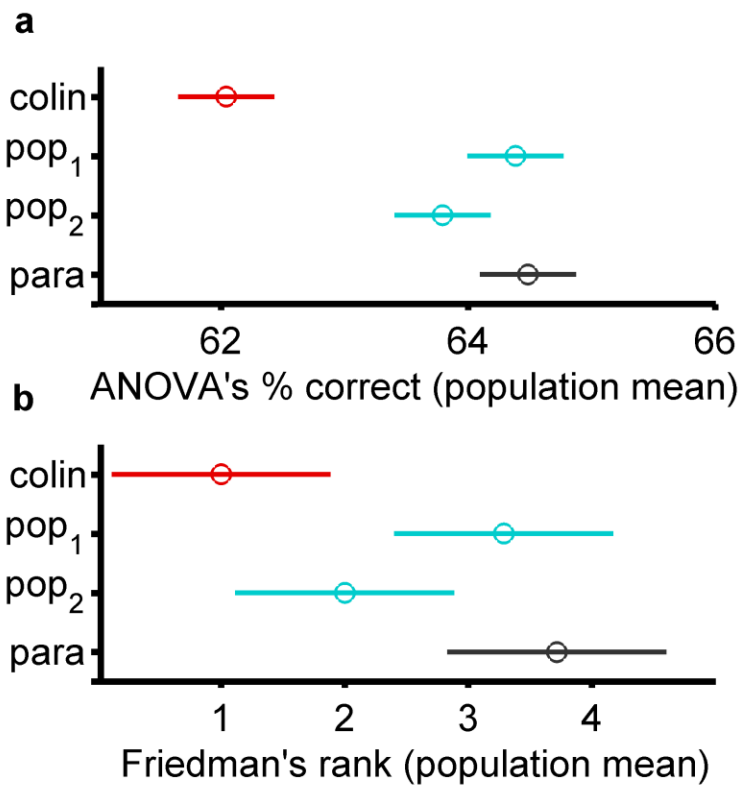

Figure 3.S6. Multiple comparison adjustments. a) A visual summary the 2-way ANOVA after having adjusted for multiple comparisons. The x-axis is the marginal means of the percent correct performance from a population of seven rats. The error bars represent the half width of the critical value for multiple comparisons using Tukey's Honestly Significant Difference at $p<0.05$. If the tails of two error bars do not overlap each other on the vertical axis, then the difference those two conditions are significantly different at $p<0.05$. b) A visual summary the Friedman's test after having adjusted for multiple comparisons using the same method. Friedman's is a non-parametric test that does not assume Gaussianity; it tests for a shift in the location of a probability distribution by analyzing the relative rank of performance across rats and conditions. Friedman's has less power to reveal an effect, especially with the small N=7. The collinear condition is reliably the 1st rank (worst performance) for all rats. The pop-out$_1$ and parallel conditions have higher performance. Interestingly, pop-out$_2$ is reliably the 2nd rank, but this test could not resolve a significant difference from the collinear condition. We do not interpret this strongly because the absence of observing an effect is inconclusive, and the population marginal means shown in a), suggests more strongly that pop-out$_2$ groups with pop-out$_1$ and parallel more than with collinear. Additionally, the significance of the each of the rats individually (Fig. 3.S5) supports the conclusions of the 2-way ANOVA. All tests that are significant in this figure (both a and b) are also significant at $p<0.01$ (not shown).

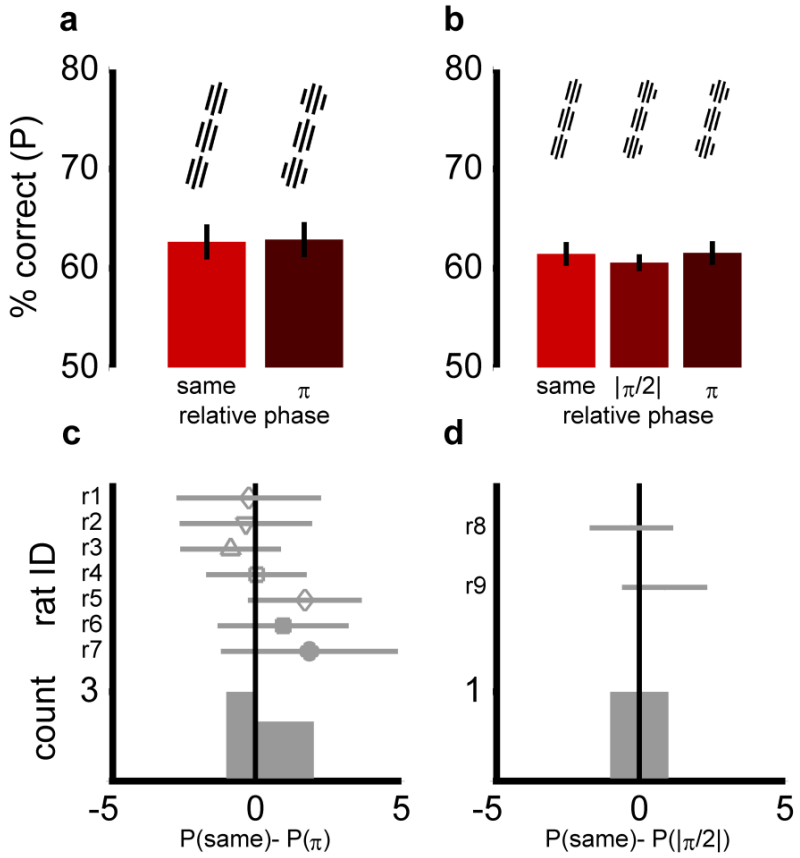

Figure 3.S7. Phase of grating has no influence on detection with flankers. a) A single rat's performance on collinear stimuli that were either phase-aligned or shifted by pi (sign-reversed). This is the same data as Fig. 3.4, reanalyzed to show that there is no effect of phase. b) A single rat in a separate pilot study in which four target phases and four flanker phases were used. Conditions are grouped by the phase difference between target and flanker. c) The difference in percent correct between phase-aligned and phase-reversed trials for the seven rats in the main study. No individual rats show significant differences. The population as a whole shows no indication that the relative phase of the flankers influences the rats detection performance ($p>0.05$, ANOVA; $p>0.05$, Friedman's test). d) The same is true for the pilot study with more phases. All three pair-wise comparisons are insignificant for each individual rat ($p<0.05$, Agresti-Caffo interval). Shown is the difference between phase aligned and stimuli in which the phase was shift $\pm\pi/2$. No population statistic was performed because N=2.

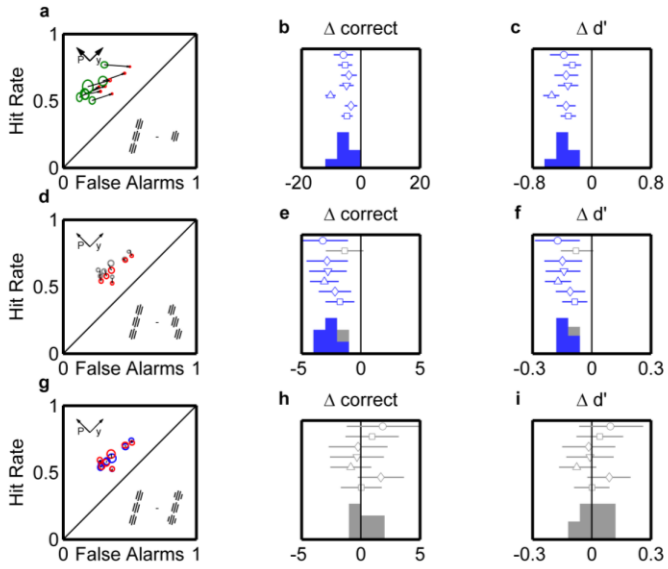

Figure 3.S8. Percent correct and d′ yield the same conclusions. The raw data are presented for three experiments: the influence of flankers (a,b,c), the influence of collinearity (d,e,f), and the influence of phase (h,i,j). The purpose is to validate that the conclusions are the same using percent correct or d′ as a performance metric. a) Same data as in Figure S3, comparing the difference between two conditions summarized by the icons lower right of the panel. The ellipses indicate performance of each rat on interleaved trials with the flanker present (red) or absent (green). Arrows indicates a single subject's change in hit rate and false alarm rate. The axis of hit rate and false alarm can be rotated to indicate percent correct performance (P) and percent yes responses (Y), as indicated by the inset axis in the upper left. All rats have a decrease in the percent correct when flankers are present. b) Performance reduction quantified by percent correct using Agresti-Caffo 95% confidence interval (identical to Fig. 3.S3c). Each horizontal line is the confidence interval for a single rat. The histogram of the mean change in performance color coded blue for each rat that is individually significant, and grey otherwise. c) Performance reduction quantified for d′ using the 95% confidence interval of an MCMC simulation. Panels (d,e,f) employ the same convention as (a,b,c), except comparing collinear flanker condition (red) to the parallel flanker condition (gray). Data are the same as Figure 5. e) difference by percent correct metric (identical to Fig. 3.5d). f) difference by d′ metric (identical to Fig. 3.5e). Panels (h,i,j) show the lack of an observed difference between collinear stimuli that where phase aligned (red) or phase reversed (blue). The data are the same as Figure 3.S7 a,c. The effects in panels b, c, e and f were all significant by ANOVA ($p<0.01$ and by Friedman's test ($p<0.01$), adjusted for multiple comparisons where required. There was no significant difference in phase (panels h and i) at the population level ($p>0.5$ both ANOVA or Friedman's test).

Table 3.1. Shaping sequence and training steps. Step numbers correspond to the numbers in the colored bars in the Figure 1c training timeline. All rats progressed from step 1 to 9 in increasing order. Two rats (r1, r2) performed contrast and spatial frequency varying psychometric curves before step 8 (not listed in this chart, results in Fig. 3.S4a). Another two rats (r4, r8) performed the flanker task with varying target contrast (not listed in this chart, Fig. 3.S4c). Other supplementary tests either used data from the training sequence (flanker contrast, step 9, Fig. 3.S4b), the main testing step (influence of luminance sign, Fig. 3.S7a,c), or after the main testing step (flanker distance, Fig. 3.S4d).

Step	**Description**	**Visual Stimulus**	**Graduation Criterion**	**Goal**	**Duration days ± std**
1	Free drinks	none (mean gray)	4 trials in 1 min	get water from port	1 ± 0
2	Earned drinks	none (mean gray)	5 trials/min for 2 min	alternate responses	2 ± 2
3	Faster drinks	none (mean gray)	6 trials/min for 3 min	sustain interest	2 ± 2
4	Side rewards	big grating	5 trials/min for 5 min	2AFC trial structure	8 ± 5
5	Easy detection	same big grating	85% correct	1st visual learning	17 ± 13
6	Linearized luminance	lower contrast	85% correct	new stim	9 ± 9
7	Thinner	double target cyc/deg	85% correct	new stim	4 ± 3
8	Smaller	target mask 1/3rd size	85% correct	new stim	35 ± 45
9	Add flanks	slowly raise flank contrast	80% correct	new stim	25 ± 13
Test	Flankers	See Fig. 3.4b	none	collect data	98 ± 36

Chapter 3, in full, appears as it was published in the *Journal of Vision*, 2011, Meier, P. M., Flister E.D., Reinagel, P. Philip Meier and Erik Flister contributed equally to this work. Each made essential contributions to the conception, design, and implementation of the general method and technology for high-throughput automated training and testing of rats in 2AFC visual tasks, with help and critical discussion from Pamela Reinagel. Philip Meier was responsible for the conception of the flanker experiment, design and implementation of the task, training and data collection, data analysis, interpretation of the results, and writing of the manuscript, with help and critical discussion from Erik Flister and Pamela Reinagel.

Chapter 4:

Rat performance on visual detection task modeled with divisive normalization and adaptive decision thresholds

Abstract

Performance on any perceptual task depends on both the perceptual capacity and the decision strategy of the subject. We provide a model to fit both aspects, and apply it to data from rats performing a detection task. When rats must detect a faint visual target, the presence of other nearby stimuli ('flankers') increases the difficulty of the task. In this study we consider two specific factors. First, flankers could diminish the sensory response to the target via spatial contrast normalization in early visual processing. Second, rats may treat the sensory signal caused by the flankers as if it belonged to the target. We call this source confusion, which may be sensory, cognitive, or both. We account for contrast normalization and source confusion by fitting model parameters to the likelihood of the observed behavioral data. We test multiple combinations of target and flanker contrast using a yes-no detection task. Contrast normalization was crucial to explain the rats' flanker induced detection impairment. By adding a decision variable to the contrast normalization framework, our model provides a new tool to assess differences in visual or cognitive brain function between normal and abnormal rodents.

4.1 Introduction

We trained rats to detect a faint visual target, tested them on variety of contrast conditions, and provide a mathematical model for their behavioral performance. The model specifies how decreasing the target's contrast impairs performance, as well as how increasing the contrast of two flanking stimuli ('flankers') impairs performance.

During natural vision, multiple stimuli compete for a limited cognitive resource. The detection paradigm determines which stimulus is relevant for behavior. In our task, the central grating patch is always the target, and its presence or absence should determine the subject's action. A perfect subject would ignore the flankers and only be influenced by the contrast of the target. Yet flankers substantially influence rats. Thus, the behavioral impairment caused by the flankers is revealing about the encoding of the visual scene, the character of the noise in the visual system, and the decision strategy of the rat.

Successfully attending to the target alone involves suppressing the influence of the flankers. This is goal-oriented selection. Additionally, natural vision involves the automatic selection of statistically relevant features. It is likely that such stimulus-driven selection of features exploits spatial contrast normalization. It has been argued that attention and contrast normalization rely on common mechanisms, (Reynolds & Heeger, 2009) such as lateral processing in inhibitory sub-networks. Therefore we are interested in developing a unified model accounting for both. Many psychophysical methods emphasize the priority of finding the minimum contrast that is required to perform correctly on a fixed fraction of trials. Contrast thresholds are useful because

they summarize a subject's aptitude with a single number, and are presumably more robust to variability in subjects' strategies, such as idiosyncratic response biases. Often determining the threshold involves measuring d' which invokes signal detection theory to disambiguate the effects of bias and aptitude. Unfortunately both d' and contrast thresholds create a layer of assumptions and abstractions between a conceptual model that explains the data, and the likelihood that the model gives rise to the data. Instead we fit a mathematical model to the subjects' raw data of hits and false alarms. We treat hits and false alarms for every stimulus condition as independent binomial proportions. We then maximize the likelihood that a given model could give rise to the observed data.

In order to explicitly model changes in bias caused by the flankers, we used a yes/no detection task. This differs from the more commonly used forced choice paradigm in which the subject chooses the stimulus at one of two spatial locations (2AFC) or one of two temporal intervals (2IFC). In these paradigms, both the target (T+) and non-target (T-) stimulus are provided on every trial. Presumably subjects performing 2AFC and 2IFC make comparisons between two representations in short term memory. We use a yes/no detection task that requires subjects to compare a single representation to an internal referent. As a result we can independently measure a subject's responses to when the target is present (T+, hit rate) and when it is absent (T-, false alarm rate). We note that yes/no detection is similar, but distinct from a go-no-go paradigm. Both are well suited for detection tasks. Our yes/no task has symmetric motor responses: a response port on the left or right side of the request

port. On the other hand, go-no-go has an inherent asymmetry in withholding and engaging in an action.

Using the framework of signal detection theory, we consider various models that differ in decision criteria, stimulus processing, and the character of the noise distribution. Our strategy is to find a single class of models that quantitatively fits the results with a small number of parameters. To adequately constrain the fit, we collected behavioral data for four target contrasts and five flanker contrasts (20 randomly interleaved conditions). Below we begin by presenting the raw data from a single subject, and the standard measure of d' and criterion for each stimulus condition independently. Then we sequentially add features to the model, explaining why parameters were added. The best model incorporates divisive normalization, and an adaptive decision criterion. Notably, subjects with very different responses can be fit by explicitly modeling the subjects' bias. Alternative inferior models are briefly considered. Finally, we consider the parameters of the model in order to explain why flankers impair detection in rats.

4.2 Methods

4.2.1 Animal training

Data were collected from two male Long Evans rats (Harlan Laboratories). All experiments were conducted under the supervision and with the approval of the *Institutional Animal Care and Use Committee* at the University of California San Diego. All subjects that were trained on these stimuli were included in the study; no

subjects were excluded. Training methods were the same as reported previously (Meier, Flister, & Reinagel, 2011). Each subject performed a 90 minute session every weekday for a total of 95 sessions. During the testing phase subjects performed about 30,000 valid trials.

4.2.2. Task

Rats initiated each trial by licking a port in the center of the behavioral training chamber. The rats' request triggered a stimulus to appear in the middle of a CRT display for 200 ms. Flankers were presented along with the target on every trial, except for stimulus conditions where flanker contrast was set to zero. A target was present on 50% of the trials. Subjects were required to select one of two response ports on the left or right hand side of the chamber to indicate that the target was either present or absent. The mapping between yes/no and left/right was opposite for the two subjects, and constant for their entire life. For correct responses rats were rewarded with a small amount of water, and for incorrect responses rats had to wait 6-10 seconds before the next trial.

A tone accompanied each trial request, providing the subject with confirmation that he had successfully initiated a trial and should proceed to ascertain and report target presence. During the time-out penalty after an incorrect response, a different tone and a flickering screen indicated to the rat that the system was non-responsive.

Stimuli were presented in blocks of 150 trials. Within a block, the flanker contrast C_F was constant and the target contrast was either zero or a fixed value C_T. This study

used four target contrasts and five flanker contrasts, for a total of twenty conditions. The stimulus with the target present is shown for all twenty conditions (Fig. 4.1c). Each block was presented once before repeating the blocks in a new random order. Subjects completed roughly 3-4 blocks per day.

Correction trials were used during initial training. After making a mistake, a correction trial began with 50% probability. During a correction trial, stimulus was presented such that the correct response was the opposite of the subject's last response. This prevented the rat from perseverating on a single port at chance reward rates. The correction trial persisted until the subject made a correct response. Correction trials were not present during the testing phase. Data was not analyzed from trials immediately following errors, allowing for the possibility that the rats were biased to switch sides after an error. This also served to avoid short-term effects of contrast normalization to the flickering visual signal presented during the time-out period (Gaudry & Reinagel, 2007).

We increased rewards for consecutive correct answers. The first correct response after an error yielded an 80 ms valve opening (approximately 16 μL). The 2nd to 4th consecutive correct responses earned 100, 150, and 250 ms rewards. Consecutive responses thereafter earned 250 ms rewards; the first incorrect response reset this schedule to the beginning value of 80 ms. Using ramped rewards discourages guessing strategies. Without a ramp, a rat performing at 75% correct only earns 50% more rewards than guessing randomly. With the ramp imposed, the same performance yields more than twice the rewards expected by chance guessing.

4.2.3. Stimuli

The target stimulus was an oriented square-wave grating with a spatial frequency of 0.22 cyc/deg, presented in a Gaussian mask with a standard deviation of 6.8°. Subjectively, this left about 3 visible periods in the grating. Previous experiments (N=2 subjects, not shown) confirmed that detection is contrast limited between contrasts of 0.25 and 1.0 at this spatial frequency, even for larger grating patches. The pixel pattern of each flanker was the same as the target. The two flanking stimuli were on the axis defined by the stimulus orientation, at a distance of 3λ, where λ (= 4.5°) is the spatial scale determined by the spatial frequency of the grating. Flankers did not spatially overlap the target, and subjectively appeared separated from the target. In order to be consistent with their previous training stimuli, subject 1 viewed gratings tilted 15°, and subject 2 viewed gratings tilted 22.5°. We do not think this small difference particularly mattered. Target and flanker stimuli were always collinear. When viewing the screen from the request port, a rat's eye is roughly 10 cm from the monitor, 10 cm below its center. At this position, the center of target grating is roughly 14 cm away. No effort was made to invert perspective of the non-tangent display screen, and thus the orientation and spatial frequency of the three stimulus patches varied slightly due to perspective.

To present visual stimuli, we used PsychToolbox (Kleiner, 2007) to control standard OpenGL capable graphics cards (Nvidia GeForce 7600) via Matlab (Mathworks, Natick, MA). Stimuli were presented on a CRT monitor (NEC FE992,

19", 100 Hz, 1024x768 resolution) with linearized luminance output. A linearization table was created for the monitor by fitting a power law with gain and offset ($y=b*x^g+m$) to photodiode measurements (Thorlabs, PDA55) of a rectangular patch in the center of the screen. The minimum, mean, and maxiumum luminance were set to 4, 42 and 80 cd/m^2, respectively (Colorvision, spyder2express).

During initial training, rats viewed stimuli with geometric configurations such that the target orientation, flanker orientation, and angular position were randomly and independently chosen to be either clockwise or counter clockwise by a fixed amount. During testing on the task reported in this manuscript, contrast was varied and geometric configurations were held constant to minimize the number of independent stimuli tested. Target orientation, flanker orientation and angular position remained constant at the same counter clockwise angle. Thus, all stimuli had a collinear configuration.

4.2.4. Model fitting

Models were created with a small number of parameters (4-7) that determined the noise distribution, the signal distribution and the decision criterion for each stimulus condition. Optimal parameters were selected by maximizing the likelihood that the model could have produced the data. If the experimental data are independent observations of a hit rate and a false alarm rate (Ogilvie & Creelman, 1968), then the likelihood is the product of binomial processes. For this experiment with 20 stimulus conditions, the likelihood can be calculated as the product of 40 binomial processes:

$$(1)\ p(data|\theta) = \prod_{i=1}^{20}\binom{H_i + M_i}{H_i} h_i^{H_i}(1-h_i)^{M_i} \prod_{i=1}^{20}\binom{F_i + C_i}{F_i} f_i^{F_i}(1-f_i)^{C_i}$$

Where *H, M, F* and *C* are integer counts of the number of hits, misses, false alarms, and correct rejections; and *h* and *f* are the rates predicted by the model. All counts and rates are indexed by the stimulus condition *i*, of which there were twenty in this study. Minimizing the negative log likelihood is numerically more stable than maximizing the likelihood. Additionally, we isolate the model-independent factors into a constant *a* that is defined purely by the data, and need not be recalculated in the optimization procedure. The log likelihood is then equivalent to the sum of products of the log probability:

$$(2)\ \mathrm{LL} = a + \sum_{i=1}^{20}(H_i \log(\mathrm{h_i}) + M_i \log(1-\mathrm{h_i}) + F_i \log(\mathrm{f_i}) + C_i \log(1-\mathrm{f_i}))$$

Each model was fit by choosing the best of 10-40 random starts of the Nelder-Mead simplex method (*fminsearch* in MATLAB) with a cost function proportional to the negative log likelihood. As some models have more parameters than others, we provide the Bayesian Information Criterion (BIC) as a measure of goodness of fit (Schwarz, 1978).

$$(3)\ BIC = -LL + \frac{k}{2}\log(n)$$

The BIC is equal to twice the negative log likelihood plus a penalty for the number of parameters *k* in the model. Each parameter is penalized by log(*n*), where *n* is the number of independent observations, equal to the number of trials performed by the rat. For our data set (*n*~30,000), each additional parameter incurred a penalty of ~10 units of a natural logarithm (nats). If a model produced a perfect fit to our data and had no parameters, the BIC would equal the negative likelihood, which is about 137 nats. Our best model has a fit around 238 nats, as compared to inferior explanatory models that have fits over 500 nats.

The variability of the fitting procedure is characterized by refitting the model 50 times to data that is sampled from the posterior of the model (Goris, Zaenen, & Wagemans, 2008). Assuming the best fit model is a true model of the data, this provides the variability in the parameters expected from limited data and the fitting procedure.

In this results section of this manuscript, models will be increased in complexity by adding more parameters. The final model contains seven parameters that were fit; they are summarized in Table 1 at the end of the results. Some parameters were constrained to a limited range by passing the fit value through a sigmoid: $\alpha, \lambda \in [0\ 1]$; or an exponential function: $\mu_T, \mu_F, \gamma, c_{50} \in [0\ \infty]$. The bias parameter is defined as a log ratio, and thus was not constrained, $b \in \Re$.

4.3 Results

4.3.1. Target contrast improves detection

We begin by considering the detection of a faint target when no flankers are present. If a target is so faint that it provides no information (is not visible), the rat must behave the same whether or not the target was there. In this case the hit rate and the false alarm rate would be equal (diagonal line in Fig. 4.2a). The rat is above chance performance for even the weakest contrast stimulus presented in this study. We find that increasing the target's contrast increases the rat's hit rate and decreases the rat's false alarm rate (Fig. 4.2a).

Even when the target is not present, the visual system still produces a noise response and the organism must internally process it to make a decision. On a given trial, an internal decision variable that summarizes the sensory evidence could take on a range of possible values. We adopt the framework of signal detection theory to mathematically model this decision variable (Green & Swets, 1966). We represent its possible values with a Gaussian distribution (see Discussion for alternatives). The rat shown in Fig.1a was tested on four target contrasts (Fig. 4.2a). We illustrate hypothetical underlying signal and noise distributions for a high contrast condition (Fig. 4.2b, C_T= 1) and a medium contrast condition (Fig. 4.2c, C_T= 0.5).

We refer to the distribution caused by a blank screen (T-) as the noise (black curves, Fig. 4.2b,c) and the response to a target (T+) as the signal (blue curves, Fig. 4.2b,c). The signal distribution increases with the target's contrast, which makes sense because there is increased sensory evidence that it was present. The metric d'

measures the separation between the noise and signal, and it increases with contrast (Fig. 4.2d). Notice that the noise distribution is the same regardless of the condition; it is always caused by a blank screen. All models in this paper assume that the sensory representation of a blank screen is not different by virtue of its being randomly interleaved with targets of higher contrast. We observe that the probability that a rat chooses yes when the target is absent (the false alarm rate) is dependent on the contrast of the target in that condition (Agresti-Caffo test between false alarms of C_T=0.25 and C_T=1.0, $p<10^{-4}$ for both subjects). How could the false alarm rate change if the stimulus is the same? Rats can adjust their decision threshold. For 150 consecutive trials, the rat views trials of the same condition. For blocks with low contrast targets, the subject reduces his threshold for making a yes decision (Fig. 4.2b, grey line is shifted left compared to Fig. 4.2c).

In figure 4.2e we present the bias as measured on each individual stimulus. This common measure of bias criterion is equal to the average of the normal transform of hit rate and false alarm rate: -[z(h)+z(f)]/2. This value is zero when the subject chooses "yes" and "no" with equal probability; zero corresponds to a threshold at the midpoint of two equal variance Gaussians. The positive criterion bias (Fig. 4.2e), indicates that this rat slightly favors no responses. This bias is reduced when the target is higher contrast. Note that the "decision threshold" increases with high contrast (Fig.4.2c, grey line moves right), while the "criterion bias" decreases (grey bar is closer to the intersection of the distributions). We clarify this point because some explanations of decision theory conflate these terms by translating the x-axis of

the decision variable so that the terms are equal. We avoid this translation, because it removes information about the means of the distributions. In all graphical depictions, we will preserve the raw means of the distributions because the absolute decision thresholds interact between different stimulus conditions.

4.3.2. Flanker contrast impairs detection

The presence of flankers makes the detection task difficult for the rats (Meier et al., 2011). Flankers decrease performance both by decreasing the hit rate and by increasing the false alarm rate (Fig. 4.3a). To facilitate comparison, a condition with no flanker contrast is re-plotted from figure 2 (blue symbol in 4.3a, all of 4.3b). Despite the high contrast target in all conditions, the separation between signal and noise is poor when then flanker contrast is high (Fig. 4.3c). There are reasons why the noise distribution might change in the presence of flankers. In this model, however, we apply the approximation that the noise distribution is constrained to be a normal Gaussian with zero mean and standard deviation of one.

As flanker contrast increases, detection sensitivity d' decreases (Fig. 4.3d). Flankers bias the rat to report yes more often (Fig. 4.3e). Measuring each stimulus condition separately, the bias criterion becomes more negative with flanker contrast. The impairment and the bias are consistent with theory that rats are confused about the source of the perceived contrast, causing accidental responses to the flankers alone. The reduction in sensitivity is also consistent with the theory that spatial contrast normalization decreases the effective strength of the target, rendering it hard to detect.

However, the impact of spatial contrast normalization is more difficult to intuit. To isolate these components we fit a model to data from a range of conditions that independently vary target contrast and flanker contrast.

Humans that detect targets sometimes improve their sensitivity when the oriented target is presented between collinear flankers, compared to without flankers (Chen & Tyler, 2008; Polat & Sagi, 2007). This increase in performance seems to be dependent on the target being low contrast. Here we ask, will low contrast targets also be easier for rats detect when they are flanked by collinear visual features? The hit rates and false alarms are presented for twenty conditions, including the previously described target only conditions (blue, Fig. 4.2) and data with variable flanker contrast and a high contrast target (red, Fig. 4.3). The additional observations for the remaining conditions (grey, Fig. 4.1c) follow the same trends (Fig. 4.4a): increasing flanker contrast always impoverishes performance (Fig. 4.4b,c), and it always biases the rats to report yes more often (Fig. 4.4d,e). Thus, regardless of the target contrast, collinear flankers do not facilitate detection in rats.

At higher target contrasts, each increment in flanker contrast mildly reduces the hit rate and substantially increases the false alarm rate. Interestingly, for low contrast targets, the hit rate increases with the flanker contrast. However, the rise in false alarms is still greater than the rise in hit rate, reflecting an overall loss of sensitivity (Fig 4b,c).

4.3.3 Fitting a model to the data

It is common for perceptual models to quantify the relationship between contrast and performance by accounting for the d' measurements and ignoring the decision criteria. This is sensible when the criteria vary between subjects, and a given subject has a fixed criterion for all stimuli. Our data would be poorly fit by such a model because the criteria shift between the stimulus conditions. That is, for a given the stimulus condition, the model would not be able to predict the observed hit rate and false alarm rate.

Our approach is to find a model with a small number of parameters to fit the raw data on all conditions. For each subject there are twenty stimulus conditions, each with a hit rate and false alarm rate, totaling 40 independent estimates of a binomial proportion per subject. In figure 4, these values were transformed into 20 estimates of d', and estimates of 20 criteria. This is guaranteed to "fit" the data exactly and represents no savings in parameters. To find a compact representation of the data, we begin with a very simple model of four parameters and incrementally increase its complexity, explaining the parameters as we include them. The goal of the model is to predict the signal (S) and noise (N) distribution, and the decision threshold (ζ) for each stimulus condition. This level of description requires at least three parameters per stimulus condition, and so it is important to compute these parameters from a simpler set of rules.

In our models we assume that the internal noise is stimulus independent and Gaussian; thus, the standard deviation of each distribution is fixed at 1. The means of

the distributions are calculated from a linear fit to the effective contrast, which employs a power law to characterize an accelerating non-linearity. Thus for each stimulus condition i, we obtain:

(4) $$\mu_s[i] = k_T(C_T[i])^\gamma + 2k_F(C_F[i])^\gamma$$

(5) $$\mu_N[i] = 2k_F(C_F[i])^\gamma$$

The parameters k_T and k_F are coefficients that determine the contribution of the target contrast and the flanker contrast to the decision variable. The coefficient for the flanker k_F is multiplied by two because there are two flankers. A subject ideally suited for this task would have a large k_T and a k_F of zero. Consequently, when relying on the state of the decision variable, the subject would have no confusion about whether the source of the signal was the target contrast or the flanker contrast. The ratio of k_T/k_F indicates how well the subject is attending to the target. Large values correspond to good selectivity of spatial attention. Values near 1 indicate that the rat is confused about the source of the contrast. If the ratio is 1, then all contrast is treated equally by the rat, and he is not preferentially attending to the target location. Notice that the mean of the signal distribution (μ_S) and noise distribution (μ_N) in fact use the same stimulus generating equation. The noise distribution is defined as the case where the target is not present (C_T=0), in which case eqn. 1 reduces to eqn. 2.

What remains is a specification of the decision threshold. To begin, we consider the case where the model can only select a single decision threshold to

account for the subject's performance on all stimulus conditions. The prediction of the hit rate and false alarms for all 40 observations is summarized by drawing iso-contrast curves (Fig. 4.5a). The intersections of the lines correspond to the model's prediction of the data. The model fit captures the gross topology: target contrast increases performance and flanker contrast biases. However the model fails in two respects. First, flanker contrast does not impair the models performance: shifting the signal and noise distribution by a constant amount will not change d'. Second, the model predicts that changing the target contrast will not change the false alarm rate. This is graphically apparent by the presence of vertical lines in the iso-contrast curves (Fig. 4.5a). Since rats do change their false alarm rate, even when the flanker is absent (Fig. 4.2a), we know that the class of models with fixed threshold will be a poor approximation.

It is natural to wonder if a dynamic threshold will predict the data if the rat always chooses the threshold that maximizes its rewards in the current block of trials. This model involves even one less parameter, because the decision threshold is determined by the distributions ($\zeta = \zeta_{opt}$). The results of this model are degenerate: all the outputs lie on a single line (Fig. 4.5e). Later in this paper we consider the addition of non-linear terms and a global bias. Even these parameters will not change the one dimensional organization of the iso-contrast curves; at most, these parameters serve to curve the line. In no case does the topology of the model prediction span a plane as we see in the raw data. Thus the hypothesis that rats know the signal and noise

distribution and choose the optimal threshold for each condition also fails to account for the observed data.

What decision thresholds could explain the data in this framework? It seems possible that rats strive for the optimal threshold, but do not confidently know what it is. If this were the case, the rat's choice of threshold could be explained by the combination of his prior for the optimal threshold with some weak evidence of the current optimal threshold. We speculate that this model could be formulated as the weighted combination of two thresholds: ζ_{prior} and ζ_{opt}. We allow ζ_{opt} to be defined by the signal and noise distribution, and take ζ_{prior} as the expected value of ζ_{opt} averaged over all conditions. The relative weight between the two hypothesis is determined by the parameter α which ranges from 0 to 1 and is fit empirically from the data. Thus for each stimulus condition i, we obtain a threshold $\zeta[i]$ as a weighted sum of the optimum threshold and a global prior:

(6) $$\zeta_{opt}[i] = \frac{\mu_S[i] + \mu_N[i]}{2}$$

(7) $$\zeta[i] = \alpha\zeta_{opt}[i] + (1 - \alpha)\zeta_{prior}$$

When the model includes a weighted hypothesis, it fits the data better (Fig. 4.5i). The Bayesian information criterion (BIC) is substantially reduced (see Methods for the interpretation of BIC). Notably the model is improved because it can produces diagonal lines for the curves of constant flanker contrast, even when flanker contrast is zero (Fig. 4.5i, blue line). However, as with the constant threshold model (Fig. 4.5a),

the influence of a flanker still does not impoverish performance as measured by d'. The reason is that that the only influence of flanker contrast is linearly additive to the mean of the decision variable. Next we incorporate the non-linearity of spatial contrast normalization, which can be described by a Naka-Rushton equation. This will cause flankers to also influence the decision variable indirectly by reducing the effective contrast of the target Rather than using a simple power law, we use the following equations to calculate the effective contrast of the target (C_T') and the flanker (C_F'), for each stimulus condition i.

(8) $$C_T'[i] = \frac{(c_T[i])^\gamma}{c_{50}^\gamma + (c_T[i] + 2\lambda c_F[i])^\gamma}$$

(9) $$C_F'[i] = \frac{(c_F[i])^\gamma}{c_{50}^\gamma + (c_F[i] + \lambda c_T[i] + \lambda^2 c_F[i])^\gamma}$$

The two terms in the denominator include the semi-saturation point (C_{50}) and the normalization pool. In other models, the normalization pool is sometimes taken to be the sum of the activity in a local neighborhood (Foley, 1994; Heeger, 1992; Tolhurst & Heeger, 1997). Here we estimate this activity by taking the weighted sum of the target and flanker contrast. The parameter λ represents the decrease in contribution to the normalization pool at the spatial separation between target and flank that was used in the experiments. In equation 8, the normalization pool includes the terms $2\lambda C_F$ because there are two flankers, and each one contributesλC_F. In equation 9, the flankers effective contrast is only influenced by one adjacent stimulus,

the target, which is captured by the term λC_T. Contrast normalization is a locally weighted phenomena, and so the other flanker, which is farther away, should not have the same influence. Because the flanker is twice as far away, we square the coefficient; this operation assumes and exponential decay of influence. Because the value of λ is always less than 1, the squared term will always have less of an effect. In pilot tests, we removed the influence of one flanker on the other, and the models were not qualitatively different. We kept the term because the simplicity of treating all contrast features equally (all stimuli in the display had the potential to contribute to the normalization pool) seemed more important than the parsimony of simply removing a term (insisting that the far flank could not contribute to the normalization pool).

When applying the divisive normalization, the calculation of the means of the signal and noise distribution remains the same. This is the same as equation 4 and 5 except that the effective contrasts are now calculated with equation 8 and 9.

(10) $$\mu_S[i] = k_T C_T'[i] + 2k_F C_F'[i]$$

(11) $$\mu_N[i] = 2k_F C_F'[i]$$

The best model fit with divisive normalization (μ_T,μ_F,γ,α,c_{50},λ) has a spatial falloff of 0.33. This means that each flanker contributes about a third as much as the target does to the target's normalization pool. The effective contrast of the target is reduced by the presence of the flanker contrast (Fig. 4.6a). A weaker contribution from the surrounding flankers would have less of an effect on the target's contrast

(Fig. 4.6b). From this experiment, it is unclear what the shape of the spatial profile is for divisive normalization. If we assume that the contribution from spatial neighbors falls off with a Gaussian profile, we can calculate the size of the normalization pool (σ_{DN}) that corresponds to the best fit model (Fig. 4.6c). The contribution is the spatial integral of the product of the contrast with the normalization pool, $A_F=\iint C_F(x,y)*DN(x,y)\delta x\delta y$. The flanker's contribution to this pool, was represented in units normalized to the strength of the target's contribution, which is equal to the parameter that we fit ($\lambda = A_F/A_T$). The best fit model has a normalization pool (σ_{DN}) about 4 times as large as the size of the target stimulus (Fig. 4.6c, red circle is the optimal fit). The corresponding model fits the data quite well, especially the reduction of performance caused by increasing contrast (Fig. 4.6e). Notably, the curvature of the iso-target contrast curves is correct. If the parameters are held constant, except for halving the influence of the flankers on the normalization pool, then the curvature of the iso-contrast contours change (Fig. 4.6f).

The model fits surprisingly well, considering that there is no parameter that explicitly accounts for the rat's bias. We consider the bias to be a stimulus independent reflection of the rat's preference for one port over the other. The model fits surprisingly well, considering that there is no parameter that explicitly accounts for the rat's bias. We consider the bias to be a stimulus independent reflection of the rat's preference for one port over the other. A bias could be caused if the rats had a prior expectation that a target would be present or absent. However the target probability was 50% across their entire training and testing experience, and so we do not think

this played a role. Alternatively a bias could be cause if the rat expected a larger reward or smaller penalty on one side than the other. In our experiment, penalties and rewards were symetric, nonetheless rats still displayed some bias.

We note that animals using the same apparatus with the same reward schedule displayed different biases. Thus we suspect that subjects' idiosyncratic strategies may affect their perceived utility of a left or right response. It is useful to have a free parameter to fit this bias. There is only a single parameter for the rat's overall bias and it is applied to all the stimulus conditions. The bias is determined by finding a threshold in the log likelihood ratio of the noise and signal distributions that best fits all of the observed data.

For equal variance Gaussian distributions the bias (b) affects the threshold by translating it proportional to the distance between the signal and noise distributions.

$$\zeta_{bias}[i] = \zeta_{opt}[i] + \frac{b}{\mu_S[i] - \mu_N[i]} \quad (12)$$

The effect of the adaptation is appropriately updated from equation 7 to use the rat's biased threshold:

$$\zeta = \alpha\zeta_{bias} + (1 - \alpha)\zeta_{prior} \quad (13)$$

Including a bias term improves the fit of data (μ_T,μ_F,γ,α,c_{50},λ,b) from subject 1, reducing the BIC from 257 to 234 (Fig. 4.7a,b). The bias is only a slight preference

to respond yes more often. On the other hand, the fit to the data from subject 2 is very poor without the bias term, and the model improves substantially from BIC of 2351 to 246 (Fig. 4.7c,d). The bias parameter captures the rats' tendency to favor no responses. In fact, without the bias term, the previous divisive normalization model completely failed to account for the data from subject 2 (Fig. 4.7c).

4.4 Discussion

We measured rats' ability to detect a dim target at a known location and onset. Additionally, we quantified the rats' performance as a function of the contrast of the target and the contrast of nearby flanking stimuli. Using an automated training and testing procedure, we observed 30,000 binary choices from each of two subjects. The experimental paradigm was a yes/no forced choice, which enables us to explicitly observe changes in the rats' bias per stimulus condition. We find that the most parsimonious description of rats' detection abilities requires three components: a decision criterion that is influenced by the ensemble of stimuli the rat views, spatial contrast normalization, and a global bias acting as a threshold of the log likelihood ratio.

Having accounted for the rats overall bias, the two sources of stimulus specific bias are a linear additive term that we use to model source confusion (k_F) and a nonlinear divisive term that we use to model the flanker's contribution (λ) to spatial contrast normalization. Collectively the models provide a quantitative description for how the flankers bias and impair the subject's performance. According to our models,

flankers impair detection in two ways: flanker contrast biases rats to report the target is present (source confusion); and flanker contrast reduces the effective contrast of the target (contrast normalization). Rats also adapt their decision criterion to the stimulus distribution.

4.4.1 What's the meaning of the parameter k_F?

The target contrast had more impact on the rat's decision than the flanker contrast. This is summarized at the ratio of k_T to k_F, which was larger than one in the best model for each of the two subjects (Fig. 4.8). If these values were equal it would indicate that the rat was not selective as to the source of the contrast that influenced the decision variable to increase, and thus caused more yes responses. The fact that k_F was lower indicates a selectivity, but it is mild. Thus, a substantial amount of the time, the rat is confusing the contrast from the flankers "as if" it came from the target.

Why do rats respond to flanker contrast as if it were target contrast? One simple possibility is that the size of the rats' receptive fields is so large that the relevant neurons span both the target and the flanker. Thus the target's signal is contaminated -correlated with the flanker contrast – at earliest stages of encoding. However, target stimuli were at least 13° and receptive fields of the rat ganglion cells are about 5° wide (Anishchenko et al.). This estimate involves a schematic model eye (Hughes, 1979; Hughes & Wassle, 1979) and a conservative approximation that the receptive field falls within 3 standard deviation of a Gaussian fit:

$3\sigma_{RF} * \frac{100\mu m}{\sigma_{RF}} * \frac{1^{\circ}}{59\mu m} = 5^{\circ}$. Some of the cell classes may vary in size, but are roughly

4.5-6° wide. The points spread function caused by the optics of the rat's eye may add some blur to the retinal image, and increase the receptive field size (Artal, Herreros de Tejada, Munoz Tedo, & Green, 1998), but this increase in width would be well less than a degree. Thus, retinal receptive fields measured as linear transfer functions do not require correlation between the target and flankers. If contrast signals are pooled due to large receptive fields, this occurs in later stages of processing.

An alternative possibility is that the rat does not know which neurons to monitor to make a decision, despite the consistent location of the target on the screen. This would be true if the rats' gaze varied from trial to trial and the decision process did not have access to the information about the gaze. Thus, any contrast signal in a visual area like V1 has multiple possible sources: it could have come from the target or the flanker. Given such a correlated noisy encoding, attributing target evidence to flanker contrast may be the result of optimally decoding a correlated signal. In other words, if a learning algorithm assigned weights to V1 neurons at a slow time scale spanning hundreds to thousands of trials, it should choose high weighs for neurons that most likely to respond to the target, and weak but non-zero weights for the neighbors. The width of these weights would reflect the variability of the stimulus encoding including behavioral variability.

Finally, it is possible that lateral interactions in the visual system cause the presence of the flankers to "fill in" (Chen & Tyler, 2008; Polat & Sagi, 2007). This is a conceptually interesting possibility, but with this data set we cannot distinguish it from the other possible lateral interactions in the visual system or the uncertainty

about the target's location (Wu & Chen). The strength of our model is that we explicitly represent the influence of source confusion and thereby isolate it as distinct from the effects of contrast normalization.

4.4.2 What's the meaning of the parameter λ?

For both subjects, the best fit of the flanker's contribution (λ) to spatial contrast normalization was about 66% of the target's contribution. Specifically, each flanker contributed about a third as much as the target. The Naka-Rushton equation has effectively modeled many saturating components of perceptual systems. Extending the normalization to include a spatial pool of local neural responses has proven to be an effective model for the nonlinearities of circuit level neural processing in the early visual system (Carandini et al., 1997; Geisler & Albrecht, 1992; Heeger, 1992). It also provides compact descriptions of behavioral data (Boynton, Demb, Glover, & Heeger, 1999; Chen & Tyler, 2008; Foley, 1994). This pooling may be selective to a specific visual channel at a given orientation or spatial frequency (Parkes, Lund, Angelucci, Solomon, & Morgan, 2001; Watson & Solomon, 1997), though this is not known for rats. Rats have orientation tuned neurons in V1 (Girman, Sauve, & Lund, 1999); but lack smoothly varying orientation maps (Ohki et al., 2005).

The model we present is a mathematical description that does not implicate a specific part of the brain. It is interesting that the same model can describe the activity of neurons in the primary visual cortex. Consequently, it is possible that neurons responsible for the interaction target and the flanker reside in the primary visual

cortex. Future experimental studies could confirm this hypothesis. However, it is also possible that the behavior we observe, which is consistent with normalization, is the result of computations in earlier visual areas, subsequent visual areas, or even part of a non-visual decision making stage.

It has been argued the control of attention is mediated by engaging the same mechanisms that are used in contrast normalization (Grossberg & Raizada, 2000; Reynolds & Heeger, 2009). The neurodynamics responsible for the divisive normalization are not known, but there are candidate mechanisms for multiplicative gain changes in individual neurons and population responses (Blomfield, 1974). Gain control mechanisms include shunting inhibition from GABA mediated conductance changes (Dreifuss, Kelly, & Krnjevic, 1969), feed-forward inhibition from fast spiking basket cells (Pouille, Marin-Burgin, Adesnik, Atallah, & Scanziani, 2009), coherent oscillations (Lakatos, Karmos, Mehta, Ulbert, & Schroeder, 2008), rapid conductance changes with counterbalanced excitation and inhibition (Atallah & Scanziani, 2009; Okun & Lampl, 2008), and stochastic fluctuation of the input variance (Shu et al., 2003). The common theme to these various circuit mechanisms is recruitment of sub-networks of inhibitory neurons that coordinate population activity and/or regulate the conductance or "leakiness" of individual cells. Much work is needed to understand how these circuit level mechanisms relate to perceptual phenomena like spatial contrast normalization and cognitive phenomena like attention. It is possible they have a common instantiation at the circuit level, sharing the same competitive and normalizing attributes of the network.

Assuming that contrast and attention share the same normalization pool, the behavioral model developed and fit to our data could also account for attention. The influence of attention could be modeled as an additional gain control *preceding* the normalization. This could competitively outweigh the flanker's contribution (λ) to the normalization pool. We note that this is different from the coefficient that we fit *after* the normalization (k_T), which takes on the more cognitive role of weighting evidence. This experiment did not include an independent manipulation of the subject's spatial attention, and so we cannot isolate the influence of attention from the influence of contrast. Thus we speculate that the relative weight of the flanker contrast is actually slightly higher, because the parameter (λ) simultaneously models the impact of flanker contrast (increase λ relative to target) and attention (decrease λ relative to target) on the normalization pool.

4.4.3 Theory

We have considered two ways the presence of flanking gratings could impair target detection performance: contrast normalization and source confusion.

How might contrast normalization impair? Flankers with a common onset should drive synchronous activity in the neurons surrounding the ones that encode the target. This is likely to induce contrast normalization: the divisive suppression of the spikes that encode the target (Carandini et al., 1997; DeAngelis et al., 1992; Heeger, 1992). Thus contrast normalization should impair detection because it weakens the signal encoding the target. Because suppression is divisive, it would have little or no

effect when target is absent. How might source confusion impair? One possibility is that visual neurons lack the spatial resolution to separate flanker from target contrast. In this case, source confusion can arise from an optimal decision strategy, given that visual detection of the signal involves some contamination from nearby regions of visual space. Alternatively, source confusion could arise at a cognitive level, for example, poor spatial resolution in the rats' representation of what screen location qualifies as a valid target. Among models we tested, the best fit to our data was achieved by models with both contrast normalization and source confusion.

Interestingly, the presence of flanking gratings never facilitated the rats' ability to detect the target. This is contrary to the observation that flanking gratings can sometimes facilitate human detection performance. Such facilitation may be caused by decreasing the uncertainty about the targets location in space and time. This would enable the decision process to attend to a smaller number of sensory neurons, thus avoiding the false positives caused by noise from irrelevant sensors. Alternately, flanking stimuli could provide a pedestal of excitation that would enable weak signals to cross a threshold for activation, similar to stochastic resonance (Goris et al., 2008; Wiesenfeld & Moss, 1995). Both theories suggest that facilitation may occur for low target contrasts. However, facilitation was not found for rats at any combination of target and flanker contrast. It is possible that there was very mild facilitation with a flanker contrast of 0.75 and a target contrast of 0.25 (see small bump in d' in Fig. 4.4 b,c). However, this did not raise performance about the level of the target-alone condition.

4.4.4 Future recommendations

The model presented here represents the groundwork for a core model of visual detection in the presence of distracters. We recommend four ways a future model could be improved: by adding priors for the models parameters, by considering alternate decision rules, by analyzing trial by trial fluctuations in behavior, and by using non-Gaussian models of internal noise.

In this paper, the model provided was fit to the likelihood, assuming a flat prior for all parameters. A Bayesian approach that provides a prior for the parameters would result in more stable fits. Providing priors, even just for incidental parameters like γ and c_{50}, can decrease the variance in estimates of parameters of interest, such as λ. Including priors may be especially useful if taking multiple measurements over time, or measuring a larger population of subjects. We provided a decision rule as the linear combination of two thresholds because it is the simplest mathematical form that fit data well. Intuitively it has some appeal because the weighted combination of the two may correspond to the limited evidence the subject has for the current optimal threshold. The model we provide does not have a temporal component of the decision process and cannot explain why the rats' bias is correlated with their reaction time on a trial to trial basis. We note that a simple diffusion model of decision making does not fit the reaction time data because accuracy does not decrease with reaction time (data not shown).

Another possible improvement for future models is to consider the trial to trial changes in behavior. The current model treated each stimulus conditions as a constant binomial proportion for all trials with the same stimulus. However, we know that rats change their bias slightly based on the context of their previous response and whether it was correct or incorrect. Specifically, rats are more likely to choose the opposite response immediately after an error. This is similar to the well known tendencies of human subjects to alternate and perseverate responses (Cho et al., 2002). It is possible that the rats increased their alternation strategy due to our use of correction trial in initial training (see Methods). The consequence of this is that animals may appear to be slightly less biased on hard stimuli where they make a large number of errors. In our data set, we avoided the variability due to post-error switching by only analyzing trials in which the previous trial was correct. Future models could explicitly model the switching tendencies. A simple example is to add or subtract a fixed switching bias (b_s) to the bias parameter (b) after errors; $+b_s$ after false alarms and $-b_s$ after misses. This shift in bias is equivalent to context specific adjustment to the prior odds.

Finally, it is possible that future models may describe decision variables with probability functions that are more interpretable with physiological parameters. For example, a Gaussian distribution is capable of having negative values; this does not make sense for firing rates, which must be positive. Additionally, it seems that noise increases with signal strength in neural representations (Shadlen & Newsome, 1998), especially for low contrast signals that are not saturated.

We performed preliminary tests on three families of distributions that are bounded at zero, and thus change their shape as the decision variable increases: the exponential distribution, the gamma distribution and the lognormal distribution. When comparing the fits with the BIC, the gamma distribution performs slightly better, but the normal distribution performs quite well, especially when the variance of the noise distribution scales with the signal such that the variance is twice the mean (data not shown). In the long run, other probability models may add explanatory power and provide a better fit to the data in terms of a lower BIC. However, for this data set we the improvements from non-Gaussian probability functions was modest. Thus, for the purpose of presenting this model, we focused on the Gaussian distribution, which has a clear relationship to d'.

4.4.5 Applications

The behavioral paradigm presented in this paper may prove useful for high throughput screening of perceptual and cognitive defects in rodents. For example, schizophrenic patients exhibit abnormal visual processing to high contrast visual stimuli (Brenner et al., 2009; Chapman & Mc, 1962) and atypical performance in the visual tasks with flankers (Gooding, Braun, & Studer, 2006; Keri, Kelemen, & Benedek, 2009) presumably due to abnormal function of GABA-ergic interneurons (Bullock, Cardon, Bustillo, Roberts, & Perrone-Bizzozero, 2008; Cruz, Weaver, Lovallo, Melchitzky, & Lewis, 2009; Lewis, Hashimoto, & Volk, 2005). A task and model like ours can be used to separately measure the influence of spatial contrast

normalization and cognitive confusion in animal models of schizophrenia. Other disease states implicated with abnormal lateral processing in inhibitory networks may also have unique behavioral signatures that could be discovered and quantified by estimating the strength of the contrast normalization in basic detection tasks.

Applying an automated method to train animals (Meier et al., 2011), we provide here a framework to fit models to the likelihood of the raw data, and examples of key parameters required to fit the bias and deficits caused by flanking stimuli. Within the model provided, the contribution of the flankers to non-linear divisive normalization (λ) is separated from the linear contribution of the flanker contrast (k_F) to the decision variable. Measuring both of these parameters, as well as nuisance parameters like the animal's overall bias, provides an assay that is sensitive to changes in the spatial contrast normalization.

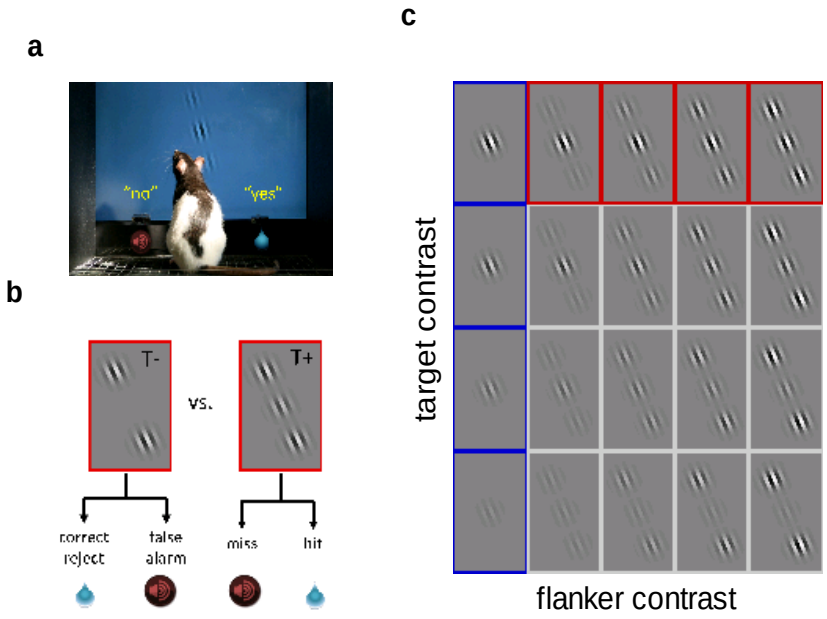

Figure 4.1. Experimental design for yes/no 2AFC target detection. a) A rat in the training chamber viewing a stimulus. The rat's task is to report the presence or absence of the central grating (the "target"). The rat is forced to choose either "yes" or "no" before proceeding to the next trial. The top and bottom gratings ("flankers") contain no information about the correct response. In panel a), the target is present in the stimulus, and thus a response to the right port would be rewarded with a drop of water (a "hit"). A response to the left port would be punished with a timeout period and an aversive tone (a "miss"). For photographic purposes, the stimulus was left on indefinitely, but during the task, stimuli persisted for 200msec. b) The target was present on 50% of the trials (T+) and absent on other trials (T-). The rat was appropriately rewarded and punished for correct rejections and false alarms, just like hits and misses. c) In the testing period, rats performed a block of 150 consecutive trials during which the target contrast and flanker contrast were held constant. Every 20 blocks visited each combination of target contrast (C_t=[0.25, 0.5 ,0.75 ,1]) and flanker contrast (C_f=[0, 0.25, 0.5 ,0.75 ,1]), in a random order. The probability of a target being absent (C_t=0) was always 50%, and thus we refer to the C_t of a block as the contrast that a target would be if it were present on a trial.

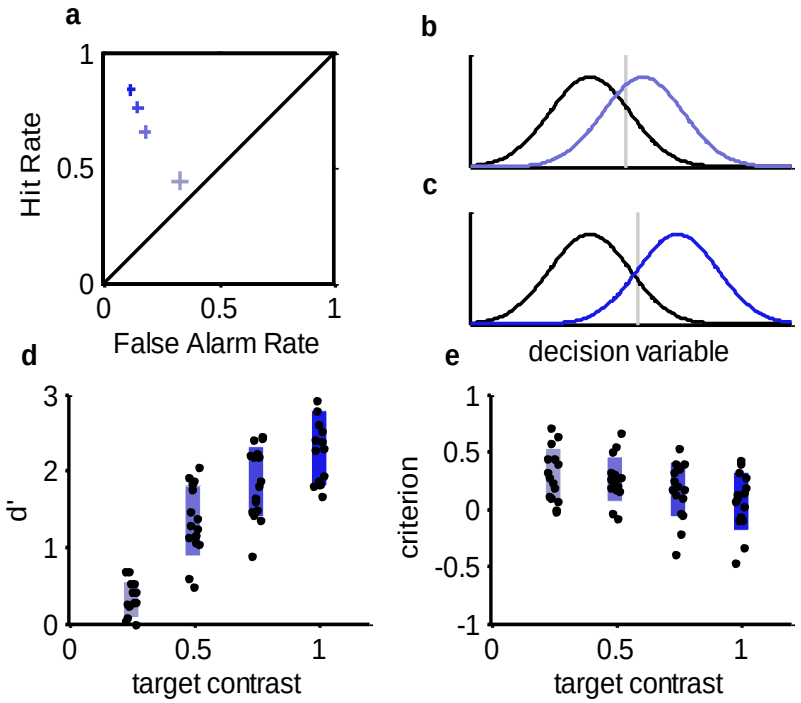

Figure 4.2. Higher target contrasts increase detection performance and bias rats to choose yes. a) The probability that a rat responds "yes" given that the target was absent (false alarm rate) or present (hit rate). If the data fall on the diagonal line, the animal's response contains no information about the target. Perfect performance is in the upper left. The four data points indicate performance from blocks where the flanker was absent and the target varied in contrast. The horizontal and vertical components of each plus symbol indicate the 95% confidence interval of a binomial proportion, for the false alarms and hit proportions respectively. b) A representation of the medium contrast target condition, assuming that a rat uses a single decision variable which has a Gaussian distribution with equal variance for the noise (T-, black) and signal (T+, blue). If a decision variable is greater than threshold criterion (gray line), then the model produces a "yes" response. The means of the distributions are separated by 1.3σ. Thus, the measure d'=1.3. c) For a higher contrast (C_t=1.0), the distributions overlap less, and the target is easier to detect (d' = 2.2). Notice that the threshold criterion is greater (grey line shifted right), resulting in fewer false alarms. d) Higher contrast targets have a larger d'. Each data point represents an estimate from a single block of trials. The vertical bars cover +/- 1 std. The variability of data is larger than expected from limited sampling; other factors beyond the stimulus also also affect performance. e) The average bias criterion is greater than zero, indicating that this rat favored "no" responses on all of these conditions. The subject's mild bias for "no" responses was reduced as the target contrast increased.

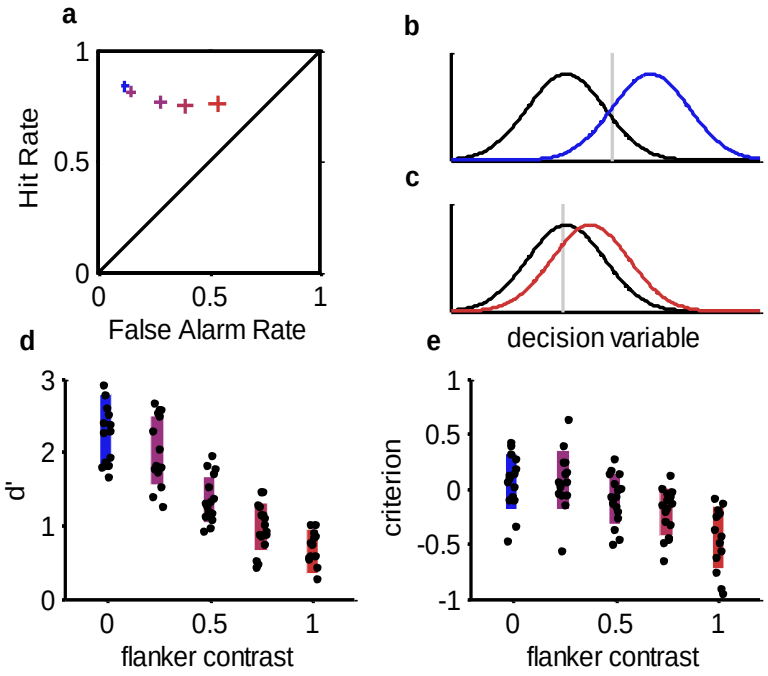

Figure 4.3. Higher flanker contrasts decrease detection performance and bias rats to choose yes. a) A scatter plot of the false alarms and hits for five stimulus conditions, varying in flanker contrast. Symbols are 95% confidence interval of a binomial proportion, as in 4.2a. All stimuli have a target contrast of 1. The blue symbol indicates zero flanker contrast; increasing flanker contrast is indicated by redness. b) The discriminability of the target when there is no flanker present. c) The discriminability when the flanker contrast is increased. The decision criteria is represented by a vertical grey line and it fit to exactly match the data with no error. Note that we plot the means of the noise distribution at zero. In subsequent models, the means of the signal and noise distribution may both be non-zero. Notably, the flanker contrast can make the discriminability go down by increasing the noise distribution, even if the signal distribution were unchanged. d) The discriminability decreases with flanker contrast. e) The subjects bias criterion shifts with flanker contrast. For consistency with the past literature, we present the criterion here as (z(cr) + z(fa))/2. T he change in the parameter reflects the rat's increasing bias to say yes as flanker contrast increases, which is also evident in the raw data in 4.3a.

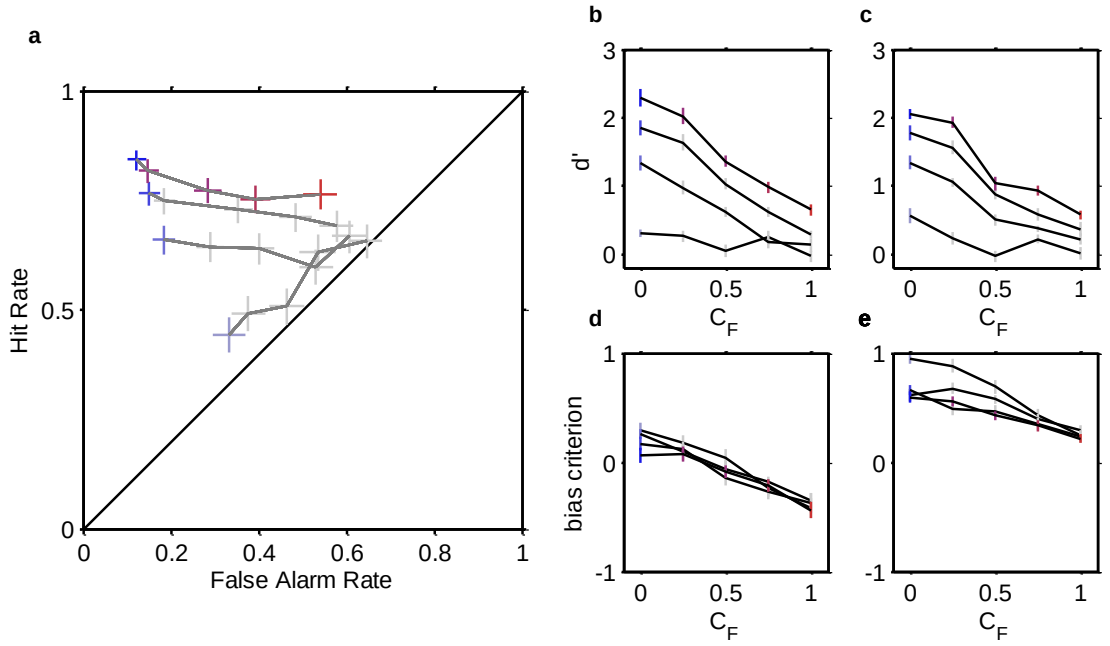

Figure 4.4. The impairment and bias caused by flankers is present for all target contrasts. a) The false alarms rate vs. the hit rate for twenty conditions including all combinations of four target contrasts (C_T=[0.25, 0.5, 0.75. 1]) and five flanker contrasts (C_F=[0, 0.25, 0.5, 0.75, 1]). Error bars are binomial confidence intervals. b) Increasing flanker contrast impairs performance. Lines connect stimuli with constant target contrast. Error bars are ±1 s.e.m. from n=14-20 measurements. c) Same as panel b, but for subject 2. d) Increasing flanker contrast increases the probability to say yes, which is a decrease in bias criterion. Lines connect stimuli with constant target contrast. e) Same as panel d, but for subject 2.

Figure 4.5. Various decision criteria. a) The best fit model ($\mu_T,\mu_F,\gamma,\alpha=0$) assuming that the subject only chooses a single fixed decision criterion for all stimuli ($\alpha=0$). The blue vertical line indicates the results of a pure change in target contrast. Other black lines indicate the influence of target contrast when the flanker contrast is higher. The red curve represents a change in the flanker contrast for a target contrast of one. Other black curves represent pure changes in flanker contrast if the target contrast is lower. If the data perfectly fit the model, the intersections of the lines would match the observed data. Blue crosses from Fig. 4.1a would fall on the blue line, and red crosses from figure 4.2b would fall on the red line. Grey crosses indicate all possible combinations of the target and flanker contrast (see Fig. 4.1c) and should be located at the intersection of the black lines. The gray contour indicates d' for each model's best fit to the high contrast condition; the curve spans all possible decision criterion thresholds. The signal and noise distribution of the model are displayed for three representative stimuli: b) a low contrast target alone, c) a high contrast target alone and d) a high contrast target with a high contrast flanker. e) The best fit model ($\mu_T,\mu_F,\gamma,\alpha=1$) if the subject chooses the optimal decision threshold for each stimulus condition ($\alpha=1$). Notice that allowing for the optimal choice for any symmetric distribution results in data that falls on a single diagonal line that extends from pure chance behavior to perfect performance. The data from the rat does not fall on a line; it is spread over a plane. This model is clearly wrong. The poor fit is reflected in the substantial rise in the Bayesian information criterion (BIC). The decision criterion are presented in (f,g,h) using the same stimulus conditions as before. Note that the entire model is refit, and so the signal and noise distribution may vary slightly as well. i) The best fit model ($\mu_T,\mu_F,\gamma,\alpha$) assuming that the subject's decision criterion is the weighted average between a single criterion (a=0), and the optimal criterion for that stimulus condition ($\alpha=1$). The best relative weight (α) is fit to the model. The value of $\alpha=0.79$ indicates that the decision criterion is close to the optimal, but ~20% influenced by a global criterion which is modeled as the average of criterion across all conditions.

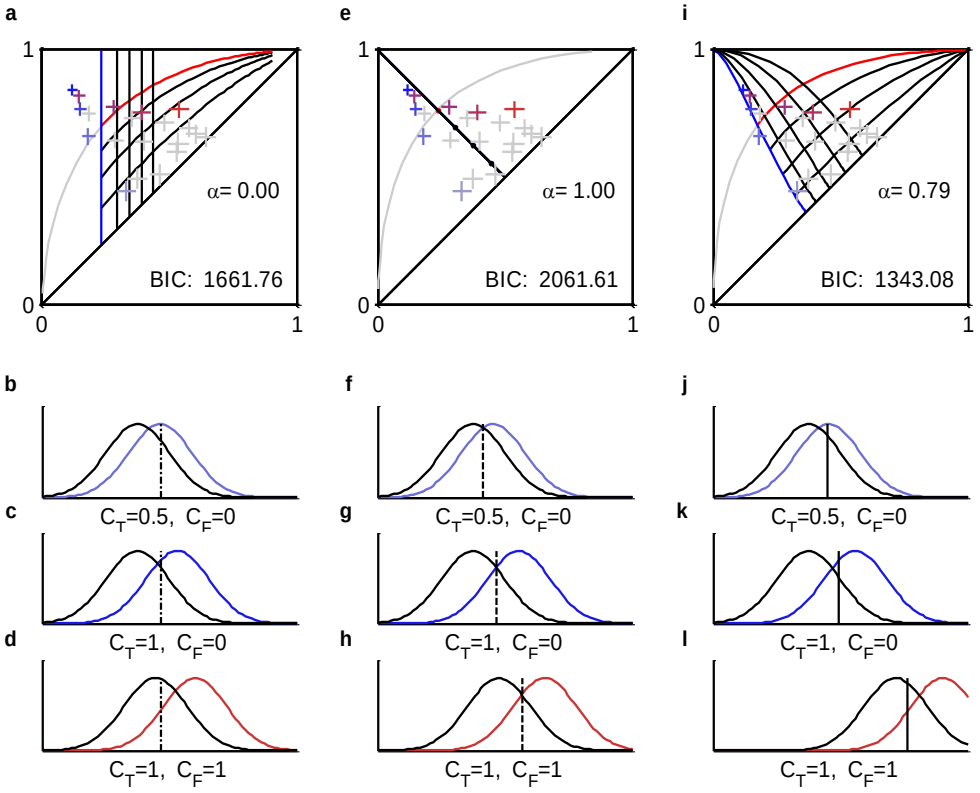
a
1
0
0
1
α= 0.00
BIC: 1661.76
e
1
0
0
1
α= 1.00
BIC: 2061.61
i
1
0
0
1
α= 0.79
BIC: 1343.08
b
c
C_T=0.5, C_F=0
f
g
C_T=0.5, C_F=0
j
k
C_T=0.5, C_F=0
d
C_T=1, C_F=0
C_T=1, C_F=1
h
C_T=1, C_F=0
C_T=1, C_F=1
l
C_T=1, C_F=0
C_T=1, C_F=1

Figure 4.6. Spatial contrast normalization improves the model fit. Spatial contrast normalization characterizes the non-linear contrast response by dividing the target's actual contrast by a normalization term to yield the effective contrast (C_T'). The full equation involves a semi-saturation constant and a power law (see eqn. 8). The appropriate nonlinearity is also applied to the flanker contrast (eqn. 9), but is not shown here. a) The effective vs. actual contrast of the target is colored blue for the condition where no flanker is present. The effective target contrast is reduced as the flanker contrast is increased (indicated by increasing redness). The curves display the nonlinearity for the best fit model (μ_T,μ_F, γ,α,c_{50}, λ) in which each flanker contributes 39% as much as the target to the normalization pool (λ=0.39). b) A smaller suboptimal parameter setting (λ=0.2) is displayed to facilitate intuitions. The flankers contribute less to the normalization pool and do not reduce the effective contrast of the target as much. c) The relationship between the contribution of the flankers (λ) to the normalization pool and the spatial extent of the normalization pool, assuming a Gaussian profile. Here lambda is calculated as the ratio of the contrast contribution from a single flanker (A_F) to the contribution from the target (A_T), see Methods. The size of divisive normalization pool (σ_{DN}) is plotted in units scaled to the size of the stimulus (σ_{stim}). The optimal fit of the model had a flanker contribution (λ=0.33, red line) which corresponds to a normalization pool about four times the radius of the stimulus. The suboptimal setting (yellow line) is also displayed. d) A schematic representation of a stimulus where in the intensity represents the contrast from the Gaussian masked gratings. The three gray contours represent the 2 std boundary of the contrast for the target and flanker patches. The red contour indicates the 2 std boundary for the optimum spatial region fit by the model. The yellow contour indicates the suboptimal region that is too small. e) The best fit for the model. f) A suboptimal model with all parameters the same expect for λ. The red contour is curved in such a way that faint flankers do not sufficiently impair detection.

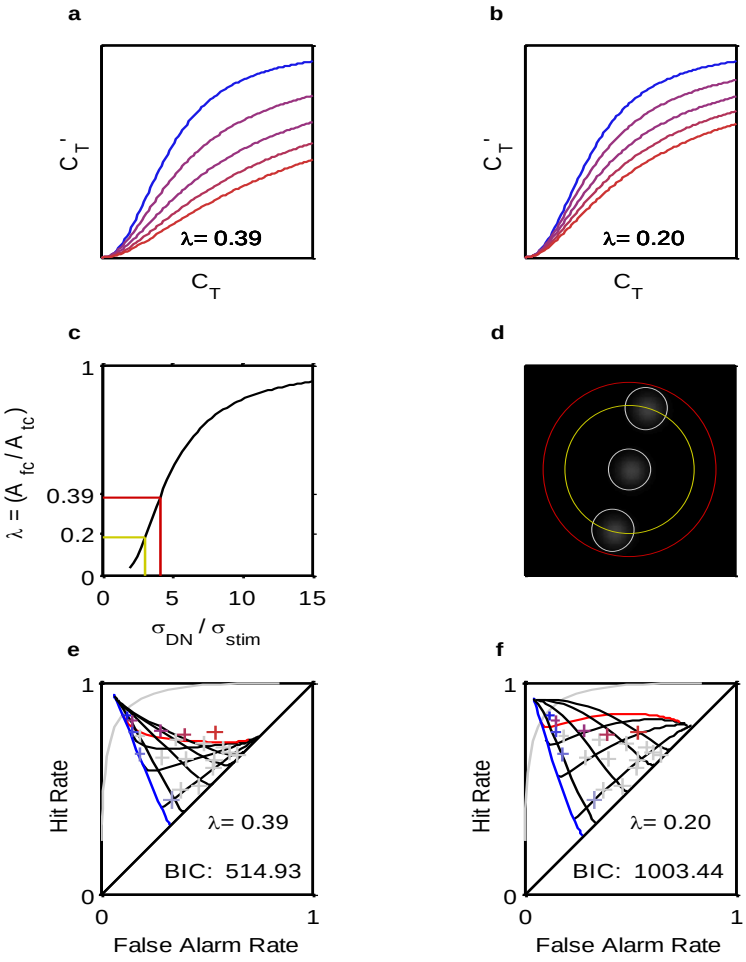
a
b
C_T'
C_T
λ= 0.39
λ= 0.20
c
d
λ = (A_fc / A_tc)
0.39
0.2
σ_DN / σ_stim
e
f
Hit Rate
False Alarm Rate
λ= 0.39
BIC: 514.93
λ= 0.20
BIC: 1003.44

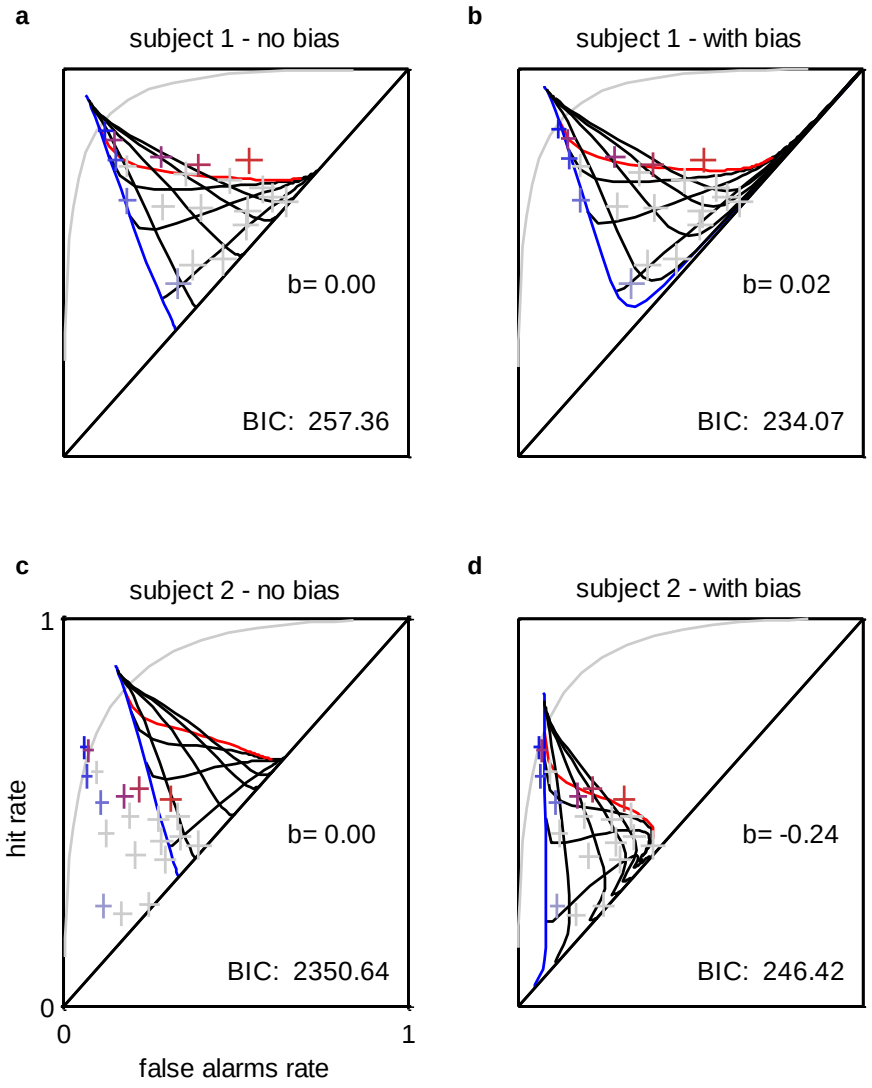

Figure 4.7. Bias allows model to generalize to different subjects. a) The best model for subject 1 when the utility for a correctly rejected trial is equal to the utility for a hit. In this panel, the log of the utility ratio is zero, b= log(util(cr) / util(hit)) =0. b is a single bias term that effects all 40 stimulus conditions. b) Subject 1 displays a modest improvement when a bias parameter is added, because he reports "yes" slightly more often. The decrease in the Bayesian Information criterion is ~ 23 nats, despite the penalty of ~ 10 nats per parameter. c) Subject two favors "no" responses, which is fit poorly by a model without a bias term. d) The addition of bias term enables the model to be fit quite well.

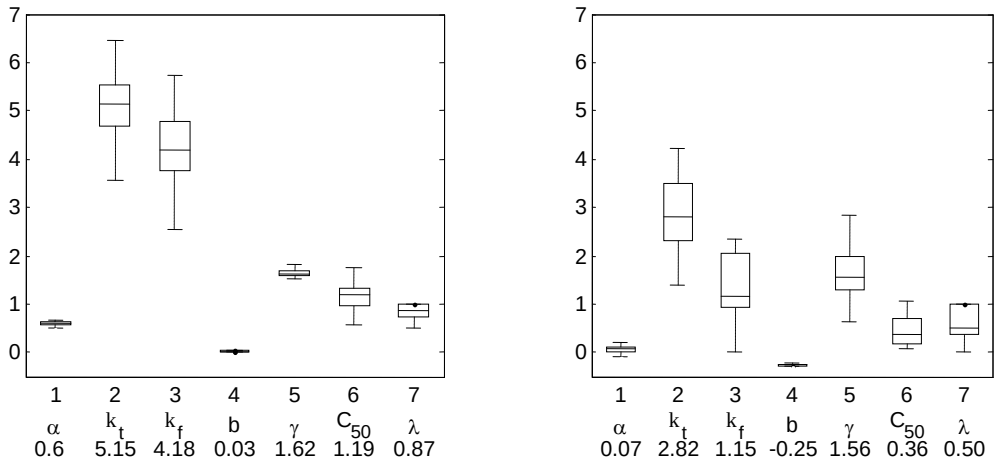

Figure 4.8. Confidence intervals for the model parameters. a) The parameter values for subject 1. b) Subject 2. Parameter distributions represent the variability due to limited data and the estimation procedure, assuming the best fit model was a true model of the data. For each parameter estimate, 30,000 trials of data were sampled from the posterior of the model, matching the total amount of real data collected from each subject. The fitting procedure was applied to this synthetic data. This process was repeated 50 times. The height of each box indicates the 1st and 3rd quartile and the slash in the middle is the median. If parameter values were distributed Gaussian, the whiskers would extend 2.7 sigma, covering 99.3% of the values. However, the data are not Gaussian, and the actual coverage shown here is less. The seven parameter symbols are displayed on the bottom, along with the median value. Notably, the target's contribution is larger than the flanker's contribution.

Table 4.1. Summary of the final model. There are seven free parameters which are fit from 40 data points.

Stage 1: effective contrast

γ acceleration of the contrast response nonlinearity

C_{50} contrast at 50% response saturation

λ space constant of normalization pool

→ *these determine effective contrasts of target and flanker,* C_T' *and* C_F' *(Eqn. 7,8)*

Stage 2: decision variable

k_T weight of target contrast C_T' input to decision variable

k_F weight of flanker contrast C_F' input to decision variable

→ *these determine the means of the signal and noise distributions,* μ_S *and* μ_N *(Eqn. 10,11)*

Stage 3: decision threshold

b response bias of subject → with μ_S and μ_N, determines ζ_{bias}

α relative weight of global vs. local threshold

→ *these determine the 20 condition-specific decision thresholds,* ζ *(Eqn. 13)*

Chapter 4, in full, appears as it was submitted for publication in the *Journal of Vision*, Meier, P. M., Reinagel, P. The dissertation author was the primary investigator and author of this paper and the co-author listed in this publication directed and supervised the research which forms the basis for this chapter.

Chapter 5:

Neural Recordings from the Lateral Geniculate Nucleus of the Rat

In this chapter I present preliminary results on the physiological properties of neurons in the rat's visual thalamus. My goal was to establish the feasibility of recording physiological data capable of supporting or rejecting the hypothesis that the correlates of pattern sensitive surround processing are present in the visual thalamus of the rat. The future study is feasible when one can reliably record single units from the visual thalamus of rats, record the position of the rats' eyes, localize the receptive field, and measure enough repetitions of a stimulus to have statistical power that quantifies the magnitude of surround modulation. My primary interest is the character of the signal from flashed stimuli similar to the ones viewed by the awake behaving rats (please see chapter 2 and 3). However, to address the animal's overall brain state under anesthesia, I also use stimuli with other properties as well: flickering noise patterns and drifting gratings.

Behavioral evidence does not implicate a particular brain region in generating collinear specific processing, so anywhere along the chain of processing could demonstrate the effect. At one extreme, the behavioral response of the animal contains the signature of interactions between the target and the flanker. Thus the motor cortex would contain a correlate of the behavior… but would not be very revealing about *why* the behavior was generated. It would not explain why collinear

flankers impair performance. My prior belief is that the interaction between the configuration of the flankers and the target is caused by mechanisms within the visual system, at early stages, and that the subsequent brain regions would simply carry the signature of a computation from an earlier stage. Thus localizing the flanker phenomena requires finding the *first* region in the brain with a signature of the interaction. Observing the phenomena in cortex would not rule out that it could have arrisen from an earlier stage. It would be particularly informative to know that contextual effects of flashed gratings are *not* present in the thalamus and that they *are* present in the visual cortex. Of course, there is always the possibility of finding an effect earlier in the visual processing stream than the traditional story suggests. The retina can display surprisingly complex phenomena (Gollisch & Meister), and some LGN cells of rats demonstrate orientation tuning.

Many experimental challenges are shared for recordings performed in V1 and the lateral geniculate nucleus, including the challenges of anesthesia, head-fixing and eye tracking. I performed a few experiments in V1 and confirmed that recording methods were also applicable to cortical experiments. In the following chapter, however, all the physiological data will come from thalamic recordings.

5.1 Choosing a preparation

Earlier in on in my thesis I performed recordings in awake free moving animals. The free moving neural recordings involved a head mounted recording device and a tether to record the data. The head mounted drives were appealing

because they were compatible with our behavioral paradigm. Unfortunately, the paradigm was not compatible with presenting controlled spatial stimuli and the cell yield was low. The animal's body position and eye position are constantly changing and unknown; therefore, it is only really feasible to reliably present temporal stimuli. If one insisted on trying to present spatial stimuli, the following challenges were prohibitive: 1) often the eye would be not be pointed at the stimulus display, 2) if it was pointed at the display, it would not be clear which region was driving visual responses, 3) the majority of variability in the data from a single neuron would be dominated by movements of the animal's gaze. I was motivated to find a paradigm where similar experiments could be performed as the ones that are possible with non-human primates. In such experiments monkeys perform a range of task, their eyes are free to move, and the stimulus within the receptive is known at every moment in time. This requires eye tracking which requires the rat's head to be still and conflicts with the free moving behavioral paradigm. Additionally, the yield per tetrode was particularly low (one cell per three trodes across an entire penetration). Since the neural drive was mounted on the animal, it was not possible to move the tetrodes to a new location. The low yield motivated me to switch to a paradigm where it was possible to move the recording device to a new location along the rostrocaudal or mediolateral axis. As a result, it was possible to search for a cell across multiple penetrations, greatly increasing the yield per subject, and the probability of finding a good cell in a given session.

Ultimately, I chose to work with head fixed animals with a well cap that allowed for access to the brain during experiments, but could easily be sealed shut between recording sessions. The well cap was prepared during the initial surgery. A thin 1-2mm layer of silicone over the brain surface reduced desiccation, and permitted a externally mounted electrode to pass though it without difficulty. An experiment could be performed either for anesthetized rats or for awake rats that had been accommodated to the rig. Rats would lie still in a soft fuzzy "sock" sandwiched between two Plexiglas surfaces. Rats were free to move their legs, but if they did, they would typically stop after a few movements. The low friction interface between the sock and Plexiglas limited the amount of force they could generate. Additionally the Plexiglas plate above the rat prevented him from arching his back and pushing down firmly with his leg muscles. Both these features reduced the strain the rat would generate on the one fixed point on the headmount.

Initially it was challenging to create a lasting head mount that could be used to hold a rat's head still during experiments. We hoped that titanium screws would create a strong connection to the bone because titanium has osseoinegrative properties. However, this alone was not enough. Over a few weeks to months, the connection between the mounting surface and the skull failed. This problem was fixed by including a luting as an improved adhesive (Rely-X) between the skull surface and dental cement and by providing daily oral antibiotics (Keflex).

5.2 Eye tracking

With the animal's head in place, it was possible to track its eyes with an infrared camera. More precisely, we were able to observe the changes in the difference between the center of the pupil and the position of the reflection of an infrared light source on the corneal surface. This technique can reliably estimate the absolute gaze angle of the eye if the subject performs a calibration procedure by fixating on known locations. Unfortunately, in rodents, we cannot not elicit saccades to a particular location, and thus it is not possible to use this calibration procedure. Other methods exist for back calculating the gaze based on the infrared data and geometric model of the eye (Stahl, 2004). However, these methods make considerable assumptions about the shape of the eye, and small errors in the model can produce substantial variation on the absolute estimated gaze angle. Fortunately, it is possible to provide valuable experimental information based on the relative changes in gaze direction, despite not knowing the absolute location. In other words, if we could confirm that the eyes were not moving, or at least that they returned to a home location for the majority of the experiment, we can simply remove the trials in which the eyes were pointed in a different direction.

Under anesthesia, the eyes would not move for minutes on end. A thirty second trace is shown in Figure 4.1a,b. However, occasionally the eyes would drift downwards for a period of five to ten seconds (Fig. 4.1c). Typically they would return to where they were before in a similar amount of time. These lolling eyes are relevant because it is not common practice for many rodent vision research labs to record the

eye position. Presumably it is assumed the eye will remain stable under anesthesia. As a result, variability of neural recording may be falsely attributed to variability in the neural processing as opposed to its true source in the eyes. Anecdotally, the probability of the eye drifting is higher when the animal is very lightly anesthetized and almost about to wake up. In deeper planes of anesthesia, the eyes are less likely to move.

Rats do not saccade under anesthesia. But they do saccade when they are awake (Fig. 4.1d,e). It is not clear if rats use saccades in order to improve their sampling of the world. Rats do not have a fovea, but they do have a circular region in the back of their eye with increased density of photoreceptors (Heffner & Heffner, 1992). The increase in photoreceptor density is broad and shallow compared to the fovea of cats or primates. One possibility is that rats use saccades as compensatory adjustments paired with head movements in order to improve image stabilization (Fuller, 1985). This could explain why rats do not saccade towards the presentation of novel localized stimuli. Saccades typically have more horizontal power. Rats do not often saccade up or down. Rather, when horizontal translations occur there is sometimes a vertical shift along with the horizontal component.

When rats are awake and head fixed, the variability in eye movements before and after the saccade is comparable to variability during anesthesia. Under isoflurane the standard deviation of the eye position in the horizontal direction is 0.06 deg (Fig. 4.1b). For an awake animal in stable period not including a saccade, the standard

deviation of the eye position is 0.08 degrees in the horizontal direction (Fig.4.1e, after the saccade).

Saccades were not obviously correlated with any feature of the experiment. It appeared that eye movements might be correlated with other body movements. However, we do not have a record of the body movements to be able to measure this directly.

Eye movements pose challenges for estimating the variability of the neural response. The two main challenges follow the categories of eye movements: saccades and drift. In order to accommodate for saccades, one strategy is to present stimuli that repeats on a time scale that is short compared to the intersaccadic interval. Since the interval is often larger than five seconds, and stimulus presentations last for 200 milliseconds, it is reasonable to display at least 10 repetitions that include a interstimulus interval of 200 milliseconds. While better than a single repetition, this is still far too low to randomly interleave different stimulus conditions in order to measure the influence of the surround.

Drift in the eye position poses different challenges. For example, the peak phase in a drifting grating would appear to shift as the eye changes position. Averaging over a long time sequence could produce a "blurred" response profile in cells which spike selectively to a particular phase of a drifting grating. For this reason, the total spike count ("F0") may be more appropriate for estimating responses to gratings when eye movements are present. However, we primarily chose to present gratings that drove the neurons to spike more; they were typically low spatial

frequencies and the small amount of drift in eye position would only minimally influence the preferred phase. Additionally, we inspected the variability of the preferred phase over time, and found it to be stable for most cells. These observations justify the metric "F1": the modulation of the spiking response at the preferred frequency of the stimulus. A more substantial change in response would occur if the entire receptive field drifted out of the region of stimulation. Thus, unless the experiment required otherwise, drifting gratings were presented on the whole screen, and the receptive field was positioned near the center of the monitor. Estimating the influence of the surround requires confidence about the location of the target with respect to a measured receptive field. This was achieved for anesthetized animals (Figure 4.1 b). However, the eye movements of awake animal prevented me from having confidence that the stimulus remained in the receptive field. Thus, the only flashed stimuli I presented to awake rats were gratings that spanned the full field. This could accommodate the changes in eye position, and enabled me to characterize the nature of the neural response at different contrasts. However it prevented me from measuring the impact of the surround on the neural response of the target.

One might expect drifting gratings to induce an optokinetic response (OKR), in which the animal smoothly follows the grating with the velocity of the drift speed (Fuller, 1985). We did not observe this with our frontally presented gratings, irrespective of their contrast, or drift speed. Nor did drifting natural images induce an OKR. Maybe stimuli spanning the lateral visual world are required to induce such a response.

5.3 Isolating single cells

Since our cell yield was fairly low, we actively choose cells, introducing a selection bias. Our selection criterion was primarily based on cells with a good signal to noise. That is, we sought spikes with large amplitudes compared to other spikes in the background. This may select for morphological properties that generate large dipoles. Additionally, we would fail to notice cells that did not elicit spikes in response to our search stimuli. We were more likely to record from cells with higher firing rates, especially if the high rates were evoked by the stimulus, and not merely a tonic background rate. We searched using either a hand held flashlight or a 2-4Hz drifting grating.

One argument for using a multichannel probe is to minimize the sampling bias in recorded cells. Since each of the recording locations cannot be moved individually, it reduces the possibility of optimizing a desired recording criteria and more fairly represents the distribution of signals in the thalamus. Unbiased sampling comes at the cost of obtaining only a small fraction of useful data. If we searched for a visual response on the tip of the multichannel probe, we could typically record 1-4 visual cells on the other 16 channels. Often the other electrodes would only contain multiunit visual activity that could not be well isolated as a single spike. Thus we favored single channel tungsten electrodes because their signal to noise was slightly better for single cell isolation, because their yield was not much lower, and their reduced cost enabled replacement.

When detecting spikes in the extra cellular voltage trace we typical set a threshold on the negative going deflection because it was larger than the positive lobe (Fig. 4.2b). Additionally, negative thresholds were useful because the presence of the CRT artifacts was restricted to positively deflecting features in the voltage trace (Fig. 4.2a). Thus, the choice of a negative threshold prevented the possibility of including "false" spikes. CRT artifacts could often be reduced with appropriate grounding, though not always. In the example presented, substantial CRT noise appears in the waveform traces, but the negative going spike is sufficiently large that the challenge of sorting spikes can be achieved with the appropriate selection of waveform features. In fact, the sorting could be accomplished with the appropriate choice of threshold, but here I use an inclusive threshold to be able to demonstrate the quality of the separation of the waveforms. The presence of noise waveforms is useful to visualize their separation from the isolated spikes (Fig. 4.2c).

By default we analyze spike waveforms using principle components. However, using principle component analysis (PCA) ends up representing features in a space that includes useless variability: for example, CRT noise in the waveform after the spike occurred. This is wasting feature space to represent variability that does not have discriminative power. Upward deflections occur at random locations from both the isolatable spike and other waveforms. In this particular case it is easier to demonstrate the quality of a sort using hand-picked features that focus on the variability of the negative going lobe. The example presented here is the exception and not the rule: most of the time, the automated clustering of klustaKwik (Harris, Henze,

Csicsvari, Hirase, & Buzsaki, 2000) works fine simply using projections of the first ten principle components as features.

5.4 Localizing individual receptive fields

In order to localize receptive fields I used a reverse correlation technique. By independently varying the stimulus over space and time, it is possible to determine which spatial location on the screen is responsible for causing a neuron to spike.

Spatial arrays of flickering stimuli did a poor job of eliciting spikes from the neurons of interest, even when it was clear that a reliable spike triggered average could be measured with full field flickering stimuli. Typically binary distributions of luminance were better at driving spikes than Gaussian distributions. Presumably this is because binary distributions have higher contrast. Additionally larger stimulus patches could elicit spikes, while very finely binned flickering squares were not good drivers of spiking activity. It was particularly challenging to drive strong responses under anesthesia. Roughly one in five visual cells that I recorded could be localized with a spatial stimulus, and typically knowledge of its position was very crude. It was less challenging to elicit visual responses to spatially independent flicker when recording from awake animals.

An example of a localized receptive field is shown in Figure 4.3a. The stimulus presented is a 12x16 grid of squares, each of which was independently drawn from a binary distribution of luminance outputs: either completely dark or completely light. The receptive field was well localized, despite eye movements. Indeed, if eye

movements had not been present in the recording, then the "size" of the receptive field would be smaller. Unfortunately, the quality of the eye recording was low during the stimulus presentation, and we could not reliably find a single region during which we were confident the eyes were not moving. Additionally, it is possible that the receptive field was sensitive to a region that spanned multiple adjacent independent squares. Thus, the center of the receptive field might be much smaller than 10 degrees to a side. Since the LGN inherits the majority of its spatial processing from the retina, receptive field sizes should be comparable to the ganglion cell receptive field centers of rats: about 2 degrees for a single standard deviation of the center width (Anishchenko et al.). In conclusion, the size of the measured region of luminance sensitivity is an upper bound for the size of the neuron's receptive field center. The increase in the measured RF size is probably due to eye motion and aliasing, though it may also be due to operations that pool laterally.

While the temporal signature of the center is immediately apparent in the spike triggered average, the influence of the surround is not (Fig. 4.3b). To get a better estimate of both the center and the surround, the average stimulus in the central four squares is pooled, as well as the average stimulus from the surrounding region immediately adjacent to the center. The pooling emphasizes the weak signal that was buried in the noise (Fig. 4.3c). The surround has a similar time course as the center, but it is inverted in sign. The peak surround is delayed one frame (10 msec) with respect to the time when the center has maximal sensitivity.

5.5 Influence of the surround under anesthesia

It was rare to obtain all the conditions necessary to perform an experiment that controlled for the placement of the stimuli in the surround. In one case, all the necessary components obtained: a cell was visually response, had good signal to noise, was stable for a long recording period, the signal from the eye was stable, the eye was not moving, the receptive field was localized in space, and the monitor had been centered upon the receptive field. Throughout a forty five minute session, the eye did not move more than 2 degrees.

The stimuli presented had the same spatial pattern as the stimuli viewed by the rats which performed the detection task in the presence of flankers. The stimuli did not drive a response, even at high contrast. I was able to confirm that the cell was still present and responsive to both white noise flicker, and a drifting grating. Second, I changed the Gaussian profile of the target, and used a circular stimulus patch. This manipulation would increase the contrast of the stimulus, especially if the stimulus were not perfectly centered. However using a circular aperture still did not drive responses. Third, I decreased the spatial frequency of the stimulus, and cell remained unresponsive. Finally, I did something more drastic; I modified the stimuli so that the dark stripes were not present. As a result, each stimulus patch was not merely a change in local contrast, but also in the average luminance. This is a different stimulus than the one used in the behavioral experiments. However, the stimulus did elicit reliable responses, so I recorded the data set. Because the stimulus is different (luminance and contrast vs. contrast alone) and the conditions are different

(anesthetized vs. awake behavior), we should be careful about extending any conclusions to the inferred neural responses during the awake behavior in Chapter 3 and 4. None the less, the data offer a useful point of reference about the neural response to stimuli that are flashed for 200 ms within the receptive field. The data set varied in two key parameters, the contrast of the target and the contrast of the flankers in the surround. In all cases the flankers were collinear and had the same phase as the target. Please keep in mind, in this experiment, changes in contrast were coupled with changes in luminance, and so it is probably better to think of the variables as "stimulus intensity" even though the graphs are labeled by "contrast." An important implication of this is that the extra classical divisive surround which is modulated by contrast could be conflated with a subtractive classical surround. This preliminary study does not distinguish between the two.

The raster plot displays forty eight repetitions of each of the 25 stimulus conditions: all combinations of 5 target contrasts and 5 flanker contrasts (Fig. 4.4). The target and flankers were present for 200 milliseconds and mean screen was present for two seconds between every presentation. Two seconds is probably longer than necessary for a cell to return to a baseline state between each presentation. In other experiments I would use 400 ms between stimulus presentations because it yielded more data within the typical 15-30 minute duration that a cell could be recorded from.

On some presentations the neuron would not respond at all. When it did, it would typically respond with a burst of about 12 spikes (Fig. 4.5a). Burst sustained

high firing rates typically between 400-500 Hz (Fig. 4.5b). The delay to the onset of the burst was about 100 ms. This delay ranged 65-120 ms, as indicated by the empirical 90% of the data. The average latency to first spike was shorter for full contrast stimuli (79 ms) than for stimuli with half that contrast (104 ms). The distributions of latencies are significantly different between the three highest contrast conditions: 1, 0.75, 0.5 (Two-sample Kolmogorov-Smirnov, $p<10\text{-}5$, $p<10\text{-}3$, $p<0.05$). Stimuli with little or no contrast ([0 0.25]) generated fewer than five instances of bursts, and there was not enough power to reveal a difference in burst latency. The relationship between contrast and latency (Fig. 4.5c) suggests that lower contrasts take longer to raise the neuron from its resting membrane potential to a spiking threshold. It is likely that reduced resting membrane potentials caused by isoflurane would also influence the latency of neural responses in the retina. It is unknown whether trials without spikes were caused by a failure to elicit a response in ganglion cells or a failure of the thalamic cells to respond to retinal activity. Likewise, without knowing the activity of the ganglion cell input, we cannot know if the burst originated in the retina or the thalamus. While thalamic relay neurons may be following rapid spiking that originates in the retina (Mitra & Miller, 2007), it seems more likely that the bursts are the result of the T-type Ca channel in the thalamus (X. Wang et al., 2007). This channel is de-inactivated when the membrane potential remains low for an extended period of time. Then, the next threshold spiking event, the typical sodium action potential is accompanied by a Ca spike with a longer time constant. The persisting calcium spike enables multiple sodium spikes in rapid succession. The

number of spike in a burst was not influenced by contrast; once a burst was initiated, low contrasts did not produce fewer spikes (Fig. 4.5d).

To characterize the bursting responses, I consider the interspike interval distribution (Fig. 4.6). More than 40% of the intervals are less than 5 msec, suggesting that a substantial fraction of the spikes are likely to belong to a burst response, and much fewer are part of tonic activity. A particular criteria for indentifying bursts has been validated by intracellular recording within the cat thalamus, in response to a slow 1Hz drifting grafting (Lu, Guido, & Sherman, 1992). If the interval before a spike is greater than 100 milliseconds and the interval after than spike is less than 4 milliseconds, then there is a high probability that the spike times will identify a calcium burst. I display this criterion, by plotting all of the spikes that fit the criteria in red in a two dimensional interspike interval plot (Fig. 4.6b). The figure shows the interval before and after each spike on a log temporal scale. A more inclusive criterion for bursts includes spikes when the interspike interval is less than 8 milliseconds (Fig. 4.6b). This criterion includes the other spikes in the cloud of data in the upper left of the plot. One concern is that this may also include high tonic firing rates that are not calcium spikes, especially given the nature of the stimulus, which includes long periods of no stimuli, and then the sudden onset of a strong driving stimulus. To address this possibility I look at the subsequent intervals after the first and second spike in the putative burst (Fig.4.6c). The intervals between the 2nd and 10th spike have a similar profile for spikes in which the first interval was slower: initially the frequency of spikes accelerates, and then it decelerates. This is true for

both the spikes that were included in the strict criteria as well as the more liberal criteria. However, there is a difference in the liberal bursts: subsequent intervals tend to be longer. This suggests that whatever state determines the sluggishness of the second spike persists throughout the burst. In conclusion, the liberal criterion probably selects real calcium spikes based on the following evidence: there are many spikes in the burst and the profile of the subsequent intervals contains the appropriate shape.

As expected, increasing the intensity of the target stimulus increased the average number of spikes (Fig. 4.4 reports the firing rate for each condition, Fig. 4.7 plots this organized by flanker contrast). However, the intensity of the flankers do not influence the firing rate. There is variability at some target intensities (such as C_T=0.75), but there is no consistent trend and differences are not significant. If there were small differences, they would be hard to resolve due to the bursting nature of the responses. In other words, a given stimulus either elicited no spikes or a whole burst. A clear example of this is the binary distribution of spike counts in response to a strong target stimulus. On 68.8% of the presentations there we no spikes at all, while on 31.8% there were more than 10 spikes. A high contrast target was presented 240 times; only twice, did a high contrast target generate between 1 and 9 spikes during the 200ms following the onset of the stimulus.

The measure of firing rate is not a particularly appropriate metric due to the non-poison nature of the response. The Fano factor measures the variance of the neural response over the mean of the response. Poisson stimuli would have a Fano

factor of 1. The neuron analyzed in Figure 4.7 has a Fano factor of about 20 for low contrast conditions and 10 for high contrast conditions. Each burst has a fairly constant number of spikes and each stimulus condition was presented the same number of times. As a consequence, plotting the firing rate looks the same as plotting the fraction of trials with at least one spike.

Since the firing rate is a fairly noisy measure for a process with a high Fano factor, I also considered the latency to the time to the first spike. This metric has the potential to provide more information because the latency is an analog value that might carry more information than the binary value that indicates the presence or absence of a burst. I restrict this analysis to conditions in which the target contrast generated enough spikes to analyze. This excludes 3 of the 5 target contrasts. The distribution of latencies for each flanker contrast is no different than the distribution of latencies with no flankers present (Two-sample Kolmogorov-Smirnov, $p>0.05$ in all eight comparisons considered, spanning two target contrasts and four flanker contrasts).

Why did the presence of additional stimuli in the surround fail to influence the response to the target? The presence and the intensity of collinear flankers did not influence the average number of spikes, the probability of a burst or the timing of the first spike. Notably, all of these metrics were influenced by the intensity of the target. Indeed, the stimuli used had the capacity to engage luminance normalization and contrast normalization. Having observed suppression would not have indicated

whether the contrast or luminance were responsible. On the other hand, observing no effect suggests that neither played a role. I consider five possible explanations below.

One possibility is that the influence of the flankers on the neural response to the target is small and I was not able to resolve it despite the large number of repetitions. This is particularly a concern because the rare responses at low contrast did not lend to analysis of firing rate changes or latency changes. Additionally, it is at the low target contrasts that we might expect to find a larger influence of the surround.

Another possibility is that other cells would demonstrate the effect. The cell that I focused on in this section is only a preliminary data point.

A third possibility is that flashed stimuli will not induce surround effects. Typically physiological effects of the surround are explored with drifting gratings (Cavanaugh et al., 2002; Levitt & Lund, 1997; Sillito et al., 1993) and perceptual effects of contrast perception are based on a static grating (Cannon & Fullenkamp, 1996) or sustained viewing contrast reversing stimuli (Xing & Heeger, 2001). It is possible that the motion signal plays an important role in establishing the suppressive surround. However physiological surround effects have been observed with contrast reversing stimuli too (Polat et al., 1998). It is possible that a sustained driving stimulus (like motion, or rapid periodic reversal) allows the network to settle into a steady state response. This is different from a briefly flashed stimulus that generates more of an impulse response. Intuitively, it may take some time for the appropriate mechanisms of lateral inhibition to influence the target response. However, it is noteworthy that many perceptual effects of the surround can be demonstrated with

flashed stimuli that only appear once. For example, flanker detection experiments often use stimuli presented for about 100 milliseconds (Chen & Tyler, 1999; Polat & Sagi, 1993; Zenger-Landolt & Koch, 2001; Zenger & Sagi, 1996).

A fourth possibility is that rats to not have the same kind of surround processing as reported by the literature in other mammals. This seems unlikely given the similarity of the mammalian visual system across species, and the robustness of surround processing. Indeed, in at least one case, our lab has recorded area summation curves that show that increasing the spatial window occupied by drifting grating initially increases the response of a cell, but eventually decreases the response for large sizes (data not shown). Some, but not all, cells demonstrate this effect.

A fifth possibility is that surround processing is disrupted by the anesthesia. One objection to this hypothesis is that surround suppression has been more often than not been documented in anesthetized animals. However, much of the previous research in the cat visual thalamus studied cats using pentobarbital, pentothal or urethane, (Bonin, Mante, & Carandini, 2005; Denning & Reinagel, 2005; Hubel & Wiesel, 1961; Lesica & Stanley, 2004) and in this study I used isoflurane as an anesthetic to study the LGN in rats. Light isoflurane anesthesia is generally touted as preferable by rodent researchers because it more resembles the activity of animals that are not anesthetized. In the following section I question the strength of this claim.

5.6 The influence of isoflurane anesthesia

How different are the response properties of the early visual system under anesthesia? Historically, much research has been motivated by the observation that visual responses remain intact under anesthesia (Hubel, 1960; Hubel & Wiesel, 1961). Many properties of the visual system may be derived from the anatomical pattern of connectivity, and one would not expect this connectivity to change with anesthesia. However, it is also possible that many visual phenomena depend on the balance of excitation and inhibition, which may be shifted substantially under anesthesia. Additionally, the functional connectivity may differ with anesthesia, even when the anatomical connectivity is the same. To assess the influence of anesthesia, I characterize visual responses of thalamic neurons in rats that are anesthetized with isoflurane as well as rats that are awake and passively viewing stimuli. I consider responses to drifting gratings of varying spatial frequency, flickering full field stimuli and flashed gratings of varying contrast. The data presented were recorded collaboratively with Balaji Sriram.

Thalamic neurons are visually responsive under anesthesia, but their firing rates are lower. I consider the visually evoked responses across three different stimulus categories. Combining the data from all stimulus categories, the neurons in awake rats fire spikes at an average of 9 Hz. Under isoflurane anesthesia this average evoked rate is four times less. Notably, under anesthesia there are many cells that have average firing rates lower than 1 Hz. It is important to consider the rates for each stimulus category, since the number of cells recorded for each category are not

matched. For example, the observed difference in firing rates between anesthetized and awake conditions could have been caused if more cells recorded under anesthesia were presented with a stimulus class that drives cells weakly. Indeed, there are different firing rates for the various stimulus types. Analyzing each class of stimulus separately confirms that the reduction of firing rate is not an artifact of the distributions of stimuli rendered: in each case, firing rates are lower under anesthesia (Fig. 4.8b,c,d).

It is possible that the overall firing rates have been reduced but that neurons maintain their selectivity under anesthesia. To test this we consider neural responses to drifting gratings of varying spatial frequency (Fig. 4.9a). Thalamic neurons recorded in awake rats are typically band pass (Fig. 4.9b). Their peak spatial frequency is an intermediate spatial frequency greater than 0.007 cycles per degree and less than 0.1 cycles per degree. The highest spatial frequencies always have a diminished response. The lowest spatial frequency typically also has a diminished response, but some cells (2/7) are low pass. On the other hand, most cells (7/10) recorded under anesthesia are low pass (Fig. 4.9c). The shift in the spatial sensitivity is comparable to blurring of the stimulus under anesthesia, as if large spatial regions of the retina were pooled. Lower firing rates could have explained a change in the width of the tuning curve, but cannot account for a shift in the location of the peak sensitivity. In conclusion, the presence of anesthesia appears to dramatically change the response properties of neurons.

An alternative explanation that could account for this observation is if the band pass cells recorded from awake rats were selectively silenced when under anesthesia. In this case, we would never record from them because they would not fit our search criteria. Thus, maybe individual cells do not lose their selectivity, but low pass cells were simply more likely to be found under anesthesia. In either case, the recordings made under anesthesia did not have properties that were representative of the awake neural data. This poses a challenge when relating the neural responses under anesthesia to the decision process of behaving animals.

Full field stimuli were also used to characterize neural responses in the awake and anesthetized condition. The response of a cell to flickering stimuli can be used to calculate the spike triggered average (STA) – the average stimulus value at each point in time preceding a spike. For every frame, the luminance of the stimulus was chosen from Gaussian distribution of intensity (Fig. 4.10a). If the stimulus is randomly and independently chosen on each frame, then the STA of an inhomogeneous Poisson process is equivalent to the transfer function that is the optimal linear filter. Both the awake and anesthetized conditions displayed a diverse range of STA waveforms of varying amplitude, sign and number of lobes. The amplitude is indicative of the fraction of spikes that are stimulus driven. The sign of the most recent peak in the transfer function indicates the cell type: positive lobes correspond to ON cells that respond to increases in luminance, negative lobes correspond to OFF cells that respond to decreases in luminance. As the shapes of the waveform were variable within conditions, many of the features did not display trends between the awake and

anesthetized recording conditions. One feature that was different between the conditions is that STAs recorded in awake rats appeared compressed and faster. We captured this with a single summary metric: the duration of time between the spike and the peak or trough with the largest deviation. The average duration from peak sensitivity to spike was 29 ms shorter for neurons recorded in awake rats (Fig. 4.10b,c). In the awake condition the average latency was 83 ms, while in the anesthetized condition the latency was 54 ms.

Why do the STAs from awake rats have a shorter duration between their maximum deviation and the time of the spike? One hypothesis is that the isoflurane reduces the resting membrane potential of cells, resulting in a longer duration of time between the primary driving stimulus and the time of the spike. This is consistent with the mechanism of action of isoflurane. Isoflurane depolarizes cells by acting on GABA receptors (Jones, Brooks, & Harrison, 1992) as well as causing a persistent leak of potassium (Ries & Puil, 1999). Another hypothesis is that the difference in the STA can be attributed to non-liner dynamics that induce correlations in the neurons spiking response. In other words, since both the stimulus and the network influence the spiking activity of the neuron – a change in the shape of an STA does not necessarily indicate that the neurons are sensitive to different features of the stimulus. For example, if a neuron bursts, each spike is not independently triggered by the stimulus 54 ms before it. Another member of the lab, Erik Flister, is working on methods to determine the statistical significance of differences in spike-triggered

stimuli. It would be fruitful to apply such methods to differences in responses to white noise stimuli, with and without an anesthetic agent.

In one case, I was able to record a rich dataset from a single neuron under isoflurane (Fig. 4.11a) and after the isoflurane was turned off (Fig. 4.11b). When isoflurane was removed the animal began to wake up. The animal was not fully alert in the un-anesthetized state due to the recent removal of the anesthesia. However the rat entered a cognitive state which was much more awake, and the changes in the neural response from the previous anesthetic plane is informative.

Gratings were repeatedly flashed during and after anesthesia. The gratings were not spatially restricted to the neurons receptive field; they filled the entire monitor screen (Fig. 4.11). Each presentation had a randomly selected contrast from one of five values: [0, 0.25, 0.5, 0.75, 1]. Every flash persisted for 200 ms, and was separated by 400 ms of a blank gray screen. The neuron responded in two volleys: once after the grating appeared, and once after the grating was removed. The second response contained more spikes than the first. Thus, I focus on the second response for characterizing differences in contrast sensitivity, though the basic observations also hold for the initial onset response.

Upon waking from anesthesia, the baseline firing rate of the neuron increased from less than 1Hz to more than 10Hz. This sustained baseline rate is visible in the raster plot in the period of time before the stimulus is presented (time<0), as well as the entire duration of the zero contrast condition (Fig. 4.12a,b top of raster, blue).

Under anesthesia, the neuron was unresponsive to a contrast of 0.25 and only very weakly responsive to a contrast of 0.5. After waking from anesthesia, the neuron responded to both of these contrasts with an evoked rate that was greater than the baseline rate. Interestingly, the response to increasing contrast was linear in the range that we studied. Up to the full contrast of the monitor, we did not observe saturation of the neural response.

A deviation from linearity was present at the low range of contrasts. Under anesthesia, there was no response to low contrasts, and the difference from linearity can be explained by the floor: a neuron cannot spike less than 0 Hz. In the awake context the low visual response was also non-linear at low contrasts, but for a different reason. The floor was not at 0 Hz, but at the baseline firing rate of 10Hz.

The influence of waking the rat is consistent with an additive pedestal of excitation. The addition of excitation reveals the responses to low contrast that were previously below the spiking threshold. At higher contrasts, all responses increased with approximately four extra spikes in the response. If these four spikes were to be smoothed over the 200 ms analysis window, they would correspond to an increase in the firing rate by 20 Hz. However, by inspecting the peri-stimulus time histogram (Fig. 4.12b), we can see that the addition of extra spikes is not uniform in the time range. Nor is it restricted to a time-locked range. Rather, higher contrasts induce more spikes for two reasons: they increase the peak firing rates and they extend the duration in which visually induced spikes are generated.

If one characterizes sensitivity by the slope of the contrast response function, then the awake state does not increase sensitivity. However, if one considers the minimally detectable contrast, then the pedestal of activity of the waking brain state increases contrast sensitivity. An ideal observer with access to a single neural channel would perform at chance detecting contrast of 0.25 using anesthetized responses. In the awake state, the ideal observer would could detect these gratings. On the other hand, performance on higher contrasts might decreased due to the addition of the baseline firing rate that increased the "noise" that persists even in the absence of a stimulus. It is unclear how much noise an increase in the baseline rate of a thalamic neuron would actually generate, given that subsequent cortical processing might adapt.

5.7 Future

I recommend recording from awake rats instead of anesthetized ones. They provide better estimates of the neural responses of subjects that were performing the detection task. Notably, in neurons recorded under anesthesia, the firing rates are lower, the selectivity for spatial stimuli is reduced, and low contrast stimuli are less detectable. Additionally, in the one case that flankers were reliably presented outside the neuron's receptive field, they had no influence on the neural response. Possibly flashed stimuli do not elicit the spatial normalization in the LGN of rats that has previously been reported in response to drifting gratings in the visual system of other mammals. It would be valuable to observe classical spatial contrast normalization to both drifting stimuli and flashed stimuli in the same neuron in awake rats.

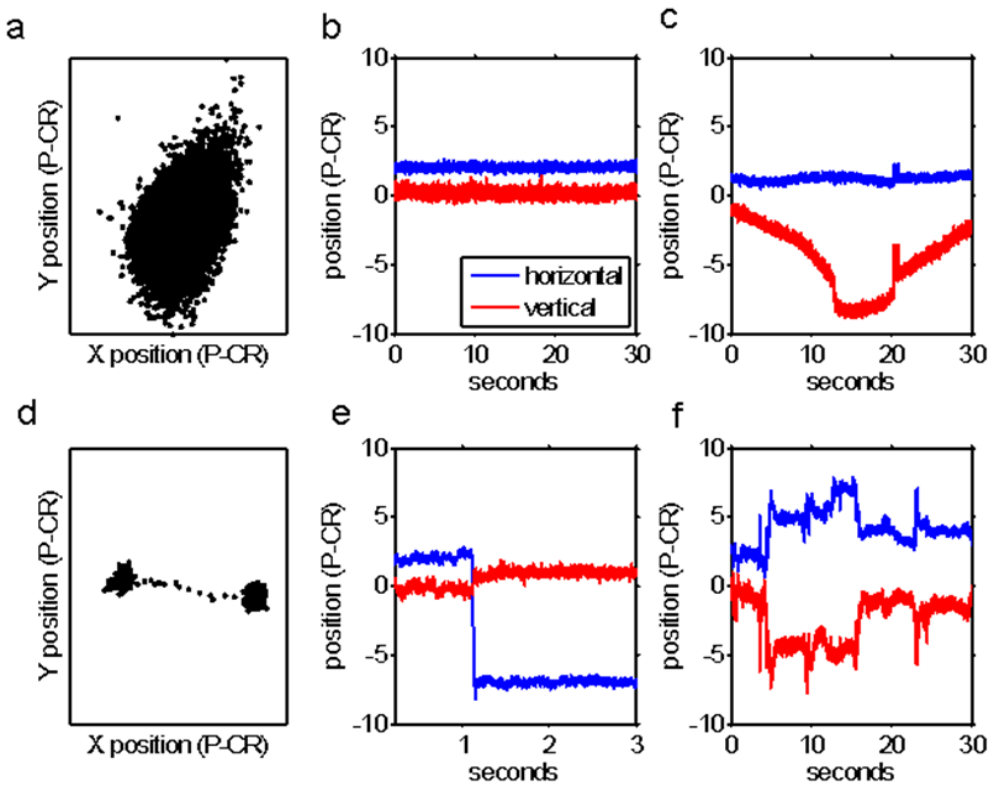

Figure 5.1. Eye movements in anesthetized and awake rats. a) A cloud of eye positions sampled for a 30 second duration while the rat is anesthetized. All the data points are bounded by a box defined by the axis limits that is 1 degree in X and Y directions. Variability includes natural movement of the eye and noise in the estimation method. b) The same data as in a) showing the horizontal and vertical position over time. Traces have been translated to reduce overlap. The Y axis indicates the difference in position between the center of the pupil (P) and the corneal reflection (CR) in units of camera pixels. These values are monotonically related to the gaze angle of the eye. An estimate of the local linear relationship between camera pixels and gaze angle is 2:1. c) An example of the eye drifting during anesthesia. The eye drops downward more than four degrees over the period of 15 seconds, before slowly returning. d) A horizontal saccade about five degrees. e) The same data as in d) showing the horizontal and vertical position over three seconds. f) An example where the eye position of an awake rat moves fairly often during a thirty second period. The sampling rate of all data in this figure is 1000 Hz.

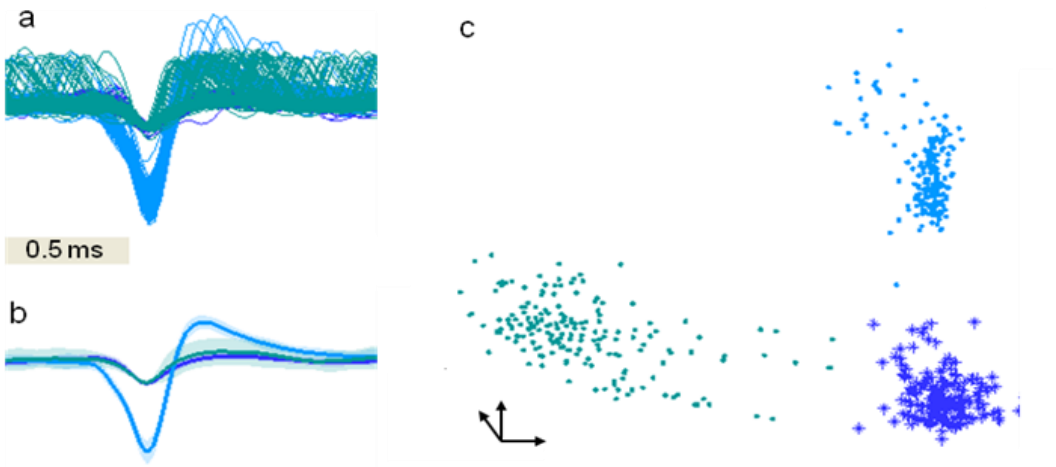

Figure 5.2. Detecting and clustering spikes. a) An example of the waveforms of an action potential recorded extracellular and bandpass filtered [200Hz 10kHz]. Each waveform is 1.6 ms long. The selection criterion for including a waveform was typically a negative going waveform. The threshold was adjusted for the conditions of each recording; in the panel shown, the threshold is -50mV. b) The average voltage of each spike cluster is displayed along with +/- 1 standard deviation. Clusters are initially determined automatically using KlustaKwik , which assigns each spike to a different Gaussian cluster using a variant of Expectation Maximization. Clusters are manually merged if they are nearby in feature space and they do not violate the refractory period as evidenced by the interspike interval histogram. c) Each spike is plotted as a single point in a three dimensional feature space. Often the features displayed are the first three principle components, though the algorithm typically is run with ten principle components of the waveforms. In the example displayed, the features of the waveforms were manually reselected for localized features that were less influenced by the variability of the CRT noise. For example, the y-axis corresponds to the peak amplitude of the spike. The green and blue clusters are deemed "noise" not belonging to the single isolated cell. The cyan cluster is considered a well isolated spike and is included for further analysis.

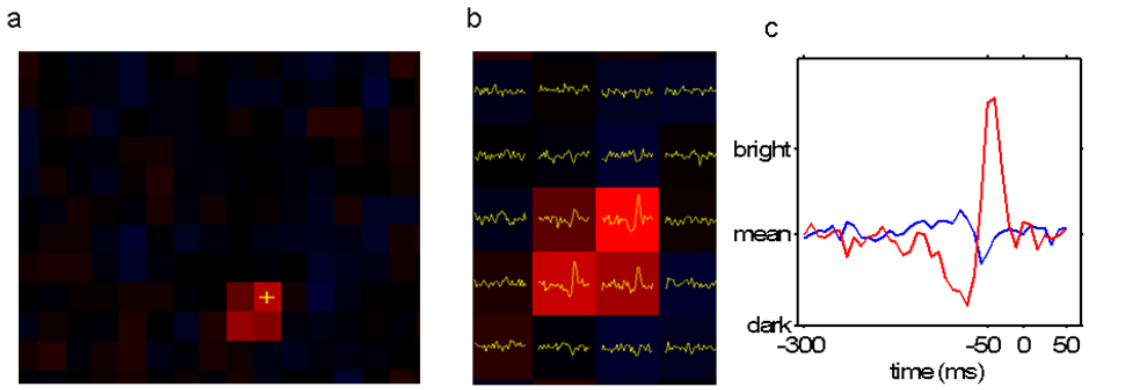

Figure 5.3. A localized LGN receptive field from an awake rat. a) A pseudo color image representing the average luminance of the stimulus 50 ms before a spike. The red region indicates that spikes were driven by an increase in luminance from a small patch on the screen. A flickering grid of squares was presented across the entire screen. Each square region subtended about 5 degrees to an edge. The visual cell had a high average firing rate of 22 Hz. The spike triggered average was generated with one minute of stimulus presentation that contained 1306 spikes. b) A detail of the responsive region including overlays of the temporal traces of the average stimulus for each location. The four central regions correspond to luminance regions that drive the ON center of the neurons receptive field. The surrounding region is weakly suppressive. c) The average temporal luminance calculated by averaging across the four squares driving the central region (red) and the 12 surrounding adjacent squares (blue). Data recorded by Balaji Sriram.

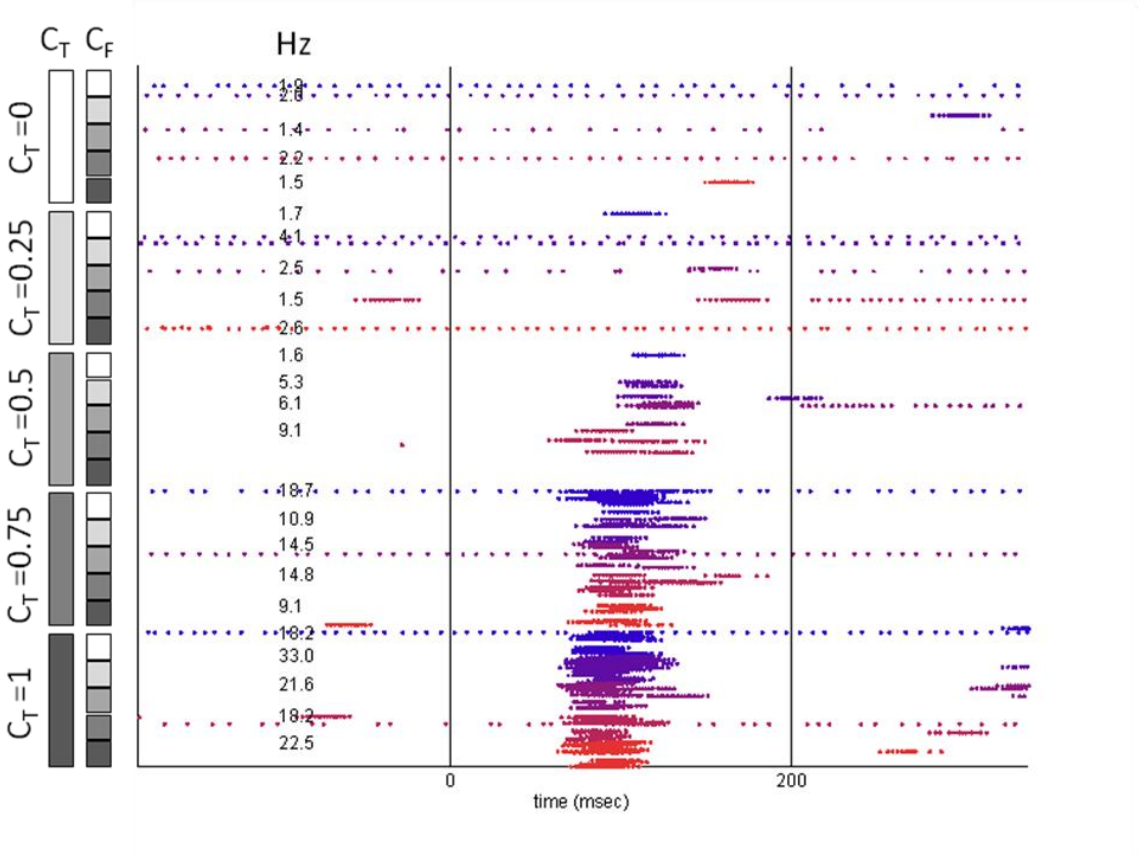

Figure 5.4. Response of an LGN cell while flashing gratings with varying target and flanker contrast. A single horizontal line of the raster corresponds to a single presentation. The vertical axis contains 48 repetitions of 25 conditions. Contrast conditions were randomly interleaved and are grouped for display purposes. The conditions are grouped by target contrast (C_T) at a coarse scale and flanker contrast (C_F) at a finer scale. Within a contrast condition, repetitions are sorted by temporal order of presentation. Additionally, to facilitate the boundary between contrast conditions, spikes are color coded to indicate flanker contrast; blue indicates zero flanker contrast and increasing redness indicates increasing flanker contrast. The stimuli onset is indicated by a vertical line at time zero and the offset is indicated by the vertical line at 200 ms. To summarize the volume of spikes elicited in each condition, the average firing rate during the stimulus is displayed for each condition. A graphical summary of the same data is available in Figure 5.7.

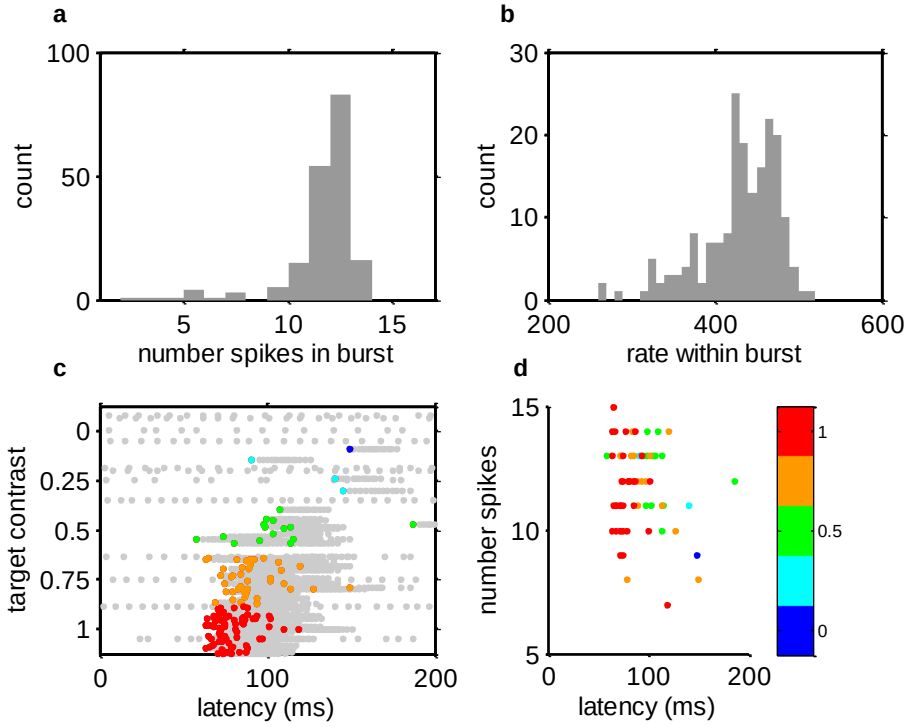

Figure 5.5. Properties of putative-bursts in an LGN cell in response to flashed gratings with varying target and flanker contrast. Data collected from a single cell in an anesthetized rat. a) A histogram of the number of spikes within a single burst. b) A histogram of the firing rate per burst, restricted to the time range of the first and last spike of each burst. c) The latency to burst onset sorted by trial condition. Each grey dot is a single spike. Each colored dot is the first spike in a burst. Time zero corresponds to the onset of the stimulus. The stimulus disappears at 200 msec. After the stimulus offset there was no additional visually locked response. Contrast conditions were randomly interleaved, and sorted for viewing purposes. The trials are color coded by the contrast of the target. Each target contrast condition also includes five sorted conditions of varying flanker contrast. This is not indicated by the color to avoid visual clutter, and also because it does not explain the variability of the data. d) The latency of the burst onset is plotted with respect to the number of spikes. Data is colored by contrast, same as figure c. While the latency of the burst changes significantly with each contrast, the number of spikes in a burst does not.

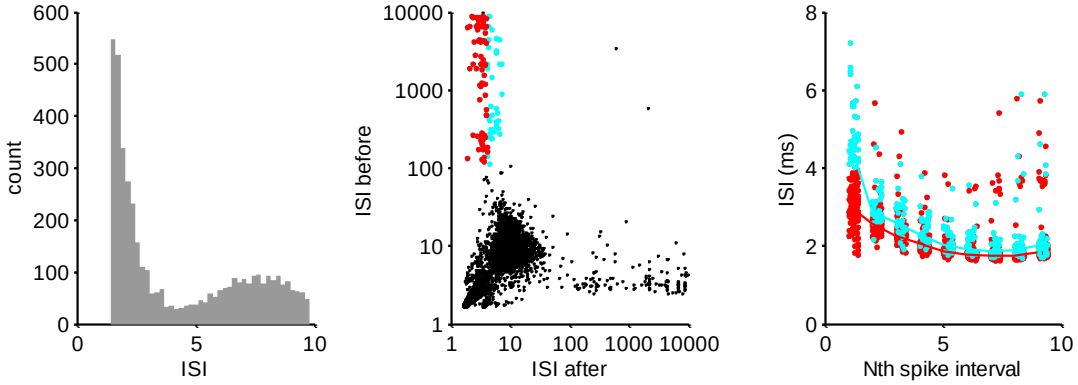

Figure 5.6. Interspike intervals of cell with putative Calcium bursts. a) Histogram of the interspike intervals. The panel focuses on intervals between 0 and 10 ms to clearly demonstrate the absence of refractory violations. The range of the plot causes 1.3% of the spike intervals to be excluded from view. b) A scatter plot of the same cell indicating the interval to the subsequent spike (ISI after on the x-axis) and the interval from the preceding spike (ISI before on the y-axis). Occasionally the cell fires tonically around 100 Hz (the dark cloud of points located near [10,10]). After long periods of silence, the cell typically responds with a rapid succession of spikes. Putative bursts are colored in red with the more stringent criteria (>100ms ISI before and <4ms ISI after) and colored in cyan for the additional spikes that are added using a more inclusive criteria(>100ms ISI before and <8ms ISI after). c) The latency between the Nth and the N+1th spike in the burst. Data were only included if they contained 10 spikes in a burst, which removed 7% of the bursts. Each spike is a single point. Horizontal jitter was added to the x-axis to assist in viewing regions of high spike density. Each spike interval was coded by the inclusion criteria: red indicates that the first ISI was less than 4ms, while cyan indicates that the first interval was between 4ms and 8ms. The red and cyan lines plot the median length for each interval in the burst.

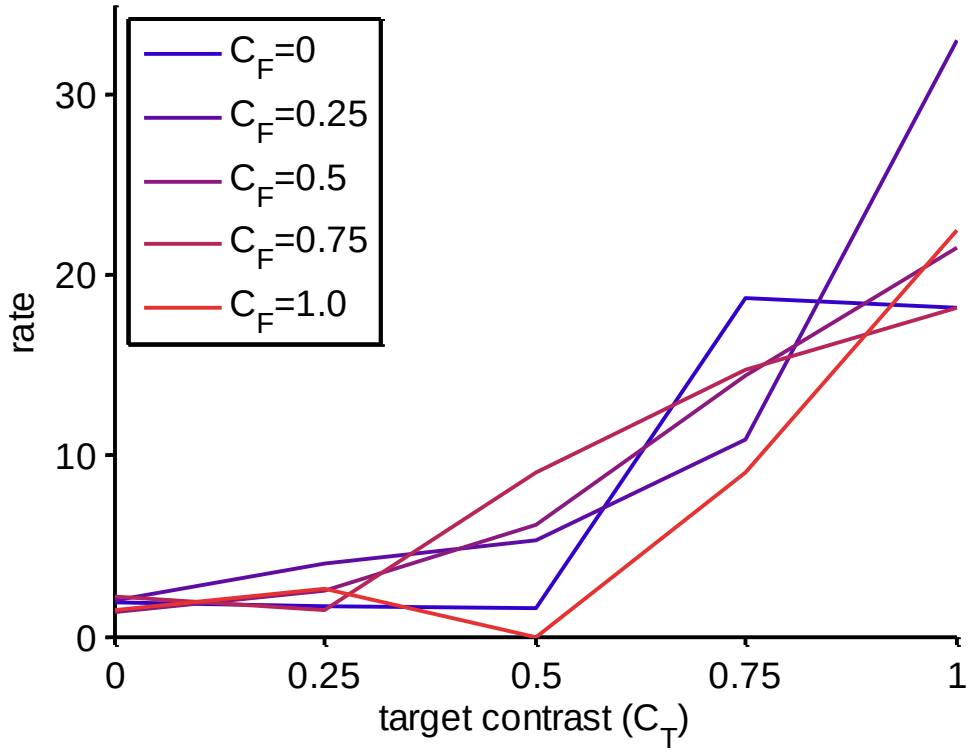

Figure 5.7. Sensitivity to the contrast of a target centered over the receptive field. The average firing rate is plotted against the contrast of the target. The contrast of the flanker is indicated by the color of each line. Error bars have been omitted to reduce clutter. Error bars are large due to the fact that some presentations elicit zero spikes, and other presentations elicit a burst of more than ten spikes.

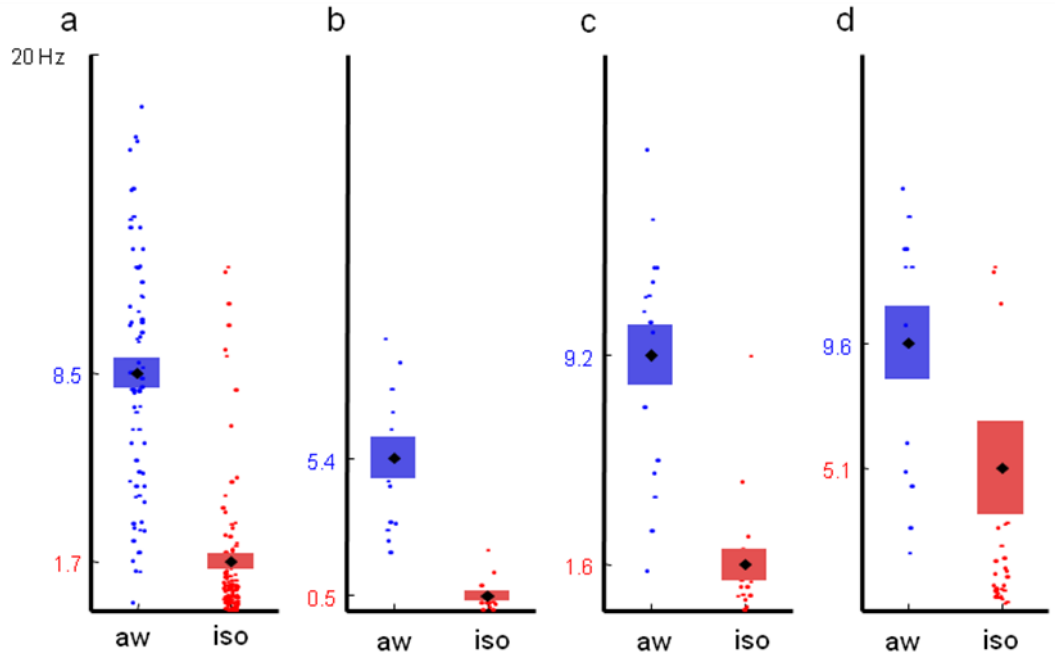

Figure 5.8. Average evoked firing rate is lower under isoflurane anesthesia compared to no anesthesia. Data were collected from two populations of LGN cells: one population recorded in awake rats (aw) and other from isoflurane anesthetized rats (iso). a) The evoked response combining all measurements. Each point represents the analysis of a stimulus type on a single cell. If multiple stimuli were presented to a cell, multiple data points are included. 34 cells were recorded from cells when the rat was anesthetized; 24 were recorded when the rat was not. Horizontal jitter is added to the data to better visualize dense regions. Error bars are one standard error of the mean. b,c,d) same convention, but only viewing the data from stimuli that were b) drifting gratings of varying spatial frequency c) spatiotemporal noise patterns of binary checkerboard, d) temporal noise of a full field Gaussian flickering stimulus.

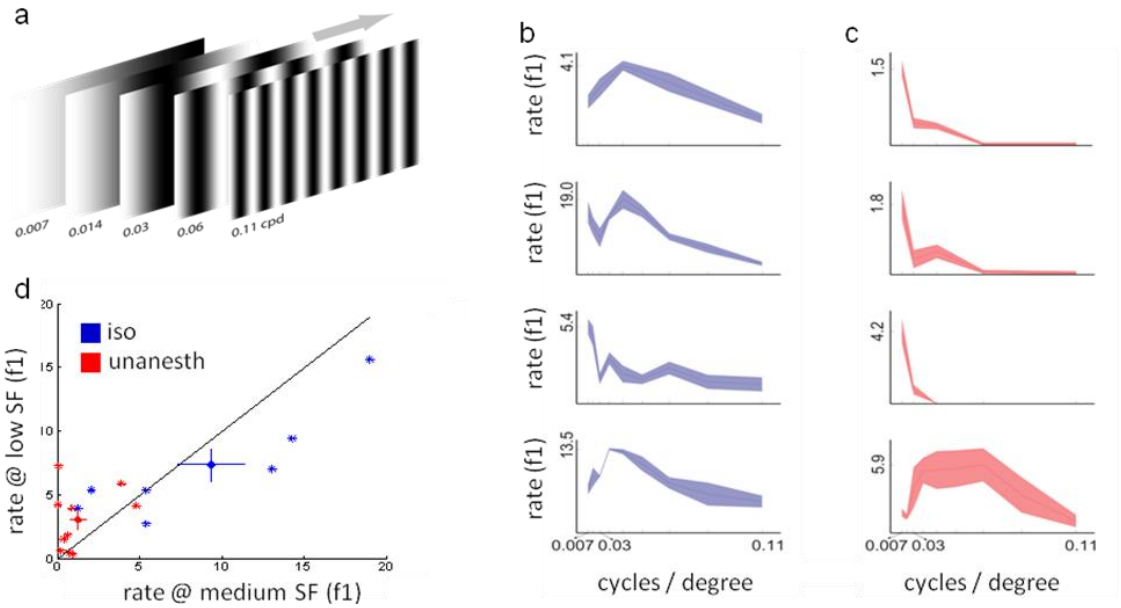

Figure 5.9. LGN cells recorded under isoflurane have less spatial selectivity. a) Stimuli consisted of vertical sinusoidal gratings that drift horizontally at 2Hz. Stimuli varied in spatial frequency including from 0.007 to 0.1 cycles per degree. Stimuli were presented repeatedly as 3 second long clips that were randomly interleaved. b) Visual response properties of four representative cells recorded in the LGN of awake rats. The vertical axis indicates the number of spikes at the fundamental frequencies of the stimulus (f1, 2 Hz). Most cells are bandpass as demonstrated by the peak response at intermediate spatial frequencies. Some cells (2/7) are low pass . c) Under isoflurane anesthesia, most cells are low pass and only a few are weakly band pass (3/10). A scatter plot indicates the rate of the response for low and intermediate spatial frequencies. Each symbol indicates the response of a single cell. The cross hairs demark the mean of the population and the standard error.

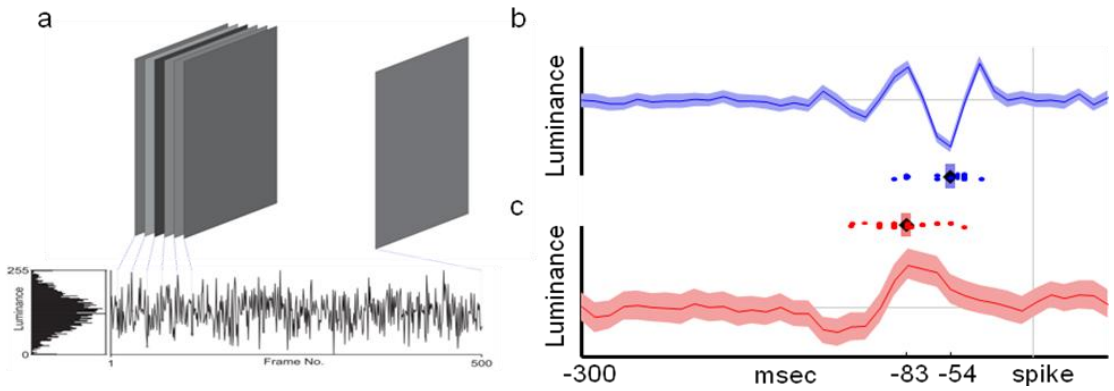

Figure 5.10. The peak response of the spike-triggered average has a longer delay for cells that are recorded under isoflurane anesthesia. a) The stimulus presented was a homogenous luminance across the entire CRT monitor. This luminance was updated every frame (100Hz) with an intensity proportional to a random draw from a Gaussian distribution. b) The average luminance preceding a spike, calculated for a single cell in an awake animal. The peak sensitivity of the cell occurs 60ms before the spike. This time is close to the population average in awake rats (54ms), indicated by the black dot below the waveform. The peak sensitivity of other cells occurred at times indicated by blue dots. The width of the bar indicates the standard error of the mean. c) The spike triggered average, as in b), calculated for a representative neuron in the isoflurane anesthetized animal. The peak sensitivity occurred at -80ms before the spike, which is close to the population average time (-83ms) of STAs recorded in anesthetized rats.

.

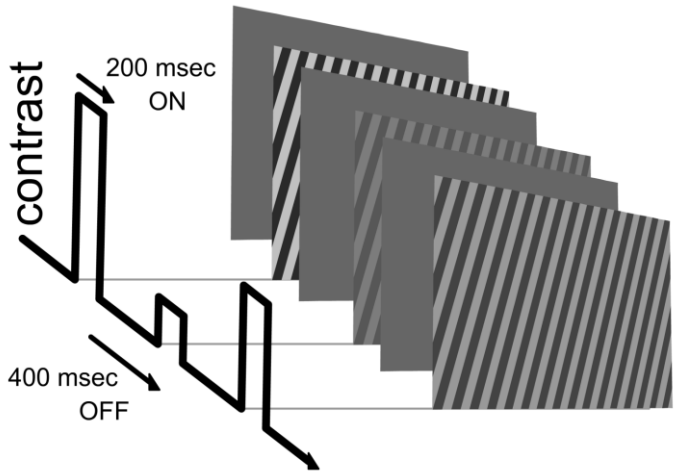

Figure 5.11. Flashed gratings that approximate the stimulus shown to other rats performing a behavioral detection task. Stimuli were matched for spatial frequency and duration but not for position. The contrast was varied randomly on each 200 millisecond presentation.

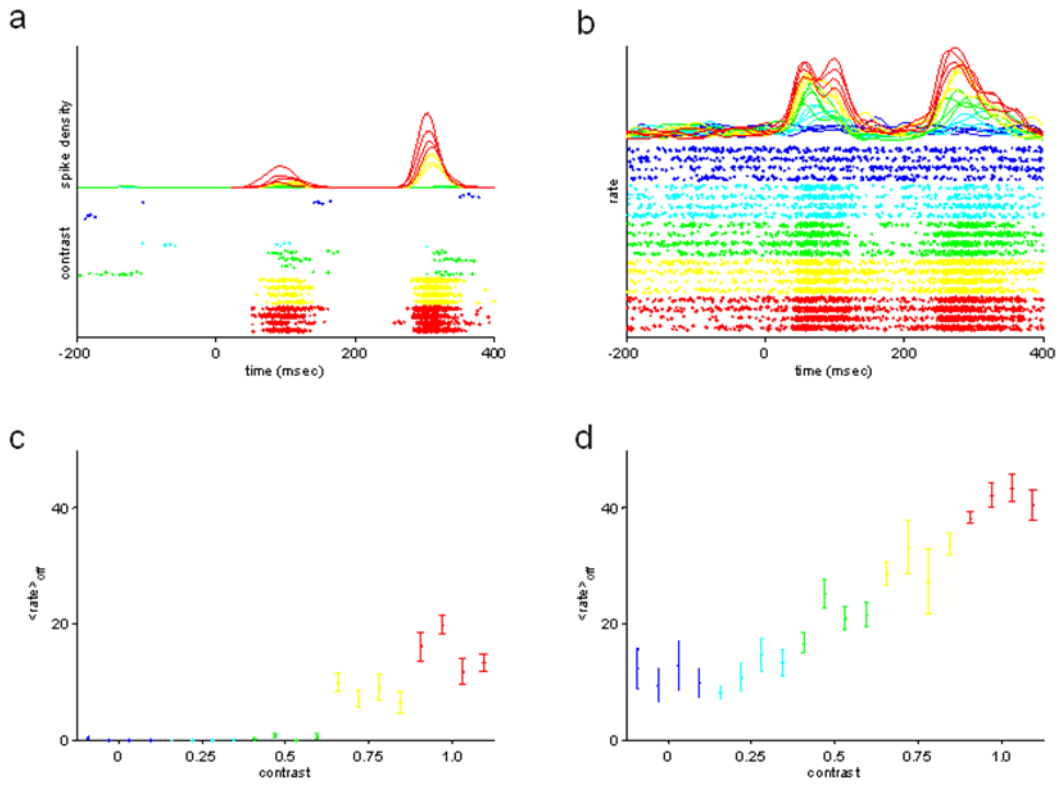

Figure 5.12. The firing rate increases when the rat exits anesthesia, increasing sensitivity to low contrasts. a) A spike raster sorted by contrast is presented in the lower portion of the panel. Above that is the peri-stimulus spike histogram calculated by smoothing the total number of spikes within a condition using a Gaussian window (σ=10 ms). b) The response of the same cell to the same stimulus after anesthesia was removed and the animal was waking. c) The average firing rate of the cell during anesthesia using the same data as in a). The firing rate was calculated during the 200 ms after the grating was turned off. d) The average firing rate of the cell while the rat was waking, using the same data as in b). The firing rate was calculated during the 200 ms after the grating was turned off.

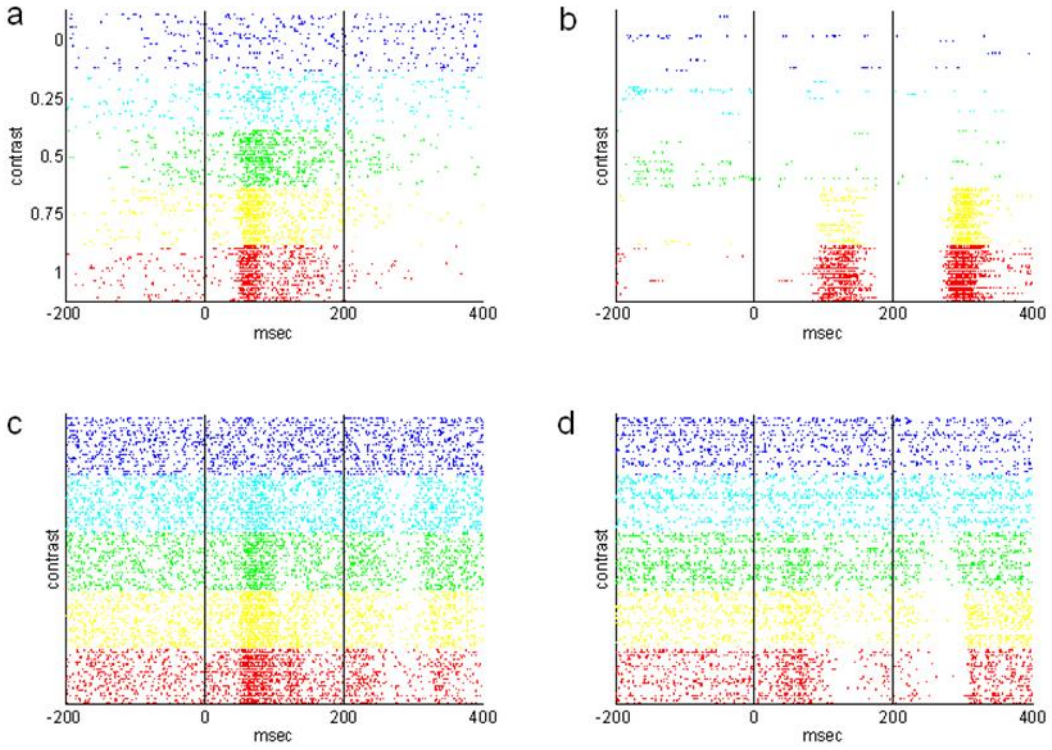

Figure 5.13. The response of four different cells in the LGN of an anesthetized rat to flashed gratings of varying contrast. a) A raster plot of multiple repetitions of the flashed grating to a single cell. Five contrasts were presented, with each contrast indicated by a different color. The stimulus duration is indicated by the black vertical lines. The grating is presented at time zero and removed at time 200ms. This cell has a response localized within 200 msec after the grating is presented. b) A different cell presented with the same stimulus produces a second volley of responses after the grating is removed. C) An example cell with a higher background tonic firing rate. The background firing rate is suppressed, especially after the stimulus is removed. d) An example cell that has a very mild increase in spikes for increased contrast. The primary influence is the reduction of spikes, both during the grating's presence and after its removal.

Chapter 5, in part, contains data presented at the conference Computational Systems Neuroscience (Cosyne), 2011, Sriram B.S., Meier, P. M., Reinagel, P. The dissertation author was a co-first author of this poster, the other author contributed equally, and the last co-author listed on the poster directed and supervised the research which forms the basis for part of this chapter.

Chapter 6:

Discussion

In this thesis I have asked and answered the question: do rats show behavioral evidence for pattern vision? I take "pattern vision" to be processing that is both sensitive to the relationship between features (like orientation) and the geometric relationship between features (like the relative position of the receptive field centers). I chose to study a flanker detection task because it is a minimal paradigm that has revealed pattern-specific effects in human psychophysics and in neurophysiology of early visual neurons in primates and cats. I chose to study the phenomenon in rats because of the long-term need to develop a small animal preparation amenable to experimental manipulations that are less feasible in large animal models. I found that rats do show a pattern-specific (collinearity) effect at the level of visual behavior.

6.1 Summary of findings

Rats are adept at a variety of visual tasks. I have described methods for training rats, and reviewed major experimental considerations for developing a behavioral training paradigm. Rats perform many trials and sustain stable performance within and across days of training.

Rats' ability to detect visual stimuli is impaired by collinear stimuli. The presence of any flanking stimulus impairs a rat's detection performance, but the impairment is greater for stimuli that are collinear. This impairment holds when comparing to flankers that have a different orientation than the detected target. The

collinear impairment also holds when comparing to parallel conditions in which flankers have the same orientation as the target, but collinearity has been disrupted by changing the angular position of the flankers. The rat's visual system is sensitive to the arrangement of stimuli near the location of a target, even when the surrounding stimuli are irrelevant for the subject's task.

Rats' performance on a detection task is influenced by the contrast of the target as well as the contrast of collinear flankers. Collinear flankers impair the detection task more at high contrast. This impairment is present regardless of the contrast; the presence of collinear stimuli did not facilitate the detection of the target at any contrast. A model based on signal detection theory captures the influence of target contrast and flanker contrast on the hit rates and the false alarm rates of the rats performing the detection task. The best fit models account for the rats' bias, use adaptive decision thresholds, quantify how much the effective flanker contrast influences the decision, and invoke nonlinear interactions between the flanker and target by means of a common normalization pool.

Visual stimuli evoke lower firing rates from thalamic neurons while the rats are under isoflurane anesthesia compared to rats that are awake and passively viewing a stimulus. Under anesthesia spatial frequency tuning is more likely to be lowpass. Rats saccade when awake. They also move their eyes under anesthesia; however, movements are slower and less frequent. Both kinds of eye movement provide challenges for localizing receptive fields in order to present stimuli beyond the classical receptive field. Preliminary results suggest that flashed stimuli in the

surround may not engage divisive normalization. However, results are inconclusive due to the low yield, and the differences in the awake and anesthetized conditions.

6.2 Significance of collinear effect

What is the significance of the fact that rats are selectively impaired by collinear stimuli? The specificity of the impairment provides evidence for the specificity of the rat's visual system. It provides a clue about how the visual system of the rat processes patterns that span multiple receptive fields in primary visual cortex. I consider two possibilities for how the flankers have an impact of target detection.

1. The flankers change the activity of the neurons that represent the target as localized image features. The rats solve the task by using the intensity of these local features. Flankers change this intensity. I refer to this as a "low level explanation."

2. The flankers and the target combine into a higher order patterns. Rats solve the task using their higher order representation which need not be spatially localized. The presence of flankers influences which higher order responses are engaged, and not all are as sensitive to the presence or absence of the target. I refer to this as a "high level explanation."

The difference between these two explanations depends on which neurons are used by the decision process. Explanation of flanker effects tend to focus on low level vision. Surround processing phenomena observed in low level neurons are similar to the nature of human psychophysics. It's possible they might be sufficient to explain effects of pattern processing on detection. Additionally the features involved in low level vision are better understood. There is more data and more parsimonious testable models. It is more tractable. However, it is possible that low level interactions are not the only reason why nearby stimuli impact target detection in a flanker experiment. This is particularly true in experiments with trained subjects that are not instructed to use a low level cue. It might be that higher level representations are more readily learned.

The behavioral data does not allow us to determine if the specificity of the collinear impairment is caused because of low level visual processing or high level visual processing. Physiological experiments are required to address whether the low level explanation are adequate to account for the data.

6.3 Low level visual explanations

Flankers may activate the surround processing of neurons, which would alter the neuron's response. Under this conception, the neural representation of the target is treated as an information channel for the subject's behavior, and the presence of a collinear flanker corrupts this information channel more. For example, during the detection task the response of the neurons encoding the target could be more

suppressed by collinear stimuli, increasing the probability that the rat would confuse the weak neural response for background neural activity. In this case, the specificity of behavioral responses must be shared by the surround processing. That is, the functional connectivity from nearby neurons it the visual system should be stronger not only because of shared orientation preferences, but also because of a particular geometric relationships of the receptive field centers with respect to the orientation preference. Collinear relationships should particularly influence surround processing. Which neural connections would have this specificity is not clear: in principle it could be horizontal connections within a visual area (Chisum et al., 2003), feedback connections from higher cortical areas (Angelucci, Levitt, Walton et al., 2002), or feedback connections from cortex to thalamus that modify a neurons support (W. Wang, Jones, Andolina, Salt, & Sillito, 2006). For the behavioral effect to be explained by surround processing in the early visual system, the relevant synaptic contacts would also be in the primary visual cortex or earlier.

This is an intriguing hypothesis because the rodent visual system lacks smooth orientation domains and the interneurons in primary visual cortex of rats may lack orientation sensitivity. Many species, including primates and cats, have neurons with orientation response properties that are spatially organized (Gilbert & Wiesel, 1989; Hubel & Wiesel, 1968). Thus, two neighboring cells have a high probability of encoding similar orientations. The function of orientation columns is not known, but they are many theories that explain how they might be useful (Ben-Shahar, Huggins, Izo, & Zucker, 2003; Worgotter, Niebur, & Koch, 1991). For example, consider the

challenge of wiring neurons with similar orientation preferences together into a sub-network. A neuron's proximal dendrites could automatically select for similar orientation properties, without any additional targeting mechanism. Rats lack smoothly varying orientation domains in primary visual cortex (Ohki et al., 2005). Thus, if rats' neurons make preferential contact between other neurons of shared orientation tuning, this requires some additional mechanism to establish the connection. Evidence from a small network of anatomically traced cells in mouse cortex reveals that at least some inhibitory neurons are non-selective for the orientations of which excitatory neurons drive them (Bock et al.). This anatomical evidence is consistent with physiological reports that interneurons in mouse visual cortex are activated by drifting gratings, but are not very selective to the orientation of the gratings (Kerlin, Andermann, Berezovskii, & Reid). The lack of orientation selectivity in interneurons may also be true for other rodents like rats. Future evidence about interneuron specific neural responses will likely continue to be advanced in mice due to genetic labeling of interneuron types.

Could surround processing maintain selectivity in such a circuit? Yes, it could by virtue of selective connections formed by orientation-tuned excitatory inputs, and the strength of un-tuned or weakly-tuned inhibition. Interestingly, a plausible way to connect such a network would be based on the statistics of the input: connections between excitatory neurons could be learned by their co-activation. A single mechanism that might establish orientation specific processing would also be sensitive to collinear specific processing because collinear contexts are more frequent in natural

scenes than other comparable spatial patterns. Despite their lack of orientation columns, squirrels have single cells in primary visual cortex that display length-summing and end-stopping similar to other visual mammals (Van Hooser et al., 2006). It is less clear how interneurons could adapt their weights to be context sensitive. It will be important to understand the specificity of interneuron's selection of excitatory targets. In summary, it is possible that rats maintain context sensitive surround processing in primary visual cortex, despite lacking orientation columns and possibly orientation tuned inhibitory neurons. Center-surround interactions have been observed in the visual cortex of anesthetized rats which including cells that have orientation independent surround suppression, and other cells where the suppression that is tuned to the disparity between center and surround orientation (Girman et al., 1999). These recordings involved drifting gratings with differing motion between center and surround. The surround was not investigated with the positional sensitivity to evaluate the role of collinear specific processing. Future physiological work is needed to address the specificity of the observed behavioral impact of collinear flankers and to better understand surround processing in the rodent visual system.

6.4 Higher level visual explanations

Alternative explanations can account for collinear specific deficits in visual processing without appealing to the early visual system. Early in the visual system, spatially separated responses are likely to engage different neural populations. However, at later stages of processing, receptive field sizes increase. Consider a

hypothetical neuron that sums the nonlinear outputs of the many neurons in the early visual system. While this neuron contains information about the target's presence, it is intertwined with information about the flanker's presence. If the rat's decision system limited to only have access to such higher order cells, then the ability to detect a target would be better described as the ability to discriminate between two different patterns of activation. The flankers alone activate one such pattern, and the flankers plus the target activate a different pattern. Using this framework there are two core explanations to why a particular discrimination is harder: either the responses are more similar, or the responses are more variable within a condition.

Not all patterns are of equal importance to discriminate between. The visual system has particular tasks to solve: object identification, visual navigation and so on. It is useful to remove the image variability that is irrelevant for behavior, enabling higher order features to be more invariant. This mapping from lower order to higher order features will not treat all differences the same.

If we speculate for a moment that a good higher order representation is one that groups stimuli with similar underlying causes, then it is possible the collinear condition is most difficult to detect because the stimulus is likely to have come from the same underlying cause, whether or not the target is present. Other conditions, contingent on the target, might have greater differences in the underlying causes that would typically explain the stimuli in a natural setting. Granted, this explanation is speculative, as we do not know the mapping from causes to stimulus patterns, nor that

this is a valid description of the similarity space for higher order representations, nor that rats decisions are limited to acting on higher order representations.

Possibly the presence of a collinear stimulus increases the noise in a higher order representation. Flankers that have a collinear arrangement may be more likely to be bound together. The binding of visual features is not well understood; there are many ways it might be implemented. Let us consider a concrete example. Suppose that the target and flanker are bound together by a process that entrains their spiking response with a mild shared inhibitory stimulus. Even if this process serves a useful function for interpreting a scene, the additional computations would modify the neural response and act as a noise source from the point of view of detection. A similar explanation applies to edge specific processing. This line of reasoning suggests that specialized processing for structure in an image will degrade weak visual signals. Most visual stimuli are higher contrast and could tolerate some amount of noise to enable the benefits of the computation for grouping, segmentation, and other subtask which may be present in the visual processing of scenes.

6.5 Flankers and divisive normalization

The response of neurons in primary visual cortex have non-linear activation functions that can be modeled by normalizing the output of each neuron to its neighbors activation. That is, the response magnitude of a neuron can be fit by a linear filter response that is divided by the summed activity of other linear filters that represent the input drive to other neurons (Heeger, 1992). This phenomenological

description of neural responses is consistent with an optimality principle in which the visual system reduces the redundancy between neural responses. Theoretical approaches that impose sparseness or maximize the independence of neuronal responses to natural images learn linear weights with properties that are similar to receptive fields (Bell & Sejnowski, 1995; Olshausen & Field, 1996). These optimality principles yield linear filters that are bandpass and oriented. However, the linear filter responses are far from independent. Notably, the variance of any two nearby filters are correlated, sometimes referred to as a "bowtie" dependency. This conditional dependence in the variance could be reduced by allowing the activity of a neuron to impact the gain of other neurons in the network (Simoncelli & Schwartz, 1998). In other words, the theory predicts the relative weights contributing to a neurons "normalization pool." A neuron that is more predictive of it's neighbor, for example, by virtue of its proximity or by sharing a common orientation, would have a stronger weight assigned to it by the optimization procedure. The weights that are learned by optimization procedure are consistent with physiological measurements of the contrast independence of orientation tuning, spatially overlapping masks, and masking stimuli in the surround of a central driving stimulus (Schwartz & Simoncelli, 2001).

If the goal of divisive normalization is to reduce the redundancy in the neural code, what does this predict about neural responses to a target with flankers? First, the presence of any flanker should engage the normalization pool by activating other neurons in the network. According to the theory, this is sensible because the targets response would have been redundant. Independent signals are useful because they

maximize the amount of information being transferred in a channel for a given set of constraints, like energy consumption. Also, independent signals may be computationally useful for discovering the independent features of a visual scene. Another way of looking at it is that a strong activation in the surround indicates that there is a high probability that the response in the center will be saturated, and thus the neural encoding of the target will be less informative. Reducing the gain of the target region would shift the response curve such that the most likely contrasts are aligned with the steepest slope of the transfer function. Any deviation about the expected value will be encoded with higher fidelity. Collinear flankers should contribute a greater response to the normalization pool than other flanker conditions because they are more predictive of oriented power at the target's location. Notice that such a non-linear lateral interaction would be useful on average in natural scenes. However, under experimental conditions, the spatial patterns and the contrast of the flanker were independent from one another. Thus, the presence of a flanker, rather than increasing the information about the target, could decrease it. A target that was difficult to detect, and below perceptual threshold on a fraction of the trials, would be perceived less often.

A theoretical account of divisive normalization, motivated by increasing the independence of the neural responses, is consistent with the impairments observed in rats behavioral performance. However, the same theory is hard pressed to account for phenomena observed in human observers. Namely, for some species and contrast conditions, the presence of flankers can improve detection of low contrast stimuli.

Additionally, at least some neurons demonstrate increased firing rates when collinear stimuli are presented outside of their receptive fields (Chen et al., 2001; Kapadia, Ito, Gilbert, & Westheimer, 1995).

The proponents of divisive normalization have argued that facilitory effects may be the result of the release of inhibition. In other words, a given cell may have a tonic level of suppression from other cells, and the suppressive effect on these other cells would yield an observed increase in responsiveness for a given neuron. While this is technically true, and a way that facilitory effects could be explained, it misses the point of the criticism. Divisive normalization does not proclaim to be only the result of monosynaptic effects in the network: it is a mathematical description of physiological responses that contain the influence of the whole network. Neural evidence of increases in the firing rate cannot be explained by the given mathematical description, suggesting that there are also other forces at play.

One possibility is that a fully descriptive theory will have to include a range of effects for different kinds of lateral interactions. For example, particular lateral interactions may be sustained or transient, target distal or proximal dendrites, or have different ratios of excitation and inhibition. The normalization model summarizes the aggregate effect of many different mechanisms. Certainly the models of individual cell responses, and possibly psychological models, will have to account for the variability in the types of lateral interaction. In short, not every contextual effect may be explained by a suppressive phenomenon. Some cell types or stimulus events may lead to excitatory interactions. Efficient information processing is a computational

principle that influences the structure of early sensory processing, but in some instances the value of a salient response, or rapid computation, may outweigh the benefit of independent information channels.

6.6 Why flankers might facilitate

A classical example of facilitation is when the addition of a pedestal of activity increases the sensitivity to the intensity of particular stimulus. Traditionally, facilitation was accounted for by the non-linearity of neural responses to local contrast (Foley, 1994; Legge & Foley, 1980; Nachmias & Sansbury, 1974). The non-linearity was invoked to explain how detection of a target can be facilitated by weak mask contrasts and impaired by stronger mask contrast. At low contrasts, the response function is accelerating, and additional contrast will increase the response to region of greater sensitivity. The greatest sensitivity occurs when the transfer function has maximal slope. Beyond this contrast, the slope decreases, and increasing the contrast of the mask will reduce the sensitivity to the target. Models that get their explanatory power from the shape of the non-linear function are collectively referred to as transducer models because they transduce contrast into an internally represented decision signal.

Traditionally transducer models were applied to masks or pedestals that spatially overlapped the location of the target. A lateral stimulus could produce the same effect if the stimuli in the mask and the target contributed to the same non-linear response function. Possibly a high contrast flanker could mimic the impact of a low

contrast pedestal. A large receptive field could produce such an effect (Solomon, Watson, & Morgan, 1999). However, experiments in humans that combine pedestals with flankers show that the function of threshold detection versus contrast does not saturate as expected if the pedestal effect were the same as the flanker effect (Chen & Tyler, 2002). The authors suggest that the influence of flankers is better explained as multiplicative gain change acting upon the target response.

Finally, a very different reason has been offered for why flankers may facilitate: their presence can reduce the uncertainty of the spatial location of the target (Petrov et al., 2006; Wu & Chen). A detection task requires more than monitoring the intensity of a single channel; it also involves knowing which channel to monitor. Since visual channels are distributed over space and the stimulus is localized, the detection task will be more difficult if the subject had more uncertainty about the location of the target. If the target could appear anywhere, then the noise form a large number of channels can result in an increase in false positives. In the particular task that my rats were performing, the target always appeared in a reliable location, reducing the spatial uncertainty a large amount. However, the rats head position and gaze will always have some variability and the exact neurons encoding the response will differ in each trial. The remaining uncertainty could be further reduced if the flankers were used as a spatial cue to indicate the location of the target. If the rats had demonstrated better performance in the presence of flankers, this might have played a role. However, since rats never exhibited an improvement in performance with

flankers, we cannot tell if the beneficial effect of uncertainty reduction were merely outweighed by other effects, or if they played no role at all.

6.7 Bursts in detection

The relay cells of the lateral geniculate nucleus have two modes firing: one the produces tonic spikes and one the produces burst of action potentials. Burst were observed in Hubel and Wiesel's initial recordings from the thalamus (Hubel & Wiesel, 1961). Burst are prevalent during slow wave sleep when the relay cells are hyperpolarized (Livingstone & Hubel, 1981). Sufficient hyperpolarization de-inactivates Ca-channel, enabling a burst on subsequent depolarization (Jahnsen & Llinas, 1984).

It is unknown whether or not bursts play a role in active perception. It has been argued that burst are not useful for visual perception (Steriade, 2001). Burst are present in awake animals, but are rare compared to the tonic spike response (Weyand, Boudreaux, & Guido, 2001). Studies in anesthetized animals show that visual stimuli can control the priming of T-type calcium channels (Denning & Reinagel, 2005), suggesting that burst could be visually informative in awake animals. Additionally the amount of information convey by a single burst is higher than a single tonic spike (Reinagel, Godwin, Sherman, & Koch, 1999). Encoding stimuli with burst is less metabolically efficient, but it may be a good strategy for stimuli when the signal to noise is particularly low.

Detecting a low contrast target may be easier for an organism when there is a presence of bursts in the visual response. Neurons were not recorded while a rat was performing a detection task, so it is not known if an alert rat performing a task will have neurons that will respond in a burst mode. Here I speculate about the conditions requiring de-inactivation of the T-type calcium channels. The period preceding a target's onset was a mean luminance screen, pending of course the rat's eye and head movements. For free moving animals it is likely that a mean screen could precede the stimulus for hundred of milliseconds, but probably not more than half a second. Thus it is possible that contrast modulation was not recently present in the neuron's receptive field, and an extended period of time might have accrued since the most recent spike. These conditions favor the generation of a burst. Of particular importance is the fact that the rat initiates an action which delivers the stimulus. Thus, the subject is in a state of expecting a possible stimulus. It would be informative if future experiments tested for the arrival of inhibition in the thalamus at this time, possibly from the perigeniculate or the brainstem (Murphy, Uhlrich, Tamamaki, & Sherman, 1994; Sherman, 2005). Expectation triggered inhibition could allow the thalamus to prime calcium channels and transmit more information by virtue of a burst event. Bursts elicit post-synaptic action potentials more effectively in layer 4 of cortex, though in part this may be due to releasing corticothalamic synapse from sustained depression during the interval preceding the burst (Swadlow, Bezdudnaya, & Gusev, 2005).

Could flashed gratings both prime and trigger bursting in the thalamus? This seems unlikely due to the temporal requirement of the Ca-channels and the transient nature of the stimulus. However a closer consideration of the neural response reveals that it is possible: a flashed stimulus produces a volley of neural activity for both the onset and offset of the grating; both are potentially useful for detecting the grating. Inhibition from the onset could be responsible for hyperpolarizing a cell and enabling a burst response to the offset. For humans performing detection tasks, the stimuli are often presented for less than 100 ms. In the experiments I performed, they were present for exactly 100 ms. Rats viewed stimuli for 200 ms in most of the experiments I performed. In both cases, there is time for the onset and the resulting network effect to potentially de-inactivate Ca-channels. Some individual cells produced reliable burst following onset and burst again following the offset. However, in these recordings the Ca-channels may have been primed non-visually by virtue of the anesthesia. This uncertainty about the capacity for the LGN to prime bursts emphasizes the importance of recording neurons from rats that are awake, ideally performing the perceptual task. Even passive viewing rats and alert behaving rats may have very different baseline states, which would affect the resting membrane potentials of thalamic cells, the burst probability and the contrast sensitivity of neural responses.

Appendix A:

Natural scene statistics

A1.1 A theoretical interpretation of collinear impairment

I speculate that the performance impairment observed in rats is the consequence of visual processing that would optimize performance for natural scenes, at the expense of impairing performance in tasks that violate natural scene statistics. Specifically I suggest that the visual system computes a prediction of target probability based on some function of the image at other locations, and this prediction is used to adjust the local representation of target. The theory implies that those surrounds that make the strongest predictions about the target in natural images should influence the representation most, and thus impair performance most in our task. This theory is neutral about whether predicted targets should be suppressed (reducing redundancy) or enhanced (propagating beliefs) at the level of early vision. The framework of normalization developed above implies suppression of predicted features.

In principle, every feature at every location in an image contains some information about the target probability, and complex higher order interactions may be required to fully describe the dependence of target probability on the surrounding image as a whole. Here I consider a much simpler relationship. Suppose the visual system only has access to the separate pair-wise correlations in strength of nearby local oriented features (e.g. the correlated firing of V1 neurons). Further suppose that

each local oriented feature's representation is normalized by the activity of all nearby local oriented features, in proportion to their statistical co-activation in natural images. Such connectivity could plausibly be learned by activity-dependent mechanisms, and does not presume orientation columns (Ohki et al., 2005). This simple pair-wise first-order correlation scheme is unlikely to capture all the informative statistics of natural images (Zetzsche & Nuding, 2005). Yet, perhaps surprisingly, it is sufficient to explain the unique status of collinear flankers in our task. In the statistics of natural images, contrast at one location is correlated with high contrast nearby (Ruderman & Bialek, 1994). The orientation at one location is correlated to the orientation nearby, and collinear features co-occur most often (Geisler et al., 2001; Sigman et al., 2001), suggesting that collinear features should normalize the most (Geisler et al., 2001; Schwartz & Simoncelli, 2001).

To illustrate this point, we show the probability of a feature matching our target in a natural image database, either at random locations, or in locations adjacent to a feature matching one flanking patch (Figure A1.1). As expected, the presence of any of our flanking features increases the probability of the target being present. Collinear flankers consistently increased the probability more than other flankers, regardless of the arbitrary threshold for considering a feature "present". This analysis did not disambiguate among the non-collinear flankers because their relative ranking depended on choice of threshold. All these observations hold true if we conditioned on both flanking patches (not shown).

Thus one can think of the normalization strength in our model (E) as representing a predicted contrast at the target location ($\hat{C}_T$) estimated on the basis of the contrast in the flanking region (C_F). This suggests that the function of divisive normalization is to reduce the effective contrast for expected features, and amplify unexpected features, thereby maximizing information transfer for natural scenes (Ruderman & Bialek, 1994).

Future studies could further test this ecological interpretation by correlating behavioral impairment with natural co-occurrence statistics of flanking features at other positions and orientations. The more specific hypothesis of normalization makes the direct prediction that flanker features that are correlated with the target in natural images should reduce the firing rate of neurons that respond to the target early in the visual system (such as V1 or LGN).

A1.2 Methods used for natural scene statistics

For the analysis shown in Figure A1, I used images of natural scenes from Correl images, as selected by the first two hundred in the Berkley scene segmentation database. Color images where transform to grayscale with Matlab's rgb2gray, but otherwise un-normalized. The estimate of the local oriented target contrast was calculated as the dot product between the image (I) and an oriented filter (f) matched to the stimuli viewed by the rats. The squared output values were used as the feature contrast:

$$C_{feature} = (I_{x,y} f_{feature})^2$$

Target contrast locations were sampled from an evenly spaced grid spanning the center of each image. This allowed all center and flanker samples to be free from edge artifacts, and avoided the oversampling that would be induced by including immediately adjacent pixels. This resulted in two million samples of the target center and two million estimates for each of the four flanker locations tested. A target was deemed present if it was in the top 10% of target contrasts. This defined the arbitrary threshold contrast C_{thresh}. The same threshold was used for defining the presence of a flanker feature.

$$p(target \mid flanker) = p(C_{target} > C_{thresh} \mid C_{flanker} > C_{thresh})$$

I do not mean to imply the binary threshold of presence or absence is employed by the visual system. This simplified analysis parallels the logic of experimental paradigm, and is sufficient to communicate the trend that collinear flankers more strongly predict the target than the other flanker types I studied.

For Figure A1.1, I selected the image and target location that best matched the collinear flanker stimulus based on the strength of both flanker positions. Specifically, I multiplied the amplitude of the filter responses from the two flanker locations, resulting in a comparable distribution of values to the target contrast (square of the filter output at one feature location). I also repeated the analysis in Figure A1.a and

A1.c conditioning on both flanker positions by requiring the product of the two flanker feature amplitudes to be above C_{thresh}. Results were qualitatively similar: the conditional target distribution was shifted to higher contrast by any flanker condition compared to the unconditioned distribution; the collinear flankers shifted the distribution more than any other flanker type regardless of choice of threshold; the other three flanker conditions had similar effects and their rank order depended on choice of threshold.

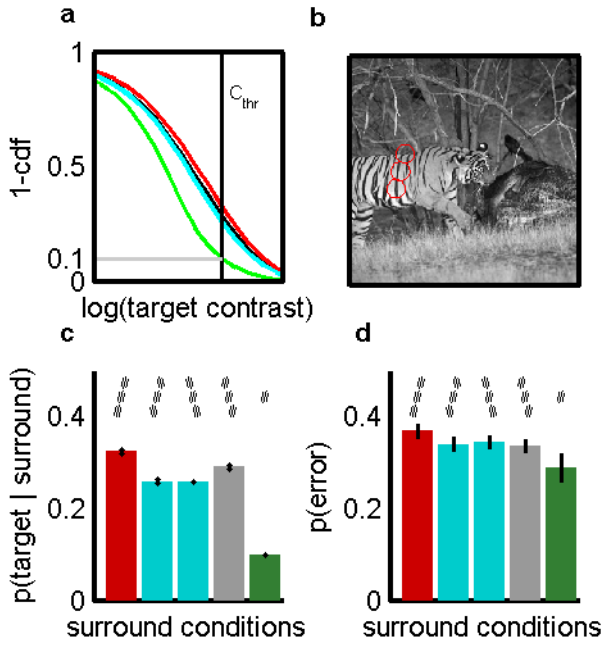

Figure A1.1. Co-occurrence of flanker and target features in natural images. I analyzed natural images with grating patches identical to the target and flanker features seen by rats. The match (dot product) between a local image patch and a feature is taken as the "contrast" of that feature at that image (See Appendix A1.2). a) Negative cumulative distributions of feature contrast for a clockwise-tilted target. The green curve shows the fraction of image positions with target contrast greater than a given contrast. For the purpose of this illustration I set an arbitrary contrast threshold (C_{thresh}) such that $C > C_{thresh}$ for 10% of this unconditioned target distribution. Contrasts higher than threshold were categorized as "feature present" and lower contrasts "feature absent". Remaining curves show the distribution of target contrast conditioned upon the presence (above-threshold contrast) of one flanker patch at the lower flanker position. All conditioned curves are shifted right, confirming that the presence of any flanker increases the probability of the target feature. Collinear flanker patches (red) shift the distribution more than pop-out (cyan) or parallel (black). b) The image in the database containing the strongest match to both collinear flanker features (top and bottom circles) also contains an above-threshold collinear target (central circle). c) The probability that a target is present, given that the surround feature was present. Symbols indicate result for the upper and lower flanker positions; bars indicate the average of these. d) Errors made by a representative rat on the different conditions tested. The worst performance is observed on the collinear condition, which is most correlated to the target's presence in natural images.

Appendix B:
Human psychophysics consistent with literature

It is natural to wonder if rats display the same behavioral responses to flanker stimuli as humans. Interestingly, humans also display collinear specific processing, but the influence of collinear stimuli is different. Compared to other flanker conditions, collinear flankers facilitate target detection by humans. In rats, collinear flankers impair target detection.

Is it possible that the difference between the phenomena I observed in rats and the literature reported in humans was based on the different experimental paradigm? Unlike most studies that use gratings with sine waves, I used square wave gratings. Also, I did not repeatedly test the same spatial pattern in isolation: all spatial patterns were randomly interleaved. This control prevents subjects from exploiting feature based attention to isolate the target because the orientation of the target is unknown on each trial. Also, the variable angular position of the flankers might prevent the flankers from acting as a useful spatial cue. Finally, rats learned the task without instructions, while humans usually receive a verbal description of the task.

I repeated the experiment in humans, matching the experiment performed in rats as close as possible. The stimuli were square wave, all spatial patterns were randomly interleaved, and the humans received no verbal instruction about the task other than to try and maximize the number of correct trials. Thus, humans also learned the task.

Rats and humans have different absolute thresholds, so it was necessary to change some of the parameters to make the task harder for humans. Compared to the

rats, humans viewed stimuli at a farther distance (2m vs. 10cm), lower target contrast (2.5% vs. 100%), and shorter duration (100ms vs. ~1 second). With these conditions, both humans and rats saw stimuli near their optimal spatial frequency sensitivity (Fig. A2.1). A single contrast was used within species. This brought all subjects into a range where behavioral differences were observable.

The previous literature predicted that performance in humans would be higher for collinear stimuli than non-collinear ones. My results agreed with this prediction (Fig. A2.2b, A2.3b). Here I show the result comparing collinear to pop-out. The effect also holds in the comparison to the parallel condition (Fig. A2.4). A summary of performance and bias yield the same qualitative results using either percent correct (Fig A2.2) or the signal detection theoretic measures of d' and criterion (Fig A2.3). Collinear stimuli facilitate detection in humans and impair detection in rats.

Humans display large differences in performance that required fewer trials to observe. Also learning was much faster. Rats required more trials to be able to observe the difference in performance between the stimulus conditions. While the changes in performance provide a useful summary for the impact of the stimulus (Fig. A2.2b, A2.3b), the absolute performance provides useful context (Fig A4).

The difference in the influence of collinear stimuli in rats and humans cannot simply be explained by difference between experimental paradigm or stimulus pattern, because both were matched across species. One objection is that the contrasts were different between the two species. This difference was necessary to measure threshold perceptual sensitivities. If the contrast were the same, then either human performance

would saturate near perfect or rat performance would be remain at chance. The alternative is to explore a range of contrast within the visible range of a species. Many target contrasts were presented to the rats, and there was never evidence of collinear facilitation with respect to pop-out. Additionally, collinear stimuli that varied both in target contrast and flanker contrast were tested in rats (Chapter 4). In no case did the addition of flankers improve performance with respect to the target alone.

In conclusion, the visual system of rats display differences in contextual processing of collinear stimuli compared to humans, and this is likely to be a difference between species. Rats and humans have different densities and ratios of cones and rods. Humans have a fovea, and rats only have a slight increase in density in their central vision. Rats have smaller eyes, fewer ganglion cells and lower acuity. Like all mammals, both species have a visual cortex with six layers. Rats have fewer higher cortical areas, and differ in the details of cortical micro-circuitry. Rats have a visual cortex that lacks smooth regions of orientation columns in which nearby cells have similar orientation preference. It is unknown which, if any, of these differences are responsible for the differences in collinear processing between the species. An alternative possibility is that the natural distribution of scenes sampled by rats have different characteristics, thus inducing different neural wiring. Or possibly the attention systems of rats are different and this interacts with spatial processing of visual features.

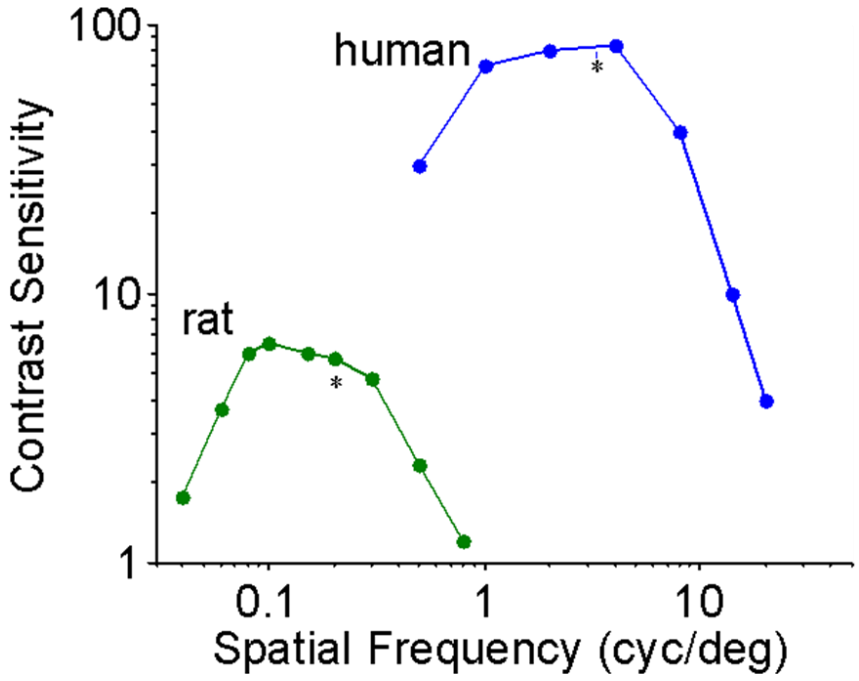

Figure A2.1. A comparison of the contrast sensitivity in humans and rats. Contrast sensitivity is measured as one over the contrast threshold. Both species performed a detection task near their peak spatial frequency sensitivity (indicated by the asterisk). Rats were tested with stimuli that were 0.22 cyc/deg while humans were tested with stimuli that were 3.3 cyc/degree. Data reproduced from Keller, 2000.

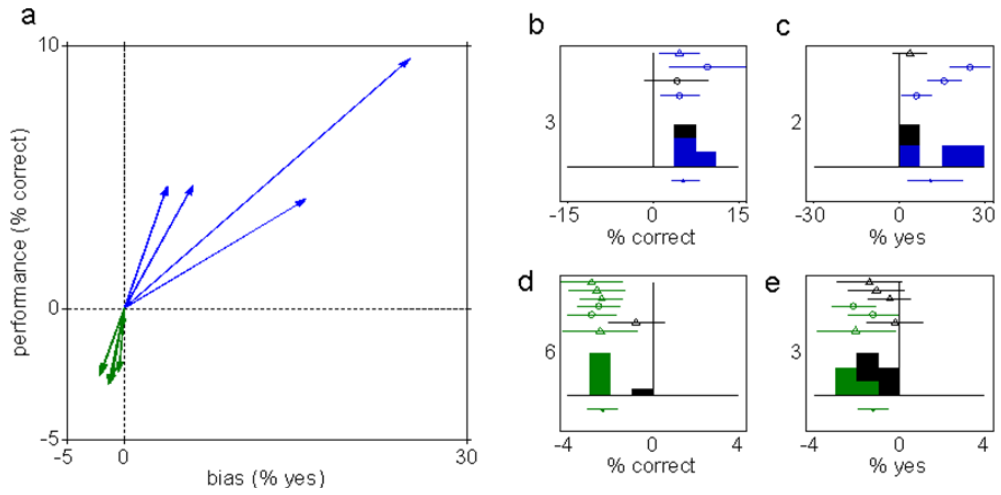

Figure A2.2. The influence of collinear stimuli compared to pop-out stimuli in humans and rats. a) Arrows indicated changes in performance for each subject. Human subjects are coded in blue, and rat subjects are coded in green. The x-axis indicates changes in bias, and the y-axis indicates changes in percent correct. b) Human subjects improved performance on collinear conditions compared to pop out conditions. Each horizontal line corresponds to a single human subject. Error bars indicate the 95% confidence of a Wald interval modified by the Agresti-Caffo method. Histogram below is color coded blue to indicate effects that were statistically significant, and black otherwise. The average population change is plotted below the histogram. c) The change in the percent of yes responses for humans viewing collinear conditions compared to pop-out conditions. Panels c d and e use the same display conventions as panel b. d) The change in percent correct performance for rats. Rats perform worse on collinear conditions than pop-out conditions. e) The change in the percent of yes responses for rats. Rats produce fewer yes responses for collinear conditions than pop-out conditions.

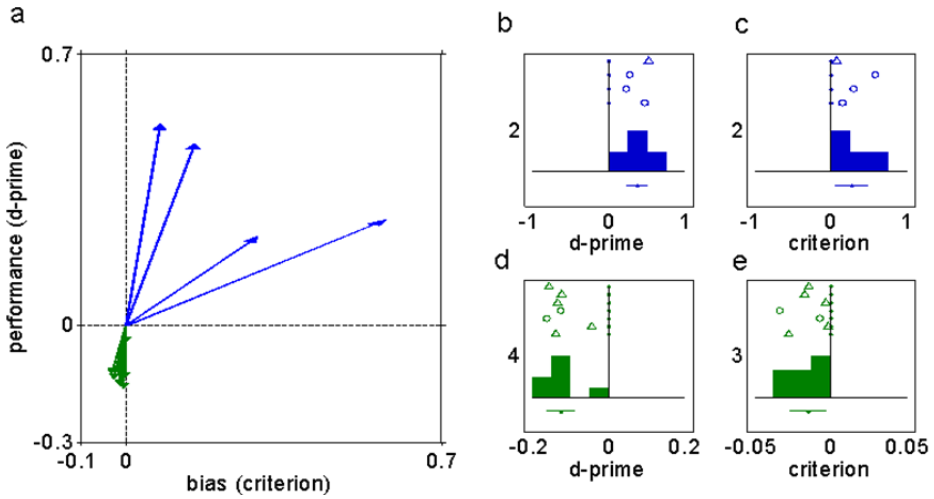

Figure A2.3. The influence of collinear stimuli compared to pop-out stimuli in humans and rats have the same interpretation with signal detection theory metrics. Arrows in panel a) and the x-axis in panels b)-e) indicated the difference between collinear and pop-out conditions. All panels a-e follow the same convention as Figure A2, expect that performance is measured as d-prime instead of percent correct and bias is measures a criterion instead of fraction of yes responses. b) In humans, d-prime is greater for collinear stimuli than pop-out stimuli. c) In humans, the criterion shift indicates an increase probability of reporting yes for collinear stimuli than pop-out stimuli. d) In rats, d-prime is lower for collinear stimuli than pop-out stimuli. d) In rats, the criterion shift indicates a decrease in the probability of reporting yes for collinear stimuli than pop-out stimuli.

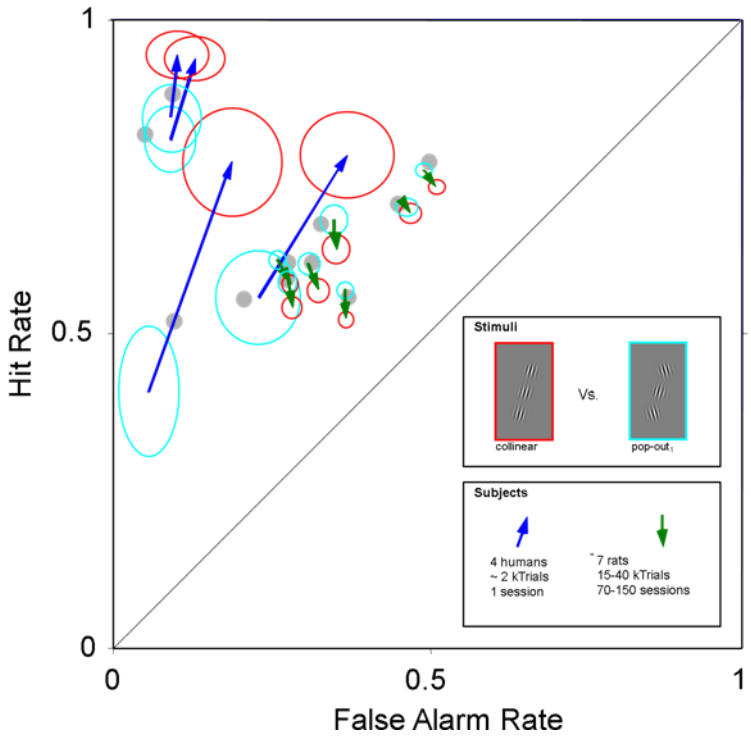

Figure A2.4. Behavioral performance of humans and rats on a target detection task in the presence of flankers. The hit rates and false alarms for each subject on a condition are indicated by the center of an ellipse. The width of each ellipse denotes the 95% confidence interval for a binomial proportion. For each subject, an arrow in this space summarizes the difference between the collinear condition (red) and randomly interleaved trials of the non-collinear reference condition (pop-out$_1$, cyan). The arrow points to the change in performance induced by the collinear feature (with respect to the reference condition). Blue arrows are used for human subjects. Green arrows are used for rats. An alternative reference condition (parallel) was also included for each subject as a grey dot. The dot indicates the mean performance. Additional arrows and ellipses were not included for the parallel condition to avoid clutter.

References

Agresti, A., & Caffo, B. (2000). Simple and effective confidence intervals for proportions and differences of proportions result from adding two successes and two failures. *American Statistician, 54*(4), 280-288.

Angelucci, A., Levitt, J. B., & Lund, J. S. (2002). Anatomical origins of the classical receptive field and modulatory surround field of single neurons in macaque visual cortical area V1. *Prog Brain Res, 136*, 373-388.

Angelucci, A., Levitt, J. B., Walton, E. J., Hupe, J. M., Bullier, J., & Lund, J. S. (2002). Circuits for local and global signal integration in primary visual cortex. *J Neurosci, 22*(19), 8633-8646.

Anishchenko, A., Greschner, M., Elstrott, J., Sher, A., Litke, A. M., Feller, M. B., et al. Receptive field mosaics of retinal ganglion cells are established without visual experience. *J Neurophysiol*.

Artal, P., Herreros de Tejada, P., Munoz Tedo, C., & Green, D. G. (1998). Retinal image quality in the rodent eye. *Vis Neurosci, 15*(4), 597-605.

Atallah, B. V., & Scanziani, M. (2009). Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron, 62*(4), 566-577.

Bakin, J. S., Nakayama, K., & Gilbert, C. D. (2000). Visual responses in monkey areas V1 and V2 to three-dimensional surface configurations. *J Neurosci, 20*(21), 8188-8198.

Barlow, H. (2001). Redundancy reduction revisited. *Network, 12*(3), 241-253.

Bell, A. J., & Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Comput, 7*(6), 1129-1159.

Ben-Shahar, O., Huggins, P. S., Izo, T., & Zucker, S. W. (2003). Cortical connections and early visual function: intra- and inter-columnar processing. *J Physiol Paris, 97*(2-3), 191-208.

Birch, D., & Jacobs, G. H. (1979). Spatial contrast sensitivity in albino and pigmented rats. *Vision Res, 19*(8), 933-937.

Blomfield, S. (1974). Arithmetical operations performed by nerve cells. *Brain Res, 69*(1), 115-124.

Bock, D. D., Lee, W. C., Kerlin, A. M., Andermann, M. L., Hood, G., Wetzel, A. W., et al. Network anatomy and in vivo physiology of visual cortical neurons. *Nature, 471*(7337), 177-182.

Bonds, A. B. (1989). Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Vis Neurosci, 2*(1), 41-55.

Bonin, V., Mante, V., & Carandini, M. (2005). The suppressive field of neurons in lateral geniculate nucleus. *J Neurosci, 25*(47), 10844-10856.

Bonin, V., Mante, V., & Carandini, M. (2006). The statistical computation underlying contrast gain control. *J Neurosci, 26*(23), 6346-6353.

Bosking, W. H., Zhang, Y., Schofield, B., & Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J Neurosci, 17*(6), 2112-2127.

Boynton, G. M., Demb, J. B., Glover, G. H., & Heeger, D. J. (1999). Neuronal basis of contrast discrimination. *Vision Res, 39*(2), 257-269.

Brenner, C. A., Krishnan, G. P., Vohs, J. L., Ahn, W. Y., Hetrick, W. P., Morzorati, S. L., et al. (2009). Steady state responses: electrophysiological assessment of sensory function in schizophrenia. *Schizophr Bull, 35*(6), 1065-1077.

Brown, R. O., & MacLeod, D. I. (1997). Color appearance depends on the variance of surround colors. *Curr Biol, 7*(11), 844-849.

Bullier, J. (2001). Integrated model of visual processing. *Brain Res Brain Res Rev, 36*(2-3), 96-107.

Bullock, W. M., Cardon, K., Bustillo, J., Roberts, R. C., & Perrone-Bizzozero, N. I. (2008). Altered expression of genes involved in GABAergic transmission and neuromodulation of granule cell activity in the cerebellum of schizophrenia patients. *Am J Psychiatry, 165*(12), 1594-1603.

Bussey, T. J., Muir, J. L., & Robbins, T. W. (1994). A Novel Automated Touchscreen Procedure for Assessing Learning in the Rat Using Computer Graphic Stimuli. *Neuroscience Research Communications, 15*(2), 103.

Bussey, T. J., Padain, T. L., Skillings, E. A., Winters, B. D., Morton, A. J., & Saksida, L. M. (2008). The touchscreen cognitive testing method for rodents: how to get the best out of your rat. *Learn Mem, 15*(7), 516-523.

Cannon, M. W., & Fullenkamp, S. C. (1996). A model for inhibitory lateral interaction effects in perceived contrast. *Vision Res, 36*(8), 1115-1125.

Carandini, M. (2004). Receptive fields and suppressive fields in the early visual system. In M. Gazzaniga (Ed.), *The Cognitive Neurosciences* (3 ed., pp. 313-326). Cambridge, MA: MIT Press.

Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci, 17*(21), 8621-8644.

Cavanaugh, J. R., Bair, W., & Movshon, J. A. (2002). Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *J Neurophysiol, 88*(5), 2547-2556.

Chapman, J., & Mc, G. A. (1962). A comparative study of disordered attention in schizophrenia. *J Ment Sci, 108*, 487-500.

Chen, C. C., Kasamatsu, T., Polat, U., & Norcia, A. M. (2001). Contrast response characteristics of long-range lateral interactions in cat striate cortex. *Neuroreport, 12*(4), 655-661.

Chen, C. C., & Tyler, C. W. (1999). Spatial pattern summation is phase-insensitive in the fovea but not in the periphery. *Spat Vis, 12*(3), 267-285.

Chen, C. C., & Tyler, C. W. (2002). Lateral modulation of contrast discrimination: flanker orientation effects. *J Vis, 2*(6), 520-530.

Chen, C. C., & Tyler, C. W. (2008). Excitatory and inhibitory interaction fields of flankers revealed by contrast-masking functions. *J Vis, 8*(4), 10 11-14.

Chisum, H. J., & Fitzpatrick, D. (2004). The contribution of vertical and horizontal connections to the receptive field center and surround in V1. *Neural Netw, 17*(5-6), 681-693.

Chisum, H. J., Mooser, F., & Fitzpatrick, D. (2003). Emergent properties of layer 2/3 neurons reflect the collinear arrangement of horizontal connections in tree shrew visual cortex. *J Neurosci, 23*(7), 2947-2960.

Cho, R. Y., Nystrom, L. E., Brown, E. T., Jones, A. D., Braver, T. S., Holmes, P. J., et al. (2002). Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cogn Affect Behav Neurosci, 2*(4), 283-299.

Chubb, C., Sperling, G., & Solomon, J. A. (1989). Texture interactions determine perceived contrast. *Proc Natl Acad Sci U S A, 86*(23), 9631-9635.

Cook, R. G., Geller, A. I., Zhang, G. R., & Gowda, R. (2004). Touchscreen-enhanced visual learning in rats. *Behav Res Methods Instrum Comput, 36*(1), 101-106.

Cowey, A., & Franzini, C. (1979). The retinal origin of uncrossed optic nerve fibres in rats and their role in visual discrimination. *Exp Brain Res, 35*(3), 443-455.

Cruikshank, S. J., Landisman, C. E., Mancilla, J. G., & Connors, B. W. (2005). Connexon connexions in the thalamocortical system. *Prog Brain Res, 149*, 41-57.

Cruz, D. A., Weaver, C. L., Lovallo, E. M., Melchitzky, D. S., & Lewis, D. A. (2009). Selective alterations in postsynaptic markers of chandelier cell inputs to cortical pyramidal neurons in subjects with schizophrenia. *Neuropsychopharmacology, 34*(9), 2112-2124.

Das, A., & Gilbert, C. D. (1995). Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature, 375*(6534), 780-784.

Das, A., & Gilbert, C. D. (1999). Topography of contextual modulations mediated by short-range interactions in primary visual cortex. *Nature, 399*(6737), 655-661.

DeAngelis, G. C., Robson, J. G., Ohzawa, I., & Freeman, R. D. (1992). Organization of suppression in receptive fields of neurons in cat visual cortex. *J Neurophysiol, 68*(1), 144-163.

Delahunt, P. B., & Brainard, D. H. (2004). Color constancy under changes in reflected illumination. *J Vis, 4*(9), 764-778.

Denning, K. S., & Reinagel, P. (2005). Visual control of burst priming in the anesthetized lateral geniculate nucleus. *J Neurosci, 25*(14), 3531-3538.

DiCarlo, J. J., & Maunsell, J. H. (2003). Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. *J Neurophysiol, 89*(6), 3264-3278.

Douglas, R. M., Neve, A., Quittenbaum, J. P., Alam, N. M., & Prusky, G. T. (2006). Perception of visual motion coherence by rats and mice. *Vision Res, 46*(18), 2842-2847.

Dreifuss, J. J., Kelly, J. S., & Krnjevic, K. (1969). Cortical inhibition and gamma-aminobutyric acid. *Exp Brain Res, 9*(2), 137-154.

Dubin, M. W., & Cleland, B. G. (1977). Organization of visual inputs to interneurons of lateral geniculate nucleus of the cat. *J Neurophysiol, 40*(2), 410-427.

Ejima, Y., & Takahashi, S. (1985). Apparent contrast of a sinusoidal grating in the simultaneous presence of peripheral gratings. *Vision Res, 25*(9), 1223-1232.

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A, 4*(12), 2379-2394.

Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: evidence for a local "association field". *Vision Res, 33*(2), 173-193.

Finn, I. M., Priebe, N. J., & Ferster, D. (2007). The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron, 54*(1), 137-152.

Fitzpatrick, D. (2000). Seeing beyond the receptive field in primary visual cortex. *Curr Opin Neurobiol, 10*(4), 438-443.

Foley, J. M. (1994). Human luminance pattern-vision mechanisms: masking experiments require a new model. *J Opt Soc Am A Opt Image Sci Vis, 11*(6), 1710-1719.

Freeman, E., Sagi, D., & Driver, J. (2004). Configuration-specific attentional modulation of flanker- -target lateral interactions. *Perception, 33*(2), 181-194.

Fukushima, K. (1980). Neocognitron: a self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol Cybern, 36*(4), 193-202.

Fuller, J. H. (1985). Eye and head movements in the pigmented rat. *Vision Res, 25*(8), 1121-1128.

Garcia-Perez, M. A., & Peli, E. (2001). Luminance artifacts of cathode-ray tube displays for vision research. *Spat Vis, 14*(2), 201-215.

Gaudry, K. S., & Reinagel, P. (2007). Benefits of contrast normalization demonstrated in neurons and model cells. *J Neurosci, 27*(30), 8071-8079.

Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annu Rev Psychol, 59*, 167-192.

Geisler, W. S., & Albrecht, D. G. (1992). Cortical neurons: isolation of contrast gain control. *Vision Res, 32*(8), 1409-1410.

Geisler, W. S., & Perry, J. S. (2009). Contour statistics in natural images: grouping across occlusions. *Vis Neurosci, 26*(1), 109-121.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res, 41*(6), 711-724.

Gilbert, C. D., & Wiesel, T. N. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J Neurosci, 9*(7), 2432-2442.

Girman, S. V., Sauve, Y., & Lund, R. D. (1999). Receptive field properties of single neurons in rat primary visual cortex. *J Neurophysiol, 82*(1), 301-311.

Gollisch, T., & Meister, M. Eye smarter than scientists believed: neural computations in circuits of the retina. *Neuron, 65*(2), 150-164.

Gong, S., Zheng, C., Doughty, M. L., Losos, K., Didkovsky, N., Schambra, U. B., et al. (2003). A gene expression atlas of the central nervous system based on bacterial artificial chromosomes. *Nature, 425*(6961), 917-925.

Gooding, D. C., Braun, J. G., & Studer, J. A. (2006). Attentional network task performance in patients with schizophrenia-spectrum disorders: evidence of a specific deficit. *Schizophr Res, 88*(1-3), 169-178.

Goris, R. L., Zaenen, P., & Wagemans, J. (2008). Some observations on contrast detection in noise. *J Vis, 8*(9), 4 1-15.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York,: Wiley.

Grossberg, S., & Raizada, R. D. (2000). Contrast-sensitive perceptual grouping and object-based attention in the laminar circuits of primary visual cortex. *Vision Res, 40*(10-12), 1413-1432.

Hansen, T., & Neumann, H. (2008). A recurrent model of contour integration in primary visual cortex. *J Vis, 8*(8), 8 1-25.

Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H., & Buzsaki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol, 84*(1), 401-414.

Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Vis Neurosci, 9*(2), 181-197.

Heffner, R. S., & Heffner, H. E. (1992). Visual factors in sound localization in mammals. *J Comp Neurol, 317*(3), 219-232.

Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biol Cybern, 82*(2), 111-121.

Helmholtz, H. v. (1866). *Treatise on physiological optics* (Dover ed.). Mineola, NY: Dover Publications.

Howe, C. Q., & Purves, D. (2005). Natural-scene geometry predicts the perception of angles and line orientation. *Proc Natl Acad Sci U S A, 102*(4), 1228-1233.

Huang, X., Albright, T. D., & Stoner, G. R. (2007). Adaptive surround modulation in cortical area MT. *Neuron, 53*(5), 761-770.

Hubel, D. H. (1960). Single unit activity in lateral geniculate body and optic tract of unrestrained cats. *J Physiol, 150*, 91-104.

Hubel, D. H., & Wiesel, T. N. (1961). Integrative action in the cat's lateral geniculate body. *J Physiol, 155*, 385-398.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J Physiol, 195*(1), 215-243.

Hughes, A. (1979). A schematic eye for the rat. *Vision Res, 19*(5), 569-588.

Hughes, A., & Wassle, H. (1979). An estimate of image quality in the rat eye. *Invest Ophthalmol Vis Sci, 18*(8), 878-881.

Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Res, 40*(10-12), 1489-1506.

Jacobs, G. H., Calderone, J. B., Fenwick, J. A., Krogh, K., & Williams, G. A. (2003). Visual adaptations in a diurnal rodent, Octodon degus. *J Comp Physiol A Neuroethol Sens Neural Behav Physiol, 189*(5), 347-361.

Jacobs, G. H., Fenwick, J. A., & Williams, G. A. (2001). Cone-based vision of rats for ultraviolet and visible lights. *J Exp Biol, 204*(Pt 14), 2439-2446.

Jacobs, G. H., Williams, G. A., Cahill, H., & Nathans, J. (2007). Emergence of novel color vision in mice engineered to express a human cone photopigment. *Science, 315*(5819), 1723-1725.

Jahnsen, H., & Llinas, R. (1984). Ionic basis for the electro-responsiveness and oscillatory properties of guinea-pig thalamic neurones in vitro. *J Physiol, 349*, 227-247.

Johnson, A. W., Crombag, H. S., Smith, D. R., & Ramanan, N. Effects of serum response factor (SRF) deletion on conditioned reinforcement. *Behav Brain Res, 220*(2), 312-318.

Jones, M. V., Brooks, P. A., & Harrison, N. L. (1992). Enhancement of gamma-aminobutyric acid-activated Cl- currents in cultured rat hippocampal neurones by three volatile anaesthetics. *J Physiol, 449*, 279-293.

Kapadia, M. K., Ito, M., Gilbert, C. D., & Westheimer, G. (1995). Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys. *Neuron, 15*(4), 843-856.

Keller, J., Strasburger, H., Cerutti, D. T., & Sabel, B. A. (2000). Assessing spatial vision - automated measurement of the contrast-sensitivity function in the hooded rat. *J Neurosci Methods, 97*(2), 103-110.

Keri, S., Kelemen, O., & Benedek, G. (2009). Attentional modulation of perceptual organisation in schizophrenia. *Cogn Neuropsychiatry, 14*(2), 77-86.

Kerlin, A. M., Andermann, M. L., Berezovskii, V. K., & Reid, R. C. Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. *Neuron, 67*(5), 858-871.

Kleiner, M., D.H. Brainard , D.G. Pelli (2007). *What's new in Psychtoolbox-3?* Paper presented at the 30th European Conference on Visual Perception.

Kryger, Z., Galli-Resta, L., Jacobs, G. H., & Reese, B. E. (1998). The topography of rod and cone photoreceptors in the retina of the ground squirrel. *Vis Neurosci, 15*(4), 685-691.

Kurki, I., Hyvarinen, A., & Laurinen, P. (2006). Collinear context (and learning) change the profile of the perceptual filter. *Vision Res, 46*(13), 2009-2014.

Lakatos, P., Karmos, G., Mehta, A. D., Ulbert, I., & Schroeder, C. E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science, 320*(5872), 110-113.

Lashley, K. S. (1912). Visual discrimination of size and form in the albino rat. *J. Anim. Behav., 2,* 310-331.

Lashley, K. S. (1931). The mechanism of vision. I. A method for rapid analysis of pattern- vision in the rat.

Lechner, H. A., Lein, E. S., & Callaway, E. M. (2002). A genetic method for selective and quickly reversible silencing of Mammalian neurons. *J Neurosci, 22*(13), 5287-5290.

Lee, M. D. (2008). BayesSDT: software for Bayesian inference with signal detection theory. *Behav Res Methods, 40*(2), 450-456.

Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *J Opt Soc Am, 70*(12), 1458-1471.

Lesica, N. A., & Stanley, G. B. (2004). Encoding of natural scene movies by tonic and burst spikes in the lateral geniculate nucleus. *J Neurosci, 24*(47), 10731-10740.

Levitt, J. B., & Lund, J. S. (1997). Contrast dependence of contextual effects in primate visual cortex. *Nature, 387*(6628), 73-76.

Lewis, D. A., Hashimoto, T., & Volk, D. W. (2005). Cortical inhibitory neurons and schizophrenia. *Nat Rev Neurosci, 6*(4), 312-324.

Li, W., Piech, V., & Gilbert, C. D. (2006). Contour saliency in primary visual cortex. *Neuron, 50*(6), 951-962.

Livingstone, M. S., & Hubel, D. H. (1981). Effects of sleep and arousal on the processing of visual information in the cat. *Nature, 291*(5816), 554-561.

Lu, S. M., Guido, W., & Sherman, S. M. (1992). Effects of membrane voltage on receptive field properties of lateral geniculate neurons in the cat: contributions of the low-threshold Ca2+ conductance. *J Neurophysiol, 68*(6), 2185-2198.

Matsubara, J., Cynader, M., Swindale, N. V., & Stryker, M. P. (1985). Intrinsic projections within visual cortex: evidence for orientation-specific local connections. *Proc Natl Acad Sci U S A, 82*(3), 935-939.

McCormick, D. A., Shu, Y., Hasenstaub, A., Sanchez-Vives, M., Badoual, M., & Bal, T. (2003). Persistent cortical activity: mechanisms of generation and effects on neuronal excitability. *Cereb Cortex, 13*(11), 1219-1231.

Meier, P., & Reinagel, P. (2010). Model of visual target detection applied to rat behavior. *Front. Neurosci. Conference Abstract: Computational and Systems Neuroscience 2010.*

Meier, P. M., Flister, E. D., & Reinagel, P. (2011). Collinear features impair visual detection by rats. *Journal of Vision*.

Minini, L., & Jeffery, K. J. (2006). Do rats use shape to solve "shape discriminations"? *Learn Mem, 13*(3), 287-297.

Mitra, P., & Miller, R. F. (2007). Normal and rebound impulse firing in retinal ganglion cells. *Vis Neurosci, 24*(1), 79-90.

Morris, R. (1984). Developments of a water-maze procedure for studying spatial learning in the rat. *J Neurosci Methods, 11*(1), 47-60.

Munoz Tedo, C., Herreros de Tejada, P., & Green, D. G. (1994). Behavioral estimates of absolute threshold in rat. *Vis Neurosci, 11*(6), 1077-1082.

Murphy, P. C., Duckett, S. G., & Sillito, A. M. (1999). Feedback connections to the lateral geniculate nucleus and cortical response properties. *Science, 286*(5444), 1552-1554.

Murphy, P. C., Uhlrich, D. J., Tamamaki, N., & Sherman, S. M. (1994). Brain-stem modulation of the response properties of cells in the cat's perigeniculate nucleus. *Vis Neurosci, 11*(4), 781-791.

Nachmias, J., & Sansbury, R. V. (1974). Letter: Grating contrast: discrimination may be better than detection. *Vision Res, 14*(10), 1039-1042.

Nothdurft, H. C. (1991). Texture segmentation and pop-out from orientation contrast. *Vision Res, 31*(6), 1073-1078.

Nummela, S. U., & Krauzlis, R. J. Inactivation of primate superior colliculus biases target choice for smooth pursuit, saccades, and button press responses. *J Neurophysiol, 104*(3), 1538-1548.

Ogilvie, J. C., & Creelman, C. D. (1968). Maximum-Likelihood Estimation of Receiver Operating Characteristic Curve Parameters. *Journal of Mathematical Psychology, 5*(3), 377-&.

Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P., & Reid, R. C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature, 433*(7026), 597-603.

Okun, M., & Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci, 11*(5), 535-537.

Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature, 381*(6583), 607-609.

Otazu, G. H., Tai, L. H., Yang, Y., & Zador, A. M. (2009). Engaging in an auditory task suppresses responses in auditory cortex. *Nat Neurosci, 12*(5), 646-654.

Papagiakoumou, E., Anselmi, F., Begue, A., de Sars, V., Gluckstad, J., Isacoff, E. Y., et al. Scanless two-photon excitation of channelrhodopsin-2. *Nat Methods, 7*(10), 848-854.

Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nat Neurosci, 4*(7), 739-744.

Pelli, D. G. (1997). Pixel independence: measuring spatial interactions on a CRT display. *Spat Vis, 10*(4), 443-446.

Petrov, Y., Verghese, P., & McKee, S. P. (2006). Collinear facilitation is largely uncertainty reduction. *J Vis, 6*(2), 170-178.

Pinna, B., Brelstaff, G., & Spillmann, L. (2001). Surface color from boundaries: a new 'watercolor' illusion. *Vision Res, 41*(20), 2669-2676.

Polat, U., Mizobe, K., Pettet, M. W., Kasamatsu, T., & Norcia, A. M. (1998). Collinear stimuli regulate visual responses depending on cell's contrast threshold. *Nature, 391*(6667), 580-584.

Polat, U., & Sagi, D. (1993). Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiments. *Vision Res, 33*(7), 993-999.

Polat, U., & Sagi, D. (2007). The relationship between the subjective and objective aspects of visual filling-in. *Vision Res, 47*(18), 2473-2481.

Pouille, F., Marin-Burgin, A., Adesnik, H., Atallah, B. V., & Scanziani, M. (2009). Input normalization by global feedforward inhibition expands cortical dynamic range. *Nat Neurosci, 12*(12), 1577-1585.

Prusky, G. T., Alam, M. S., Beekman, S., & Douglas, R. M. (2004). Rapid Quantification of Adult and Developing Mouse Spatial Vision Using a Virtual Optomotor System. *Invest. Ophthalmol. Vis. Sci., 45*(12), 4611-4616.

Prusky, G. T., West, P. W., & Douglas, R. M. (2000). Behavioral assessment of visual acuity in mice and rats. *Vision Res, 40*(16), 2201-2209.

Reinagel, P., Godwin, D., Sherman, S. M., & Koch, C. (1999). Encoding of visual information by LGN bursts. *J Neurophysiol, 81*(5), 2558-2569.

Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron, 61*(2), 168-185.

Ries, C. R., & Puil, E. (1999). Ionic mechanism of isoflurane's actions on thalamocortical neurons. *J Neurophysiol, 81*(4), 1802-1809.

Robinson, L., Bridge, H., & Riedel, G. (2001). Visual discrimination learning in the water maze: a novel test for visual acuity. *Behav Brain Res, 119*(1), 77-84.

Rodieck, R. W. (1998). *The first steps in seeing*. Sunderland, Mass.: Sinauer Associates.

Roorda, A., & Williams, D. R. (1999). The arrangement of the three cone classes in the living human eye. *Nature, 397*(6719), 520-522.

Ruderman, D. L., & Bialek, W. (1994). Statistics of natural images: Scaling in the woods. *Phys Rev Lett, 73*(6), 814-817.

Rumberger, A., Tyler, C. J., & Lund, J. S. (2001). Intra- and inter-areal connections between the primary visual cortex V1 and the area immediately surrounding V1 in the rat. *Neuroscience, 102*(1), 35-52.

Saito, T., & Nakatsuji, N. (2001). Efficient Gene Transfer into the Embryonic Mouse Brain Using in Vivo Electroporation. *Developmental Biology, 240*(1), 237-246.

Salin, P. A., Girard, P., Kennedy, H., & Bullier, J. (1992). Visuotopic organization of corticocortical connections in the visual system of the cat. *J Comp Neurol, 320*(4), 415-434.

Schwartz, O., Sejnowski, T. J., & Dayan, P. (2009). Perceptual organization in the tilt illusion. *J Vis, 9*(4), 19 11-20.

Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nat Neurosci, 4*(8), 819-825.

Schwarz, G. (1978). Estimating Dimension of a Model. *Annals of Statistics, 6*(2), 461-464.

Series, P., Lorenceau, J., & Fregnac, Y. (2003). The "silent" surround of V1 receptive fields: theory and experiments. *J Physiol Paris, 97*(4-6), 453-474.

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci, 18*(10), 3870-3896.

Shapley, R. M., & Victor, J. D. (1979). Nonlinear spatial summation and the contrast gain control of cat retinal ganglion cells. *J Physiol, 290*(2), 141-161.

Shepherd, G. M. (2004). *The synaptic organization of the brain* (5th ed.). Oxford ; New York: Oxford University Press.

Sherman, S. M. (2005). Thalamic relays and cortical functioning. *Prog Brain Res, 149*, 107-126.

Sherman, S. M., & Guillery, R. W. (2002). The role of the thalamus in the flow of information to the cortex. *Philos Trans R Soc Lond B Biol Sci, 357*(1428), 1695-1708.

Shu, Y., Hasenstaub, A., Badoual, M., Bal, T., & McCormick, D. A. (2003). Barrages of synaptic activity control the gain and sensitivity of cortical neurons. *J Neurosci, 23*(32), 10388-10401.

Sigman, M., Cecchi, G. A., Gilbert, C. D., & Magnasco, M. O. (2001). On a common circle: natural scenes and Gestalt rules. *Proc Natl Acad Sci U S A, 98*(4), 1935-1940.

Sillito, A. M., Cudeiro, J., & Murphy, P. C. (1993). Orientation sensitive elements in the corticofugal influence on centre-surround interactions in the dorsal lateral geniculate nucleus. *Exp Brain Res, 93*(1), 6-16.

Simoncelli, E., & Schwartz, O. (1998). Derivation of a Cortical Normalization Model from the Statistics of Natural Images. *Investigative Opthalmology and Visual Science Supplement (ARVO), 39*, 424.

Smith, M. A., Kelly, R. C., & Lee, T. S. (2007). Dynamics of response to perceptual pop-out stimuli in macaque V1. *J Neurophysiol, 98*(6), 3436-3449.

Solomon, J. A., & Morgan, M. J. (2000). Facilitation from collinear flanks is cancelled by non-collinear flanks. *Vision Res, 40*(3), 279-286.

Solomon, J. A., Watson, A. B., & Morgan, M. J. (1999). Transducer model produces facilitation from opposite-sign flanks. *Vision Res, 39*(5), 987-992.

Stahl, J. S. (2004). Eye movements of the murine P/Q calcium channel mutant rocker, and the impact of aging. *J Neurophysiol, 91*(5), 2066-2078.

Steriade, M. (2001). To burst, or rather, not to burst. *Nat Neurosci, 4*(7), 671.

Super, H., Spekreijse, H., & Lamme, V. A. (2001). Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nat Neurosci, 4*(3), 304-310.

Swadlow, H. A., Bezdudnaya, T., & Gusev, A. G. (2005). Spike timing and synaptic dynamics at the awake thalamocortical synapse. *Prog Brain Res, 149*, 91-105.

Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature, 381*(6582), 520-522.

Tolhurst, D. J., & Heeger, D. J. (1997). Comparison of contrast-normalization and threshold models of the responses of simple cells in cat striate cortex. *Vis Neurosci, 14*(2), 293-309.

Tolhurst, D. J., Tadmor, Y., & Chao, T. (1992). Amplitude spectra of natural images. *Ophthalmic Physiol Opt, 12*(2), 229-232.

Trejo, L. J., & Cicerone, C. M. (1982). Retinal sensitivity measured by the pupillary light reflex in RCS and albino rats. *Vision Res, 22*(9), 1163-1171.

Uchida, N., & Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nat Neurosci, 6*(11), 1224-1229.

Uhlrich, D. J., Cucchiaro, J. B., Humphrey, A. L., & Sherman, S. M. (1991). Morphology and axonal projection patterns of individual neurons in the cat perigeniculate nucleus. *J Neurophysiol, 65*(6), 1528-1541.

van der Schaaf, A., & van Hateren, J. H. (1996). Modelling the power spectra of natural images: statistics and information. *Vision Res, 36*(17), 2759-2770.

Van Hooser, S. D., Heimel, J. A., Chung, S., & Nelson, S. B. (2006). Lack of patchy horizontal connectivity in primary visual cortex of a mammal without orientation maps. *J Neurosci, 26*(29), 7680-7692.

Van Horn, S. C., Erisir, A., & Sherman, S. M. (2000). Relative distribution of synapses in the A-laminae of the lateral geniculate nucleus of the cat. *J Comp Neurol, 416*(4), 509-520.

Wang, J., Borton, D. A., Zhang, J., Burwell, R. D., & Nurmikko, A. V. A neurophotonic device for stimulation and recording of neural microcircuits. *Conf Proc IEEE Eng Med Biol Soc, 2010*, 2935-2938.

Wang, W., Jones, H. E., Andolina, I. M., Salt, T. E., & Sillito, A. M. (2006). Functional alignment of feedback effects from visual cortex to thalamus. *Nat Neurosci, 9*(10), 1330-1336.

Wang, X., Wei, Y., Vaingankar, V., Wang, Q., Koepsell, K., Sommer, F. T., et al. (2007). Feedforward excitation and inhibition evoke dual modes of firing in the cat's visual thalamus during naturalistic viewing. *Neuron, 55*(3), 465-478.

Watson, A. B., & Solomon, J. A. (1997). Model of visual contrast gain control and pattern masking. *J Opt Soc Am A Opt Image Sci Vis, 14*(9), 2379-2391.

Webb, B. S., Dhruv, N. T., Solomon, S. G., Tailby, C., & Lennie, P. (2005). Early and late mechanisms of surround suppression in striate cortex of macaque. *J Neurosci, 25*(50), 11666-11675.

Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nat Neurosci, 5*(6), 598-604.

Wenderoth, P., & Johnstone, S. (1988). The different mechanisms of the direct and indirect tilt illusions. *Vision Res, 28*(2), 301-312.

Weyand, T. G., Boudreaux, M., & Guido, W. (2001). Burst and tonic response modes in thalamic neurons during sleep and wakefulness. *J Neurophysiol, 85*(3), 1107-1118.

Wiesenfeld, K., & Moss, F. (1995). Stochastic resonance and the benefits of noise: from ice ages to crayfish and SQUIDs. *Nature, 373*(6509), 33-36.

Williams, C. B., & Hess, R. F. (1998). Relationship between facilitation at threshold and suprathreshold contour integration. *J Opt Soc Am A Opt Image Sci Vis, 15*(8), 2046-2051.

Wolfart, J., Debay, D., Le Masson, G., Destexhe, A., & Bal, T. (2005). Synaptic background activity controls spike transfer from thalamus to cortex. *Nat Neurosci, 8*(12), 1760-1767.

Worgotter, F., Niebur, E., & Koch, C. (1991). Isotropic connections generate functional asymmetrical behavior in visual cortical cells. *J Neurophysiol, 66*(2), 444-459.

Wu, C. C., & Chen, C. C. Distinguishing lateral interaction from uncertainty reduction in collinear flanker effect on contrast discrimination. *J Vis, 10*(3), 8 1-14.

Xing, J., & Heeger, D. J. (2001). Measurement and modeling of center-surround suppression and enhancement. *Vision Res, 41*(5), 571-583.

Zarrinpar, A., & Callaway, E. M. (2006). Local connections to specific types of layer 6 neurons in the rat visual cortex. *J Neurophysiol, 95*(3), 1751-1761.

Zeki, S. M. (1978). Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex. *J Physiol, 277*, 273-290.

Zenger-Landolt, B., & Koch, C. (2001). Flanker effects in peripheral contrast discrimination--psychophysics and modeling. *Vision Res, 41*(27), 3663-3675.

Zenger, B., & Sagi, D. (1996). Isolating excitatory and inhibitory nonlinear spatial interactions involved in contrast detection. *Vision Res, 36*(16), 2497-2513.

Zetzsche, C., & Nuding, U. (2005). Nonlinear and higher-order approaches to the encoding of natural scenes. *Network, 16*(2-3), 191-221.

Zoccolan, D., Oertelt, N., DiCarlo, J. J., & Cox, D. D. (2009). A rodent model for the study of invariant visual object recognition. *Proc Natl Acad Sci U S A, 106*(21), 8748-8753.