

# **Spatial integration of natural scenes in the mammalian retina**

Julian Freedland

A dissertation submitted in fulfillment of the requirements for the degree of

Doctor of Philosophy

University of Washington

Reading Committee:

Fred Rieke, Chair

Greg Horwitz

John Tuthill

Program Authorized to Offer Degree:

Molecular Engineering

© Copyright 2023

Julian Freedland

University of Washington

## **Abstract**

### **Spatial integration of natural scenes in the mammalian retina**

Julian Freedland

Chair of the Supervisory Committee:

Fred Rieke

Department of Physiology and Biophysics

The human visual system relies on neural signals that travel along the optic nerve from the eye to the brain. Signals within the optic nerve are generated by individual neurons (“ganglion cells”) that reside within the retina and produce neural spikes in response to specific visual inputs. Establishing a causal link between the visual features in a single scene and the subsequent spike response of retinal ganglion cells is an important step in developing a complete understanding of human vision – and can aid the design of prosthetic eyes that restore sight.

In this dissertation, we focus our efforts towards how retinal ganglion cells in the mammalian retina encode realistic scenes from the natural world (“natural scenes”). Natural scenes are typically difficult to study in the retina due to their high dimensionality (large number of pixels) and complex statistical structure (how pixels are arranged). In Chapter 2, we gain traction on this problem by simplifying the spatial structure of natural movies and asking, using electrophysiological techniques, whether these simplifications are reflected in the spike responses of retinal ganglion cells. We find that in parasol ganglion cells, natural movies with 16 wedge-shaped pixels elicit similar responses to the original natural movie. Importantly, these simplified stimuli perform consistently across populations of neurons and only require knowledge of the neuron’s classical receptive field, enabling relatively accurate predictions of neural responses to natural scenes with minimal inputs.

In Chapter 3, we utilize a similar electrophysiological approach to investigate whether retinal ganglion cells can resolve certain spatial disruptions to natural scenes. We find that breaking the spatial structure of images (introducing “spatial discontinuities”) uniformly delays spike responses. Our work provides insight into how complex visual inputs are encoded by the mammalian visual system.

# Acknowledgments

Upon entering graduate school, I expected to live out the dramatic tales of my professors' past experiences. I would surely eat my daily lump of poison at dawn on a one-penny-a-day salary and spend each evening fluffing a pillow at the lab bench. Yet these insights couldn't have been more wrong – as I quickly learned, graduate students must now ingest *two* lumps of poison each morning. The National Institutes of Health recommends increasing daily dosages to three lumps of poison for all graduate students within the next five years.

Contrary to what many may expect, graduate school has been an overwhelmingly positive experience. Much of this positivity is attributable to my advisor and undisputed king of Slamtown; Mr. Electrophysiology himself; Fred Rieke. Fred has always provided the space and time to closely tinker with scientific questions. This research philosophy – of carefully and rigorously approaching questions from multiple angles – is deeply time consuming, yet calls for fierce creativity and the thrill of challenging every assumption. This approach to science now fuels much of my scientific thinking.

My graduate experience would amount to little without the wonderful Breena Sperry, my wife and fellow graduate student. Our life is defined by joy, laughter, and board game-derived smack talk, and all of my scientific efforts were only possible because my spirit and soul were refreshed each and every day.

In addition, I'd like to thank:

- My home department, the Molecular Engineering and Sciences Institute, for their easygoing and flexible guidance throughout my graduate years.
- My labmates: Mike Ahlquist, Todd Appleby, Jacob Baudin, Chris Chen, Shellee Cunnington, Gabrielle Gutierrez, Arthur Hong, Norianne Ingram, Mike Manookin, Phil Mardoum, Srinidhi Naidu, Greg Newkirk, Vyom Raval, Adree Songco-Aguas, and Alison Weber.
- The College of Arts & Sciences at the University of Washington for the opportunity to complete courses in orchestral conducting, tap dance, modern dance, club dance, and West African dance – experiences that broadened my perspectives of the arts and provided zest to my day-to-day routine.
- My committee: David Gire, Greg Horwitz, and John Tuthill for their kind inputs while challenging, encouraging, and pushing the limits of my intuition and knowledge.

Finally, a special word of thanks to my undergraduate research advisor, Mike Fasullo. The ease, spoils, and bakery-esque smells of working with baker's yeast will always hold a special place in my heart. Mike's consistent presence at the lab bench and kind, easy-to-talk-to approach provided such deep excitement and encouragement for a young scientist beginning his journey.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Natural and artificial stimuli . . . . .	1
1.2	Encoding of artificial stimuli by retinal ganglion cells . . . . .	3
1.2.1	Linear receptive field . . . . .	4
1.2.2	Limitations of a linear receptive field model . . . . .	9
1.2.3	Rectification . . . . .	11
1.2.4	Order of linear and nonlinear mechanisms . . . . .	12
1.2.5	Expanded linear-nonlinear models . . . . .	15
1.3	Difficulty of generalizing models to natural stimuli . . . . .	17
<b>2</b>	<b>Simplifying the spatial structure of scenes</b>	<b>19</b>
2.1	Abstract . . . . .	19
2.2	Introduction . . . . .	20
2.3	Results . . . . .	22
2.3.1	Spatial integration of natural images in the receptive field center . . . . .	22
2.3.2	Rectified spatial regions combine near-linearly to control RGC response . . . . .	26
2.3.3	Receptive field surrounds have modest impact on responses to naturalistic inputs . . . . .	29
2.3.4	Spatial integration of natural images in the receptive field surround . . . . .	32
2.3.5	Building low-dimensional reductions across the entire receptive field . . . . .	33
2.3.6	Building high-dimensional metamers across the entire receptive field . . . . .	35
2.3.7	Modeling spatial integration across the entire receptive field . . . . .	37
2.3.8	Properties of spatial integration in 16-dimensional space . . . . .	40

2.3.9	Leveraging spatial integration mechanisms with “degenerate” stimuli . . .	44
2.4	Discussion . . . . .	47
2.4.1	Spatial integration and receptive field subunits . . . . .	47
2.4.2	Model limitations . . . . .	49
2.4.3	Characterizing the relevant spatial statistics of natural images . . . . .	51
2.5	Methods . . . . .	51
2.6	Supplementary Figures . . . . .	56
<b>3</b>	<b>Distorting the spatial structure of scenes</b>	<b>58</b>
3.1	Introduction . . . . .	58
3.2	Results . . . . .	59
3.2.1	Adding circular spatial discontinuities to natural movies delay spike responses	60
3.2.2	Spatial discontinuities occur throughout the receptive field center . . . . .	65
3.2.3	Spatial discontinuities engage temporal delays in artificial stimuli . . . . .	69
3.2.4	Adding spatial discontinuities result in larger temporal delays. . . . .	72
3.2.5	Next steps . . . . .	75
3.3	Supplementary Figures . . . . .	78
<b>4</b>	<b>Future Directions</b>	<b>79</b>
4.1	Identifying a complete low-dimensional space . . . . .	79
4.1.1	Spatial integration in the receptive field surround . . . . .	80
4.2	Utilizing a high-dimensional null space . . . . .	81
4.2.1	Ethical considerations . . . . .	82
	<b>Bibliography</b>	<b>82</b>
4.3	List of Publications . . . . .	91
4.4	Vita . . . . .	91

# List of Figures

## Chapter 1

<b>Figure</b>	<b>Summary</b>	<b>Page</b>
Fig. 1.1.1	Examples of natural and artificial stimuli	1
Fig. 1.2.1	Diagram of the retinal circuit	3
Fig. 1.2.2	Measuring a classical receptive field	5
Fig. 1.2.3	Linear integration across the receptive field	6
Fig. 1.2.4	Predictions by linear models	8
Fig. 1.2.5	Identifying nonlinear mechanisms	10
Fig. 1.2.6	Rectification introduces stimulus degeneracies	12
Fig. 1.2.7	Predictions to split-field spots	13
Fig. 1.2.8	Improving predictions to split-field spots	16
Fig. 1.3.1	Nonlinearities depend on spatial scale	18

## Chapter 2

<b>Figure</b>	<b>Summary</b>	<b>Page</b>
Fig. 2.3.1	Integrated wedges in the receptive field center	24
Fig. 2.3.2	Wedge-based regions rectify inputs	28
Fig. 2.3.3	Integrated wedges in the receptive field surround	31
Fig. 2.3.4	Generating degenerate stimuli (metamers)	34
Fig. 2.3.5	Modeling spatial integration in center and surround	39
Fig. 2.3.6	Invariance of wedges to spatial location	42
Fig. 2.3.7	Tuning stimuli across populations of neurons	45
Supp. Fig. 2.6.1	Spatial scale of integrated wedges	56
Supp. Fig. 2.6.2	Low-dimensional stimuli for midget RGCs	57

## Chapter 3

<b>Figure</b>	<b>Summary</b>	<b>Page</b>
Fig. 3.2.1	Introducing circular discontinuities to natural stimuli	61
Fig. 3.2.2	Circular discontinuities evoke spiking delays	63
Fig. 3.2.3	Delays occur throughout receptive field center	68
Fig. 3.2.4	Similar delays occur after disrupting artificial spots	71
Fig. 3.2.5	Adding more discontinuities evokes stronger delays	74
Fig. 3.2.6	<i>(In progress)</i> Shifts in excitatory and inhibitory input	76
Fig. 3.2.7	<i>(In progress)</i> Functional maps of spike delays	77
Supp. Fig. 3.3.1	Effect of discontinuities on spatial integration	78

## List of Equations

<b>Figure</b>	<b>Summary</b>	<b>Page</b>
Eq.1.2.1	Linear integration across receptive field	6
Eq.1.2.2	Converting raw luminance to Weber contrast	7
Eq.1.2.3	Applying linear integration to a split-field spot	9
Eq.1.2.4	Predicting nonlinear responses using scalar weights	15

## List of Tables

<b>Figure</b>	<b>Summary</b>	<b>Page</b>
Table 2.1	Fitted parameters for spatial integration models	56

# Chapter 1

## Introduction

### 1.1 Natural and artificial stimuli

Each second, 18,000 photons from the star Polaris (our North Star) reach a single square millimeter of Earth's surface [1]. If we look towards Polaris, the collection of photons reaching our eyes provides visual information about the star – such as its brightness and size. Given that only 22% of these particles lie within the visible spectrum in humans, our perception of Polaris can be exhaustively described by the 4,000 photons that influence our vision. Collectively, these 4,000 photons reveal the particularly bright spot in the night sky that guides our path northward.

Human vision is not limited to viewing stars in the night sky. On Earth, our vision consistently resolves large assortments of complex, varied objects. Such scenes are referred to as natural scenes (i.e. viewed in the natural world). Three example natural scenes are presented in Fig. 1.1.1(A), and typically consist of fairly common sights (such as leaves and trees). In general terms, the human visual system interprets these natural scenes similar to the star Polaris: light enters the eye, is encoded by the retina, and is parsed across various stages of processing in the brain to form a perceptual image.

A. examples of natural stimuli



B. examples of artificial stimuli




Figure 1.1.1: (A) Three naturalistic images from [2]. (B) Three common laboratory stimuli [3–6].

Although a majority of our vision is tuned towards natural scenes, we lack a deep understanding of how the mammalian visual system measures and interprets these collections of photons [7]. This is because natural scenes contain several highly complex structures that are difficult to directly isolate and study. For example, each image in Fig. 1.1.1(A) contains 70,000 pixels, and each pixel could reasonably impose some effect on our vision. Additionally, these pixels are organized into discrete objects that vary in brightness (dark/bright), shape (circle/square), texture (rough/smooth), and size (small/large). Simply *describing* natural scenes can be extremely difficult – even for individual objects like trees, which consist of a very specific pattern of pixel values. A complete understanding of how these complex spatial features affect our visual system remains even more challenging.

To avoid the complicated nature of natural scenes, visual neuroscientists often prefer to study artificial stimuli. Artificial stimuli are deliberately designed to have a simpler structure than natural scenes; three examples are presented in Fig. 1.1.1(B). Although artificial stimuli typically aren't encountered in the natural world, artificial stimuli strongly stimulate key visual mechanisms (such as the perceptual effects of a sudden flash of light). Studying artificial stimuli is essential for identifying key individual visual mechanisms that, in unison, describe our vision under natural conditions. Retinal models built using artificial inputs, however, have traditionally struggled to consistently describe responses to natural scenes. [7, 8].

In this dissertation, we bridge the gap between artificial and natural stimuli by investigating the key spatial features that elicit responses in retinal ganglion cells. Retinal ganglion cells generate the neural signals that travel along the optic nerve and serve as a direct link between the eye and the brain. In brief, our approach uses electrophysiology to measure the neural signals of retinal neurons in response to a variety of artificial and natural stimuli. By identifying similarities in neural responses shared across both artificial and natural stimuli, we can effectively isolate the spatial features that mediate responses within the mammalian visual system.

## 1.2 Encoding of artificial stimuli by retinal ganglion cells

The retinal neurons relevant to this study do not directly measure light from the natural world. When humans initially view a scene, light is encoded by photoreceptors (i.e, cones) using a light-reactive pigment [1] (Fig. 1.2.1). The primary purpose of photoreceptors is to transform incoming light into a biochemical neurotransmitter. Output signals from photoreceptors are then passed to horizontal and bipolar cells (“Synapse A”). Afterwards, output signals from bipolar cells are passed to amacrine and retinal ganglion cells (“Synapse B”). Retinal ganglion cells, the focus of our studies, subsequently use inputs at Synapse B to generate neural spikes along the optic nerve into the brain. Fig. 1.2.1 depicts a simplified version of the retina (compared to anatomical models which contain a broader variety of pathways and cell types).




Figure 1.2.1: Cross-section of the mammalian retinal circuit (associated with daylight vision).

Since neural signals must first pass through multiple retinal layers before reaching retinal ganglion cells, each retinal layer has the opportunity to affect our perception of visual scenes. If Synapse B deleted all neural activity within the circuit, for instance, retinal ganglion cells would never detect any nonzero inputs regardless of how strongly photoreceptors are stimulated. Thus, our understanding of ganglion cell responses, as well as our perception, is inextricably tied to our

understanding of encoding across multiple stages within the retina. Luckily, we can rely on decades of research with artificial stimuli to closely investigate how visual inputs are shaped across multiple retinal layers.

In the following sections, we will directly measure neural responses to classic artificial stimuli in visual neuroscience. These stimuli, which were developed over decades of research, are still highly relevant for studying neuronal behaviors today. By measuring and analyzing neuronal behaviors in our own hands, we will (1) introduce the exact tools and analyses that inform our subsequent work with natural scenes and (2) build intuitions for how neurons in the retina encode visual inputs.

### **1.2.1 Linear receptive field**

We first investigated how a single retinal ganglion cell responds to a flash of light. We began by recording the neural spikes from an On-parasol retinal ganglion cell. (A full description of methods are available in Section 2.5). On-parasol ganglion cells are a common retinal neuron that are larger in size (relative to other types of retinal ganglion cells) and typically respond to increments of light (“On-”). We identified the center of the neuron and flashed bright spots of varying sizes, counting how many neural spikes were produced by each spot. The results of this experiment are in Fig. 1.2.2.

As we increased the diameter of our flashed spot from 0 to 240  $\mu\text{m}$ , we observed that the total spike count increased (Fig. 1.2.2, left). Increasing spot diameter beyond 240  $\mu\text{m}$ , however, resulted in a decrease in total spike count. Importantly, this experiment measures integrated behavior. Photoreceptors that lie in the absolute center will be stimulated by every spot, regardless of size, and will contribute to each spot’s spike response. If we instead subtract the spikes count between two spot sizes, we isolate the responses uniquely stimulated by the bigger spot. Repeating this operation at an infinitesimally small scale is mathematically described as a derivative and reveals an important intrinsic property of most neurons: their classical receptive field [4,9,10] (Fig. 1.2.2, right).




Figure 1.2.2: (Left) Total number of neural spikes (“spike count”) elicited from an On-parasol retinal ganglion cell (25° temporal) after flashing various uniform, circular spots (at +90% contrast) for 500 milliseconds. Spot size refers to the diameter of each spot (in microns). A difference-of-gaussians fit is provided in red. (Right) Differentiating our difference-of-gaussians fit (with respect to space) reveals receptive-field properties of the neuron.

Our measurements in Fig. 1.2.2 reveal an On-parasol with a receptive field that spans two key spatial regions. In one region ( $< 240 \mu\text{m}$ ), increasing the size of a bright spot elicited more spikes, forming a classical receptive field center. Spots larger than the receptive field center ( $> 240 \mu\text{m}$ ) reached into a second region, the receptive field surround, where larger spot sizes resulted in fewer total spikes. Together, the receptive field offers a two- or three-dimensional model (applying weight  $w$  over one or two spatial dimensions) for approximating how an On-parasol retinal ganglion cell integrates inputs across space (Fig. 1.2.3(A)). First, a stimulus is convolved (multiplied) by the neuron’s receptive field. The resulting stimulus is then integrated (summed) across space to produce spike counts. Performing this operation across flashed bright spots of different sizes will return the same measured results as Fig. 1.2.2 and is referred to as linear integration.




Figure 1.2.3: (A) 3-dimensional model of a classical neuronal receptive field. (B) To predict spike counts, a spatial stimulus is first convolved (multiplied) by the neuron's receptive field, then integrated (summed).

Mathematically, linearly integrating a visual stimulus is equivalent to calculating the weighted average (using the receptive field as a weight over space). For each pixel ( $p$ ) in a region ( $r$ ) with a light intensity ( $I$ ), the linearly integrated luminance ( $\mathcal{L}$ ) is calculated via,

$$\mathcal{L} = \frac{\sum_p^r [I_p \times R_p]}{\sum_p^r [R_p]} \quad (1.2.1)$$

where  $R$  is the receptive field of the neuron. A variation of Eq. 1.2.1 without the denominator is recommended if the receptive field center and surround are perfectly balanced ( $\sum_p^r [R_p] = 0$ ), where  $\mathcal{L} \propto \sum_p^r [I_p \times R_p]$ .

Linear integration offers important notions about how neurons encode visual inputs from the natural world. Within the receptive field center of an On-parasol retinal ganglion cell, larger uniform spots produced more spikes, providing some sensitivity to the *size* of a visual input (Fig. 1.2.3). Additionally, since receptive fields vary their weights over space, spike responses depend on *location*: a bright spot in the absolute center of a receptive field is weighted very differently

to an identical spot hundreds of microns away. As a result, linear integration offers a (limited) connection between the size and location of a spatial feature and a retinal neuron's resulting spike response.

Another important factor that affects retinal encoding is the relative *brightness* of a feature ( $I$  in Eq. 1.2.1), which we quantify using a scale known as Weber contrast [11]. Weber contrast is a useful metric because it translates well to the retinal circuit. To calculate the Weber contrast ( $c$ ), the light intensity of a single pixel ( $I_p$ ) is compared to the average light intensity across the scene ( $I_0$ ):

$$c = \frac{I_p - I_0}{I_0} \quad (1.2.2)$$

By identifying the brightness of a pixel ( $I_p$ ) relative to a background light intensity ( $I_0$ ), the Weber contrast mimics the response functions of individual photoreceptors.

In the spirit of Weber contrast, consider savoring the printed text of a cherished book. Ambient light ( $I_0$ ) sits at exactly 0% contrast. Printed black words appear relatively dark ( $c < 0\%$ ) pressed against a highly reflective white page ( $c > 0\%$ ) and remain constant regardless of whether the book is read under direct sunlight (large  $I_0$ ) or beside a dim bedside lamp (small  $I_0$ ). This largely matches expectations of our visual system, which can perceive the same black-and-white page across a variety of lighting conditions. Since converting light intensity to Weber contrast is a linear (reversible) process for a fixed  $I_0$ , we can describe the linearly integrated luminance  $\mathcal{L}$  as either a pixel intensity ( $I$ ) or as a Weber contrast ( $c$ ).

To better understand how the relative brightness of a visual input affects encoding, we flashed uniform spots with varying Weber contrasts onto the receptive field center of an On-parasol retinal ganglion cell (Fig. 1.2.3, black curve, bottom x-axis). We observed that spike counts increased in response to higher contrast spots, similar to how spike counts increased in response to larger spot sizes (Fig. 1.2.3, blue curve, top x-axis). This result supports our equation for linearly integrated luminance  $\mathcal{L}$  (Eq. 1.2.1) by observing that the left- and right-hand sides of our equation can be

linearly mapped to be approximately equal (Fig. 1.2.3, right).

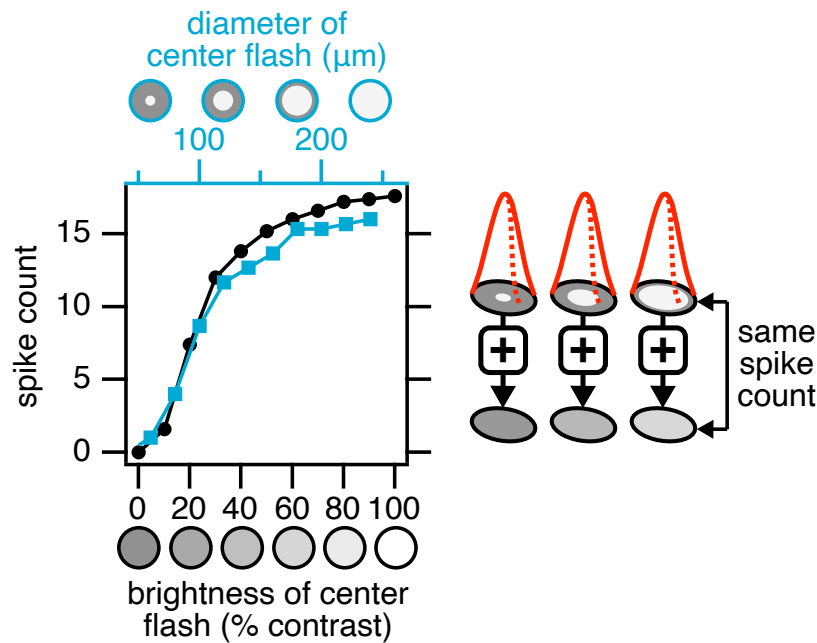


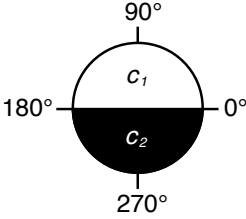
Figure 1.2.4: (Left) Total number of neural spikes (“spike count”) elicited from an On-parasol retinal ganglion cell ( $25^\circ$  temporal). The top x-axis (blue curve) measures spike count to circular spots with a single, uniform luminance (+90% contrast) at different sizes (60 to  $240 \mu\text{m}$ ). The bottom x-axis (black curve) measures spike count to circular spots with varying luminances (0% to +90% contrast) at a single size ( $240 \mu\text{m}$ ). All stimuli were flashed for 500 milliseconds. (Right) As uniform spots grow larger in size, their integrated luminance  $\mathcal{L}$  increases, providing the ability to map stimuli in the top x-axis to stimuli in the bottom x-axis.

In the context of human vision, linear integration offers a startling prediction: a large, dim spot (bottom x-axis in Fig. 1.2.3) can generate the same spike count as a smaller, brighter spot (top x-axis in Fig. 1.2.3). In other words, two completely different visual inputs with the same integrated luminance  $\mathcal{L}$  will elicit the same spike count. Generalizing this concept introduces a class of stimuli that we will consistently use throughout our work: degenerate stimuli, or collections of different visual inputs that elicit identical spike responses in retinal neurons.

The results in this section highlight the importance of a linear receptive field model and its success in predicting spike counts to certain artificial stimuli.

## 1.2.2 Limitations of a linear receptive field model

Certain artificial stimuli within the receptive field center, however, elicit spike responses that strongly diverge from a linear integration model [3, 12]. Consider a grey (0% contrast) spot that spans the receptive field center of an On-parasol retinal ganglion cell. If we replace half of the grey circle with a brighter region (contrast  $> 0$ ) and the other half with a darker region (contrast  $< 0$ ), the resulting split-field stimulus will resemble a half-moon. For some angle ( $\theta$ ), we can define the Weber contrast for each pixel ( $c$ ) within a split-field stimulus ( $s_\theta$ ) as,

$$s_\theta = \begin{cases} c_1 & 0^\circ < \theta \leq 180^\circ \\ c_2 & 180^\circ < \theta \leq 360^\circ \end{cases}$$


where  $c_1 > 0$  (light intensity of the brighter half) and  $c_2 < 0$  (light intensity of the darker half). Using the linearly integrated luminance (Eq. 1.2.1), the spike response of a split-field spot can be described by a simple linear sum:

$$\begin{aligned} \mathcal{L} &= \frac{\sum_{\theta>0^\circ}^{180^\circ} [c_1 \times R_\theta]}{\sum_{\theta>0^\circ}^{360^\circ} [R_\theta]} + \frac{\sum_{\theta>180^\circ}^{360^\circ} [c_2 \times R_\theta]}{\sum_{\theta>0^\circ}^{360^\circ} [R_\theta]} \\ &= c_1 \left( \frac{\sum_{\theta>0^\circ}^{180^\circ} [R_\theta]}{2 \times \sum_{\theta>0^\circ}^{180^\circ} [R_\theta]} \right) + c_2 \left( \frac{\sum_{\theta>180^\circ}^{360^\circ} [R_\theta]}{2 \times \sum_{\theta>180^\circ}^{360^\circ} [R_\theta]} \right) \\ &= \frac{c_1 + c_2}{2} \end{aligned} \tag{1.2.3}$$

Visual neuroscientists realized they could strongly challenge a linear model using a specific split-field spot where  $c_1 = 90\%$  and  $c_2 = -90\%$  contrast [12]. When linearly integrated, the

resulting split-field spot is expected to produce no spike response ( $\mathcal{L} = \frac{1}{2} ([+90] + [-90]) = 0$ ). To test this prediction, we flashed this exact split-field spot onto the receptive field center of an On-parasol retinal ganglion cell (Fig. 1.2.5(A)). We observed a consistent, nonzero spike response during both the onset and offset of the split-field spot, an observation that diverges strongly from expectations of  $\mathcal{L}$ .

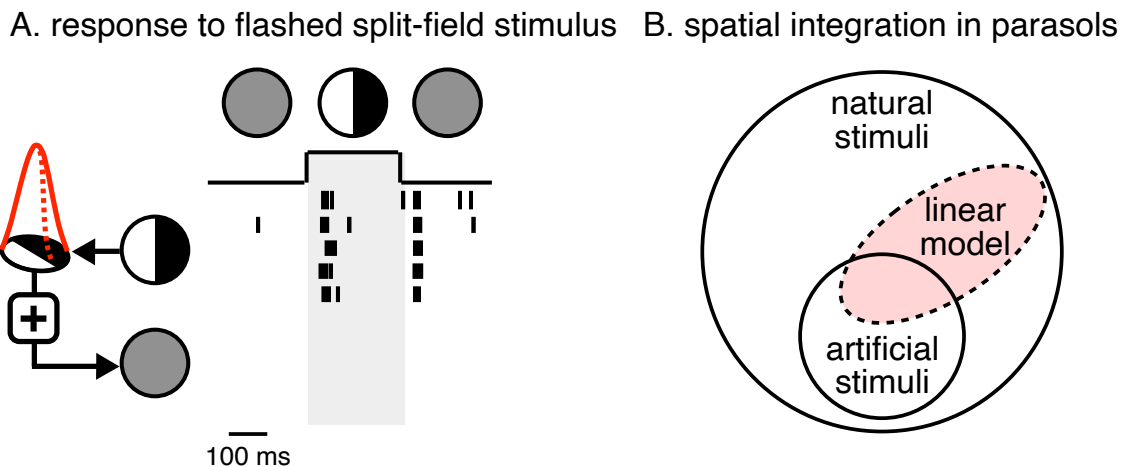


Figure 1.2.5: (A) Spike train (5 repeats) for an On-parasol retinal ganglion cell ( $25^\circ$  temporal) after flashing (for 250 ms) a split-field stimulus (top) and linearly-integrated split-field stimulus (bottom). The bright/dark half of the spot stimulus contains a luminance of  $+90\%/ -90\%$  contrast. (B) Linear integration describes neural responses to some artificial and natural stimuli, but does not ubiquitously describe all stimuli.

Our experiments thus far highlight how linear integration describes neuronal responses to some stimuli (uniform spots) but not others (split-field spots). Since spike responses to uniform and split-field spots were recorded over the exact same region of space in the exact same On-parasol ganglion cell, a single neuron can exhibit both linear and nonlinear behavior depending on the specific spatial properties of the stimulus.

To resolve observations of both linear and nonlinear behavior within a single neuron, we can hypothesize that spatial integration in retinal ganglion cells is globally nonlinear, but can appear

linear under certain cases (Fig. 1.2.5(B)). This is consistent with our upcoming observation that linear integration describes neural responses to some natural scenes but not others (see: Section 2). Identifying the specific stimulus characteristics that mediate linear and nonlinear integration regimes is essential for understanding the full range of human vision under natural conditions.

### 1.2.3 Rectification

If we assume that spatial integration in retinal neurons is globally nonlinear, what specific nonlinear mechanisms (or nonlinearities) within the retinal circuit are engaged by certain stimuli (split-field spots) but not others (uniform spots)?

A nonlinear split-field spot contains three key features: two contrasts (+90%, -90%) separated by a single edge (or spatial discontinuity, to be discussed in Section 3). We began by removing our split-field edge to see whether we could engage nonlinear behavior using only uniform luminances. We flashed uniform spots with a wide variety of spot contrasts across the receptive field center of an On-parasol retinal ganglion cell and counted the number of resulting spikes (Fig. 1.2.6). We observed that an On-parasol produced spikes to brighter-than-grey spots (contrasts  $> 0\%$ ), and as the spot became dimmer, the cell produced fewer spikes. Once the spot became darker than the grey background, no spikes were produced. In other words, a grey flash (0% contrast) – as well as *any* darker-than-grey flash ( $< 0\%$  contrast) – produced no spike response from an On-parasol. This violates predictions by a linear model, which states that brighter flashes should always produce more spikes than darker flashes.

Our experiment in Fig. 1.2.6 reveals that all darker-than-grey flashes are encoded similarly in On-parasol retinal ganglion cells. In the context of human vision, this suggests that On-parasols can detect if a feature is dark ( $< 0\%$  contrast) but cannot discern *how* dark that feature is (0% vs. -90% contrast). This general behavior falls into a class of rectifying nonlinearities. Rectification arises from within the retinal circuit, occurring at the bipolar cells that form inputs to retinal ganglion cells at Synapse *B* (Fig. 1.2.1) [13, 14]. Additionally, another rectification-like effect, saturation,

also occurs in response to extremely bright spots ( $I \gg I_0$ ) (Fig. 1.2.6).

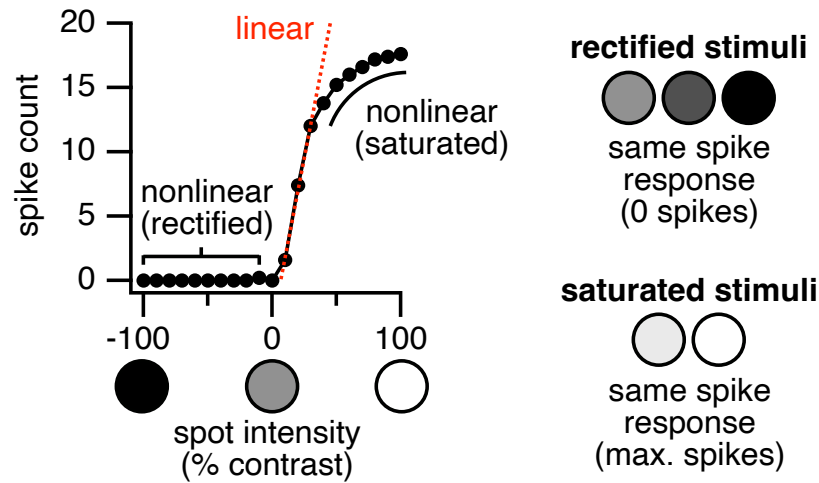


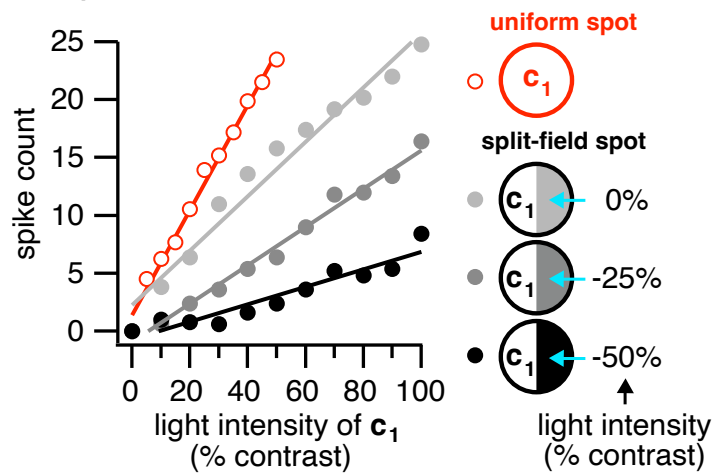
Figure 1.2.6: Total number of neural spikes (“spike count”) elicited from an On-parasol retinal ganglion cell ( $33^\circ$  temporal) after flashing (500 ms) circular spots with uniform luminances (0% to +90% contrast) across the receptive field center.

Our results highlight how certain “linear-nonlinear” models (which assume spatial integration is linear) may not describe all spike responses due to nonlinear processes like rectification and saturation [15].

#### 1.2.4 Order of linear and nonlinear mechanisms

We then investigated how to add nonlinear processes to a linear spatial integration model. We first flashed a uniform spot of varying brightnesses across the receptive field center (red curve) (Fig. 1.2.7(A)). We then set one half of the spot ( $c_2$ ) to either 0%, -25%, and -50% contrast and flashed split-field spots with varying brightnesses in the other half ( $c_1$ ). We observed that as we held  $c_2$  at increasingly darker contrasts, the total number of spikes across all  $c_1$  brightnesses decreased (Fig. 1.2.7(A)).

### A. spike counts to flashed stimuli



### B. model predictions

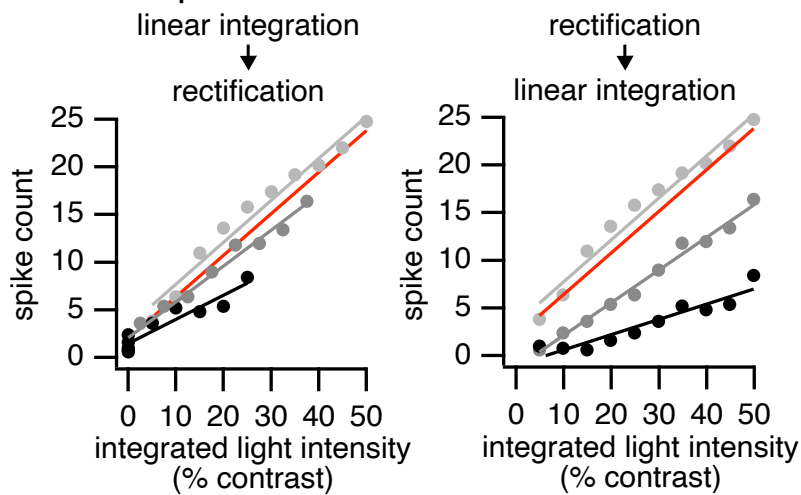


Figure 1.2.7: (A) Total number of neural spikes (“spike count”) elicited from an On-parasol retinal ganglion cell ( $33^\circ$  temporal) after flashing uniform spots (red curve) or split-field spots (grey curves) in the receptive field center. For split-field spots, one half of the receptive field center was kept either 0% (light grey curve), -25% (dark grey curve), or -50% contrast (black curve). The other half varied in light intensity according to the x-axis. (B) Spikes count as predicted by models in which a split-field spot is (left) integrated and then rectified, or (right) rectified and then integrated.

We asked whether different models for spatial integration could successfully predict spike responses to split field spots (Fig. 1.2.7(B)). We built two models that contain the same operations (linear integration and rectification) but only differ in which operations occur first. One model linearly integrates each split-field spot and then rectifies the result (Fig. 1.2.7(B), left). Another model rectifies each split-field spot and then linearly integrates the result (Fig. 1.2.7(B), right). We observed that both models made different predictions, highlighting how the precise order of various linear and nonlinear operations can affect how spatial models perform.

A model with accurate predictions would map each stimulus (data point) onto the red line in Fig. 1.2.7(B). We observed that a model that linearly integrates and then rectified the result (Fig. 1.2.7(B), left) performed better than a model that rectified and then integrated each split-field spot (Fig. 1.2.7(B), right). The relatively poor performance of a rectification-first model occurred because all  $c_2$  contrasts ( $\leq 0\%$  contrast) were rectified to the same value – and as a result, spike counts were only dependent on the value of  $c_1$  (Fig. 1.2.7(B), right).

Our results indicate that while linear-nonlinear models can broadly predict spike responses to split-field spots, such models do not fully capture all spike behavior (i.e, decreasing slopes in Fig. 1.2.7(B) as  $c_2$  became darker). Since model performance cannot be resolved by simply switching the order of linear and nonlinear operations (Fig. 1.2.7(B), right), additional parameters are required to improve spike predictions to split-field spots. As we will discuss in Section 2, adding local nonlinearities at specific spatial scales to models will improve spike predictions considerably.

## 1.2.5 Expanded linear-nonlinear models

The partial success of linear-nonlinear models at predicting spike responses to split-field spots indicate that small adjustments may improve model performance (Fig. 1.2.7). We began by testing whether the inclusion additional linear terms may meaningfully improve performance. Specifically, we introduced a new variable ( $w$ ) that mathematically weights  $c_1$  prior to linear integration such that,

$$\mathcal{L} = \frac{wc_1 + c_2}{2} \quad (1.2.4)$$

Weighting  $c_1$  against  $c_2$  is similar to models that weigh excitatory inputs against inhibitory inputs [16,17]. Eq. 1.2.4 can also partially capture nonlinear processes: if  $|wc_1| \gg |c_2|$ ,  $c_2$  will negligibly impact spike predictions during integration, similar to if  $c_2$  was rectified.

Using our On-parasol retinal ganglion cell from Fig. 1.2.7, we identified a scalar value,  $w$ , that maps each stimulus (data point) onto the red line in Fig. 1.2.7(B). We observed that while the inclusion of  $w$  in Eq. 1.2.4 united uniform and split-field spots with the same spot count, the exact value of  $w$  varied considerably across different split-field spots (Fig. 1.2.8).

The variability of  $w$  in Fig. 1.2.7 indicates that weights applied to  $c_1$  are not appropriately represented as a single scalar value. Such observations offer two potential insights into improving spatial models. Firstly, models that include weight ( $w$ ) as a function that varies with spot brightness (i.e, similar to excitatory or inhibitory inputs) may improve spike predictions. Alternatively, the inability to strongly improve models with a single linear term ( $w$ ) indicate that a nonlinear approach may be better suited for spike predictions.

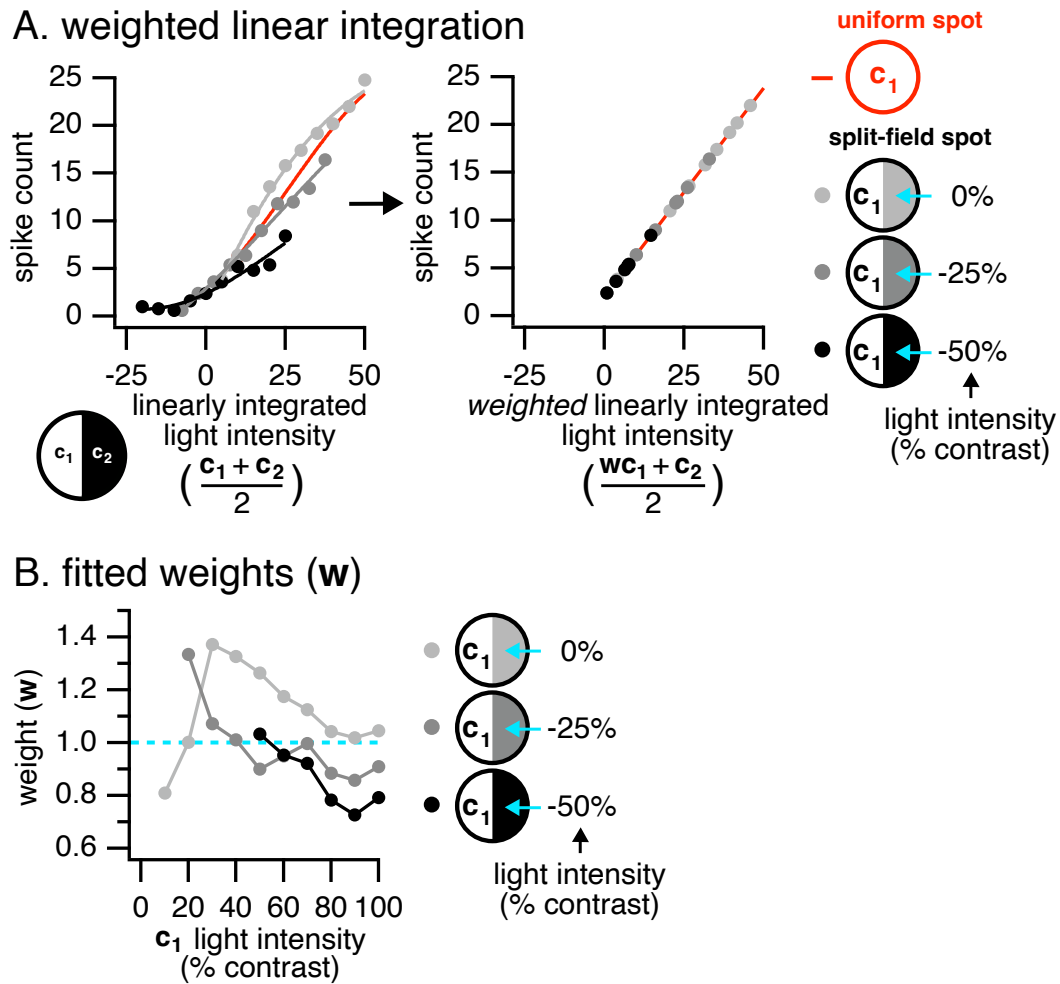


Figure 1.2.8: (A) Total number of neural spikes (“spike count”) elicited from an On-parasol retinal ganglion cell ( $33^\circ$  temporal) with respect to the (left) linearly integrated light intensity ( $\mathcal{L}$ ) or (right) *weighted* linearly integrated light intensity of flashed uniform (red) and split-field spots (grey) across the receptive field center. (B) Individual weights ( $w$ ) applied to  $c_1$  such that the weighted luminance matches a uniform spot with the same spike count (A, left). Points above the blue line ( $w > 1$ ) denote that  $c_1$  is weighted more heavily than  $c_2$  during integration.

### 1.3 Difficulty of generalizing models to natural stimuli

Improving linear-nonlinear models (by weighting inputs as a function of spot brightness or by including spatial nonlinearities) in the context of natural scenes, however, remains a challenge.

Split-field spots consist of two half circles, each with a different contrast value ( $c_1$ ,  $c_2$  from Eq. 1.2.7). Introducing linear weights ( $w$  in Eq. 1.2.4) to split-field spots effectively weighs one contrast value ( $c_1$ ) against another ( $c_2$ ). Natural stimuli, however, contain a wide variety of contrasts over space. Identifying appropriate weights for each individual pixel is a computationally-intensive effort that often yields poor predictions [8].

Spatial models can instead leverage local nonlinearities that account for mechanisms like rectification. Nonlinear models built for split-field spots, however, are limited by the spatial scale of each half circle, which provide similar inputs to retinal ganglion cells regardless of interneuron size, location, or density (Fig. 1.3.1, top row). Stimuli with varied spatial scales are more strongly shaped by individual interneurons, and typically evoke more nonlinear activity (Fig. 1.3.1, bottom row). As a result, nonlinear models built for split-field spots often neglect many of the nonlinear mechanisms stimulated by natural stimuli at very fine spatial scales.

In Section 2, we circumvent this issue by building spatially-nonlinear models using exclusively natural movies. We find that spike predictions can be improved by integrating inputs into approximately 8 wedge-shaped regions within the receptive field center. These spatial regions appear to be a sufficient spatial scale for account for the nonlinear rectification of functional subunits, a key nonlinearity within the retinal circuit.




Figure 1.3.1: Two sets of stimuli (left) linearly integrated by populations of interneurons (top row, in green).

# Chapter 2

## Simplifying the spatial structure of scenes

This text appears in: J. Freedland, F. Rieke, Systematic reduction of the dimensionality of natural scenes allows accurate predictions of retinal ganglion cell spike outputs. *Proceedings of the National Academy of Sciences* 119, e2121744119 (2022). [18]

### 2.1 Abstract

The mammalian retina engages a broad array of linear and nonlinear circuit mechanisms to convert natural scenes into retinal ganglion cell (RGC) spike outputs. Although many individual integration mechanisms are well understood, we know less about how multiple mechanisms interact to encode the complex spatial features present in natural inputs. Here, we identified key spatial features in natural scenes that shape encoding by primate parasol RGCs. Our approach identified simplifications in the spatial structure of natural scenes that minimally altered RGC spike responses. We observed that reducing natural movies into 16 linearly integrated regions described  $\sim 80\%$  of the structure of parasol RGC spike responses; this performance depended on the number of regions but not their precise spatial locations. We used simplified stimuli to design high-dimensional metamers that recapitulated responses to naturalistic movies. Finally, we modeled the retinal computations that convert flashed natural images into 1-dimensional spike counts.

## 2.2 Introduction

Visual scenes in the natural world contain complex spatial features that must be reliably encoded by the mammalian visual system to guide behavior [19, 20]. Although many of the individual physiological mechanisms that enable vision are well understood, robust predictions of how these mechanisms collectively encode natural scenes remain elusive [7, 21]. This discrepancy originates at least in part because natural inputs engage multiple interacting mechanisms in visual neurons differently than typical laboratory stimuli.

The dimensionality of natural inputs is reduced as signals are integrated across space to control retinal outputs. Except in the primate fovea, photoreceptors considerably outnumber downstream ganglion cells [22]. Linear integration across a large number of photoreceptors, as suggested by classical models, predicts that encoded scenes lose much of their fine spatial detail [23, 24]. Indeed, the responses of some ganglion cells to natural scenes are often well-predicted by linear spatial models [25]. Such linear models greatly reduce the dimensionality of scenes by identifying key operations (spatial integration across a receptive field) that drive retinal ganglion cell spike responses.

However, a rich history of work identifies situations in which linear spatial integration fails, revealing receptive field subunits that provide sensitivity to spatial structure [3, 12–14, 26, 27]. Subunit behaviors can be difficult to predict; for instance, subunits in the receptive field center are sensitive to luminances in the receptive field surround [28, 29], highlighting how multiple retinal mechanisms interact to shape natural encoding. Such findings have motivated approaches to identify subunits and incorporate them into encoding models, resulting in improved predictions to natural scenes [8, 30–35]. Functional subunits identified via these approaches are tailored to a specific cell and can reveal the retinal circuitry that shapes inputs to that cell, such as the bipolar cells that provide inputs to ganglion cells [32, 33]. A key question is whether such specificity is required for encoding models and for the interpretation of retinal output signals.

Here we explore an alternative approach: we systematically reduced the complexity of natural stimuli and quantified changes in the resulting spike responses of primate parasol retinal ganglion cells (RGCs). Much of the structure of the spike response was retained when natural movies were simplified into eight linearly-integrated wedge-shaped regions in both the receptive field center and surround. Importantly, our low-dimensional representations only require knowledge of the RGC's classical receptive field. We confirmed the completeness of our dimensional reduction approach by designing novel high-dimensional "metamers": stimuli that share the same reduced dimensional representation but diverge in other higher-order properties. Metamers generated responses very similar to the original natural stimuli. We also observed that RGC responses were largely invariant to the specific locations of wedge-shaped regions and correlations between them. Insensitivity to subunit location further simplifies encoding models and should similarly simplify models for decoding parasol responses. These results highlight the ability of simple dimensionality-reduction approaches to capture the relevant spatial features of natural inputs that govern RGC responses.

## 2.3 Results

Our goal was to determine how On- and Off-parasol retinal ganglion cells (RGCs) respond to spatial structure in natural images. We found that we could reduce natural movies to 16 spatial regions (i.e. representing a point in a 16-dimensional space) while retaining the majority of structure in the RGC spike response. We identified this space by dividing both the center and surround of the classical receptive field into subregions. We found that 8 regions were sufficient to capture the key structure of natural movies in both the center and surround. We then determined how spatial regions in the receptive field center and surround interact to shape spike counts in response to flashed natural images. Finally, we show that parasol responses are insensitive to specific rearrangements of spatial subregions.

### 2.3.1 Spatial integration of natural images in the receptive field center

We started by investigating how the spatial structure of naturalistic scenes in the receptive field center of parasol RGCs could be simplified while minimizing the impact on spike responses. The RGC receptive field is classically modeled as the difference between center and surround components, with the center dominating over some regions of space and the surround others (Fig. 2.3.1(A)) [4, 9, 10]. We measured the center-dominated region of each recorded neuron by first locating the center of the receptive field and then flashing spatially uniform circular disks of several sizes centered on this location. Responses initially increased with increasing disk size and then decreased as the disks began to impinge on surround-dominated regions. The disk that produced the largest response defined the bounds of the center-dominated region of the neuron's receptive field.

Naturalistic stimuli were delivered with a circular aperture that restricted them to the receptive field center of the neuron. To sample a large number of image patches, we recreated the scenes that fall on a single neuron using an online database of eye movement trajectories of human

observers viewing still naturalistic images [36]. Such movies characteristically consist of small eye movements that produce correspondingly small shifts in spatial input, interrupted several times a second by saccades that produce large input shifts. Using eye movements across several still images produced a diverse library of naturalistic stimuli for use in this study.

To simplify naturalistic stimuli, each movie frame was divided into discrete wedge-shaped regions with each wedge containing a uniform luminance. The luminance value of each wedge was determined by linearly integrating the pixel values of the original movie frame weighted by the estimated receptive field (Fig. 2.3.1(B); see Eq. 1.2.1). Since our spatial wedges were calculated independently across each movie frame (without consideration of any prior or subsequent movie frame), collections of wedges represent a purely spatial simplification of the original stimulus. Because we repeat this procedure for each movie frame, light intensities within a single subregion follow a similar temporal trajectory to the light intensities encountered while viewing natural stimuli.

Dividing an image into wedge-shaped regions clearly differs from the receptive fields expected for the retinal interneurons that form anatomical subunits. However, preliminary experiments suggested that tiling the entire receptive field center was more important for encoding than subregion shape. Wedges offer a convenient approach for dense tiling compared to more sparsely-packed geometries (e.g. packed discs) and conveniently preserve the radial curvature (from the center outwards) of classical receptive field weights. We return to the (lack of) correspondence between the functional regions from our simplified stimulus and anatomical subunits in Fig. 2.3.6.

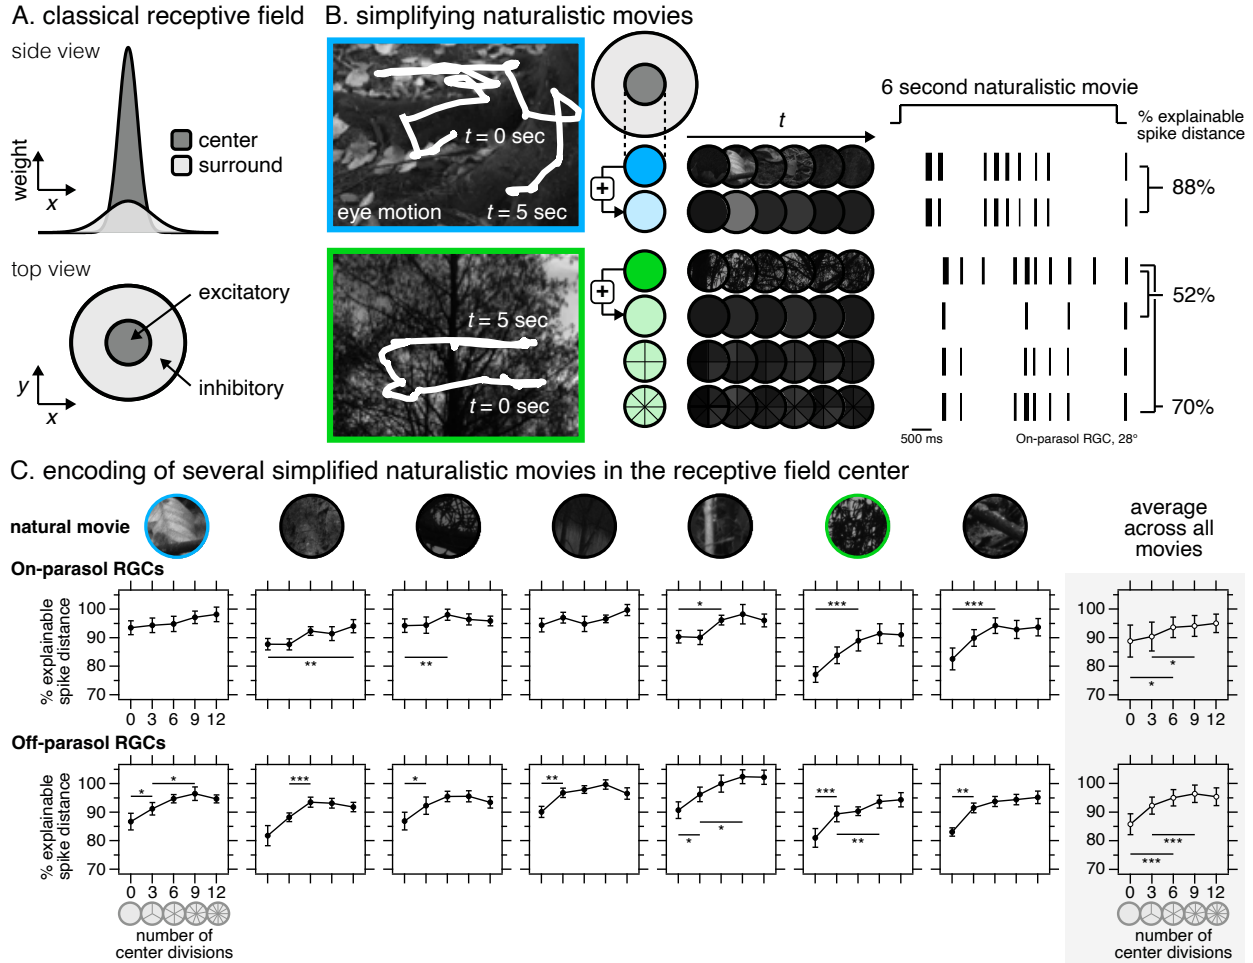


Figure 2.3.1: (A) Side- and top-view of a classical neuronal receptive field. (B, left) Eye-tracking observers during naturalistic image viewing, courtesy of [36], enable the creation of movies that sample various natural image patches over time. (B, middle) Each movie frame can subsequently be simplified to a series of linear-equivalent regions. (B, right) Spike traces for a single On-parasol retinal ganglion cell (“RGC”, 28° temporal) in response to various 6 second naturalistic and simplified movies as presented to their receptive field center. Encoding similarities between simplified movies and their naturalistic counterparts are quantified via explainable spike distance (see Section 2.5; [37]). (C) Encoding properties in the receptive field center of individual On- and Off-parasol RGCs ( $n \geq 12$  across 5 retinas) as simplified movies are divided into increasing numbers of linear-equivalent regions. Seven unique naturalistic movies, each derived from a unique image, are shown with (right) an average across all moves. Error bars for individual movies are SEM; error bars across all movies are S.D. Statistical significance is determined via paired Student’s t-tests: (\*)  $p < 0.05$  (\*\*)  $p < 0.1$ , (\*\*\*)  $p < 0.001$ .

To compare the performance across various simplified spatial architectures, we quantified the similarity of spike responses to original and simplified naturalistic movies using the explained spike distance (see Section 2.5; [37, 38]). The spike distance measures the similarity of two spike trains by computing a cost to convert one to another. Upper and lower bounds allowed us to normalize this measure. The upper bound used random spike trains to calculate the spike distance in the absence of any temporal correspondence between predicted and measured spike trains. The lower bound used responses across repeated trials of the same natural stimulus to determine the (unpredictable) spike distance associated with variability in the neural responses. Figure 2.3.1(B) illustrates this process for an example On-parasol RGC. For the top image, the response to the original movie is well approximated by a single spatially uniform region (88% explained spike distance). For the bottom image, the response to the reduced image becomes more similar to that of the original as more regions are added, and correspondingly the explained spike distance increases.

Figure 2.3.1(C) collects the results from this approach across populations of recorded On- and Off-parasol RGCs for seven different natural movies. For Off-parasol RGCs, the explained spike distance consistently increased as the receptive field was subdivided into 6 to 9 regions. Movies with 6 regions performed better across neurons than movies with a uniform center, and movies with 9 divisions performed better than movies with 3 divisions. When averaged across all seven images, movies with 6 divisions performed 5% and 9% better than movies with uniform centers for On- and Off-parasols, respectively (far right panel in Fig. 2.3.1(C)). Increasing the number of regions from 6 to 9 or 12 resulted in little or no change in average performance in either cell type. The specific number of spatial divisions with the best performance varied across individual natural images, suggesting that the spatial features of a scene may affect the extent to which specific naturalistic movies can be simplified.

Performance improvements for incorporating 6 divisions into our reduced stimuli were smaller and more variable for On-parasol RGCs (5%) than for Off-parasol RGCs (9%). This is consistent with previous work that found that On-parasol responses to natural movies were better approximated

by a spatially-linear receptive field than Off-parasol responses [25]; our results here indicate that the spatial linearity of On-parasol responses holds for some images, while others (e.g. the 6th from the left in Fig. 2.3.1(C)) evoke clear nonlinear spatial integration.

In addition to adjusting the number of linear wedges in our reduced stimuli, we were curious whether spatial integration would benefit from subdividing the receptive field center into additional circular regions. Across a large population of Off-parasol RGCs ( $n = 23$ ), simplified movies with eight wedge-shaped regions across the receptive field center performed similarly to movies with eight wedges and an additional circular subdivision (creating a total of 16 regions) (Supp. Fig. 2.6.1).

The results illustrated in Figure 2.3.1 show that natural movies restricted to the receptive field center of On- and Off-parasol RGCs can be reduced to a 6 to 9 dimensional space with minimal loss of structure relevant for the cells' spike responses. Some movies can be captured with fewer regions. Our observed number of required regions is considerably less than the  $\sim 200$  cone photoreceptors and the  $\sim 30$  cone bipolar cells that fall within the parasol receptive field. Further, the specific shape and position of the regions do not correspond to expectations from anatomy [22], an issue that we return to in Fig. 2.3.6.

### **2.3.2 Rectified spatial regions combine near-linearly to control RGC response**

We next investigated how luminances across discrete wedge-shaped regions within the receptive field center are spatially integrated to produce parasol spike outputs. We opted to exclusively model spike counts to flashed natural images. Our intent was to computationally probe a simpler (purely spatial) stimulus space without the additional challenges associated with modeling the array of temporal effects that likely arise from movies. The ease of investigating purely spatial models became particularly important in scaling our model to include more complex integration regimes

(e.g. center-surround integration).

To probe nonlinear behavior, we flashed 48 naturalistic images for 250 ms in the receptive field center of On- and Off-parasol RGCs and counted the total resulting spikes (Fig. 2.3.2(A)). As expected, the relationship between spike count and the integrated luminance within the receptive field center was nonlinear for both cell types (Fig. 2.3.2(B), top). We then computed a predicted luminance value by reducing each natural image into 8 subregions, rectifying the values in each subregion, and summing. (Without applying rectification, the sum across 8 subregions is identical to a spatially-linear model). Spike counts showed a linear or near-linear relationship with our model output (Fig. 2.3.2(B), bottom), indicating that rectification across 8 subregions is a good predictor of RGC spike counts for these stimuli.

To further test this conclusion, we compared the quality of predictions made by four models (Fig. 2.3.2(C), left): (1) linear summation; (2) weighted linear summation, in which optimal weights were placed on each of the 8 receptive field regions; (3) rectified summation, in which signals in each subregion were rectified and then summed; (4) weighted and rectified summation, in which subregion signals were rectified and then optimally weighted. Optimal weights were determined by performing numerical optimization to identify parameters that maximized the correlation ( $r^2$ ) between the model output and spike count (see Section 2.5). Additionally, we prevented nonlinear models from bypassing their nonlinearity and converging to linear models by requiring that rectification operate within the range of luminance values across our images (see Section 2.5).

Figure 2.3.2(C) (right) shows the explained variance of the responses across 48 flashed images for each model. Models incorporating rectification performed considerably better than those that did not, and applying individual weights to subregions offered little additional improvement despite an increase in the number of model parameters. Similar performance across models with a single uniform weight and models with eight individually-fitted weights indicates that different wedge-shaped regions of space are weighted similarly during spatial integration.

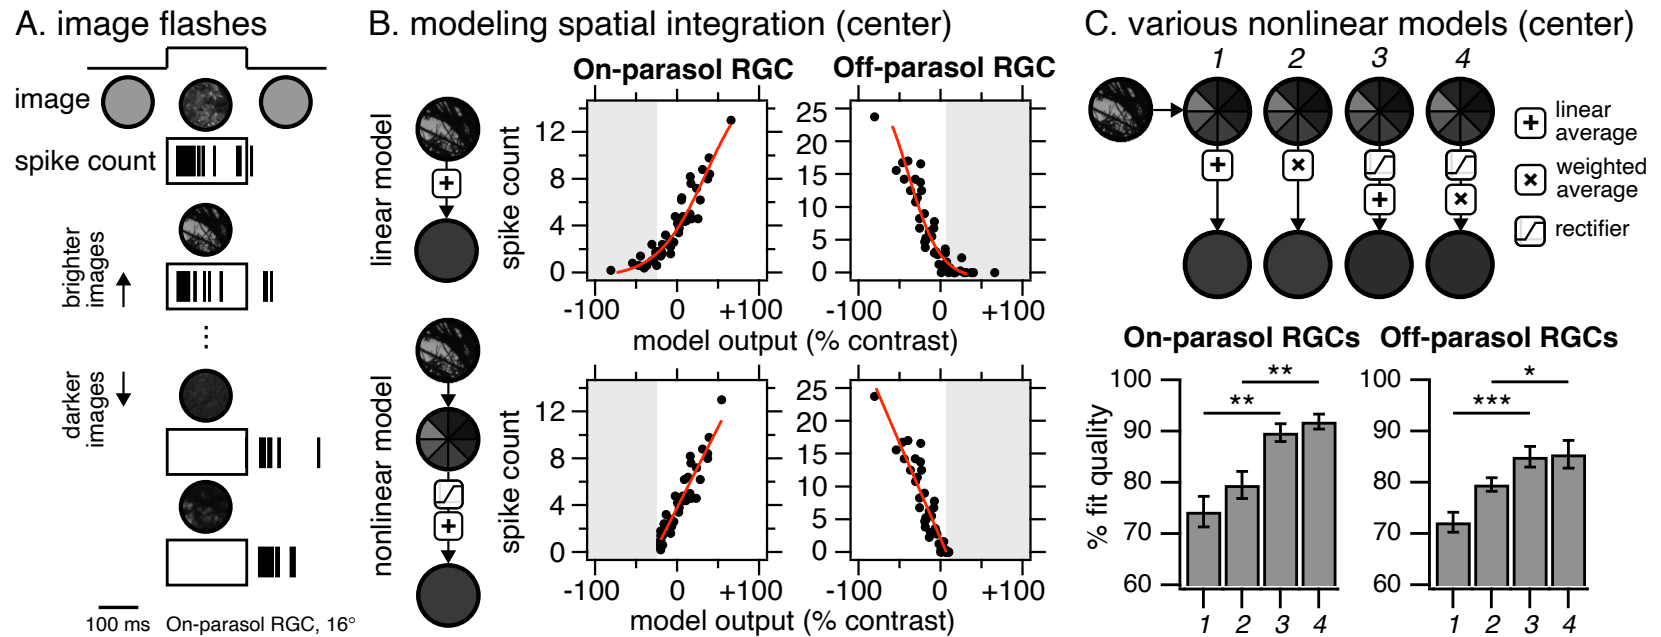


Figure 2.3.2: (A) 48 naturalistic images, spanning an array of luminances, were flashed for 250 ms onto the receptive field center of an On-parasol RGC. Spike count was determined by counting the total number of spikes that occurred within a 250 ms window (boxed region). (B) All 48 naturalistic images were individually integrated through linear and nonlinear models and correlated with total spike count. Each point represents the total number of spikes in response to a single naturalistic image flash for one On- and Off-parasol (47°, 16° respectively). (C) Four downstream integration architecture models were fitted to predict spikes responses to 48 naturalistic image flashes in the receptive field center of individual On- and Off-parasol RGCs ( $n = 9$ ,  $n = 10$ , respectively, across 4 retinas). Averaging operations [+] compute the raw average luminance over the receptive field center; weighted averaging [×] assigns each wedge a weight and then computes the weighted average across the receptive field center. Percent fit quality was determined by optimizing fit parameters for the best linear correlation ( $r^2$ ) between the model output and measured spikes. Statistical significance is determined via paired Student's t-tests: (\*)  $p < 0.05$  (\*\*)  $p < 0.1$ , (\*\*\*)  $p < 0.001$ .

The results illustrated in Figures 2.3.1 and 2.3.2 support a model in which receptive field subregions are first rectified and then summed to predict the RGC response. This is broadly consistent with past work identifying bipolar cells as the origin of receptive field subunits [13], with little or no nonlinear interaction between subunits. These subregions, however, are larger than expected for bipolar subunits and need not be aligned spatially with bipolar cells.

### **2.3.3 Receptive field surrounds have modest impact on responses to naturalistic inputs**

The ability to identify a low-dimensional representation for naturalistic spatial inputs in the receptive field center encouraged us to test a similar empirical approach in the receptive field surround. We began by quantifying how strongly spatial features in the receptive field surround affected spike outputs of On- and Off-parasol RGCs. We compared spike responses to naturalistic movies restricted to the receptive field center with responses to the same movie spanning the entire receptive field (Fig. 2.3.3(A)). Across seven naturalistic movies, responses to receptive field center stimulation alone reached  $\sim 70\%$  explained spike distance when compared to responses to stimuli across the full receptive field for both On- and Off-parasol RGCs. This provides a bound ( $\sim 30\%$ ) on the impact of the receptive field surround.

The limited impact of the surround could reflect either modest activation by natural inputs or a generally weak role of the surround. To distinguish these possibilities, we measured responses (from the same RGCs as Fig. 2.3.3(A)) to movies with a naturalistic image in the receptive field center and temporally-modulated uniform disk in the receptive field surround (Fig. 2.3.3(B)). To strongly stimulate the receptive field surround, the luminance of the surround disk inverted polarity after each saccade. We observed that the circuitry of the receptive field surround is capable of imposing a far greater effect on encoding, nearly doubling its impact during full-field encoding (a  $\geq 57\%$  change in explained spike distance) at  $\pm 90\%$  contrast modulation (Fig. 2.3.3(B)). The effect

of receptive field surround became less pronounced as we decreased the contrast of the surround disk (Fig. 2.3.3(B)). This indicates that the modest impact of the surround during naturalistic movies occurs because such movies do not exercise the surround strongly.

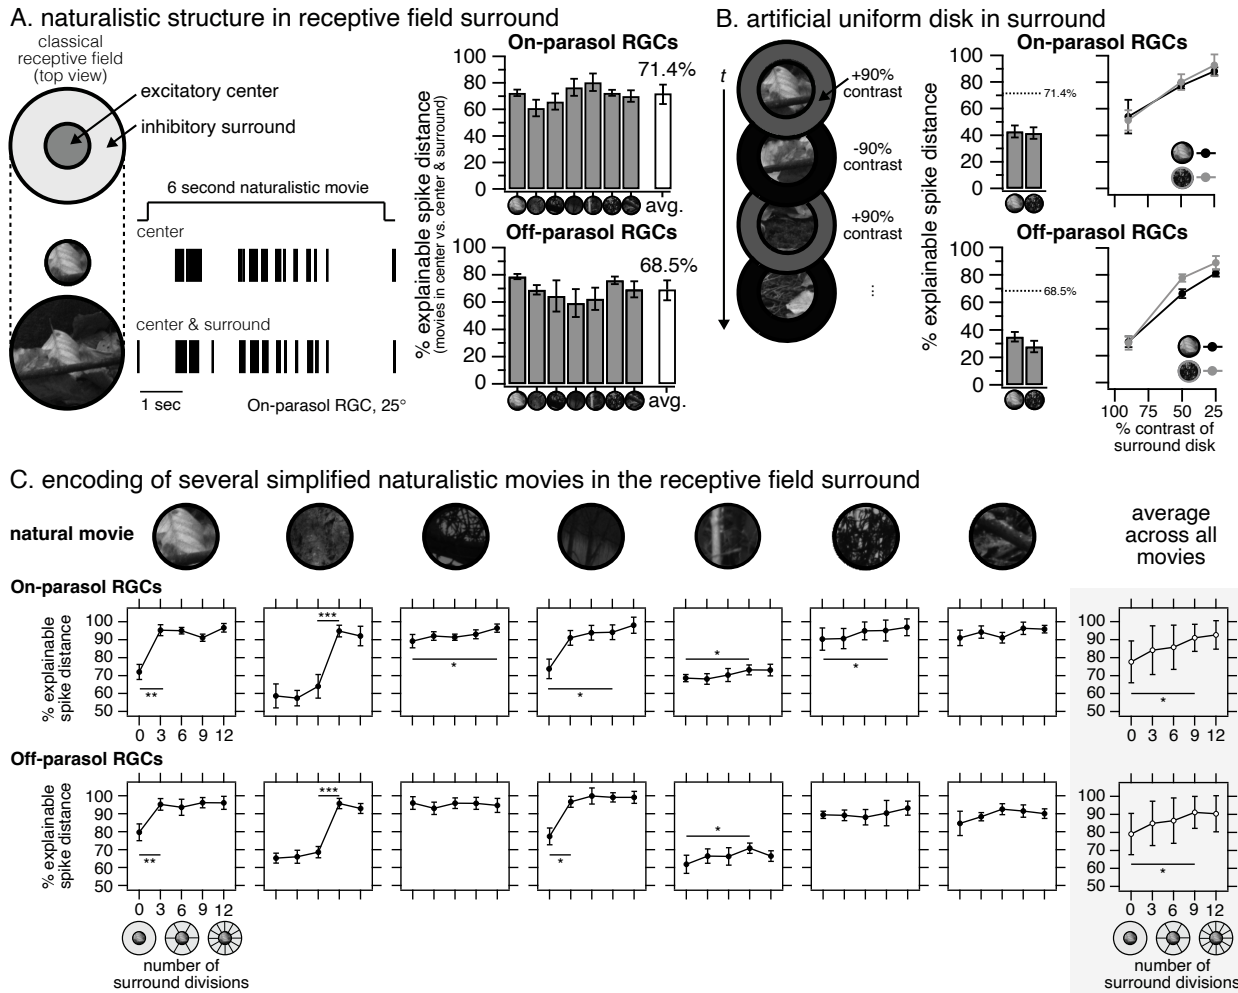


Figure 2.3.3: (A) Spike responses to naturalistic movies presented to the receptive field center and full receptive field of an On-parasol RGC (25° temporal). (A, right) Encoding properties of individual On- and Off-parasol RGCs ( $n \geq 8$  across 5 retinas) comparing movies restricted to the receptive field center relative to the same movies spanning the entire receptive field. (B, left) Across the same set of On- and Off-parasol RGCs, we presented an artificial movie with a naturalistic center and an artificial uniform receptive field surround that inverts polarity after every saccade at  $\pm 90\%$  contrast. (B, right) Encoding effects as the contrast of the uniform surround was altered for individual On- and Off-parasol RGCs ( $n = 4$  across 2 retinas). (C) Encoding properties of individual On- and Off-parasol RGCs ( $n \geq 6$  across 4 retinas) to movies with naturalistic centers and reduced linear-equivalent surrounds divided into increasing numbers of regions. Error bars for individual movies are SEM; error bars across all movies are S.D. Statistical significance is determined via paired Student's t-tests: (\*)  $p < 0.05$  (\*\*)  $p < 0.1$ , (\*\*\*)  $p < 0.001$ .

### **2.3.4 Spatial integration of natural images in the receptive field surround**

Figure 2.3.1 shows that the spatial structure in the receptive field center that controls parasol responses is largely captured by dividing the center into eight or fewer subregions. The experiments described below indicate that a similar dimensional reduction is possible in the receptive field surround.

We compared responses to naturalistic movies that spanned the entire receptive field with responses to movies with linear wedges in the surround. Movies with surround wedges retained the original naturalistic movie in the receptive field center. The number of surround wedges that performed best varied considerably across images (Fig. 2.3.3(C)); in some cases, dividing the surround into regions had little effect (e.g. rightmost panels), while in others, 6 to 9 regions improved the ability to recapitulate responses to the full natural movie. When averaged across all images, dividing the receptive field surround into 6 regions improved full-field encoding by 8% and 7% in On- and Off-parasol RGCs, respectively, compared to a uniform surround. Increasing the number of surround wedges from 6 to 9 resulted in another 5% increase across both cell types, and 12 divisions provided no indication of better or worse performance than 9 divisions (Fig. 2.3.3(C), right).

Similar to our center studies, we were curious whether adding circular subdivisions to the surround might improve performance. In Off-parasols, simplified stimuli with eight wedges across the receptive field surround performed similarly to the same stimuli with an additional circular division in the surround (creating a total of 16 regions across the same region of space) (Supp. Fig. 2.6.1, right).

The results in Figure 2.3.3 indicate that we can replace naturalistic inputs in the receptive field surround with a set of subregions, as is the case for the receptive field center. We observed that dimensional reduction had more varied effects across images in the surround compared to the center, a point that we return to in the Discussion.

### **2.3.5 Building low-dimensional reductions across the entire receptive field**

The work described above suggests that 6 to 9 spatially-uniform wedge-shaped regions in both the receptive field center and surround describe a significant portion of parasol responses to a collection of naturalistic stimuli. Together, these studies indicate that many nonlinear aspects of spatial integration can be captured in a  $\sim 16$ -dimensional space, with fewer dimensions needed for some images. To confirm these findings, we united our center and surround studies to reduce naturalistic movies that span the full receptive field into 16 wedges (8 in the center, 8 in the surround). Movies with 16 regions consistently performed better than movies with a uniform center and surround (Fig. 2.3.4(A)). Across two naturalistic movies, our 16-D representation captured  $\sim 80\%$  of full-field naturalistic spiking in both On- and Off-parasol RGCs (Fig. 2.3.4(A)).

We then investigated whether simplified stimuli could identify relevant spatial integration mechanisms in other cell types in the primate retina. We introduced natural and reduced stimuli to On-midget RGCs, which are classically described to integrate linearly over space [39] (Supp. Fig. 2.6.2(A)). While linear representations (with a uniform center and surround) described 87% of On-midget spike responses to one natural movie (black trace), linear representations for a different movie (red trace) only captured 47% of the overall spike structure. For the latter movie, introducing 16 wedges across the full receptive field improved the explainable spike distance by 19%. When isolated to the receptive field center, the same simplified movie with 8 wedges improved spike distance by 13% compared to a uniform center.

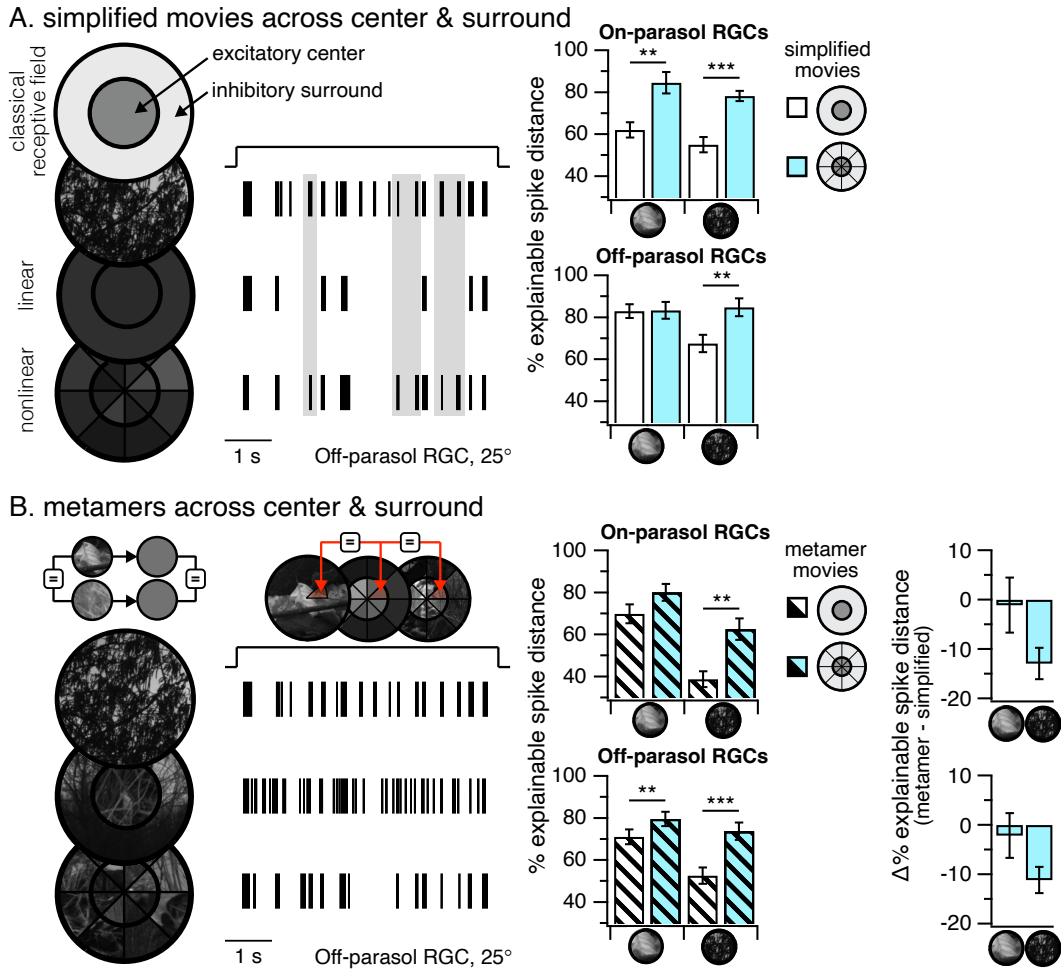


Figure 2.3.4: (A) Spike responses to a naturalistic movie (top) and the same movie simplified into a linear representation (middle), and 16-D nonlinear simplification (bottom). Each movie was presented across the receptive field center and surround of an Off-parasol RGC (25° temporal). Gray highlights indicate specific spiking events captured by both a natural movie and its nonlinear form (but not by movies with a linear center and surround). (A, right) Encoding properties of populations of On- and Off-parasol RGCs ( $n = 16$  across  $\geq 4$  retinas) for linear and nonlinear simplifications of two different naturalistic movies across the full receptive field. (B, left) For each simplified region, we identified sets of naturalistic images that share the same local integrated luminance. Each simplified region was then replaced with its analogous image to produce a metamer. Metamers built from linear (middle) and nonlinear (bottom) representations were then presented to On- and Off-parasol RGCs ( $n \geq 15$  across  $\geq 4$  retinas) and compared directly with the original full-field naturalistic movie (top). (B, right) Change in performance between 16-D metamers and 16-D simplified movies as averaged across individual cells. Statistical significance is determined via paired Student's t-tests: (\*)  $p < 0.05$  (\*\*)  $p < 0.1$  , (\*\*\*)  $p < 0.001$ .

Similar to our studies in On- and Off-parasols, we were curious whether subdividing both the center and surround into additional circular regions would improve performance in On-midgets (Fig. 2.6.2(A)). Although representations with uniform center and surround regions did not improve with additional circular subdivisions, applying circular subdivisions to 16-D representations resulted in a 22% increase over a 16-D representation alone, reaching 88% explainable spike distance. This result was consistent with experiments isolated to the receptive field center: 8 wedges with an additional circular subdivision captured 92% of overall spike behavior and performed 26% better than an 8 wedge representation alone.

The experiments described in this section indicate that our dimensional reduction approach can account for how the entire parasol receptive field integrates spatial inputs, and that similar low-dimensional geometries can generalize to other cell types in the mammalian retina.

### **2.3.6 Building high-dimensional metamers across the entire receptive field**

We have shown that low-dimensional representations of natural stimuli are largely suitable for describing a significant portion ( $\sim 80\%$ ) of the structure of spike responses. An extension of this description predicts that any two natural images that share the same low-dimensional representation (i.e. identical luminances across all wedges in 16-D space) should stimulate the same linear and nonlinear spatial integration mechanisms and therefore produce near-identical spike responses.

To test this prediction, we divided each naturalistic full-field movie frame into 16 wedge-shaped regions. For each wedge, we searched across  $\sim 1800$  natural image patches and identified an entirely different image with a wedge that shares the same integrated value. We then placed the newly-identified natural image wedge into our movie. Repeating this procedure for each spatial region produced a high-dimensional “metamer” consisting of 16 stitched-together image patches from a variety of different images that reduce to the same point in 16-dimensional space as the original image (Fig. 2.3.4(B)). We repeated this approach for each movie frame, and ensured that no two wedges within a given frame were sampled from the same image.

Although naturalistic and metameric movies share a similar low-dimensional representation, their high-frequency features diverge wildly. Additionally, metamers introduce new spatial discontinuities (hard edges that lie between each inserted wedge) and temporal discontinuities (the diversity of images sampled within a given wedge over time) not present in the original naturalistic stimulus. Thus, a comparison of spike responses elicited by naturalistic movies and their metamers is a strong test of whether a 16-dimensional representation sufficiently captures relevant image features.

We observed that responses to metamers across the full receptive field captured spike responses to the original naturalistic movie nearly as well as simplified movies (70 – 80%) for both On- and Off-parasol RGCs (Fig. 2.3.4(B)). Metamers derived from a 16-wedge representation performed better than metamers derived from a linear center and surround, consistent with the performance of their corresponding low-dimensional forms (Fig 2.3.4(B)). These results indicate that the additional structure added in constructing the metameric movies has little impact on a cell's response, confirming that 16-D reductions capture a large portion of the relevant stimulus structure. Metamer performance also highlights that discontinuities in spatial structure at the edge of the wedges do not strongly impact parasol encoding.

We also tested metamer performance in On-midget RGCs (Supp. Fig. 2.6.2(B)). Surprisingly, metamers derived from models with a linear center and surround performed similarly to metamers derived from 16-D representations, even though simplified 16-D forms outperformed simplified linear forms (Supp. Fig. 2.6.2(A)). Adding additional circular subdivisions to a 16-D description increased performance across both metamer movies by  $\sim 18\%$ ; this is consistent with the increase in performance by adding the same circular subdivisions to the low-dimensional representations (22%). Similarly, the performance of metamers restricted to the receptive field center also increased (13%, 28% for both movies) when further subdividing an 8-wedge representation into 2 circular subregions. Since the addition of circular subregions only improved the performance of metamers built from representations with wedges (and not metamers built from linear center and surround

representations), our metamers appear to identify a relevant spatial scale in which some On-midget nonlinearities operate.

### 2.3.7 Modeling spatial integration across the entire receptive field

Encouraged by the simplicity of low-dimensional scenes and their ability to elicit spikes from nonlinear processes, we next probed how simplified 16-dimensional stimuli are converted to spike outputs. In other words, what nonlinear functions operate in 16-D space? We restricted our analysis to purely spatial models by flashing  $> 60$  full-field naturalistic images and measuring the resulting spike counts in individual On- and Off-parasol RGCs. We then expanded our best spatial model in the receptive field center (Fig. 2.3.2(C), model 3) to include spatial integration across the entire receptive field.

Our first question was how to integrate luminances in the receptive field center and surround. Leveraging the simplicity of a purely spatial model, we investigated an array of possible center-surround integration architectures (Fig. 2.3.5). Our resulting models (*a* through *l*) all integrate across center wedges identically, but diverge in how surround luminances exert an effect. For instance, we tested cases where surround luminances only affect model outputs before (*a* through *d*) or after (*e* through *h*) the center undergoes rectification. Similarly, we tested the effect of first applying various operations to surround luminances, such as integration (models *b*, *f*, *j*) or rectification (models *c*, *g*, *k*), prior to integration with the center (Fig. 2.3.5(A)). Given that linear surround models (such as *a*) are subsets of nonlinear models (such as *c*), we also constrained nonlinearities to require that rectification operate within the range of intensities sampled by our images (similar to our center modeling; see Section 2.5).

For each neuron, we fit each model across all  $> 60$  images and estimated performance across smaller subsets of images. Subsets of images were selected prior to testing by reducing a large library of images into 16-D space and identifying images with wedge luminances that were expected to undergo rectification (i.e. luminances outside the expected linear range of the cell). Depending

on the location of predicted rectified wedges, images were classified as either “center-rectified”, “surround-rectified”, or “center- and surround-rectified” (Fig. 2.3.5(B)). Fitting each model to our full collection of > 60 images allowed us to reliably compare model architectures while minimizing concerns about overfitting, while the use of image subsets allowed us to closely monitor each model’s performance across collections of stimuli that might better differentiate between models. Identical model architectures were used for On- and Off-parasol RGCs, differing only by inverting the polarity of their input. Similar to modeling in the receptive field center, model performance was defined as the correlation ( $r^2$ ) of the optimal linear fit between the model output and measured spike counts.

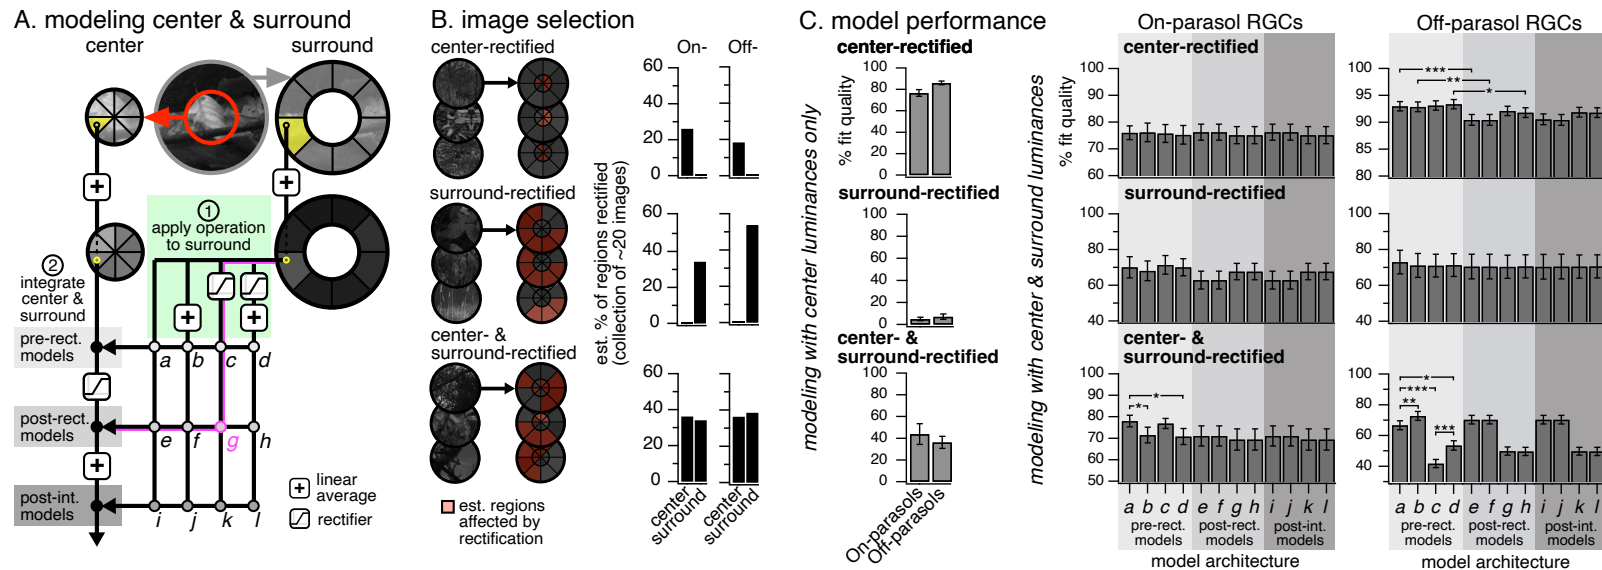


Figure 2.3.5: (A) To model spatial integration across the receptive field center and surround, we tested 12 possible model expansions to our best center model (model 3, Fig. 2). Each model applies one of four sets of operations (in green) to surround luminances. Surround luminances are then integrated with center luminances before rectification (light gray), after rectification (medium gray), or after integration (dark gray). For example, model *g* (in magenta) rectifies surround luminances and integrates them with rectified center luminances. (B) We then probed low-dimensional 16-D forms of natural stimuli to cluster images into three different collections (16 – 25 images per collection) depending on the expected location of rectification in the receptive field center and/or surround (for On- and Off-pathways). (C) All three collections of naturalistic images were flashed for 250 ms onto individual On- and Off-parasol RGCs ( $n = 7$ ,  $n = 9$ , respectively) and quantified via spike count. (Left) Our best center model (“modeling with center luminances”) was used to predict spike counts to flashed images across the full receptive field, effectively ignoring all luminances in the surround. (Right) “Modeling with both center and surround luminances” used each full-field model architecture to predict spike counts to full-field image flashes. Statistical significance is determined via paired Student’s t-tests: (\*)  $p < 0.05$  (\*\*)  $p < 0.1$ , (\*\*\*)  $p < 0.001$ .

Every model architecture performed well across our full image set, but diverged when analyzing performance across smaller subsets of images. Across center-rectified images, models in which Off-parasol RGCs integrated center and surround luminances prior to rectification performed better than those in which center-surround integration occurred after rectification, consistent with previously reported mechanisms [28, 29]. For center- and surround-rectified images, Off-parasol RGC model performance was best when surround wedges were combined prior to center-surround integration whereas On-parasol RGC models had the opposite preference. In center- and surround-rectified images, rectifying surround wedges decreased performance in Off-parasol RGCs (but not in On-parasol RGCs), with a majority of our Off-parasols containing nonlinear bounds expanded to the allowable rectification limits in our model (and hence minimizing the impact of rectification). These apparent differences in how On- and Off-parasol RGCs integrate inputs across space adds to literature that suggests significant asymmetries between On and Off pathways [40–42].

Across our tested architectures, model *a* had consistently strong performance across both On and Off pathways. Other models varied strongly across pathways; for instance, model *b* performed consistently better in Off-parasols but also consistently worse in On-parasols. Model *a* also contained the fewest fitting parameters of our tested models (4 parameters) and produced an  $r^2$  of 75% and 82% in On- and Off-parasol RGCs, respectively. Model *a* operates by reducing each image into 16 linearly-integrated wedges, weighting the center (1 parameter) and surround (1 parameter), combining each surround wedge with its nearest center wedge, and then rectifying (2 parameters) and averaging all remaining values. Table 2.1 contains averaged parameters for model *a* as fitted across populations of On- and Off-parasol RGCs.

### **2.3.8 Properties of spatial integration in 16-dimensional space**

The ability of low-dimensional stimuli to preserve most of the structure in responses to natural movies is curious since 16 wedge-shaped regions do not correspond in shape, number, or location to anatomical features of the retinal circuitry. An invariance to subunit location is consistent with

subunits providing sensitivity to the presence, but not specific location, of spatial structure within the receptive field. Equipped with (1) a spatial model that appears to describe spike counts to flashed natural images, and (2) a methodology for using naturalistic movies to test integration behaviors, we sought to ask two further questions about the properties of subunits.

First, we were curious about the effect of disrupting (decorrelating) specific pairs of center-surround wedges. In our spatial model (model *a*), each center wedge is integrated with its nearest surround wedge, and perturbing this process may affect performance. We computationally rotated the center or surround of each flashed image and asked whether the performance of our model changed (Fig. 2.3.6(A)). Rotating the receptive field center by  $180^\circ$  decreased performance by only 4% and 1% in On- and Off-parasol RGCs, respectively.

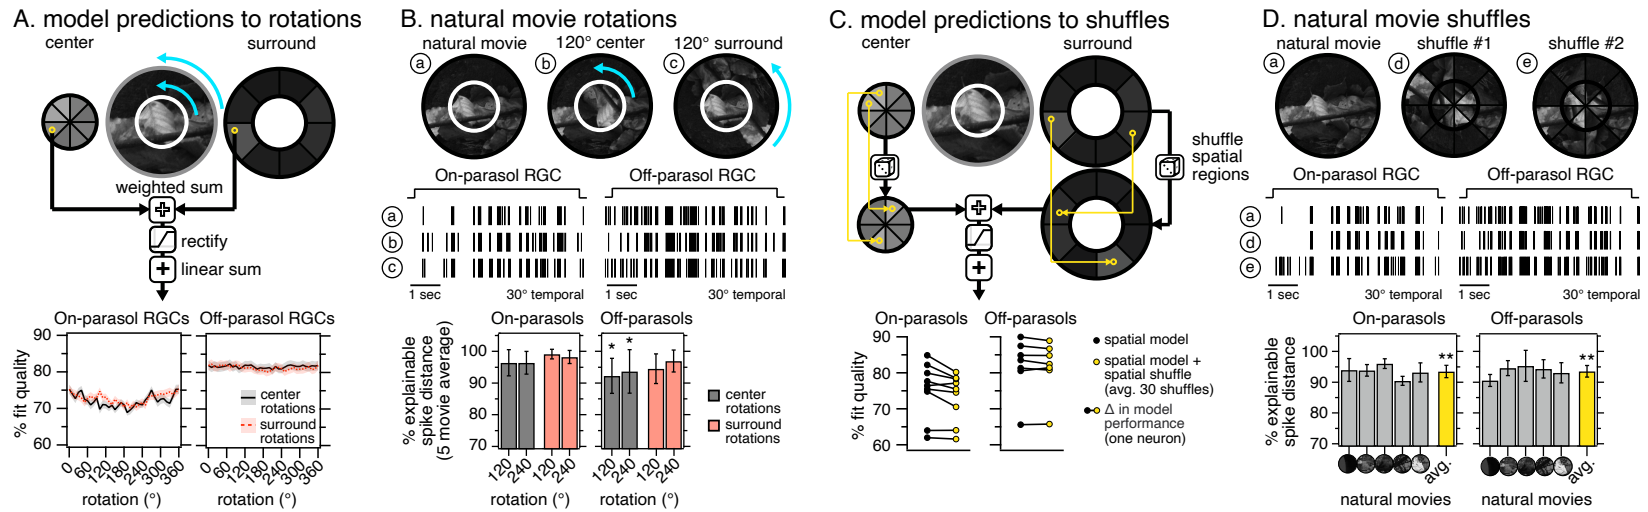


Figure 2.3.6: (A) Model performance (model a in Fig. 5) after computationally rotating the receptive field center or surround of all flashed images for On- and Off-parasol RGCs ( $n = 7$ ,  $n = 9$ , respectively). (B) For every frame of a given natural movie, the receptive-field center or surround was rotated counter-clockwise (by either  $120^\circ$  or  $240^\circ$ ). On- and Off-parasol spike responses ( $n = 5$  cells across 2 retinas) to rotated stimuli were compared to their non-rotated counterparts and averaged across 5 different naturalistic movies. Spike traces for two such movies (one movie for On-parasols, another movie for Off-parasols) are provided. (C) Model performance after randomly redistributing (“shuffling”) low-dimensional luminance wedges in both the receptive field center and surround in On- and Off-parasol RGCs. Fit quality for each neuron (paired black/yellow dot) was averaged across 30 unique shuffles. (D) In a single frame of a natural movie, spatial features were divided into 16 natural wedges, shuffled, and placed in a new location. Shuffling defined in our first frame was then identically repeated for each subsequent movie frame and resulted in “kaleidoscope” stimuli. On- and Off-parasol spike responses ( $n = 5$  cells across 2 retinas) to kaleidoscope stimuli were compared to their non-shuffled counterparts, with spike traces to two different movies (one movie for On-parasols, another movie for Off-parasols) provided. Error bars across populations of neurons are SEM (A, C, individual movies in D); error bars across movies are S.D (movie averages in B, D). Statistical significance is determined via Student’s t-tests (relative to 100% spike distance): (\*)  $p < 0.05$  (\*\*)  $p < 0.1$ , (\*\*\*)  $p < 0.001$ .

We then empirically tested the effect of center-surround decorrelations using naturalistic movies. We rotated the receptive field center or surround of each movie frame and asked whether spike responses from On- and Off-parasol RGCs changed relative to the original naturalistic movie (Fig. 2.3.6(B)). The same rotation applied to one movie frame was applied across the entire movie to avoid introducing motion artifacts. When averaged across 5 full-field naturalistic movies, movies with rotated centers accounted for 96% and 93% of full-field spikes for On- and Off-parasols, respectively. Only rotations to the receptive field center of Off-parasols performed consistently below 100% spike distance across five natural image movies.

Second, the relatively small impact of decorrelating features between the receptive field center and surround encouraged us to ask whether parasol RGCs can detect decorrelated features within the receptive field center or surround. We returned to our spatial model (model *a*) and computationally integrated each image into uniform wedges. We then swapped the location of any two center wedges and repeated this operation -- swapping center wedges within the center and surround wedges within the surround -- until each image was maximally disrupted (i.e, no wedge remained in its original location and a supermajority of wedges had different neighbors). In the context of our spatial model, each pair of center-surround wedges are integrated independently of each other, and are thus only sensitive to pairing a center wedge with a different surround wedge. Disrupting intra-center and intra-surround structure, on average, resulted in 2% worse performance in On-parasols and negligibly worse performance in Off-parasols (Fig. 2.3.6(C)).

To test whether invariance to swapping individual wedges was a failure of our spatial model (which does not consider correlations between subunits) or a property of spatial integration in parasol RGCs, we turned to empirical measurements with naturalistic movies. We divided each movie frame into 16 natural image wedges, and, similar to our computational work, repeatedly swapped each naturalistic wedge -- keeping center wedges within the center and surround wedges within the surround -- until the original image was maximally disrupted. The precise set of shuffling operations applied to one movie frame was repeated for each subsequent movie frame

and produced “kaleidoscope” stimuli. We presented two to five shuffled variants for each of five different natural image movies and observed that, on average, responses to kaleidoscope stimuli were 94% similar to responses to the original full-field movie in both On- and Off-parasol RGCs (Fig. 2.3.6(D)). The decrease in performance with shuffling of regions was small but significant in both On- and Off-parasol RGCs. While the shuffling of spatial regions alters the specific point in 16-dimensional space that natural movies lie in, our observation of preserved spike responses suggest that kaleidoscope stimuli all reduce to a similar point in 1-dimensional space. Such experiments further support a model in which subunits are integrated independently of each other.

### **2.3.9 Leveraging spatial integration mechanisms with “degenerate” stimuli**

Our work serves as the basis for a general class of “degenerate” stimuli: unique images and movies that leverage the integration strategies of single neurons to elicit similar spike responses (Fig. 2.3.7). One version of degenerate stimuli, subunit-targeted degenerate stimuli, are designed to activate a set of subregions similarly (Fig. 2.3.7(A)). Our metamers are one example of subunit-targeted degenerate stimuli, and operate by tuning subunit inputs to stimulate similar linear and nonlinear integration mechanisms, resulting in similar spike outputs. Subunit-targeted degenerate stimuli share the same 16-dimensional representation (i.e. reside at the same point in 16-dimensional space).

Other versions of degenerate stimuli can similarly be designed to capture integration across subunits. Nonlinear subunits are summed to produce a single integrated luminance value. Different stimuli with the same integrated value may produce a similar spike response even if their 16-dimensional representations diverge (Fig. 2.3.7(B)). Such stimuli would be characterized as “parasol-targeted degenerate stimuli.” The diversity of stimuli that fall into this category is greatly increased by observations of (1) location invariance of individual subunits and (2) the lack of correlations between subunits, both of which create additional degeneracies in how subunits are integrated into spike outputs (Fig. 2.3.6).

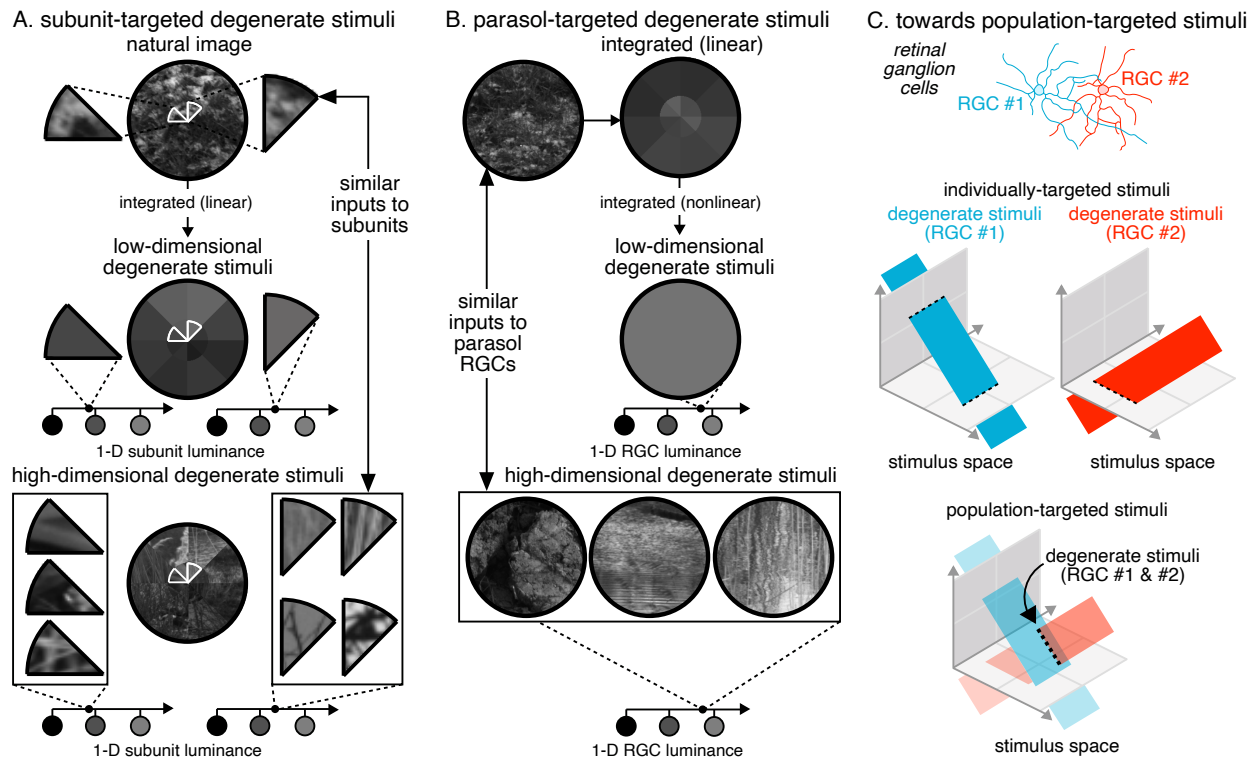


Figure 2.3.7: (A) Subunit-targeted low-dimensional degenerate stimuli subdivide a parasol RGC receptive field into 16 linearly integrated regions. To produce high-dimensional variants, each region can then be replaced with another image that reduces to a similar 1-dimensional subunit luminance. (B) Parasol-targeted low-dimensional degenerate stimuli subdivide a parasol RGC receptive field into 16 linearly integrated regions and then nonlinearly integrate to produce a 1-D luminance. This operation requires measured nonlinear parameters (see Table 2.1 for nonlinear parameters from spatial modeling), but can identify other natural scenes that are integrated across the entire receptive field similarly. (C) Theoretically, targeted stimuli can be scaled to encode similarly across populations of RGCs by matching the 1-D integrated luminance of all RGCs within a given region of space. This requires the spatial location and receptive field of each RGC as well as individual nonlinear parameters (see Table 2.1 for averaged parameters across parasols), but not the location of specific interneurons or photoreceptors.

By design, degenerate stimuli only elicit similar responses within a single neuron, but may be scalable to produce similar responses in populations of neurons simultaneously (Fig. 2.3.7(C)). The design of such images would require leveraging spatial integration mechanisms across all neurons (of the same cell type) within a given region of the retina. Given our observations that low-

dimensional stimuli can also improve encoding in On-midget ganglion cells (Supp. Fig. 2.6.2), we predict that other cell types in the retina may also contain useful degeneracies that can be similarly probed by targeted stimuli. Such stimuli would elicit degeneracies in some neurons but not others due to differences in spatial integration. Current approaches to isolate signals originating from different RGC types (e.g. parasol vs midget RGCs) are based on classical receptive field properties. An approach like that illustrated in Figure 2.3.7(C) could exploit a more complete understanding of spatial integration to help reveal which RGC types participate in specific higher visual functions.

## 2.4 Discussion

We investigated the spatial structure of natural scenes at the level of retinal ganglion cells (RGCs) by comparing the spike responses elicited by naturalistic movies, simplified (“reduced-dimensional”) naturalistic movies, and high-dimensional “degenerate” movies (“metamers” and “kaleidoscope” movies). Reducing movies to eight linearly-integrated spatial regions in the receptive field center and surround captured ~80% of spiking behavior in On- and Off-parasol RGCs. Performance generally declined for representations with substantially fewer than 16 subregions. We then used computational modeling to assess how 16-dimensional static images are integrated to control parasol spike counts to flashed full-field natural images. We consider each of these in more detail below.

### 2.4.1 Spatial integration and receptive field subunits

Spatial integration mechanisms in the retina dictate how signals from several hundred cone photoreceptors converge to control a single parasol spike output. The ability to accurately predict parasol spike responses to visual stimuli requires an understanding of the many linear and nonlinear spatial integration mechanisms that shape cone photoreceptor signals. The specific order in which integration mechanisms are applied across inputs is also important: linear integration followed by a nonlinearity is not the same as a nonlinearity followed by linear integration. Identifying the ordered set of operations that dictate parasol spike responses can highlight the key spatial features that drive encoding and other spatial features that do not.

Classical center-surround receptive field models predict that linear integration over space and time reduces photoreceptor inputs to a 1-dimensional signal. Linear-nonlinear and generalized-linear models use linear integration as an initial step followed by a single nonlinearity. Such models

have identified linear subspaces that describe neuronal responses to artificial stimuli [8]. Decades of physiological studies, however, have observed that notable nonlinearities occur in both space and time. Temporally, photoreceptor responses are strongly shaped by adaptation and hence are poorly described by linear models [43–45]. Spatially, bipolar cells rectify (a nonlinear process) the signals that provide inputs to RGCs [3, 12, 14]. Such nonlinear mechanisms are strongly stimulated by the complex visual features in the natural world, requiring improved models to describe vision under natural conditions [7].

Empirical measurements with gratings indicate that spatial nonlinearities in parasol responses have a diameter about one fifth the diameter of the receptive field center [25,26], consistent with the anatomical density of the diffuse bipolar cells that provide excitatory input to the parasol cells [22]. Such observations predict that models for RGC responses can be improved by identifying the locations of the bipolar cells that comprise nonlinear receptive field subunits. Indeed, computational approaches that leverage large, diverse sets of stimuli can identify receptive field subunits, and the inclusion of nonlinear subunits can improve model predictions [8, 31, 35]. In some cases the extracted subunits correspond closely to bipolar cells [32], suggesting that RGC responses depend on specifics of the upstream retinal circuitry. In other cases, however, the connection between bipolar cells and identified subunits is less clear. The performance of maximum likelihood models saturates when including more than four subunits [46]. In parasol cells specifically, improvements in predicted responses saturate with as few as five subunits [34], substantially fewer than the ~25 bipolar cells expected from anatomy.

Our results similarly suggest that models that account for the spatial nonlinearities present in the retinal circuitry may not need to incorporate specific subunit shapes or locations. Our empirical measurements support a surprisingly flexible description for how nonlinear subunits contribute to parasol responses that is based entirely on the classical receptive field. We find that dividing the receptive field center and surround into 8 wedge-shaped linearly-integrated regions improves our ability to describe parasol responses, including spikes evoked by spatially-nonlinear mechanisms.

Interestingly, the regions need not correspond in size, shape, or location to anatomical subunits. Instead, we only need to reach the minimum number of regions needed to describe a stimulus while fully covering the receptive field. This simplification predicts that the specific locations of subunits are not needed to understand the relationship between the spatial structure of natural images and parasol responses.

The identification of a simple reduced-dimension space enabled the design of “degenerate” stimuli that are tuned to leverage the specific spatial integration properties of each cell type. Metameric stimuli shared the same reduced representation as an original natural movie but had very different spatial structure. The similarity of responses elicited by original and metameric stimuli suggests that our dimensional-reduction approach, despite its simplicity, captures the key spatial features that control parasol responses. This strengthens the conclusion that spatial integration can be largely explained without specific information about subunit location and shape.

The independence we observe across individual wedges (see: “kaleidoscope” stimuli) offers additional flexibility for nonlinear models. We find that each wedge-shaped region of space is weighted equally in our models during integration, and correlations between adjacent wedges do not strongly affect parasol spike responses. These results support a model in which subunits integrate inputs independently of each other, and as a result, suggest parasol responses are sensitive to the presence of spatial structure but are ambiguous about the precise pattern or location. This is consistent with the spatial invariance observed in ganglion cell responses to artificial gratings [26,47]. As a direct result, models for retinal encoding may only need to consider one subunit (i.e. wedge-shaped region of space) at a time. The specific set of operations applied to one subunit can then be identically repeated for all subunits that fill the receptive field.

## **2.4.2 Model limitations**

While our 16-dimensional model captures much of the structure of parasol responses to natural movies, it does not represent a “complete” model. Here, we discuss several limitations and possible

extensions of our approach.

First, our reduced-dimensional stimuli neglect nonlinearities that operate on a fine spatial scale - such as adaptation in the cone photoreceptors, which is engaged strongly and rapidly by changes in luminance [45]. Individual cones encounter much smaller modulations in luminance in our reduced-dimensional movies compared to their original natural form, especially in movies with strong local image structure at spatial scales much smaller than subunits. As a result, cones within the same wedge of our reduced-dimensional stimulus will adapt similarly, whereas the same cones will often adapt differently in response to the original natural movie. Our high-dimensional degenerate stimuli (i.e, metamers) may help reveal the effects of such local cone adaptation. Cones are expected to encounter even more extreme changes in luminance during our metamer movies since spatial features at the scale of cones change every frame. Adaptational effects may underlie why some metamers performed  $\sim 10\%$  worse than their reduced-dimensional forms.

Second, linear wedges may not be an ideal architecture for reducing the dimensionality of input in the receptive field surround. Introducing wedges into the surround resulted in large increases in performance for some movies but not others; this differs from the consistent improvements of adding wedges of the receptive field center (Fig. 2.3.3(C)). Given the limited impact of the surround on natural movies (Fig. 2.3.3(A)), resolving this issue likely will benefit from more directed experiments investigating how spatial inputs to the receptive field surround are integrated.

Finally, our empirical approach probes dimensional-reduction in the context of movies with both spatial and temporal behavior. Our spatial modeling uses flashed images and spike counts to probe simpler (purely spatial) stimuli, and generally shows more subtle differences across models than those observed by comparing reduced and natural movies. Incorporating time into full models for parasol spike responses is an important future direction. Such temporal behaviors can be quite complex to model: for instance, both photoreceptor adaptation and dynamic shifts in bipolar cell nonlinearities cause responses to depend on stimulus history [45,48]. Models for full spatiotemporal movies may be easier to investigate by starting in our reduced dimensional space, which retains

most of the temporal nonlinearities that shape responses to full natural movies.

### 2.4.3 Characterizing the relevant spatial statistics of natural images

Visual inputs from the natural world contain a complex, high-dimensional spatial structure that stimulates multiple nonlinear mechanisms in neural circuits. Understanding the link between specific spatial features and specific nonlinear behaviors in neurons remains an open question in visual neuroscience. Our work here shows that the nonlinearities that shape the integration of spatial inputs are retained when natural images are linearly projected into a simple 16-dimensional space. The performance of high-dimensional “degenerate” metamers (Fig. 2.3.4(B)) provided a direct test of the completeness of this space: the original movie and matched metamer share little other than the same reduced dimensional representation, yet elicited very similar spike responses.

Identifying a low-dimensional space that captures the statistics of natural inputs that are relevant for parasol responses could provide a useful tool for studies of subsequent visual processing. For example, low-dimensional spaces can also be used to design stimuli that are more easily manipulated than fully-natural inputs and hence provide a bridge between artificial stimuli aimed at specific mechanisms and complex natural stimuli. Further, higher-level computations that rely on parasol responses presumably only have access to the image statistics present in reduced-dimensional space (generalized to parasol populations) (see Fig. 2.3.7). Incorporating such dimensional reduction into stimulus design and/or models for higher-level computations could help determine which visual tasks are predominantly enabled by parasol vs midget responses.

## 2.5 Methods

**Tissue preparation** We obtained retinal tissue from Macaque monkeys (*M. nemestrina*, *M. mulatta*, or *M. fascicularis*) via the tissue distribution program at the Washington National Primate Research Center. All procedures were approved by the Institutional Animal Care and Use Commit-

tee at the University of Washington. Results did not depend obviously on age (2-18 years) or sex, although our ability to discern such differences is quite limited. Dissection procedures have been described previously [25]. After enucleation, the eye was hemisected and the vitreous humor was removed mechanically, sometimes assisted by treatment with human plasmin (~50 g/mL, Sigma or Haematologic Technologies Inc.). The retina was dark adapted for ~1 hr, and all subsequent procedures were performed under infrared light using night-vision goggles. The retina and pigment epithelium were separated from the sclera and stored in oxygenated (95% O<sub>2</sub>/5% CO<sub>2</sub>) Ames bicarbonate solution (Sigma) in a light-tight container. Retinal mounts with the pigment epithelium attached were placed onto a poly-D-lysine coated coverslip (BD biosciences) with the RGCs facing up. During experiments, retinal tissue was perfused at 7–9 mL/min with Ames solution at ~32° C. Recordings were made from the peripheral retina with an average eccentricity of 25° ± 10° (mean ± S.D.).

**Generating naturalistic stimuli** Spatiotemporal stimuli were generated using eye motion trajectories of freely-moving observers viewing still naturalistic images [36]. Each frame was downsampled to 60 Hz and segmented into spatial regions. The linear-equivalent luminance,  $\mathcal{L}$ , of each region was calculated using Eq. 1.2.1.

Metamers were generated by calculating linear-equivalent luminance across a library of 1826 patches sampled from natural images. The patch which had a linear-equivalent luminance closest to that of the original movie was selected. This process was repeated independently for each spatial region in each frame. The resulting spatial metamers included a collection of patches from multiple images, with the caveat that no two image regions within a single movie frame could be sampled from the same original image. A MATLAB package for generating naturalistic movies, reductions, and metamers is available on GitHub: <https://github.com/jfreedland/reduce-natural-scenes>.

**Stimulus projection** Pre-generated .mp4 movies were focused onto RPE-attached macaque photoreceptors via an OLED microdisplay monitor (eMagin). Each monitor pixel (600 x 800 pixels)

corresponded to 1.65  $\mu\text{m}$  or 0.5 arcmin. Presented movies (300 x 400 pixels) were enlarged to match the spatial scale of physiological eye movements from the DOVES database (1 pixel per 1 arcmin) [36]. The average pixel intensity across naturalistic images used in the study are similar (16.3-18.6% of total monitor intensity); 17 monitor intensity was scaled to correspond to 5000 L/M-cone isomerizations per second. Monitor outputs were linearized via gamma correction. Stimuli presentation and data acquisition was guided by MATLAB software packages Stage (<http://stage-vss.github.io>) and Symphony (<http://symphony-das.github.io>).

**Cell identification & acquisition** Retinal ganglion cells were initially categorized by soma size, receptive field size, and spike transience in response to a step in light intensity. Cell type was confirmed by observing unique responses to naturalistic movies. Recorded cells were screened for sensitivity, and data was collected only from cells that exhibited at least 20 sp/s modulation of their firing rate in response to a uniform 5 contrast, 4 Hz-modulated stimulus at 5000 R\*/cone/s. Prior to data collection, each neuron was centered relative to our monitor via the presentation of a 4 Hz contrast-reversing split field stimulus. Specifically, this stimulus contains two halves of a circle (300  $\mu\text{m}$  in diameter) with light intensities that oscillate out-of-phase sinusoidally between  $\pm 90\%$  contrast. This grating is then continuously moved until both grating halves produce similar spike rates in our neuron over a 2-second presentation. Centering was repeated in both x- and y-directions. After centering, the receptive field size of each neuron was probed by flashing (250 ms) circular spots with sizes between 10 and 1000  $\mu\text{m}$  (see Results). We then used a difference of gaussians fit to describe spike response as a function of spot size which defined the receptive field in our calculation of linear-equivalent luminance (Eq. 1.2.1).

**Explainable spike distance** Spike responses for reduced stimuli were compared to naturalistic responses via explainable spike distance. The spike distance quantifies the cost for transforming a spike train response for one trial of an experimental stimulus (reduced movie) into a spike train response for a trial of the control stimulus (naturalistic movie) [37]. The two spike trains are matched

by a combination of shifting spikes in time, with a distance proportional to the time shift, and deleting and adding spikes. We used a time shift distance of 3 ms and the distance associated with deleting or adding a spike was 1. The algorithm then identifies the specific sequence of operations required to ensure both spike trains exactly match while maintaining the overall lowest distance. We averaged spike distances across all possible pairs of responses to the two stimuli with each spike distance calculated using publicly available code (<http://www-users.med.cornell.edu/~jdvicto/spkdm.html>).

To calculate the “explainable” spike distance, we compared the calculated spike distance with an upper bound formed from random spike trains with the same spike count. Specifically, this upper bound was calculated by generating 20 spike trains (with the same number of spikes as an experimental spike train placed at random times) and calculating the average spike distance relative to a control spike train. The lower bound was identified by calculating the spike distance across repeats of the same control stimulus in the same neuron. If the spike distance for a given reduced image equaled the lower bound (corresponding to 100% explainable spike distance), the responses to that stimulus were indistinguishable from those to the original stimulus (within a margin of error defined by the underlying noise of the cell). If the spike distance equaled the upper bound (corresponding to 0% explainable spike distance), none of the structure of the response to the original stimulus was retained in the response to the reduced stimulus (within a margin of error defined by the underlying noise of the cell).

Although the spike distance is highly sensitive to the specific time shift distance (3 ms in this study), we found that the “explainable” spike distance was largely invariant of this value because the “explainable” variance (noise) of each cell scaled proportionately with the time shift distance. We also observed that it was possible for two spike trains to have slightly over 100 explainable spike distance. This occurred if two spike trains were more similar than expected by the underlying noise of the cell (i.e., 98% similar spike trains in a cell with 5% spike variability as measured across repeated trials of the same stimulus).

**Modeling** All models in this study used numerical optimization methods to fit parameters. Our data consisted of a series of flashed images, spike counts, and receptive field information. We computationally reduced our images into low-dimensional space according to Eq. 1.2.1. Reduced luminances were then adjusted from 8-bit luminances (0 to 255) to Weber contrasts via Eq. 1.2.2.

For each model architecture, we introduced a series of nesting functions that transform fractional contrasts into spike counts via a series of fitting parameters. Nesting functions define equations with free variables, allowing for the design of basic linear and nonlinear stages that we could interweave to build more complex models. After building each model, we numerically maximized the correlation ( $r^2$ ) between fractional contrast ( $c$ ) and measured spikes using `fminsearch` in MATLAB.

For each model, numerical optimization was repeated 20 times with each subsequent optimization initialized using optimal fitting results from the previous one. Each optimization was allowed to iterate up to 10,000 times. To prevent overfitting, we modeled performance across complete batches of flashed images for a given cell. We then used the individual cell's globally-fitted parameters to calculate fit quality across smaller batches of images chosen to challenge the model in specific ways (“center-rectified”, “surround-rectified”, and “center- & surround-rectified” in Fig. 2.3.5).

Constraining certain parameters imposed a separation between linear and nonlinear models and aided in interpretation of the modeling results. For instance, a positive nonlinearity across our models could not remain linear below  $-40\%$  contrast or above  $+150$  contrast. Such bounds restricted nonlinearities from infinitely expanding into a linearity and allowed us to differentiate more strongly between linear and nonlinear models (i.e models  $a$  and  $c$  in Fig. 2.3.5).

## 2.6 Supplementary Figures

	Weight of receptive field center	Weight of receptive field surround	Negative rectification (% contrast)	Positive rectification (% contrast)
On-Parasols	87.3%	-61.2%	-17.6%	70.8%
Off-Parasols	88.9%	-56.0%	-75.5%	11.3%

Table 2.1: Fit parameters for the full-field modeling architecture presented in Fig. 2.3.5(A) (model *a*) as averaged across 9 On- and 7 Off-parasol RGCs.

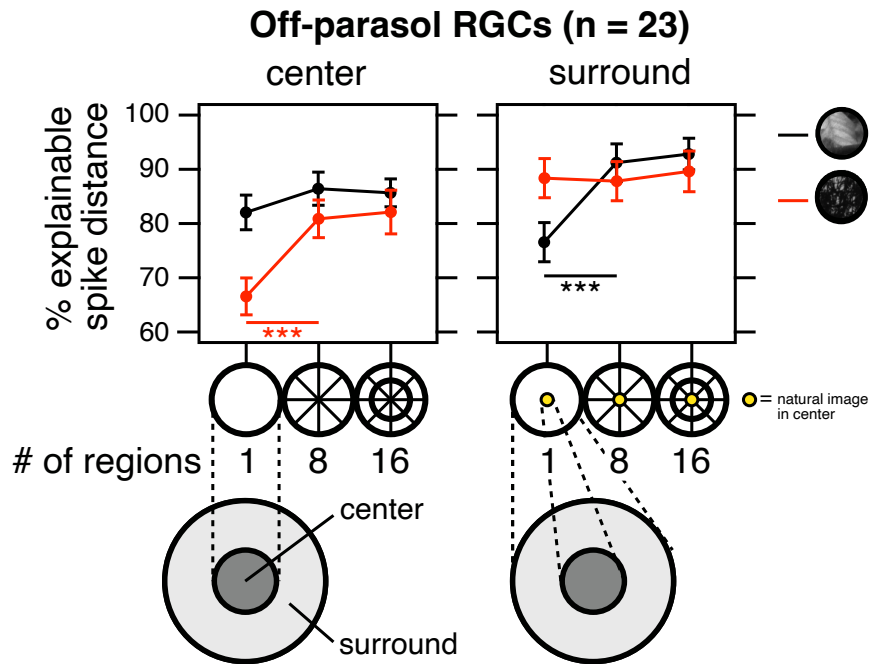
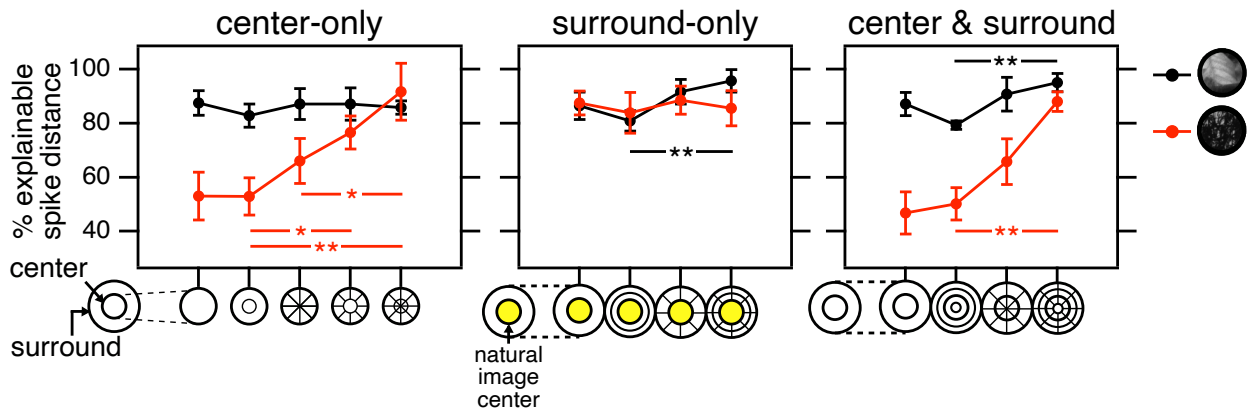


Figure 2.6.1: Explainable spike distance for Off-parasol RGCs ( $n = 23$ ) upon presenting two natural image movies with reduced linear regions in the receptive field (left) center or (right) surround. Reduced movies with 1 and 8 regions are similar to center-only stimuli in Fig. 2.3.1(C) (left) and surround stimuli in Fig. 2.3.3(C) (right). Reduced movies with 16 regions contain eight divisions further divided into a circular near- and far-center/surround. The near-center/surround is defined as a continuous spatial region that captures 50% of measured excitatory/inhibitory behavior. Error bars are SEM. Statistical significance is determined via paired Student's *t*-tests: (\*)  $p < 0.05$  (\*\*)  $p < 0.1$ , (\*\*\*)  $p < 0.001$ .

### A. reduced stimuli in On-midget RGCs



### B. metamers in On-midget RGCs

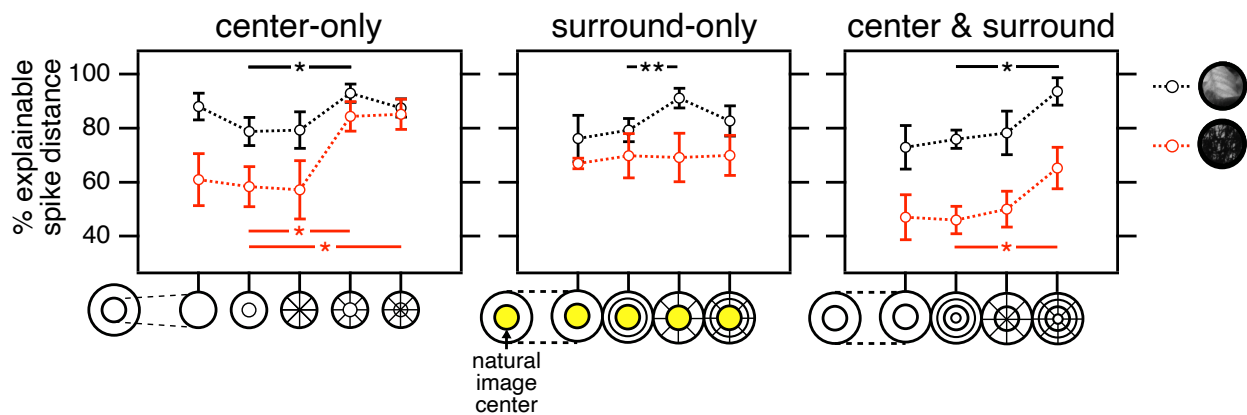


Figure 2.6.2: Encoding properties of populations of On-midget RGCs ( $n \geq 8$  across  $\geq 5$  retinas) for (A) simplified (reduced) stimuli and (B) metamers presented to in the receptive field (left) center, (middle) surround, and (right) full receptive field. Spike responses to reduced and metameric stimuli in the receptive field center were compared to natural movies presented only to the receptive field center; surround and full-field movies were compared to the original full-field natural movie. All experiments in the receptive field surround contained the original natural movie in the receptive field center. Movies with additional circular divisions within the center and surround (i.e, second from left) were placed at a location that captures 50% of measured excitatory/inhibitory behavior. Error bars are SEM. Statistical significance is determined via paired Student's t-tests: (\*)  $p < 0.05$  (\*\*)  $p < 0.1$  , (\*\*\*)  $p < 0.001$ .

# Chapter 3

## Distorting the spatial structure of scenes

This work is still in development.

**Working title** Disrupting the spatial structure of visual inputs in the receptive field center delays spike responses in parasol retinal ganglion cells.

### 3.1 Introduction

Spatial integration regimes within the receptive field center of primate parasol retinal ganglion cells (RGCs) greatly reduce the dimensionality of natural inputs. Classical models integrate the receptive field center linearly, generating spike responses by summing spatial features (as weighted by the neuron's receptive field) to produce a one-dimensional value (Eq. 1.2.1) [4, 10]. Other stimuli -- including many natural scenes -- are better described by models that sum across small numbers of nonlinear subunits (see review: [49]). As a result, the reliable prediction of spike responses to broad sets of stimuli require updated models that account for various nonlinear mechanisms within the retinal circuit [50].

Previous work from our laboratory observed that spatial models with approximately eight wedge-shaped regions tiled across the receptive field center captured the spike responses of natural movies better than linear models [18]. These wedge-shaped regions appeared to match expectations for functional subunits by rectifying inputs similarly to bipolar cells within the retinal circuit [14]. Further experiments revealed flexibility in how wedge-shaped subunits were integrated: for many natural movies, wedge-shaped regions could be freely swapped with each other or replaced

with entirely different images with minimal effects on encoding [18]. Such work highlights how disrupting the high-dimensional spatial structure of visual inputs (i.e, swapping or replacing wedges) may not affect RGC spike responses if alterations are filtered out by the retinal circuit during integration.

The insensitivity of parasol RGCs to changing wedge-shaped spatial regions encouraged us to ask whether other spatial disruptions (different locations, sizes, or shapes) exert stronger effects on encoding. More broadly, identifying particularly small disruptions to spatial structure that elicit large changes to spike response (or large disruptions that elicit negligible changes to spike response) might reflect particularly sensitive (and insensitive) regions within the receptive field. Adjusting the spatial features of low-dimensional stimuli (i.e, gratings and textures) have previously identified spatially nonlinear behaviors within the retinal circuit [26, 51]. By focusing our efforts towards high-dimensional stimuli, we hoped to identify key spatial areas within the receptive field that specifically mediate the functional encoding of complex natural scenes.

Here, we disrupted the spatial structure of natural movies and asked whether these changes were reflected in the spike responses of parasol retinal ganglion cells. We observed that rotating certain circular spatial regions globally delayed spike responses relative to the original natural movie. We observed similar delays after applying spatial distortions to artificial split-field spots due to delays in excitatory inputs and sped up inhibitory inputs. Finally, we quantified temporal delays throughout the receptive field center, building spatial maps that suggest a dependency on the size and location of perturbed spatial structure. Our results highlight how stimuli with spatial discontinuities can be used to interrupt and study important functional integration mechanisms.

## **3.2 Results**

We sought to investigate how the addition of spatial discontinuities in naturalistic stimuli shaped encoding within the receptive field center of parasol retinal ganglion cells. We observed that rotating

specific regions in both natural movies and artificial gratings globally delayed spike responses. We then used low-dimensional stimuli to map temporal delays throughout the receptive field center, and observed that delays depended on both the location and size of spatial discontinuities.

### **3.2.1 Adding circular spatial discontinuities to natural movies delay spike responses**

We began by investigating the effects of rotating spatial structure within the receptive field center of On-parasol retinal ganglion cells. We first identified a stimulus of interest. We then divided the stimulus into a series of circular spatial regions and independently rotated each region to form a “distorted” stimulus (Fig. 3.2.1(B), left). We presented both stimuli -- distorted and undistorted -- with an aperture that restricted inputs to the receptive field center of On-parasol retinal ganglion cells. We then compared spike responses elicited from distorted and undistorted stimuli to gauge the effects of added circular rotations on encoding.

The use of circular rotations to disrupt spatial structure may offer unique insights into spatial integration mechanisms. Rotated pixels always retain their original distance relative to the absolute center of the receptive field center and are thus weighted identically by a classical receptive field (Fig. 3.2.1(A)). As a result, a linear model (with a circularly symmetric weighting function) will elicit the same output from both rotated and unrotated stimuli (Fig. 3.2.1(B), left). In contrast, rotating spatial structure will severely interrupt the luminance of integrated wedges as local spatial features change (Fig. 3.2.1(B), left). Thus, stimuli with several independently-rotated regions may specifically target functional mechanisms associated with nonlinear, but not linear, integration regimes.

To investigate spatial disruptions under natural conditions, we generated high-dimensional natural movies using a database of human eye movements viewing still naturalistic scenes [36] (Fig. 3.2.1(B), right). Each six-second movie sampled a natural image as viewed over multiple

fixations and saccades. We disrupted the spatial structure of each movie by isolating a single movie frame and independently rotating various numbers of circular regions. The specific rotation applied to each circular region in one movie frame was repeated exactly in every other movie frame. Hence, while rotations alter the spatial relationship between pixels in a single movie frame, individual pixels contain the same temporal trajectory of light intensities over the course of the entire movie.

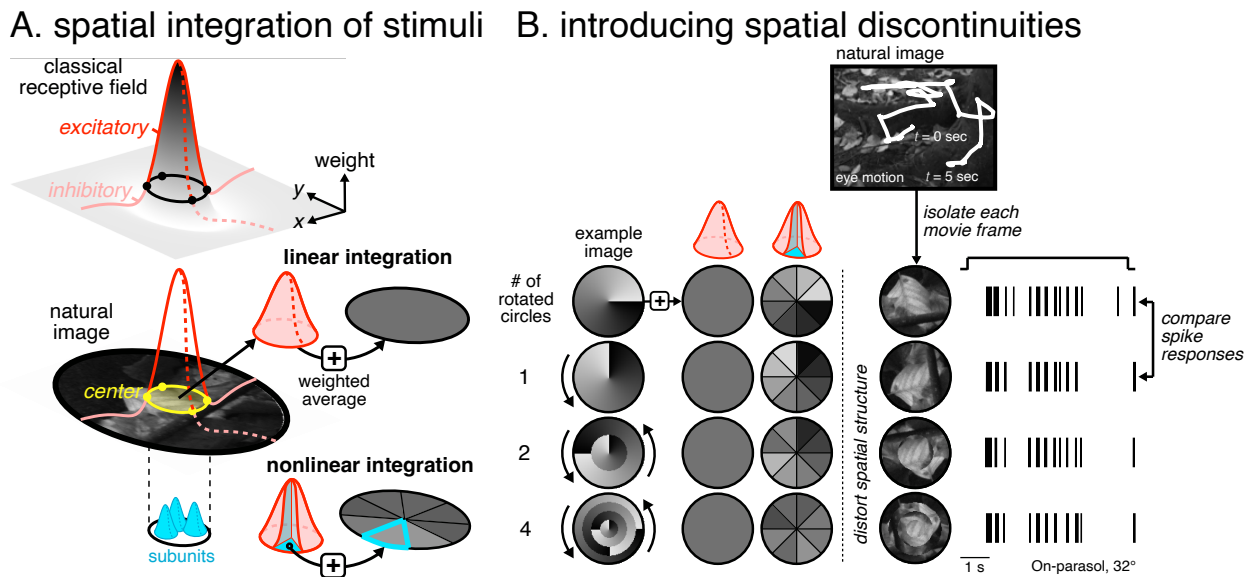


Figure 3.2.1: (A) Models for linear and nonlinear spatial integration within the receptive field center of retinal ganglion cells. (B, left) To introduce circular spatial discontinuities, an image was first divided into circular subregions. Each circular subregion was then independently rotated to produce “distorted” stimuli with similar linearly integrated values but different nonlinearly integrated values. (B, right) Using human eye movements from (9), we generated 6 second movies with a variety of naturalistic inputs typically encountered by retinal ganglion cells. Circular spatial discontinuities were applied to each movie frame and presented to the receptive field center of an On-parasol retinal ganglion cell (32° temporal). Spike responses from movies with and without circular spatial discontinuities were compared to determine whether spatial discontinuities affected encoding.

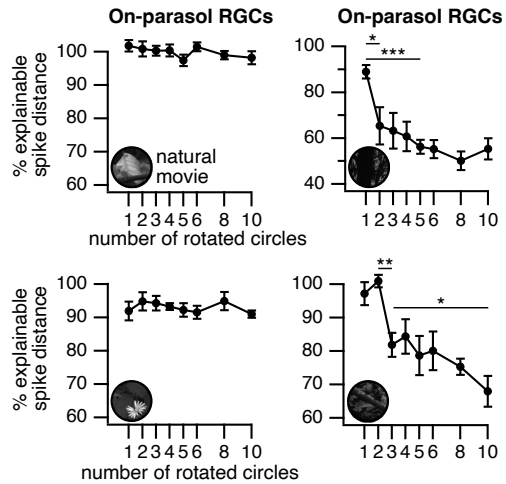

We compared spike responses elicited from distorted and undistorted movies using the explainable spike distance (see: Section 2.5). Briefly, the explainable spike distance computes the minimum cost to exactly match two spike trains by adding, removing, or shifting spikes (Fig.

3.2.2(A), top). This metric was bounded for each individual neuron. We defined that any stimulus with 0% explainable spike distance performed as well as a spike response with the same number of spikes placed at random times. Additionally, any stimulus with 100% explainable spike distance performed as closely to the original natural movie as noise allows (with noise quantified across multiple repeats of the same movie) (see: Section 2.5).

We randomly sampled a variety of natural movies and asked whether adding circular rotations affected encoding in the receptive field center of On-parasol retinal ganglion cells. (We also sampled Off-parasols, which will be included in a future draft of this manuscript). For two natural movies, rotating multiple circular rotations imposed little to no effect on the explainable spike distance, indicating that rotated movies elicited similar spike responses to the original movie. In such cases, independently rotating each of 10 spatial divisions (of equal width) throughout the receptive field center produced spike responses that were only 1% and 4% worse than spike responses to movies with a uniformly rotated center (Fig. 3.2.2(A), left column). Furthermore, disrupting any number of circular regions (between 1 and 10 regions) consistently captured >90% spike distance (Fig. 3.2.2(A)). Invariance to the rotation of multiple circular regions within the receptive field center indicates that a linear model is sufficient to capture spatial integration for certain movies.

For two other movies, however, rotating spatial structure elicited much stronger changes to the explainable spike distance, indicating that rotations strongly affected spike responses (Fig. 3.2.2(A), right column). For the movie in the top right panel, uniformly rotating the entire center captured 89% of the spike distance. Independently rotating just two circular regions, however, resulted in an additional 24% decline in spike distance relative to the original movie. Given that rotations affected the encoding of some movies but not others (in the same population of On-parasol retinal ganglion cells), certain high-dimensional spatial features may dictate whether circular spatial discontinuities strongly impact encoding or not. This observation generally agrees with observations that On-parasols integrate some (but not all) natural stimuli linearly [25, 52].

## A. distorting natural movies



## B. effects of circular rotations on encoding

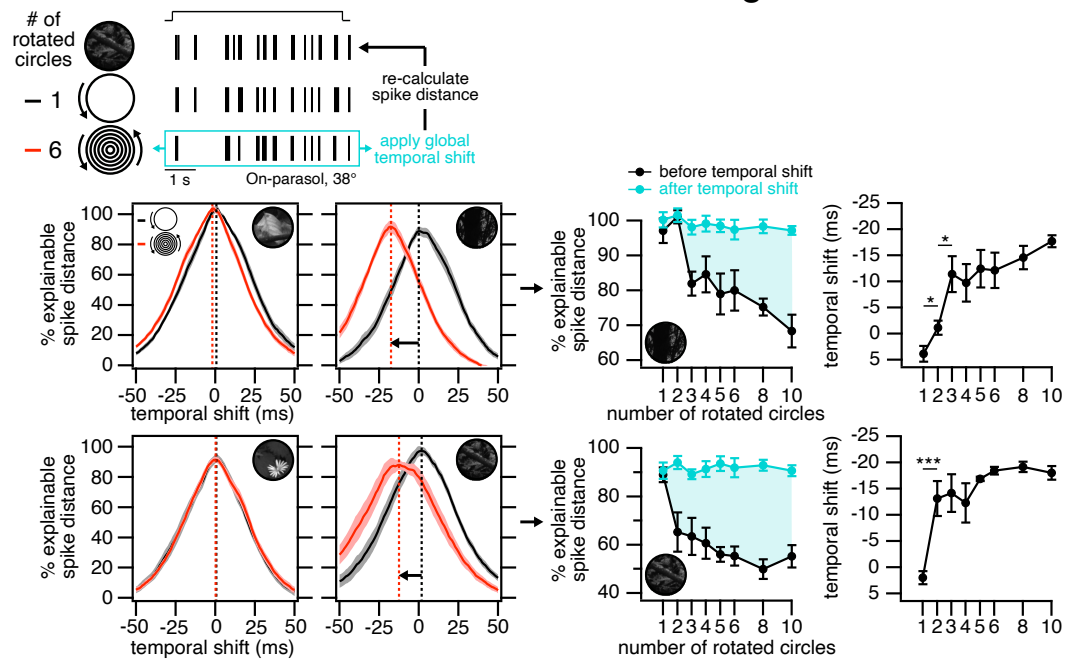


Figure 3.2.2: (A, top) Spike responses of On-parasol from naturalistic movies with added spatial discontinuities (“distorted movies”) were compared to their original counterparts using the explainable spike distance, a cost function for exactly matching one spike train to another (see: Methods). (A, bottom) Explainable spike distance for four natural movies (denoted via inset images) in the receptive field center of On-parasols with various numbers of spatial circular discontinuities (generated by rotating each circular region independently) ( $n = 7$  neurons across 3 retinas). (B, top) Each spike in response to a distorted movie was identically shifted in time and compared to the original movie’s unshifted spike response. (B, bottom) Explainable spike distance for four natural movies after globally shifting the distorted movie’s spike response. (B, right) Temporal shifts that maximize the spike distance. Negative temporal shifts correspond to initial delays in spike response. Error bars are SEM. Statistical significance is determined via Student’s t-tests (relative to 100% spike distance): (\*)  $p < 0.05$  (\*\*)  $p < 0.1$  , (\*\*\*)  $p < 0.001$ .

Rotated movies in the receptive field center largely affected encoding by imposing global delays in spike timing. We shifted every spike from a disrupted movie equally forwards or backwards in time and recomputed the explainable spike distance relative to the original movie's unshifted spike response (Fig. 3.2.2(B), top). For movies unaffected by circular rotations, any shift to the timing of spike responses consistently decreased the spike distance, indicating that spike trains were most similar with no global shifts (Fig. 3.2.2(B), left). In movies strongly affected by circular rotations, however, shifting all spikes equally forward in time (but not backward in time) improved the spike distance considerably, suggesting an initial delay in spike response (Fig. 3.2.2(B)).

Accounting for global temporal delays resolved most performance loss in movies affected by rotations (Fig. 3.2.2(B), middle). Much of the decrease in spike distance from added rotations could be eliminated by imposing an 18 ms global shift forwards in time (corresponding to a ~30% increase in spike distance), reaching 91% and 97% spike distance, respectively. The > 90% similarity of spike responses after uniform temporal shifts suggest that rotated spatial regions do not impose other large effects on encoding. Additionally, the use of global temporal shifts (applied equally to every spike) instead of local shifts (applied variably to certain spikes) indicate that observed delays are consistent throughout the entirety of the spike response.

Delays to spike response could potentially be explained using a spatial wedge-based model: for instance, this can occur if rotations interrupt large temporal changes to the integrated wedge luminance (i.e, by dampening inputs during a saccade). We isolated two different natural movies: one where added rotations elicited delays, and another where rotations did not (Supp. Fig. 3.3.1). For both movies, we observed that adding 6 rotated regions imposed no effect on the linearly integrated luminance of each movie frame, highlighting that rotating spatial structure imposes minimal effects on linear models. Interestingly, adding 6 rotated regions to both movies resulted in a similar decrease in the spatial variance across integrated wedges. Since added rotations affect integrated wedge luminances similarly in movies with and without observed delays, a spatial wedge-based model did not appear to explain observed delays in spike response.

Our results in this section indicate that certain spatial discontinuities (circular rotations within the receptive field center) introduce global delays to the spike responses to some, but not all, natural movies. The timescale of these delays (~18ms) suggest circular discontinuities may be reflected in the kinetics of excitatory and inhibitory inputs [53,54].

### **3.2.2 Spatial discontinuities occur throughout the receptive field center**

Since only a small number of rotated circular regions were required to elicit strong temporal delays in certain natural movies (2 to 3 rotated regions; Fig. 3.2.2(B)), we investigated the effect of spatial discontinuities throughout the entire receptive field center. We began by separating the receptive field center into 2 arbitrary regions: an “inner” and “outer” center (Fig. 3.2.3). For natural movies where rotations strongly affected On-parasol encoding (Fig. 3.2.2(B)), we rotated spatial structure in only the outer center, keeping all spatial structure in the inner center the same (Fig. 3.2.3(A), top gray x-axis).

For one natural movie (Fig. 3.2.3(A), left), we observed that rotating the outer 20% of the receptive field center captured 76% of the spike distance relative to the original movie (without any temporal shift). Applying the same rotation to a slightly larger outer center (40%) resolved most differences in spike responses (98% spike distance), and all larger rotated outer centers continued to elicit similar responses to the original movie (> 96% spike distance). Since rotating the outer 40% (but not the outer 20%) retained most spike distance, disruptions to encoding can be attributed to the presence of a ring-shaped discontinuity in the outer 20% of the receptive field center.

We then introduced a different type of spatial discontinuity by dividing spatial structure in the outer center into wedges and swapping (“shuffling”) regions (Fig. 3.2.3(A), bottom black x-axis). Here, each movie frame was divided into 8 wedge-shaped regions that were repeatedly swapped until every wedge was located in a new spatial position with none of its original neighboring wedges. Shuffling operations introduce two sets of spatial discontinuities: radial divisions between each wedge and ring-shaped divisions on the outer and inner edge of each wedge. Similar to rotated

movies, the shuffling operation applied to one movie frame was exactly repeated for every other movie frame.

For one movie (Fig. 3.2.3(A)), rotating or shuffling the outer 20% of the receptive field center similarly affected the spike distance (76% spike distance). This indicates that declines to the spike distance in the outer 20% of the center were largely attributable to circular discontinuities (from rotations) and not radial discontinuities between wedges (from shuffles). Interestingly, while increasing the size of the rotated outer center improved the spike distance (> 96% spike distance), increasing the size of the shuffled outer center did not (< 80% spike distance). This indicates that for larger outer centers, a decline in spike distance was attributable to radial discontinuities between wedges (from shuffles) and not circular discontinuities (from rotations).

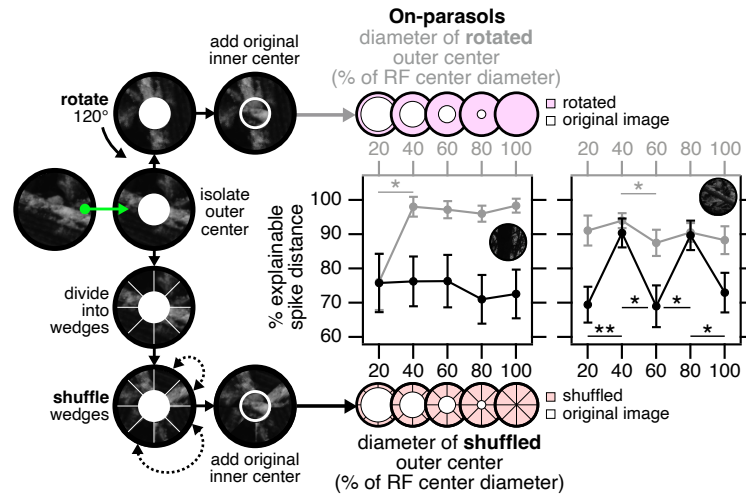
A decrease in spike distance due to certain discontinuities (circular discontinuities in the outer 20%; radial discontinuities between wedges in  $\geq 40\%$  of the outer center) highlight how spatial discontinuities can affect encoding differently depending on their spatial location within the receptive field center. Such effects could either reflect fundamental properties of the receptive field or the complex spatial structure of the movie itself, a question we return to in Fig. 3.2.4.

In a different natural movie sensitive to spatial rotations, we observed that adding circular and wedge-shaped discontinuities imposed strong effects throughout the receptive field center (Fig. 3.2.3(A), right). Shuffling the outer 20% of the receptive field center imposed a strong effect on encoding (69% spike distance), whereas shuffling a slightly larger outer center (40%) largely resolved performance (90% spike distance). Unexpectedly, shuffling an even larger outer center (60%) decreased the distance similarity again to 69%, which was resolved by again increasing the size of the outer center to 80% (90% spike distance). Shuffling the entire outer center (100%) captured only 73% spike distance. Rotated outer centers elicited similar but more subtle changes to the spike distance ( $\pm 4\%$  spike distance for rotated stimuli vs.  $\pm 20\%$  for shuffled stimuli). This indicates that while circular discontinuities drive some encoding differences, the spike distance was most impacted by the presence of radial discontinuities between wedges.

Similar to previous work with these movies (Fig. 3.2.2), we could largely resolve losses in spike distance via global temporal shifts (Fig. 3.2.3(B)). After applying temporal shifts, shuffled stimuli for both movies consistently captured >96% (Fig. 3.2.3(A), left) and >93% (Fig. 3.2.3(A), right) of the spike distance, indicating high similarity to the original natural movie.

The results in this section indicate that introducing two types of spatial discontinuities (circular discontinuities via rotations; radial discontinuities between wedges via shuffles) to natural movies can impose varying delays to spike response. Temporal delays appear to depend on both the location and shape of discontinuities as well as the underlying spatial structure of the naturalistic movie (Fig. 3.2.3(B)).

## A. rotations and shuffles in receptive field center



## B. temporal shifts in receptive field center

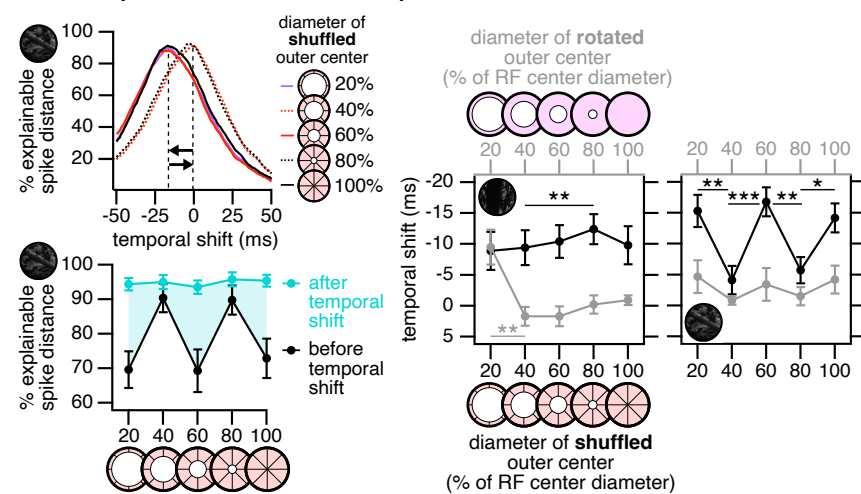


Figure 3.2.3: Temporal delays observed throughout the receptive field center. (A) Each frame of a naturalistic movie was divided into an arbitrary “inner” and “outer” center, denoting its relative location within the receptive field center. Spatial discontinuities were exclusively introduced to the outer center by rotating features (top x-axis) or dividing spatial features into 8 wedges and shuffling their location (bottom x-axis). Spike responses to distorted movies were compared to spike responses from the original movie for two different natural movies (denoted via inset images) across a variety of outer center diameters in On-parasol retinal ganglion cells ( $n = 10$  neurons across 4 retinas). (B, left) Spike distance after applying shifts for one natural movie (from A, right). (B, right) Temporal shifts that maximize the spike distance. Negative temporal shifts correspond to initial delays in spike response. Error bars are SEM. Statistical significance is determined via Student’s t-tests (relative to 100% spike distance): (\*)  $p < 0.05$  (\*\*)  $p < 0.1$ , (\*\*\*)  $p < 0.001$ .

### 3.2.3 Spatial discontinuities engage temporal delays in artificial stimuli

The effects of spatial discontinuities on the encoding of natural movies (Fig. 3.2.3(B)) were difficult to investigate due to the highly complex spatial structure of such movies. Thus, we asked whether we could measure similar delays using stimuli with simpler spatial structure. In this section, we show that added spatial discontinuities to an artificial stimulus evoke similar temporal delays as natural movies in On-parasol retinal ganglion cells.

We began by presenting a 4Hz contrast-reversing split-field spot to the receptive field center of an On-parasol retinal ganglion cell (Fig. 3.2.4(A), top). Briefly, this stimulus consists of two half-circles that sinusoidally oscillate out-of-phase between  $\pm 90\%$  contrast. We selected a classic split-field spot because it evokes spatially nonlinear behavior with periodic luminances over time that enable the close investigation of temporal delays. We then applied spatial discontinuities to spatial regions within the receptive field center and compared spike responses to the original split-field spot.

We first replaced spatial regions in the inner or outer center of a split-field spot with a uniform (zero spatial contrast) spot (Fig. 3.2.4(B), top, black trace). Replacing increasingly larger split-field spots with uniform spots strongly disrupt linear models, and we correspondingly observed a strong decrease in the explainable spike distance. Upon applying temporal shifts to spike responses, we observed that stimuli with added uniform regions elicited a faster response than split-field spots alone, contrasting the delays elicited by other stimuli (Fig. 3.2.4(B), bottom, black trace). Such observations highlight the accuracy of our temporal shift measurements and their ability to reflect key spatial properties of a stimulus.

We then rotated spatial regions in the inner or outer center of a split-field spot and observed varying temporal delays across the receptive field (Fig. 3.2.4(B), bottom, gray curve). Delays to spike responses after rotating split-field spots generally agreed with observed delays from rotated natural movies. We also observed that increasing the size of the rotated region (from 0% to 50% of

the receptive field center) gradually introduced temporal delays. Further increasing the size of the rotated region (from 50% to 100%) eliminated delays as the frequency-doubled stimulus began to resemble an out-of-phase split-field spot. This trend was roughly similar to the frequency-doubled (F2) power (Fig. 3.2.4(A)). Differences in spike responses from a split-field and fully-rotated split-field spot (68% spike distance) reflect experimental error from imperfectly centering the split-field spot within the receptive field center (see: PSTH in Fig. 3.2.4(A)).

Finally, we “shuffled” the spatial structure of a split-field spot (swapping wedge-shaped spatial regions) and observed that delays depended on both the spatial scale and location of shuffled wedges (Fig. 3.2.4(B), bottom, red curves). On average, shuffling the entire receptive field center into 4, 6, and 8 wedges delayed spike responses by 4, 6, and 8 ms, respectively, indicating a fairly linear relationship. Compared to shuffling the entire receptive field center, shuffling smaller regions (inner 20%, 40%) evoked smaller temporal shifts (90% less delayed, 50% less delayed, respectively), highlighting some correspondence with the sum of receptive field weights over space.

Leveraging the simplicity of artificial stimuli, we then observed that split-field spots replaced with uniform, rotated, or shuffled regions evoked similar temporal delays at both  $\pm 90\%$  and  $\pm 30\%$  contrast (Fig. 3.2.4(C)). Such experiments indicate that temporal delays are not restricted to high contrast stimuli (and are not the product of saturated firing rates, for instance).

The results in this section show that the same spatial discontinuities applied to natural movies (rotations and shuffles) can evoke similar temporal shifts when applied to artificial stimuli. Temporal shifts measured from disrupted split-field spots tend to be small ( $< 6$  ms for rotations;  $< 9$  ms for shuffles) but vary depending on the location and size of the disrupted region.

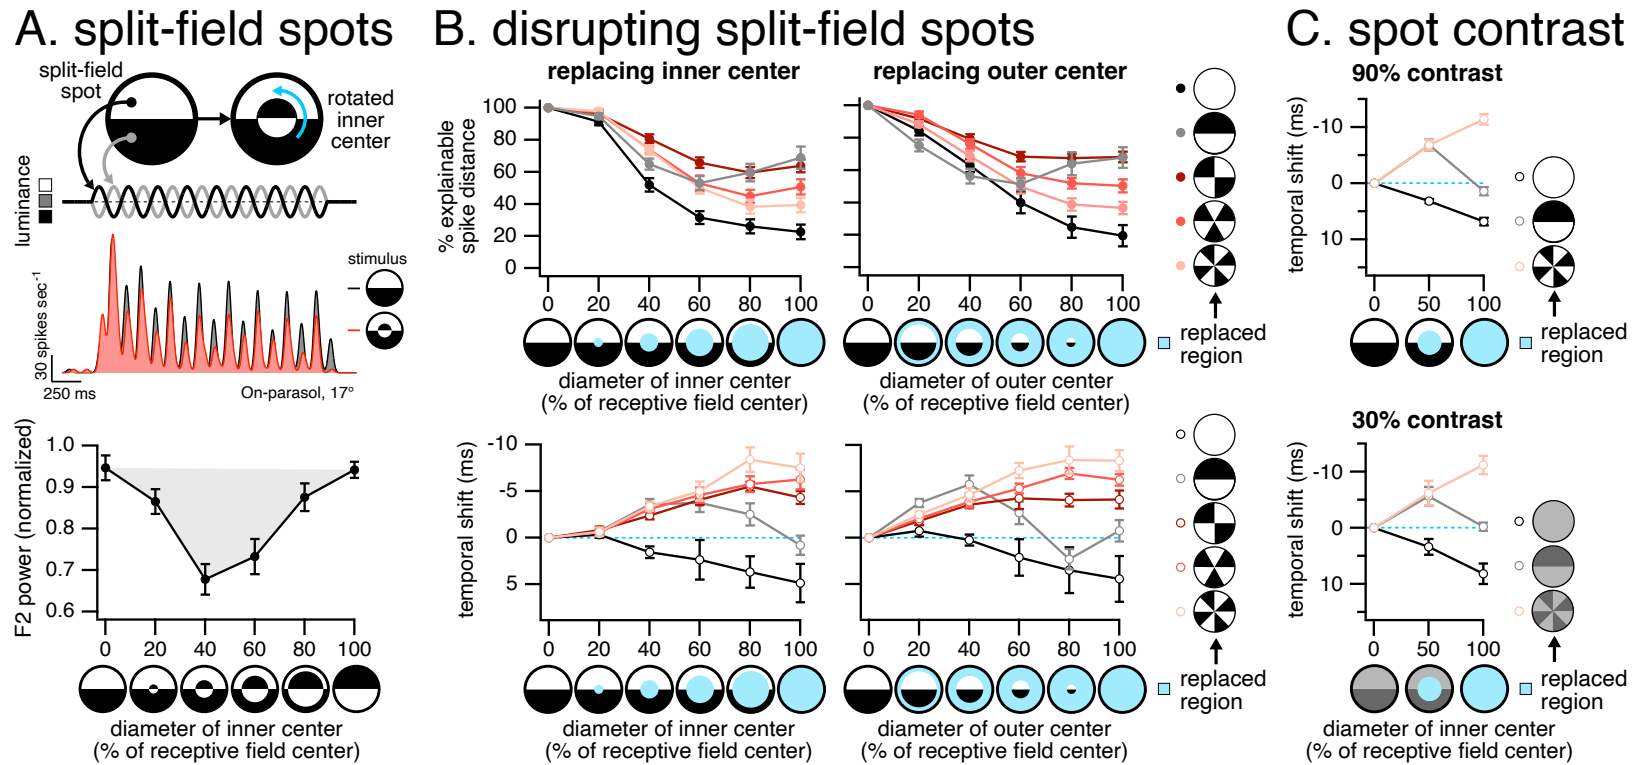


Figure 3.2.4: Discontinuities in artificial stimuli evoke similar temporal delays as natural movies. (A) Spike responses to a split-field 4Hz contrast reversing spot, with and without the outer 60% rotated, as presented to the receptive field center of an On-parasol retinal ganglion cell (17° temporal). (A, bottom) Frequency-doubled power across On-parasols ( $n = 16$  across 3 retinas) as increasingly larger inner centers are rotated. (B, top) Explainable spike distance as the inner center (left) or outer center (right) are replaced with a uniform spot (black curve), rotated split-field spot (gray curve), or shuffled split-field spot (into various numbers of wedges, red curves) for On-parasols ( $n = 16$  across 3 retinas). (B, bottom) Temporal delays that maximize the spike distance. Negative temporal shifts correspond to initial delays in spike response. (C) Temporal delays (that maximize the spike distance) for disrupted split-field spots with a maximum luminance of  $\pm 90\%$  contrast (top) or  $\pm 30\%$  contrast (bottom) ( $n = 5$  across 1 retina).

### 3.2.4 Adding spatial discontinuities result in larger temporal delays.

Equipped with a simple artificial stimulus (split-field spot) sensitive to a variety of spatial discontinuities (rotations, shuffles), we then investigated the effect of disrupting multiple spatial regions simultaneously. We first subdivided a split-field spot into 5 concentric rings and combinatorially rotated some regions but not others. We then presented each stimulus variant as a 4Hz, contrast-reversing split-field spot to the receptive field center of both On- and Off-parasol retinal ganglion cells (Fig. 3.2.5(A), top left). We included both On- and Off-parasols to build and compare computational spatial integration models across cell types by leveraging the simplicity of split-field spots.

We observed that increasing the number of circular discontinuities within the receptive field center similarly increased temporal delays in both On- and Off-parasols (Fig. 3.2.5(A), top). For instance, stimuli with two circular discontinuities in the innermost center elicited 4 and 5 ms delays in On-parasols and Off-parasols, respectively. Adding two additional circular discontinuities (to the outermost center) elicited 18 ms and 22 ms delays, respectively. Delays from stimuli with multiple circular discontinuities generally matched the sum of delays across stimuli with a single circular discontinuity placed in each of the same regions, indicating fairly linear interactions between discontinuities (Fig. 3.2.5(A), bottom left).

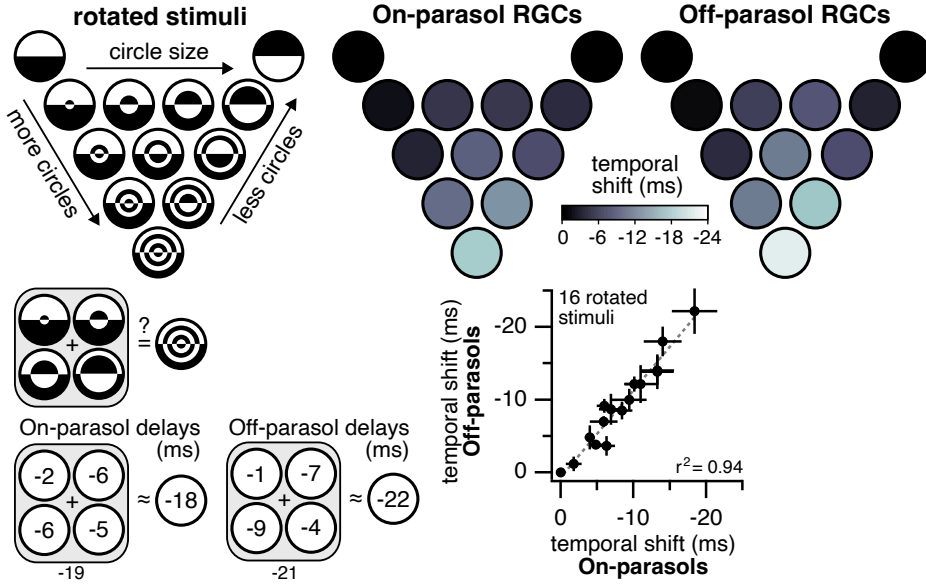
Given that our previous work was focused towards On-parasols (Figs. 3.2.1 to 3.2.4), we compared temporal delays between On- and Off-parasols for each split-field spot and observed highly correlated behavior ( $r^2 = 0.94$ ) (Fig. 3.2.5(A), bottom right). This indicates that the retinal mechanisms that underlie temporal shifts may be shared across On- and Off-pathways. However, similar delays across pathways may only occur under specific stimulus conditions: if added spatial discontinuities affect one polarity (i.e, positive contrasts) more strongly than the other (i.e, negative contrasts), temporal delays may be inconsistent between On- and Off-parasols.

We then expanded our library of contrast-reversing spots to include “shuffles”, which introduce

radial discontinuities (between each wedge) and circular discontinuities (enclosing each wedge) (Fig. 3.2.5(B)). The resulting stimulus set contained 62 variations of split-field spots that were presented to the receptive field center of Off-parasols. After observing circular discontinuities reduced On-parasol firing rates more strongly than in Off-parasols, we chose to temporarily shift our focus towards Off-parasols to mediate concerns that low firing rates might result in more inaccurate temporal shift measurements.

Across our 62 spot variants, we observed that certain spatial discontinuities evoked particularly long or short delays (Fig. 3.2.5(B), top). We quantified the number of spatial discontinuities in each split-field spot by first passing each stimulus through a 2-D high-pass filter and then quantifying the total number of pixels surrounding hard edges. We observed that the number of discontinuities were well-correlated with observed temporal delays across all 62 contrast-reversing stimuli ( $r^2 = 0.81$ ) (Fig. 3.2.5(B), bottom). Our results indicate that the increasing the number of spatial discontinuities are correlated with stronger temporal delays in Off-parasol retinal ganglion cells.

### A. rotated split-field spots



### B. shuffled split-field spots

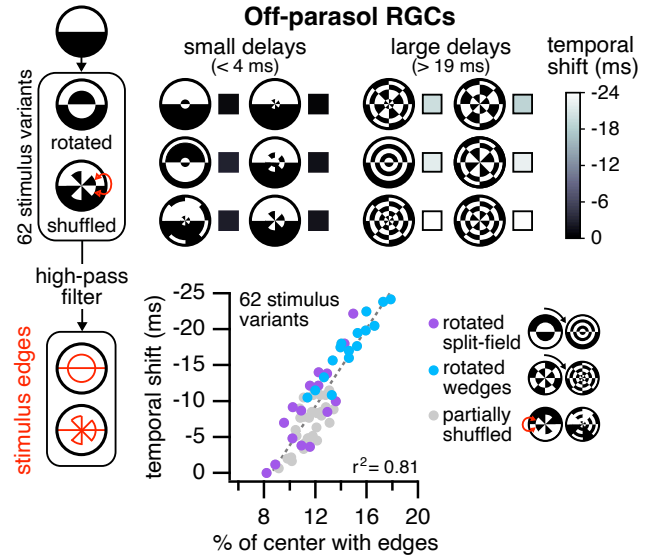


Figure 3.2.5: Adding spatial discontinuities generally elicit larger temporal delays. (A, top left) An assortment of split-field, 4Hz contrast-reversing spots, with and without certain subregions rotated, as presented to the receptive field center of retinal ganglion cells. (A, top right) Temporal shifts (that maximize the explainable spike distance between a distorted and undistorted stimulus) as measured across On-parasols ( $n = 10$  in 2 retinas) and Off-parasols ( $n = 6$  in 2 retinas). Negative temporal shifts correspond to initial delays in spike response. (A, bottom left) Temporal delays across multiple stimuli with one circular spatial discontinuity and one stimulus with multiple spatial discontinuities placed in the same locations. (A, bottom right) Correlation between temporal shifts for each disrupted spot (each dot) for On- and Off-parasols. (B, top) A larger assortment of 62 distorted stimuli (with rotations and shuffled locations of each of 8 wedges) were presented to Off-parasols ( $n = 6$  in 2 retinas). Certain stimuli that evoked minor ( $< 4$  ms) and major ( $> 19$  ms) temporal delays are provided. (B, bottom) Each stimulus was passed through a high-pass filter to identify spatial discontinuities (edges). The total number of spatial discontinuities were quantified by summing the total number of edges (red pixels) in each filtered image.

### 3.2.5 Next steps

Given that development of this work is still underway, we will discuss our planned next steps for concluding this work. We have just begun to develop our final figures which we include below as in-progress.

Our first goal is to identify, with as much specificity as possible, underlying mechanisms within the retinal circuit that contribute to temporal delays. We have conducted whole-cell recordings in response to rotated split-field spots in one On- and Off-parasol (Fig. 3.2.6). Our preliminary results indicate that temporal delays arise from either delayed excitatory inputs or quickened inhibitory inputs. Future whole-cell recordings will leverage shuffled stimuli (Fig. 3.2.5), which impose very large delays that will allow us to strongly stimulate and study such mechanisms.

Our second goal is to spatially map the functional inputs that underlie temporal shifts. We have previously observed that certain discontinuities delay spikes depending on their size and location (Fig. 3.2.4). We have begun to build such maps by passing each distorted stimulus through a high pass filter, convolving filtered images with temporal delays, and averaging across stimuli to identify spatial regions that impose stronger delays than others (Fig. 3.2.7). Interestingly, these spatial maps appear to change depending on whether we rotate a split-field spot or a spot with 8 split-field wedges (Fig. 3.2.7, bottom left and middle columns). This indicates that temporal shifts may depend on the spatial scale of stimulus features, and may provide insight into why some natural images are sensitive to rotations but others are not.

Our eventual goal is to map the functional inputs that mediate temporal shifts in parasol ganglion cells. What is their functional size, shape, and anatomical basis? Do all parasols share a common functional map, or do functional inputs differ from neuron to neuron? Lastly, do models that account for these functional inputs improve predictions to natural movies? Our hope is that the use of artificial stimuli with spatial discontinuities will provide a simple methodology for addressing these questions.

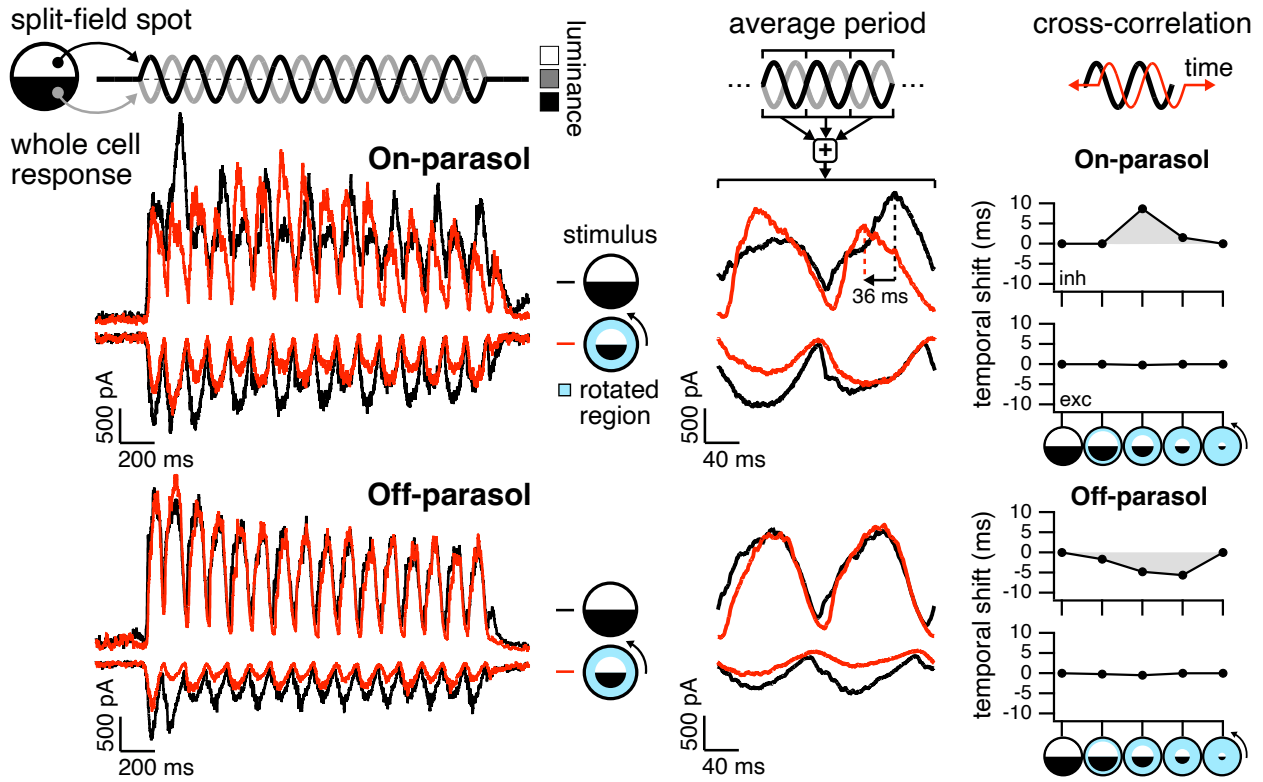


Figure 3.2.6: (in development): (Left) Whole cell recordings of an On- and Off-parasol retinal ganglion cell to a 4Hz contrast-reversing split-field spot (presented to the receptive field center) with and without the outer 40% rotated. (Middle) Average excitatory and inhibitory current across repeated sinusoidal periods. (Right) Cross correlation of whole-cell current (full traces, left) for split-field stimuli with increasingly larger outer center regions rotated. Negative temporal shifts correspond to delays.

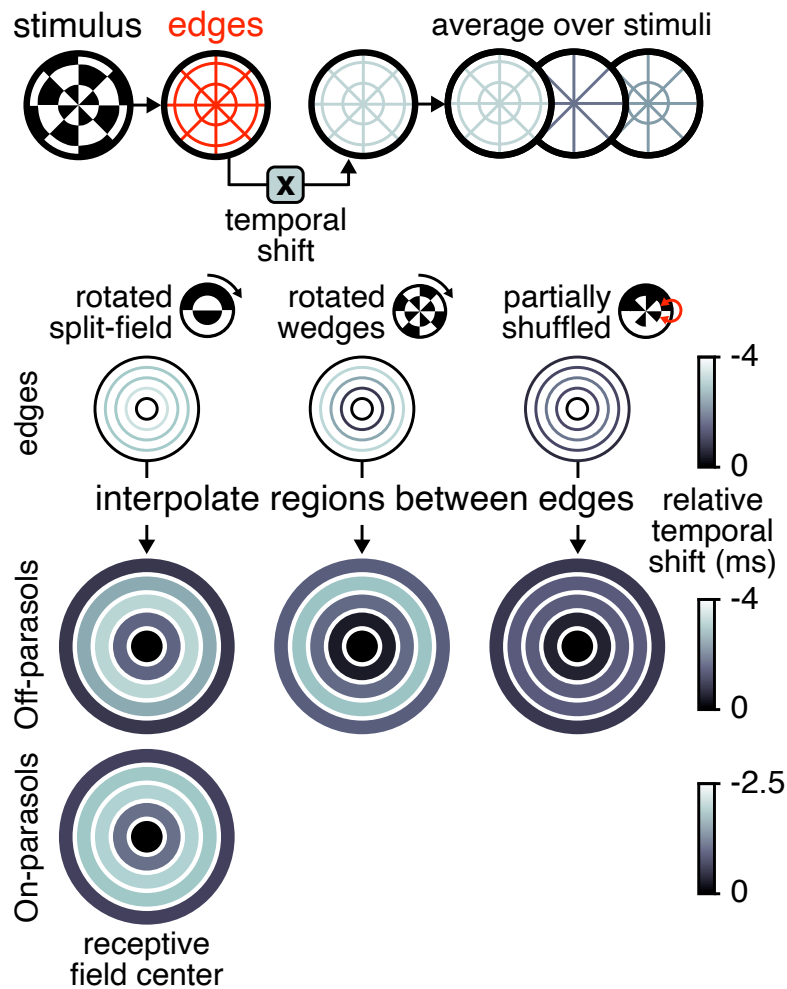


Figure 3.2.7: (in development). (A, top) Several 4Hz, contrast-reversing stimulus variants were presented to the receptive field center of On-parasols ( $n = 10$  in 2 retinas) and Off-parasols ( $n = 6$  in 2 retinas). Each stimulus was then passed through a high-pass filter and multiplied by its measured temporal shift. (A, middle) The average temporal shift for introducing a discontinuity (edge) relative to the centermost region. (A, bottom) Average temporal shift of each region (calculated via the average temporal shift between surrounding edges) relative to the centermost region. Negative temporal shifts correspond to initial delays in spike response.

### 3.3 Supplementary Figures

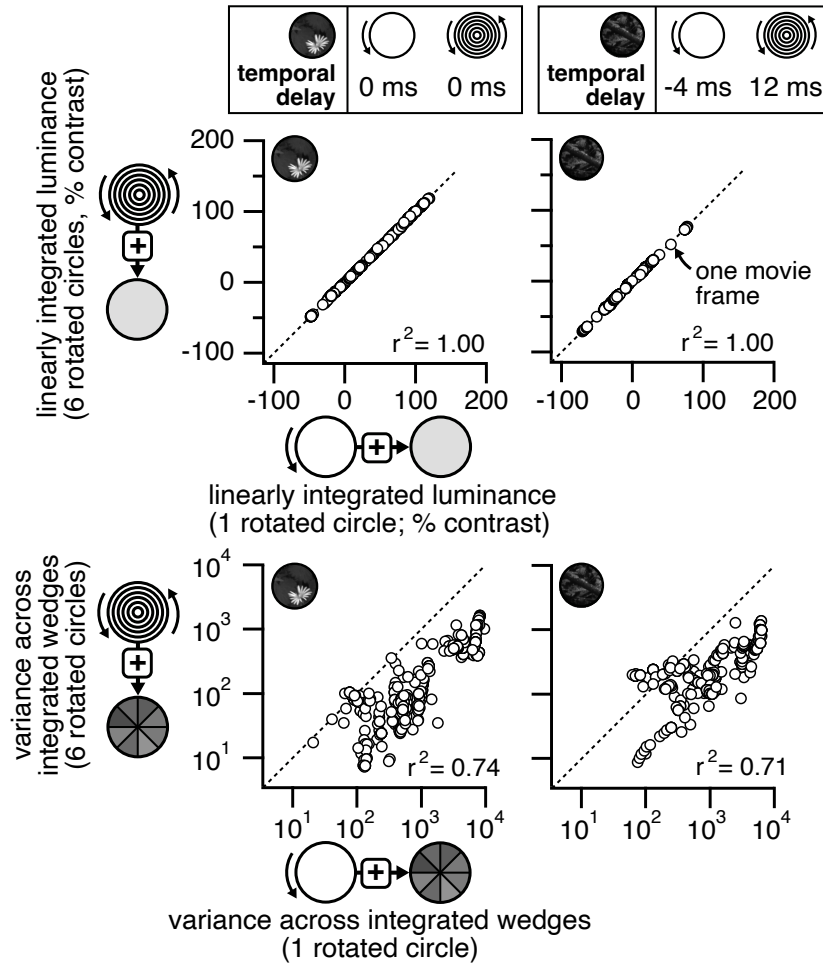


Figure 3.3.1: Figure S1. Circular spatial discontinuities affect integration similarly in movies with and without sensitivities to rotations. (Top) Average temporal delay (from Fig. 3.2.2) of two natural movies with either one uniform rotation (left icon) or 6 independently-rotated circular subregions (right icon). (Middle) Each movie frame (each dot) for movies with 1 or 6 rotated regions linearly integrated into a single luminance. (Bottom) Each movie frame (with 1 or 6 rotated regions) nonlinearly integrated into 8 wedges and correlated via spatial variance across wedges.

# Chapter 4

## Future Directions

In this dissertation, we investigated the functional properties of the retinal circuit by changing the spatial structure of stimuli. For a variety of natural movies, replacing high-dimensional spatial regions with uniform, linearly integrated wedges (forming “simplified” or “reduced” stimuli) resulted in minor changes to parasol retinal ganglion cell spike responses (Chapter 2). We then replaced uniform wedges with spatial features from other images (“metamers” or “degenerate” stimuli) and observed similarly minor changes to spike response (Chapter 2). Finally, we observed that disrupting the spatial structure of many scenes (via “rotations” or “shuffles”) imposed little effect on spike responses except for small, uniform delays ( $< 30\text{ms}$ ; Chapter 3). In this section, we briefly discuss a series of future directions for continuing this type of research.

### 4.1 Identifying a complete low-dimensional space

Our investigations broadly outline an important concept in retinal encoding: individual neurons contain strict upper limits for the amount of information they can encode. In my experience, parasol retinal ganglion cells elicit a maximum spike rate of  $\sim 300$  spikes per second, in agreement with the 1 – 2 ms refractory period immediately following a spike [55]. It would be surprising if 300 spikes could completely account for all high-dimensional spatiotemporal information encountered each second, especially given that visual inputs are constantly interrupted by eye movements [56]. In such context, it seems reasonable that spikes are shaped by a few key inputs (i.e, wedges).

Given that a few uniform wedges described most spike responses to natural movies, other low-dimensional representations may also improve predictions. We observed that a linear model

fully described responses to certain movies; in other movies, a wedge-based model improved predictions considerably. I expect that certain collections of stimuli benefit from other low-dimensional representations that more precisely reflect the functional properties of the retinal circuit. Such stimuli could be identified by painstakingly sampling a huge number of high-dimensional stimuli (and their low-dimensional forms) and isolating only the stimuli that perform poorly. However, more directed experiments could design stimuli that directly challenge spatial models (i.e, replacing wedges with sets of images with diverse of spatial structure). Additionally, as noted in the introduction to this dissertation (paired with work in Chapter 3), future efforts should not neglect the use of artificial stimuli (i.e, split-field spots) which can be used in tandem with natural stimuli to study functional integration mechanisms.

Another important question is whether our low-dimensional wedge-based models generalize to describe vision under all lighting conditions. All work in this thesis was exclusively conducted at photopic (daytime) light levels with achromatic (black-and-white) stimuli. At scotopic (nighttime) light levels, retinal ganglion cells instead receive input from rod photoreceptors and contain larger receptive fields [57] with potentially different functional integration mechanisms. In the case of color, while the use of black-and-white stimuli was appropriate for parasol retinal ganglion cells (which contain largely achromatic receptive fields), our findings may not fully describe the color-opponent receptive fields of cell types such as midget ganglion cells [39,58]. Using black-and-white stimuli, we observed that tiling 8 uniform wedges across two circular divisions was sufficient to describe 92% of spike behavior in the receptive field center of chromatic-sensitive On-midget ganglion cells (Fig. 2.6.2). An open question is whether low-dimensional stimuli change across light levels or upon the introduction of distinct color channels.

#### **4.1.1 Spatial integration in the receptive field surround**

In our investigations, we observed that integrating the receptive field surround into 8 wedges improved predictions for some, but not all, movies (Fig. 2.6.2; Fig. 2.6.2). Our results leave a

lingering feeling that we never truly solidified a low-dimensional representation of the receptive field surround.

There are several reasons why we haven't dedicated more effort towards the receptive field surround. First and foremost is the relatively subtle effect of the surround: our results from full-field naturalistic movies indicate that spatial features in the surround only describe  $\sim 30\%$  of spike behavior (Fig. 2.3.3). As a result, comparing models for the surround will be more subtle (and prone to noise) than in the center. Additionally, inhibitory surrounds only exert a measurable effect if a spike train already has  $> 0$  spikes, indicating that some center stimulation is required to study the surround. Yet the specific effects of the surround vary depending on spatial features within the center [29], complicating our ability to isolate surround-specific behaviors. I recommend leveraging our improved understanding of functional mechanisms in the receptive field center to better isolate integration mechanisms in the receptive field surround.

I also recommend ensuring that receptive field measurements in the surround are sufficiently precise. All calculations of integrated luminances ( $\mathcal{L}$ ) are limited by measurements of the receptive field ( $R$ ) (via Eq. 1.2.1) The use of circular spots to measure the receptive fields of retinal ganglion cells, which are more oblong and ovular than circular [59], may add unnecessary noise to surround-focused investigations that assume circularly symmetric surrounds.

## 4.2 Utilizing a high-dimensional null space

Our work in this dissertation not only identified the spatial features that strongly shape encoding (i.e, linearly integrated wedges), but also identified spatial features that do not (i.e, fine spatial detail within wedges). Designing stimuli with features that are not expected to affect encoding (thereby probing the null space of a neuron) directly test the completeness of low-dimensional spaces.

Characterizing a neuron's null space can greatly serve population-based stimulus design. Consider a low-dimensional stimulus that targets two adjacent On-parasol retinal ganglion cells simul-

taneously. This stimulus may contain a variety of integrated wedges (placed over each neuron) with added structure where their receptive fields overlap [59]. Presenting this population-targeted stimulus to a retina will elicit strong responses from our two On-parasols, as well as any other neuron that falls within their receptive field.

Utilizing the null space of On-parasols may allow us to further tune our stimulus to *only* stimulate two On-parasols (and silence other nearby neurons). This process consists of adding spatial structure that does not affect parasol integration (i.e, integrated wedges) yet strongly affects integration in other neurons (i.e, On-midgets). The result is a high-dimensional stimulus (similar to “metamers” in Fig. 2.3.4(B)) with low-dimensional features (wedges) that stimulate spikes in some neurons and high-dimensional features (pixels within wedges) that selectively inhibit other neurons. The resulting stimulus would push stimulus design to its most extreme.

#### **4.2.1 Ethical considerations**

The design of visual stimuli that selectively stimulate subpopulations of retinal neurons carry significant ethical considerations when generalized to the broader world. Carefully tuning visual inputs to stimulate some neurons and silence others could conceivably be applied to any visual media, including advertisements that seek to engage specific downstream behaviors. In the near future, this could manifest as complex spatial features in the periphery deliberately designed to grab one’s attention, and may grow even more disruptive depending on the selectivity of visual stimuli. I find great concern in the nefariousness of such stimuli being introduced to humans without consent, similar to hypothetical “subliminal” stimuli that lie below the detection threshold for certain neurons. Leveraging the null space of certain neurons in stimulus design requires an active, ongoing discussion on observer consent and reasonable use within the general media.

# Bibliography

- [1] Rodieck, R. W. *The First Steps in Seeing* (Sinauer, 1998). Google-Books-ID: uiBmQgAA-CAAJ.
- [2] van Hateren, J. H. & van der Schaaf, A. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society B: Biological Sciences* **265**, 359–366 (1998). URL <http://rspb.royalsocietypublishing.org/cgi/doi/10.1098/rspb.1998.0303>.
- [3] Hochstein, S. & Shapley, R. M. Quantitative analysis of retinal ganglion cell classifications. *The Journal of Physiology* **262**, 237–264 (1976). URL <http://doi.wiley.com/10.1113/jphysiol.1976.sp011594>.
- [4] Hubel, D. H. & Wiesel, T. N. Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology* **148**, 574–591 (1959). URL <http://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1959.sp006308>. \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1959.sp006308>.
- [5] Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology* **160**, 106–154.2 (1962). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1359523/>.
- [6] Hubel, D. H. & Wiesel, T. N. Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology* **195**, 215–243 (1968).
- [7] Carandini, M. *et al.* Do We Know What the Early Visual System Does? *Journal of Neuroscience* **25**, 10577–10597 (2005). URL <https://www.jneurosci.org/content/25/46/10577>. Publisher: Society for Neuroscience Section: Symposia and Mini-Symposia.
- [8] Sharpee, T. O. Computational Identification of Receptive Fields. *Annual Review of Neuroscience* **36**, 103–120 (2013). URL <http://www.annualreviews.org/doi/abs/10.>

1146/annurev-neuro-062012-170253. Publisher: Annual Reviews.

- [9] Barlow, H. B. Summation and inhibition in the frog's retina. *The Journal of Physiology* **119**, 69–88 (1953). URL <http://onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1953.sp004829>. \_eprint: <https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1953.sp004829>.
- [10] Kuffler, S. W. Discharge patterns and functional organization of mammalian retina. *Journal of Neurophysiology* **16**, 37–68 (1953).
- [11] Weber, E. H. *De Pulsu, resorptione, auditu et tactu: Annotationes anatomicae et physiologicae* ... (C.F. Koehler, 1831). Google-Books-ID: QbJUAAAACAAJ.
- [12] Enroth-Cugell, C. & Robson, J. G. The contrast sensitivity of retinal ganglion cells of the cat. *The Journal of Physiology* **187**, 517–552 (1966). URL <http://doi.wiley.com/10.1113/jphysiol.1966.sp008107>.
- [13] Demb, J. B., Haarsma, L., Freed, M. A. & Sterling, P. Functional Circuitry of the Retinal Ganglion Cell's Nonlinear Receptive Field. *The Journal of Neuroscience* **19**, 9756–9767 (1999). URL <http://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.19-22-09756.1999>.
- [14] Demb, J. B., Zaghloul, K., Haarsma, L. & Sterling, P. Bipolar cells contribute to nonlinear spatial summation in the brisk-transient (Y) ganglion cell in mammalian retina. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* **21**, 7447–7454 (2001).
- [15] Chichilnisky, E. J. A simple white noise analysis of neuronal light responses. *Network: Comput. Neural Syst.* 199–213 (2001).
- [16] Eccles, J. C. Ionic Mechanism of Postsynaptic Inhibition. *Science* **145**, 1140–1147 (1964). URL <https://www.jstor.org/stable/1714225>. Publisher: American Association for the Advancement of Science.
- [17] Cafaro, J. & Rieke, F. Regulation of Spatial Selectivity by Crossover Inhibition. *Journal of Neuroscience* **33**, 6310–6320 (2013). URL <https://www.jneurosci.org/lookup/doi/>

10.1523/JNEUROSCI.4964-12.2013.

- [18] Freedland, J. & Rieke, F. Systematic reduction of the dimensionality of natural scenes allows accurate predictions of retinal ganglion cell spike outputs. *Proceedings of the National Academy of Sciences* **119**, e2121744119 (2022). URL <https://www.pnas.org/doi/10.1073/pnas.2121744119>. Publisher: Proceedings of the National Academy of Sciences.
- [19] Ruderman, D. L. & Bialek, W. Statistics of natural images: Scaling in the woods. *Physical Review Letters* **73**, 814–817 (1994). URL <https://link.aps.org/doi/10.1103/PhysRevLett.73.814>.
- [20] Simoncelli, E. P. & Olshausen, B. A. Natural Image Statistics and Neural Representation. *Annual Review of Neuroscience* **24**, 1193–1216 (2001). URL <https://www.annualreviews.org/doi/10.1146/annurev.neuro.24.1.1193>.
- [21] Turner, M. H., Sanchez Giraldo, L. G., Schwartz, O. & Rieke, F. Stimulus- and goal-oriented frameworks for understanding natural vision. *Nature Neuroscience* **22**, 15–24 (2019). URL <http://www.nature.com/articles/s41593-018-0284-0>. Bandiera\_abtest: a Cg\_type: Nature Research Journals Number: 1 Primary\_atype: Reviews Publisher: Nature Publishing Group Subject\_term: Sensory processing;Visual system Subject\_term\_id: sensory-processing;visual-system.
- [22] Wässle, H. & Boycott, B. B. Functional architecture of the mammalian retina. *Physiological Reviews* **71**, 447–480 (1991).
- [23] Gollisch, T. & Meister, M. Eye smarter than scientists believed: neural computations in circuits of the retina. *Neuron* **65**, 150–164 (2010).
- [24] Schreyer, H. M. & Gollisch, T. Nonlinear spatial integration in retinal bipolar cells shapes the encoding of artificial and natural stimuli. *Neuron* **109**, 1692–1706.e8 (2021). URL <https://www.sciencedirect.com/science/article/pii/S0896627321001859>.
- [25] Turner, M. & Rieke, F. Synaptic Rectification Controls Nonlinear Spatial Integration of Natural Visual Inputs. *Neuron* **90**, 1257–1271 (2016). URL <http://www.sciencedirect.com>.

com/science/article/pii/S0896627316301660.

- [26] Crook, J. D. *et al.* Y-Cell Receptive Field and Collicular Projection of Parasol Ganglion Cells in Macaque Monkey Retina. *Journal of Neuroscience* **28**, 11277–11291 (2008). URL <http://www.jneurosci.org/content/28/44/11277>. Publisher: Society for Neuroscience Section: Articles.
- [27] Euler, T., Haverkamp, S., Schubert, T. & Baden, T. Retinal bipolar cells: elementary building blocks of vision. *Nature Reviews Neuroscience* **15**, 507–519 (2014). URL <http://www.nature.com/articles/nrn3783>.
- [28] Enroth-Cugell, C. & Freeman, A. W. The receptive-field spatial structure of cat retinal Y cells. *The Journal of Physiology* **384**, 49–79 (1987). URL <http://onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1987.sp016443>. \_eprint: <https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1987.sp016443>.
- [29] Turner, M. H., Schwartz, G. W. & Rieke, F. Receptive field center-surround interactions mediate context-dependent spatial contrast encoding in the retina. *eLife* **7**, e38841 (2018). URL <https://doi.org/10.7554/eLife.38841>. Publisher: eLife Sciences Publications, Ltd.
- [30] Aljadeff, J., Lansdell, B. J., Fairhall, A. L. & Kleinfeld, D. Analysis of Neuronal Spike Trains, Deconstructed. *Neuron* **91**, 221–259 (2016). URL <https://www.sciencedirect.com/science/article/pii/S0896627316302501>.
- [31] McIntosh, L., Maheswaranathan, N., Nayebi, A., Ganguli, S. & Baccus, S. Deep Learning Models of the Retinal Response to Natural Scenes. In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I. & Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 29 (Curran Associates, Inc., 2016). URL <https://proceedings.neurips.cc/paper/2016/file/a1d33d0dfec820b41b54430b50e96b5c-Paper.pdf>.
- [32] Liu, J. K. *et al.* Inference of neuronal functional circuitry with spike-triggered non-negative matrix factorization. *Nature Communications* **8**, 149 (2017). URL <http://www.nature>.

- com/articles/s41467-017-00156-9. Number: 1 Publisher: Nature Publishing Group.
- [33] Maheswaranathan, N., Kastner, D. B., Baccus, S. A. & Ganguli, S. Inferring hidden structure in multilayered neural circuits. *PLoS computational biology* **14**, e1006291 (2018).
- [34] Shah, N. P. *et al.* Inference of nonlinear receptive field subunits with spike-triggered clustering. *eLife* **9**, e45743 (2020). URL <https://doi.org/10.7554/eLife.45743>. Publisher: eLife Sciences Publications, Ltd.
- [35] Zheng, Y., Jia, S., Yu, Z., Liu, J. K. & Huang, T. Unraveling neural coding of dynamic natural visual scenes via convolutional recurrent neural networks. *Patterns* **2**, 100350 (2021). URL <https://www.sciencedirect.com/science/article/pii/S2666389921002051>.
- [36] Van Der Linde, I., Rajashekar, U., Bovik, A. C. & Cormack, L. K. DOVES: a database of visual eye movements. *Spatial Vision* **22**, 161–177 (2009).
- [37] Victor, J. D. Spike train metrics. *Current Opinion in Neurobiology* **15**, 585–592 (2005). URL <http://www.sciencedirect.com/science/article/pii/S0959438805001236>.
- [38] Victor, J. D. & Purpura, K. P. Metric-space analysis of spike trains: theory, algorithms and application. *Network: Computation in Neural Systems* **8**, 127–164 (1997). URL [https://doi.org/10.1088/0954-898X\\_8\\_2\\_003](https://doi.org/10.1088/0954-898X_8_2_003). Publisher: Taylor & Francis \_eprint: [https://doi.org/10.1088/0954-898X\\_8\\_2\\_003](https://doi.org/10.1088/0954-898X_8_2_003).
- [39] Derrington, A. M. & Lennie, P. Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. *The Journal of Physiology* **357**, 219–240 (1984). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1193256/>.
- [40] Chichilnisky, E. J. & Kalmar, R. S. Functional asymmetries in ON and OFF ganglion cells of primate retina. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* **22**, 2737–2747 (2002).
- [41] Ratliff, C. P., Borghuis, B. G., Kao, Y.-H., Sterling, P. & Balasubramanian, V. Retina is structured to process an excess of darkness in natural scenes. *Proceedings of the National Academy of Sciences of the United States of America* **107**, 17368–17373 (2010).

- [42] Ravi, S., Ahn, D., Greschner, M., Chichilnisky, E. J. & Field, G. D. Pathway-Specific Asymmetries between ON and OFF Visual Signals. *The Journal of Neuroscience* **38**, 9728–9740 (2018). URL <https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.2008-18.2018>.
- [43] Baylor, D. A. & Hodgkin, A. L. Detection and resolution of visual stimuli by turtle photoreceptors. *The Journal of Physiology* **234**, 163–198 (1973). URL <http://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1973.sp010340>. \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1973.sp010340>.
- [44] Endeman, D. & Kamermans, M. Cones perform a non-linear transformation on natural stimuli. *The Journal of Physiology* **588**, 435–446 (2010). URL <http://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.2009.179036>. \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.2009.179036>.
- [45] Angueyra, J. M., Baudin, J., Schwartz, G. W. & Rieke, F. Predicting and Manipulating Cone Responses to Naturalistic Inputs. *Journal of Neuroscience* **42**, 1254–1274 (2022). URL <http://www.jneurosci.org/content/42/7/1254>. Publisher: Society for Neuroscience Section: Research Articles.
- [46] Kaardal, J., Fitzgerald, J. D., Berry II, M. J. & Sharpee, T. O. Identifying functional bases for multidimensional neural computations. *Neural computation* **25**, 1870–1890 (2013). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3672375/>.
- [47] Shapley, R. & Victor, J. D. The contrast gain control of the cat retina. *Vision Research* **19**, 431–434 (1979).
- [48] Ozuysal, Y. & Baccus, S. A. Linking the Computational Structure of Variance Adaptation to Biophysical Mechanisms. *Neuron* **73**, 1002–1015 (2012). URL <https://www.sciencedirect.com/science/article/pii/S0896627312000797>.
- [49] Zapp, S. J., Nitsche, S. & Gollisch, T. Retinal receptive-field substructure: scaffolding for coding and computation. *Trends in Neurosciences* **45**, 430–445 (2022). URL <https://www.sciencedirect.com/science/article/pii/S0166223622000797>.

//www.sciencedirect.com/science/article/pii/S016622362200056X.

- [50] Karamanlis, D., Schreyer, H. M. & Gollisch, T. Retinal Encoding of Natural Scenes. *Annual Review of Vision Science* **8**, 171–193 (2022). URL <https://doi.org/10.1146/annurev-vision-100820-114239>. \_eprint: <https://doi.org/10.1146/annurev-vision-100820-114239>.
- [51] Schwartz, G. W. *et al.* The spatial structure of a nonlinear receptive field. *Nature Neuroscience* **15**, 1572–1580 (2012). URL <http://www.nature.com/articles/nn.3225>. Number: 11  
Publisher: Nature Publishing Group.
- [52] Yu, Z., Turner, M. H., Baudin, J. & Rieke, F. Adaptation in cone photoreceptors contributes to an unexpected insensitivity of primate On parasol retinal ganglion cells to spatial structure in natural images. *eLife* **11**, e70611 (2022). URL <https://doi.org/10.7554/eLife.70611>.  
Publisher: eLife Sciences Publications, Ltd.
- [53] Trong, P. K. & Rieke, F. Origin of correlated activity between parasol retinal ganglion cells. *Nature Neuroscience* **11**, 1343–1351 (2008).
- [54] Roska, B., Molnar, A. & Werblin, F. S. Parallel processing in retinal ganglion cells: how integration of space-time patterns of excitation and inhibition form the spiking output. *Journal of Neurophysiology* **95**, 3810–3822 (2006).
- [55] Richard, E., Goetz, G. A. & Chichilnisky, E. Recognizing retinal ganglion cells in the dark. In *Advances in Neural Information Processing Systems*, vol. 28 (Curran Associates, Inc., 2015). URL [https://papers.nips.cc/paper\\_files/paper/2015/hash/fe70c36866add1572a8e2b96bfede7bf-Abstract.html](https://papers.nips.cc/paper_files/paper/2015/hash/fe70c36866add1572a8e2b96bfede7bf-Abstract.html).
- [56] Rucci, M. & Victor, J. D. The unsteady eye: an information-processing stage, not a bug. *Trends in Neurosciences* **38**, 195–206 (2015). URL <http://linkinghub.elsevier.com/retrieve/pii/S016622361500017X>.
- [57] Troy, J. B., Bohnsack, D. L. & Diller, L. C. Spatial properties of the cat X-cell receptive field as a function of mean light level. *Visual Neuroscience* **16**, 1089–1104 (1999). URL

<https://www.cambridge.org/core/journals/visual-neuroscience/article/spatial-properties-of-the-cat-xcell-receptive-field-as-a-function-of-mean-light-18821E0D86364A0D71674840B4EA2AC2F>. Publisher: Cambridge University Press.

- [58] Wiesel, T. N. & Hubel, D. H. Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *Journal of Neurophysiology* **29**, 1115–1156 (1966). URL <https://journals.physiology.org/doi/abs/10.1152/jn.1966.29.6.1115>. Publisher: American Physiological Society.
- [59] Field, G. D. *et al.* High-sensitivity rod photoreceptor input to the blue-yellow color opponent pathway in macaque retina. *Nature Neuroscience* **12**, 1159–1164 (2009). URL <https://www.nature.com/articles/nn.2353>. Number: 9 Publisher: Nature Publishing Group.

### 4.3 List of Publications

- J. Freedland & F. Rieke. “Systematic reduction of the dimensionality of natural scenes allows accurate predictions of retinal ganglion cell spike outputs.” *Proceedings of the National Academy of Sciences U.S.A.*, vol. 119, no. 46. e2121744119. Nov. 2022. doi: 10.1073/pnas.2121744119
- J. Freedland & K. Young. “On the transformation of humans into vampires in *Buffy the Vampire Slayer*: A viral disease model.” *Slayage: The International Journal of Buffy+*. 20.1 [55] Winter/Spring 2022. ISSN 1546-9212
- N. St. John, J. Freedland, . . . , M. Fasullo. “Genome profiling for aflatoxin B1 resistance in *Saccharomyces cerevisiae* reveals a role for the CSM2/SHU complex in tolerance of aflatoxin B1-associated DNA damage,” *G3: Genes, Genomes, Genetics*, vol. 10, no. 11, pp. 3929–3947, Nov. 2020. doi: 10.1534/g3.120.401723
- J. Freedland. “Vampire slaying in *Buffy the Vampire Slayer* may result from disrupted ion signaling.” *Slayage: The Journal of Whedon Studies*, 18.2 [52], Summer/Fall 2020. ISSN 1546-9212
- M. Fasullo, J. Freedland, . . . , X. Ding, “An in vitro system for measuring genotoxicity mediated by human CYP3A4 in *Saccharomyces cerevisiae*,” *Environmental and Molecular Mutagenesis*, vol. 58, pp. 217–227, May 2017. doi: 10.1534/g3.120.401723. doi: 10.1002/em.22093
- J. Freedland, C. Cera, and M. Fasullo. “CYP1A1 I462V polymorphism is associated with reduced genotoxicity in yeast despite positive association with increased cancer risk,” *Mutation Research/Genetic Toxicology and Environmental Mutagenesis*, vol. 815, pp. 35–43, Mar 2017. doi: 10.1016/j.mrgentox.2017.02.002

### 4.4 Vita

Julian grew up in Staten Island, NY. As a child, Julian loved learning a little bit about everything. Julian graduated with a B.S. in nanoscale science and mathematics from SUNY Albany in 2017 and salivated at the thought of taking a bite out any tough problem he could get ahold of. That included publishing early works on vampire physiology in the television series *Buffy the Vampire Slayer*. Julian graduated with his Ph.D. in Molecular Engineering from the University of Washington in 2023. Julian’s happy place is spending evenings with his wonderful wife, Breena, getting his butt mercilessly whooped in board games.