Contextual effects in the early visual system and their modulation by attention

Sonia Poltoratski

Dissertation

Submitted to the Faculty of the
Graduate School of Vanderbilt University
in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY
in
Psychology

August 11, 2017

Approved:

Frank Tong, Ph.D.

Randolph Blake, Ph.D.

Sam Ling, Ph.D.

Alexander Maier, Ph.D.

Acknowledgements

Thank you to my stellar dissertation committee: Drs. Sam Ling, Randolph Blake, and Alexander Maier. You have not only guided this work to fruition, but have each made unique and lasting impressions on my scientific development at Vanderbilt. I feel immensely privileged to undergo this process with the three of you.

Of course, no part of this dissertation would have been possible without the guidance of my advisor, Dr. Frank Tong: your tenacity, insight, and true love of science have served as a guiding light in these last six years. Your impact on my development as a scientist (and a human being) is immeasurable.

Thank you to the Tong lab, present and past – so many of life's highs and lows have been lived in our space on the fourth floor of Wilson Hall. Thank you to the staff of the Vanderbilt University Institute of Imaging Science for your help with all of the data in this dissertation, and the 90% of my experiments that didn't work. Thank you to the NSF and NIH for financial support.

Tim: your support during all of this has been, in a word, superlative. My family: you are the inspiration that has taken me to this very edge of knowledge. My friends: the good times (champagne) and bad (red wine & Cheddar Rockets) were both better spent with you all.

My deepest and most sincere gratitude.

Table of Contents

Page
Acknowledgements ii
List of figures iii
Chapter
1. Introduction 1
1.1 Mechanisms of feature-tuned surround suppression 2
1.2 Balancing bottom-up salience and top-down attention 3
1.3 Figure-ground segmentation in V1 3
1.4 Functional role of the LGN in feature-tuned suppression 4
1.5 Questions for the current proposal 5
1.6 Overview of methods 5
1.7 Overview of experiments 6
2. Characterizing the effects of feature salience and top-down attention in the early visual system 8
2.1 Introduction 8
2.2 Materials and Methods 10
2.3 Results 14
2.4 Discussion 18
3. Distinct effects of boundary detection and figure enhancement in human V1 21
3.1 Introduction 21
3.2 Results 22
3.3 Discussion 25
3.4 Supplemental Materials and Methods 26
4. Automatic cortical feedback mediates figure-ground modulation in the human lateral geniculate nucleus 31
4.1 Introduction 31
4.2 Materials and Methods 32
4.3 Results 33
4.4 Supplemental Materials and Methods 36
5. Conclusions 39
5.1 Summary of findings 39
5.2 The neural basis of figure processing 39
5.3 Figure enhancement and orientation processing in the LGN 41
5.4 Recurrent processing of contextual information in the early visual system 42
References 44

List of Figures

Chapter 1: Introduction
1.1 Canonical surround suppression 1
1.2 Stimuli and model of figure-ground modulation 4
1.3 Basic implementation of pRF mapping 6

Chapter 2: Characterizing the effects of feature salience and top-down attention in the early visual system
2.1 Sample experimental displays 10
2.2 Results of Experiment 1 14
2.3 Examples of single-subject data from Experiments 1 and 2 15
2.4 Magnitude of effects plotted as a function of ROI size 16
2.5 Results of Experiment 2 17
2.6 Individual subject ANOVA results 17

Chapter 3: Distinct effects of boundary detection and figure enhancement in human V1
3.1 Sample experimental displays and results as a function of eccentricity 21
3.2 pRF estimates across regions V1-V3 22
3.3 Effects as a function of distance from the center/surround boundary 23
3.4 Average BOLD responses across ROIs within V1 and LGN 24
3.5 Average BOLD results in V2 and V3 24
3.6 pRF-based reconstruction of figure enhancement and boundary detection 25
3.7 Comparison of regression-based reconstruction and weighted averaging 26

Chapter 4: Automatic cortical feedback mediates figure-ground modulation in the human lateral geniculate nucleus
4.1 Sample displays and results from Experiment 1 33
4.2 Sample displays and results from Experiment 2 34
4.3 Results in regions V2-hV4 35

CHAPTER 1

Introduction

A primary goal of the visual system is extracting meaningful form from incredibly complex and variable 2D input to the retina. To do so, it is critical for visual representations to reflect not just the value of luminance, contrast, orientation, or other features; it is also important to determine how these features compare to the rest of the visual field. Representing visual features in this relative manner, rather than absolutely, allows the visual system to flexibly and efficiently extract important information.

A primary mechanism by which context can modulate the representation of visual features is *surround suppression*. Surround suppression is the property by which a neuron's response is decreased when it is stimulated beyond its classical receptive field, as illustrated in Figure 1.1A (Jones, Grieve, & Wang, 2001; Knierim & van Essen, 1992; Levitt & Lund, 1997). In the primary visual cortex (V1), surround suppression is feature-tuned (Knierim & van Essen, 1992; Nelson & Frost, 1978), such that maximal suppression occurs when the surround shares feature properties with the center stimulus; in the case of dissimilar center and surround features, the center will be suppressed to a lesser extent or even enhanced (e.g., Sillito, Grieve, Jones, Cudeiro, & Davis, 1995). Feature-tuned surround suppression yields a computation of local feature contrast; it can serve as a form of redundancy reduction in the coding of information across the visual field, an important step towards building a more efficient visual code (Barlow, 1961).

Computed across the visual field, surround suppression can also allow the visual system to recognize 'pop-out' elements (Knierim & van Essen, 1992) or particularly salient regions of the environment. The *visual salience* of an item or portion of the visual field describes the degree to which it can 'grab' attention; this subjective percept is strongly influenced by the local feature differences that occur within the visual scene. Behavioral studies have demonstrated that local differences in color, orientation, size, motion, or other basic features can be rapidly detected in visual search tasks (Nothdurft, 1993; Treisman & Gelade, 1980; Wolfe, 1994). Accordingly, computational models of visual salience have underscored the importance of an initial computation of local feature contrast. Models differ, however, on whether feature-selective interactions at early stages of visual processing might be sufficient to compute a salience map (Koene & Zhaoping, 2007; Z. Li, 1999; 2002; Zhang, Zhaoping, Zhou, & Fang, 2012), or whether a summation of various feature-contrast maps occurs at a higher stage of the visual pathway (Itti & Koch, 2001; Koch & Ullman, 1985). These models can predict where people are more likely to look when viewing natural scenes (Itti & Koch, 2001; Parkhurst, Law, & Niebur, 2002), though cognitive goals and top-down factors can also exert powerful influence (Henderson, 2003; Tatler, Hayhoe, Land, & Ballard, 2011). Functioning across several spatial scales, surround suppression can function locally to efficiently code information, and more broadly to enhance particularly salient regions of a visual

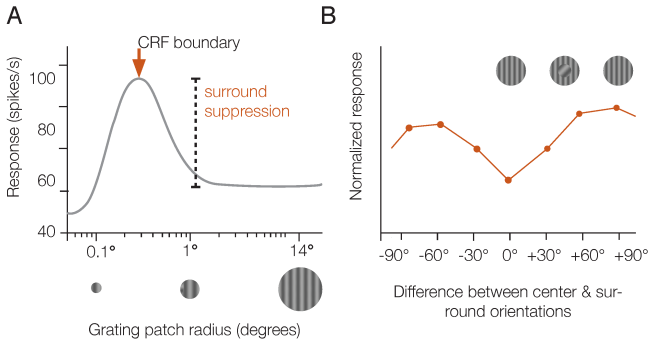

Figure 1.1. A) Illustration of canonical surround suppression. A neuron's response peaks when the stimulus fills the classical receptive field (CRF), and decreases when the stimulus extends beyond the CRF. Adapted from Nurminen & Angelucci (2014). B) Illustration of orientation-tuned surround suppression. A neuron's response is maximally suppressed when the orientation of the center matches that of the surround. Adapted from Naito et al. (2007).

scene – critical steps toward building a behaviorally relevant representation of the environment.

1.1 Mechanisms of feature-tuned surround suppression

As illustrated in Figure 1.1B, feature-tuned surround suppression (Blakemore & Tobin, 1972; Cavanaugh, Bair, & Movshon, 2002; Nelson & Frost, 1978) is a property of neuronal firing by which a surround stimulus with similar features will maximally suppress responses to a central target. In naturalistic vision – in which objects appear in a broader visual context – this has the effect of enhancing responses at image locations that differ in one or more features from their surround. This property holds true for across many features (Allman, Miezin, & McGuinness, 1985; Blakemore & Tobin, 1972; Chao-Yi & Wu, 1994; DeAngelis, Freeman, & Ohzawa, 1994; Jones et al., 2001; Zipser, Lamme, & Schiller, 1996), and has been proposed to arise from a combination of horizontal connections between feature-selective neurons within V1 and more distal interactions involving feedback from higher extrastriate areas (Adesnik, Bruns, Taniguchi, Huang, & Scanziani, 2012; Cavanaugh, 2002; Gilbert & Wiesel, 1979; Lamme, Super, & Spekreijse, 1998a); the relative contributions of these mechanisms may be determined by the spatial proximity and extent of the surround (Nurminen & Angelucci, 2014).

At more local spatial scales, suppression in the near surround may function to reduce redundancy in the visual field, as V1 neurons show evidence of much more sparse coding when they are stimulated beyond their classical receptive field with natural images (Vinje & Gallant, 2000). In support of this view, recent work examined the wide variability of surround suppression strength in response to different natural images. Coen-Cagli et al. (2015) found that the degree of statistical dependencies between neighboring regions, or image redundancy, predicted the strength of surround suppression in V1 neurons, such that images containing greater statistical redundancy elicited greater surround suppression. At a larger spatial scale, surround suppression is more broadly orientation-tuned than in the near-surround (Shushruth et al., 2013), and likely relies on feedback from higher areas rather than horizontal connections between feature-tuned neurons (Nassi, Lomber, & Born, 2013). This mechanism may subserve a fast representation of feature contrast across the visual field, which is necessary for a computation of visual salience.

How might the effects of feature-tuned surround suppression at the earliest stages of visual processing build toward a behaviorally relevant representation of a visual scene? In human psychophysics, modulatory effects of orientation-tuned suppression decrease with distance between the center and surround, but persist even when 3-5 degrees of visual angle separate the central target and modulatory surround (Cannon & Fullenkamp, 1991; Petrov & McKee, 2006). Correspondingly, fMRI BOLD responses in V1 to spatially separated oriented Gabor stimuli have been shown to depend on the orientation of neighboring items (Joo, Boynton, & Murray, 2012; Schallmo, Grant, Burton, & Olman, 2016) suggesting a contribution of surrounding context on even early visual representations of features. Existing fMRI reports differ, however, in their findings of orientation-based salience enhancement, or 'pop-out' in V1 when spatially separated stimuli are used: some report that this enhancement occurs only at higher stages of the visual hierarchy, even when stimuli are attended (Bogler, Bode, & Haynes, 2013); others find this enhancement in V1, but only for attended stimuli (Hopf et al., 2004); and an influential recent paper reported orientation-tuned enhancement even when stimuli are rendered invisible by a powerful mask (Zhang et al., 2012). This literature suggests that directed attention may selectively modulate contextual interactions in the early visual system; our understanding of salience in particular requires a concurrent examination of top-down directed attention.

1.2 Balancing bottom-up salience and top-down attention

Natural vision requires us to supplement stimulus-driven processing with voluntary selective attention, by which we can enhance the processing of stimuli corresponding to our specific goals and needs. While frontal and parietal networks have been implicated in the control of attention (for reviews: (Desimone & Duncan, 1995; Silver & Kastner, 2009; Squire, Noudoost, Schafer, & Moore, 2013), considerable work has shown that attention strongly modulates responses in early visual areas, including the primary visual cortex (V1; Gandhi, Heeger, & Boynton, 1999; Somers, Dale, Seiffert, & Tootell, 1999), extrastriate visual areas, and even the lateral geniculate nucleus (LGN; Ling, Pratte, & Tong, 2015; O'Connor, Fukui, Pinsk, & Kastner, 2002; Schneider & Kastner, 2009). In the early visual cortex, attention has recently been shown to selectively modulate several feature-based contextual interactions, including perceptual grouping (McMains & Kastner, 2011), suppression and feature-spread (Flevaris & Murray, 2015), and crowding (J. Chen et al., 2014); it is unclear, however, how directed attention interacts with the perceptual consequences of feature-tuned suppression.

We have discussed the fact that feature-tuned contextual effects can enhance salient locations or reduce local redundancy in visual input, with separate mechanisms operating at different spatial scales. How might these functional mechanisms interact with top-down, guided attention? One possibility is that once an item is attended, feature-tuned surround modulation does not provide any further enhancement. The response to an item might be "saturated" once it becomes the focus of spatial attention. Alternatively, it may be that the effects of attention and feature-based enhancement are simply summed to determine the overall prioritization of a stimulus in the visual field. Neither the Itti and Koch (Itti & Koch, 2001; Itti, Koch, & Niebur, 1998) nor Zhaoping Li's (1999) models of saliency make specific predictions about directed attention, although the former posits that top-down attention operates 'independently' from bottom-up salience. Desimone and Duncan (1995) suggested that top-down attention can differentially weight feature-contrast maps; this predicts that attention and salience are likely independent in V1, where these feature maps are first computed. While the normalization model of attention (Reynolds & Heeger, 2009) formulates separate stimulus and suppression fields, suppressive normalization occurs *after* attention modulates the stimulus signal; this allows for more complex interactions, particularly by feature-based attention and by the spatial extent of the stimulus and attention fields.

Recent work echoes multiple possible interactions between attention and bottom-up stimulus properties: for instance, while Schallmo et al. (2016) report additive effects of attention and feature-based suppression, other research describes more complex interactions between attention and feature spread spanning spatially separated items (Flevaris & Murray, 2015). Given that near- and far-surround modulations might stem from different mechanisms in early visual cortex (Nurminen & Angelucci, 2014), it is also not apparent whether attention would interact similarly with local and spatially separated contextual modulation.

1.3 Figure-ground segmentation in V1

The separation of visual objects or figures from their surround is another primary proposed function of feature-tuned surround suppression. In seminal work, Victor Lamme (1995) measured the response of V1 cells to orientation-defined figures – texture patches of lines that were oriented orthogonally to a surrounding texture, as illustrated in Figure 1.2A. He found that an equivalent texture placed in a neuron's receptive field would elicit a stronger response when it belonged to the figure than to the surround. This figural enhancement occurs for items defined not only by orientation but also by disparity, color, or luminance (Marcus & van Essen, 2002; Zipser et al., 1996), and may even modulate activity in the LGN (Jones et al., 2015).

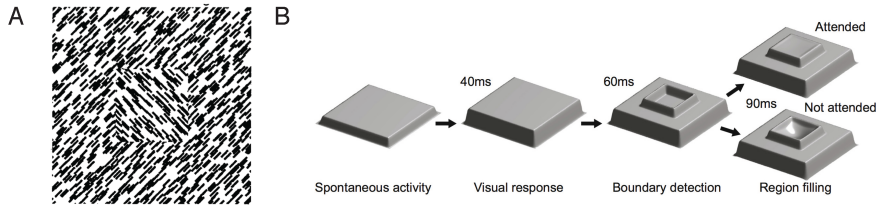

Figure 1.2. A) Sample orientation-defined texture display from Lamme (1995). B) Schematic illustration of the different phases of figure-ground modulation, taken from Poort et al. (2014).

Two separable mechanisms have been identified for the separation of figure from ground: an initial *boundary detection*, and a subsequent *region growing* that groups portions of the image that share features (illustrated in Figure 1.2B). Edge detection is thought to arise from feature-tuned suppression, and can act very rapidly (e.g. 7-10ms, Bair, Cavanaugh, & Movshon, 2003; <25ms, Nothdurft, Gallant, & van Essen, 1999). Enhancement within the figure, however, emerges at a longer latency, suggesting feedback from higher-order visual areas (Lamme, Super, & Spekreijse, 1998a; Roelfsema, Tolboom, & Khayat, 2007). In support of a dual-stage segmentation process, human performance at discriminating figures is best matched by models that instantiate both border detection and a subsequent filling-in process (Mumford, Kosslyn, Hillger, & Herrnstein, 1987; Wolfson & Landy, 1998). Other groups, however, have suggested that figure-ground modulation may be largely explained by boundary detection (Z. Li, 1999); this interpretation is consistent with the findings of Rossi et al. (Rossi, Desimone, & Ungerleider, 2001), who demonstrated that enhancement in response to a larger figure (3°) fell off with distance from the boundary. Task relevance appears to specifically modulate the filling-in component of figure-ground effects, while edge detection occurs even in the absence of attention (Poort et al., 2012).

While recent electrocorticography (ECoG) data from the early visual cortex of a human observer appears to support dual mechanisms of figure-ground modulation (Self et al., 2016), these effects are difficult to measure with traditional fMRI methods because of the difficulty of isolating responses to a small region of a stimulus. However, population receptive field (pRF) mapping, a voxel-wise encoding model that deduces a response field for each measured voxel, may allow us to distinguish enhancements at the edge from those over the entire figure.

1.4 Functional role of the LGN in feature-tuned suppression

Traditionally, the lateral geniculate nucleus (LGN) has been considered minimally selective for visual features, inheriting many of its tuning biases from cells in the retina (Leventhal & Schall, 1983). In Hubel and Wiesel's excitatory convergence model of orientation selectivity, V1 orientation is built from inputs of non-orientation selective LGN cells whose receptive fields are arranged in a row in visual space (Hubel & Wiesel, 1968). However, subsequent recordings have demonstrated orientation biases in LGN cells across several species (e.g. cat: Suematsu, Naito, & Sato, 2012; owl monkey: Cheong, Tailby, Solomon, & Martin, 2013; Xu, Ichida, Shostak, Bonds, & Casagrande, 2002; and macaque: Smith, Chino, Ridder, Kitagawa, & Langston, 1990), and it has been suggested that V1 orientation selectivity may arise directly from the biases (Vidyasagar & Eysel, 2015). Our group has also found that orientation-specific information can be decoded from the LGN and specifically modulated by top-down attention (Ling et al., 2015), suggesting a more active processing role for this structure.

Might the LGN encode representations of feature-contrast? While macaque LGN neurons exhibit non-orientation-tuned surround suppression (Alitto & Usrey, 2008; Sceniak, Chatterjee, & Callaway, 2006), positive effects of feature-tuned suppression in this structure have only been reported in the cat (Cudeiro & Sillito, 1996; Jones, Andolina, Oakely, Murphy, & Sillito, 2000; Naito, Sadakane, Okamoto, & Sato, 2007; Sillito, Cudeiro, & Murphy, 1993; C. Sun, Chen, Huang, &

Shou, 2004). Given the anatomical differences in the LGN between these species, it is unknown whether or not surround suppression in human LGN might be orientation-tuned. It may also be the case that even if LGN neurons themselves show weak tuning for orientation, corticothalamic feedback may serve to modulate responses in a feature-specific way (Briggs & Usrey, 2008). This is consistent with the findings of Ling et al. (2015), who demonstrated that attention can modulate LGN responses in an orientation-specific way. Recently, figure-ground modulation has been demonstrated in single- and multi-unit recordings in the primate (Jones et al., 2015): LGN neurons, like those in V1 (Lamme, 1995) exhibit much greater responses when their receptive field is placed over a motion-defined figure region than a matched region in the surround. Might cortical feedback to the LGN carry second-order feature information in this way? The relative paucity of neuroimaging work in this region makes these questions novel and timely.

1.5 Questions for the current proposal

The proposed work seeks to characterize how visual responses to an item depend on its surrounding context at the earliest stages of the visual system. Specifically, we will examine how feature-tuned suppression modulates fMRI BOLD responses in the lateral geniculate nucleus (LGN) and in the early visual cortex. As we consider the functional role of feature-tuned suppression, which can enhance inhomogeneous or salient regions of an image, we will also test how these contextual effects may interact with the observer's voluntary attention. Doing so will inform our understanding of how bottom-up information and top-down directed attention are processed and integrated at the earliest stages of visual processing.

We seek to address the following questions: Are feature-tuned suppressive effects limited to cortical visual processing, or might they also be represented in the LGN? How does orientation-tuned suppression function across multiple items in a display, yielding a computation of salience? How might it function to segment a particular item, enhancing either its boundary with the surround or the entire figure? Finally, do feature-tuned contextual effects interact with top-down, directed attention?

1.6 Overview of methods

We addressed these questions through several fMRI experiments in healthy adult human subjects. fMRI is a neuroimaging method that measures blood-oxygenation-level-dependent (BOLD) signal, which reflects haemodynamic changes in the brain in response to increased neuronal firing. While limited in its spatial and temporal resolution, fMRI is a non-invasive and flexible method of measuring neural activity. Moreover, responses from several regions of interest can be recorded at once, allowing for powerful comparisons of processing at different stages of the visual hierarchy. fMRI in humans is particularly amenable for answering questions about the function of attention, as human observers can also easily learn and switch between tasks and attended locations. Overall, we can use fMRI results to inform our knowledge of visual processing in conjunction with the findings of more invasive methods of recording neural firing in non-human primates and other animals, as well as psychophysics in human observers.

All scanning was conducted at the 7 Tesla magnet at the Vanderbilt University Institute of Imaging Science (VUIIS). This high field strength allows us to use a relatively small voxel size (1.5-2mm isotropic) and increases the signal-to-noise ratio in our data. Imaging the LGN poses unique challenges: this structure is small, typically encompassing less than 10% of the number of voxels in V1 in any individual subject; its position in the subcortical thalamus is also frequently affected by distortions of the magnetic field in fMRI. Our lab has worked alongside physicists from the VUIIS to

develop scan pulse sequences that minimize distortion and maximize signal strength in this part of the brain. We can now typically localize the LGN using a standard visual localizer consisting of one or more flickering checkerboards at a target location in space. We localize additional regions of interest in the cortical visual pathway (V1-hV4) by the conjunction of this independent localizer and a separate session of retinotopic mapping (Engel, Glover, & Wandell, 1997; Wandell, Dumoulin, & Brewer, 2007).

To resolve the spatial profile of surround suppression and figure-ground modulation in cortex, we used a voxelwise population receptive field (pRF) model. The pRF model was developed as an extension of visual field mapping (Dumoulin & Wandell, 2008), resolving not only the location in space to which a particular voxel responds, but also the size and shape of this response area. This voxel-wise encoding model allows us to explicitly predict neural activity from the features of visual input: based on the given set of features, the model provides a testable hypothesis of representational structure in each voxel. The basic implementation of the model, which assumes a circular 2D Gaussian pRF shape, is illustrated below.

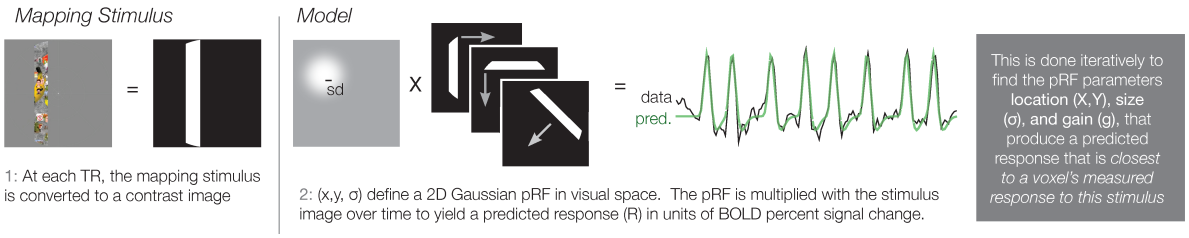

PRF modeling involves presenting a mapping stimulus that spans the visual field across time and estimating the pRF parameters that most likely produce the measured BOLD response in a particular voxel. These parameters define the pRF's location, size, and gain. PRF properties for each voxel reflect the combined RF of the neural population in a voxel, and appear well-aligned with single-neuron receptive field properties, such as contralateral preference, increasing size at greater eccentricities, and increasing size as one ascends the visual hierarchy (Dumoulin & Wandell, 2008; Wandell & Winawer, 2015). In our work, this allowed us to resolve the spatial profile of surround suppression and resulting figure-ground modulation, as we can isolate voxels that respond to the region in space occupied by a target region, its surround, or the boundary between the two.

1.7 Overview of experiments

We addressed the questions posed above in a series of fMRI experiments targeting cortical visual areas and the LGN. These are described in three chapters consisting of completed manuscripts, which have each been either published or prepared for peer-review. In Chapter 1, we describe work measuring modulatory effects of bottom-up salience, which can be conceptualized as a computation of feature contrast across the visual field, and of top-down directed attention. The computation of salience is a key functional consequence of surround-suppression, and this experiment informs how these stimulus-driven contextual interactions may interact with directed attention serving the observer's goals. Although these experiments revealed no modulatory effects of salience in the LGN, we wondered whether this structure might be sensitive to orientation-tuned suppression at more local spatial scales. In Chapter 2, we quantified the effects of near-surround suppression in the LGN and in the visual cortex. We used a voxelwise population receptive field (pRF) modeling approach to distinguish figure-ground modulations corresponding to either boundary enhancement or figural filling-in. Finally, Chapter 3 asked whether directed attention is necessary for orientation-defined figures to be enhanced in the LGN, as has been suggested by

recent work in the macaque (Jones et al., 2015; Poort et al., 2012), and if corticothalamic feedback likely drives this enhancement. Together, these experiments inform our understanding of fundamental mechanisms of contextual modulation in the early visual system, as well as how these bottom-up processes are integrated with an observer's directed attention toward a behaviorally relevant representation of the world.

CHAPTER 2

Characterizing the effects of feature salience and top-down attention in the early visual system

2.1 Introduction

At any given moment, our visual system is presented with far more information than it can process, yet we seldom feel that our visual experience is incomplete or degraded. This reflects a sophisticated balance of attentional mechanisms: the observer can willfully guide attention toward a task-relevant item, but is also predisposed to notice salient stimuli that appear outside of the current focus of attention. Our ability to navigate the dynamic visual environment critically depends on this interplay of top-down guidance of spatial attention and bottom-up processing of visually salient information.

The visual salience of an item or portion of the visual field describes the degree to which it can 'grab' attention; this subjective percept is strongly influenced by the local feature differences that occur within the visual scene. Behavioral studies have demonstrated that local differences in color, orientation, size, motion, or other basic features can be rapidly detected in visual search tasks (Nothdurft, 1993). Accordingly, computational models of visual salience have underscored the importance of an initial computation of local feature differences, although models differ on whether feature-selective interactions at early stages of visual processing sufficiently compute a salience map (Z. Li, 2002; Zhang et al., 2012), or whether a summation of various feature-contrast maps occurs at a higher stage of the visual pathway (Itti & Koch, 2001). These models can predict where people are more likely to look when viewing natural scenes (Itti & Koch, 2001), though cognitive goals and top-down factors can also exert powerful influence (Henderson, 2003).

Neurophysiological studies have found that early visual areas are strongly modulated by local feature contrast. The response of a V1 neuron, for example, is suppressed when a presented stimulus extends beyond the neuron's receptive field and into the surround; this suppression is orientation-tuned, such that greater suppression occurs when the orientation of the surround matches that of the center than when the center and surround orientations are orthogonal (Blakemore & Tobin, 1972; Cavanaugh, 2002; Nelson & Frost, 1978). This form of feature-tuned surround suppression has been shown to emerge for a variety of visual features (Allman et al., 1985; Blakemore & Tobin, 1972; Chao-Yi & Wu, 1994; DeAngelis et al., 1994; Jones et al., 2001; Kapadia, Ito, Gilbert, & Westheimer, 1995; Zipser et al., 1996). Multiple mechanisms likely contribute to feature-tuned surround suppression, including shorter-range interactions arising from horizontal connections between feature-selective neurons within V1 as well as more distal interactions that rely on feedback from higher extrastriate areas (Angelucci & Bressloff, 2006; Bair et al., 2003; Cavanaugh et al., 2002; Gilbert & Wiesel, 1979; Lamme, 1995). The computation of local feature contrast can be explained in terms of greater mutual inhibition, or divisive normalization (Carandini & Heeger, 2012) among neurons that share similar feature preferences. Thus, a local region that differs in feature content from its immediate surround should evoke a stronger visual response due to a release from feature-tuned surround suppression.

Human neuroimaging studies have investigated the effects of orientation-selective surround suppression and orientation-defined salience, using both simple and more complex displays. Studies employing large target stimuli with immediately abutting surrounds have reported powerful effects of surround suppression (Zenger-Landolt & Heeger, 2003) as well as a more modest feature-selective component of surround suppression (e.g., McDonald et al., 2009). Other studies have tested for effects of orientation-defined salience using more complex displays of multiple spatially separated gratings or lines, akin to the displays commonly used in behavioral

investigations of attention and visual search (e.g. Nothdurft, 1993). Results from these neuroimaging studies have been mixed: some studies find that a salient, uniquely oriented item evokes stronger responses in V1 (e.g., Schallmo et al., 2016; Zhang et al., 2012), while others find no reliable differences in early visual areas (Beck & Kastner, 2005; Bogler et al., 2013) or more complex interactions that depend on top-down spatial attention (Flevaris & Murray, 2015; Hopf et al., 2004). Our understanding of visual salience and its neural bases relies critically on testing with multi-item displays, which inform much of our knowledge of the mechanisms of attention and visual search. One factor to consider is that feature-tuned suppression may be more difficult to detect with widely separated items, as surround suppression effects in V1 are known to decrease as a function of retinotopic distance (Adesnik et al., 2012; Bair et al., 2003; Shushruth et al., 2013). Another consideration is that some prior studies have employed more complex visual tasks, raising the possibility that visual attention may interact with the processing of salient visual information.

Prioritization of visual stimuli is determined not only by stimulus-driven factors but also by voluntary selective attention, by which one can enhance the processing of stimuli corresponding to his or her specific goals and needs. While frontal and parietal networks have been implicated in the control of attention (for reviews: Corbetta & Shulman, 2002; Kastner & Ungerleider, 2000; Silver & Kastner, 2009; Squire et al., 2013), considerable work has shown that attention strongly modulates responses in early cortical visual areas, and even the lateral geniculate nucleus (LGN; Brefczynski & DeYoe, 1999; Gandhi et al., 1999; Ling et al., 2015; O'Connor et al., 2002; Schneider & Kastner, 2009).

Our goal in this fMRI study was to determine how directed spatial attention and feature-defined salience modulate responses at early stages of visual processing. One possibility is that once an item is attended, its salience does not provide any further enhancement. Alternatively, it could be that the effects of top-down attention and feature-based salience are simply summed to determine the overall prioritization of a stimulus in the visual field. We also sought to determine the earliest stages at which salience and top-down attention would lead to reliable modulations of visual activity. Recent work from our lab has demonstrated orientation-selective responses in the human LGN as well as modulatory effects of attention (Ling et al., 2015). Neurophysiological studies in animals suggest the presence of coarse selectivity in LGN (Cheong et al., 2013; Smith et al., 1990; Suematsu et al., 2012; Xu et al., 2002), arising primarily from elongation of the RF (Leventhal & Schall, 1983). Moreover, there is some evidence to suggest that orientation-tuned surround suppression emerges within cat LGN (Jones et al., 2000; Naito et al., 2007). Therefore, it seemed possible that the LGN might be sensitive to orientation-defined salience. However, feature-selective interactions among spatially distributed items require either long-range connections between feature-tuned neurons or feedback from higher visual regions (Nassi et al., 2013; Nurminen & Angelucci, 2014); these effects may be restricted to cortical visual areas, which have larger receptive fields and effective surrounds.

We used high-resolution fMRI at 7 Tesla to quantify the effects of top-down spatial attention and feature-defined salience at multiple levels of the visual pathway, including areas V1 through hV4 and the lateral geniculate nucleus (LGN). Observers viewed multi-item displays that contained a single salient grating with a unique orientation or motion direction, and were cued to attend to either the salient or a non-salient grating that was located immediately to the left or right of central fixation. The spatial separation of these elements emphasized longer-range feature-tuned contextual interactions. By testing both orientation- and motion-defined salience, we could evaluate the generality of the effects of feature-defined salience, as most previous research has have focused exclusively on orientation processing (Bogler et al., 2013; Flevaris & Murray, 2015; Hopf et al., 2004; McDonald, Seymour, Schira, Spehar, & Clifford, 2009; Schallmo et al., 2016; Zhang et al., 2012; but see also Harrison, Stephan, Rees, & Friston, 2007).

2.2 Materials & Methods

Participants

Six healthy adults (ages 22 to 31, one female) participated in Experiment 1, and six (ages 22 to 33, three females) participated in Experiment 2. Three subjects participated in both experiments. All participants had normal or corrected-to-normal vision, and were compensated for their role. All aspects of the study were approved by the Vanderbilt University Institutional Review Board.

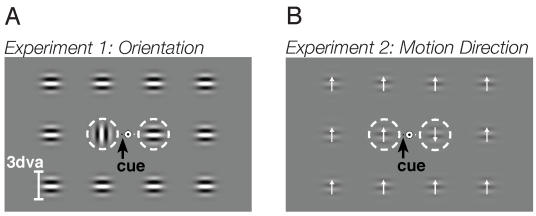

Figure 2.1. Annotated sample displays for both experiments. Feature-contrast salience is defined by orientation in A and by drifting motion direction in B (indicated by arrows, which were not present during the experiment). Dotted circles (likewise not present during the experiments) indicate the two locations at which the salient patch could appear and to which spatial attention could be directed. While the task-relevant contrast decrement occurred at both of these locations, the observer was instructed to only perform the detection task on one patch, as indicated throughout each block by a small cue. In these examples, if the participant was cued to attend to the black cue (as labeled), in A the left patch is attended/salient, while the right patch is unattended/non-salient; in B, the left patch is attended/salient, and the right patch is unattended/salient. Gabor patch edges were Gaussian blurred and also spatially separated by a gap of 0.8°,

Stimuli & Task: Experiment 1

Our experimental displays, illustrated in Figure 2.1A, were created using Matlab and Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Displays consisted of a 3 x 4 array of Gabor gratings presented at 100% Michelson - contrast, 1.5 cycles/degree, and Gaussian envelope σ= 0.4°. The stimuli measured 3.8° center-to-center. If one considers the effectively visible portions of the Gabor to span 5 standard deviations in width, where the contrast at the edge would dip to 4%, then each Gabor grating can be considered to have an effective width of 1° with a gap of 1.8° between the Gabors.

In each block, all gratings except for one appeared with a common orientation, either vertical or horizontal; the other, salient, grating had an orthogonal orientation. The salient grating appeared at one of two possible target locations, immediately to the left or right of the fixation point (0.5° in diameter with a central dot of 0.175°) – we refer to these as the left and right 'target locations' for our region of interest (ROI) analyses. All patches flickered on/off in 200ms intervals (i.e. 2.5 Hz), with the spatial phase of the Gabors randomized on each presentation.

In each experimental block, the participant was shown a cue at fixation, instructing them to covertly attend to one of the target gratings while keeping their eyes focused on the central fixation point. The attentional cue consisted of a pair of dots (0.1°) that appeared to the left and right of fixation. One of these dots was black, and the other white; each participant was told to attend to the side marked by one of these two colors throughout the experiment. The spatially balanced design of this fixation cue ensured equivalent stimulus-driven activity in each hemifield and avoided potential effects of exogenous cuing, which could occur with a single lateralized cue (Jehee, Brady, & Tong, 2011). The attended cue color was counterbalanced across participants.
Participants were asked to perform an attentionally engaging contrast-decrement detection task on the cued stimulus. We chose this task to direct covert spatial attention to the cued grating, and to minimize the potential influences of feature-based attention. In previous work, we have shown that the performance of an attentional task on the orientation of a grating leads to a strong enhancement of orientation-selective responses in V1-V4, while tasks that require attending to the contrast of a grating do not (Jehee et al., 2011). Contrast decrements occurred at independent intervals at each of the two target locations throughout the experimental block, but participants were instructed to respond to decrements only at the cued location via a button box. The

decrement occurred 8 times per 16s block at randomly determined intervals, lasting for the full 200ms duration of that 'on' interval with a minimum 800ms time difference between targets.

Stimuli & Task: Experiment 2

The stimulus parameters were similar to those of Experiment 1, except for two key differences: motion direction, rather than orientation, was used to define salience, and the drifting Gabor patches were continuously presented at 30% contrast (see Figure 2.1B). All of the patches were oriented horizontally and phase-randomized, and the gratings drifted either upward or downward at a speed of 5 degrees/second (temporal frequency, 7.5 Hz) within a stationary window. The salient patch moved in the opposite direction relative to the motion of all other Gabor gratings in the array. We chose upward-downward motion directions to minimize the likelihood of inducing involuntary optokinetic nystagmus (OKN; Honrubia et al., 1968) and to further ensure that if a small eye movement was occasionally induced, it would not be directed toward either of the lateralized target gratings. Participants were asked to perform an attentionally engaging contrast-decrement detection task on the cued stimulus. The contrast decrement lasted for 400ms, and occurred on 20% of these implicit 400ms intervals within each 16s stimulus block.
In a separate behavioral eye-tracking session, we confirmed that our participants could maintain stable fixation while viewing these arrays of moving gratings. Participants performed the same experimental task on displays that matched the stimuli and timing (16s blocks) of the fMRI experiment. Overall, our subjects kept their eyes fixed well within the 0.5° fixation dot. Horizontal eye movement deviations (reported as the standard deviation from the true fixation position), ranged from 0.07-0.33° across the 6 participants (median 0.15°); vertical deviations ranged from 0.11-0.48° (median 0.23°).

Experimental Design and Procedure

Apart from differences in the stimuli, the design of both experiments was identical. There were three experimental factors that consisted of: 1) location of the salient grating (left or right of fixation), 2) location to be spatially attended (left or right grating), and 3) the context feature (horizontal or vertical orientation in Experiment 1, upward or downward motion in Experiment 2). This 2 x 2 x 2 factorial design led to 8 experimental conditions, which were presented in a randomized order within each run. Each experimental run followed a 16s block design (272s duration), with each of the 8 experimental conditions occurring once, interleaved with 16s fixation-rest periods that also occurred at the beginning and end of each run. A fixation circle remained present throughout the experimental run; the spatial attention cue appeared 1s prior to the onset of each stimulus block, informing participants to direct their attention towards the left or right grating. Participants were informed of their performance accuracy at the end of each run, and the magnitude of the contrast decrement was adjusted between runs so that the detection hit rate for each subject fell in a range from approximately 70-90%; across subjects, the magnitude of the contrast decrement ranged from 25% to 40% of the original contrast value.

Behavioral Performance

In Experiment 1, performance at the contrast detection task led to a mean hit rate across individual subjects of 76.4% (st. dev. = 12.1%). Contrast decrements applied to the salient target in the display were detected with 77.4% accuracy, while those applied to non-salient targets yielded 75.4% correct performance. There was no significant difference in performance for salient and non-salient targets across the group ($t(5) = 0.81$, $p = 0.45$). In Experiment 2, subjects averaged 83.4% hit rate (st. dev. = 7.1%), 83.7% correct on salient targets and 82.5% correct on non-salient targets; again, there was no reliable difference in performance for the salient and non-salient

targets (t(5) = .59, p = 0.58). Our task, which requires contrast discrimination using supra-threshold contrast changes, was designed to manipulate the locus of spatial attention rather than to assess sensitivity to contrast, which has been shown to be enhanced at salient locations. (e.g., Joo et al., 2012; Kapadia et al., 1995). These behavioral results suggest that participants were able attend selectively to the non-salient grating location when it was the target, and were not distracted by the presence of the salient grating.

fMRI Data Acquisition

Data were acquired using a Philips Achieva 7-Tesla MRI scanner at the Vanderbilt University Institute of Imaging Science (VUIIS), with a 32-channel head coil equipped for SENSE imaging. The functional scans employed single-shot gradient-echo echoplanar imaging to measure BOLD activity, and were aligned approximately parallel to the AC-PC line to best capture the LGN and occipital pole. Subjects were scanned using either 1.5 mm isotropic voxel resolution in 20 slices (6 subjects in Experiment 1, 4 in Experiment 2) or a 2 mm x 1.875 mm x 1.875 mm voxel resolution in 36 slices (2 subjects in Experiment 2). The 1.5 mm thickness scans used the following parameters: 2s TR, 26ms TE, 75° flip angle, 192 mm FOV, with no gap. The 2 mm scans used identical parameters, except that the FOV was 210 mm. Twelve to 16 runs of functional data and 2-3 localizer runs were collected for each subject. The spatial extent of our slice prescription allowed us to monitor BOLD activity in the LGN and areas V1, V2, V3 and hV4. Area V3A often appeared outside of this slice prescription and was therefore not included in this study.

Functional ROI Definition

Cortical visual areas V1-hV4 were demarcated using standard retinotopy procedures, using data acquired from separate scan sessions at 3 Tesla (Engel et al., 1997; Wandell et al., 2007). We used a typical phase-encoded design in which subjects fixated while they viewed flickering checkerboards consisting of rotating wedges to map polar angle and expanding rings to map eccentricity. Retinotopy data was acquired using a Philips 3T Intera Achieva MRI scanner equipped with an 8-channel coil. Subjects were scanned using 3 mm isotropic resolution (TR 2 s, TE 35 ms, flip angle 80°, 28 slices, 192 x 192 FOV). Boundaries between retinotopic areas V1-hV4 were delineated by hand, by identifying reversals in the phase of the polar angle map measurements; the resulting ROIs were aligned to the functional space of the current experiment using FSL and Freesurfer software.

In the experimental scan session, we ran 2-3 runs of a visual localizer to identify the target regions of interest corresponding to the spatial extent of the left and right gratings. This involved presenting flickering checkerboards at full contrast within a Gaussian contrast-envelope at each of the two target locations, with alternating 16-s cycles of left or right stimulation, as well as a 16 s fixation period at the beginning and end of each run. Individual checks were 0.75° in width, and contrast-reversed at a rate of 5Hz. Subjects did not perform a task during localizer runs, and were instructed to keep their eyes on a central fixation point. Cortical ROIs were selected from the conjunction of retinotopy and a statistical map of the left vs. right contrast of our functional localizer. We report results from the 100 most functionally selective voxels as defined by the t-statistic map in each lateralized ROI in early visual areas.

The lateral geniculate nucleus was defined functionally from the same localizer contrast, using a t-value threshold of no less than 2.8; thresholds for each subject were selected to yield distinct, continuous clusters of voxels such that left and right hemisphere nuclei were generally aligned dorsally and were maximally lateralized, so as to avoid the inclusion of the pulvinar region. As the LGN cannot be localized anatomically from T1- or T2-weighted images, we cannot be fully sure that our regions of interest do not include other portions of the thalamus. However, the LGN is

more readily activated by visual stimulation than other subcortical regions, and there is evidence to suggest that functional localizers that rely on passive viewing, as ours did, do not activate pulvinar activity as strongly as the LGN (Kastner et al., 2004).

In sessions using a 1.5 mm isotropic voxel size, the bilateral LGN region of interest consisted of an average of 72.9 voxels (stdev. = 16.8); in the two participants who were tested using a larger voxel size of 2 mm x 1.875 mm x 1.875 mm voxel size, the LGN region of interest encompassed an average of 42.5 voxels.

fMRI Analysis: Preprocessing

Data were preprocessed using FSL and Freesurfer tools (documented and freely available for download at http://surfer.nmr.mgh.harvard.edu), beginning with 3D motion correction and linear trend removal, followed by a high-pass filter cutoff of 60s. Functional images were registered to a reconstructed anatomical space for each subject; this registration was first automated in FSL and then checked and corrected by hand. This allowed the alignment of the current fMRI data to the retinotopy data, which was collected in a separate session. The functional localizer was spatially smoothed using a 1-mm Gaussian kernel; no spatial smoothing was done for the experimental runs. Further analyses were conducted using a custom Matlab processing stream. For each ROI-based analysis, data were masked using the labels defined from the retinotopy and functional localizer, which corresponded to the left and right target grating locations in each visual area. Each voxel's intensities were normalized by the mean of the time series, converting to mean percent signal change within each run. Outliers were defined as time points for which the voxel's response measured more than 3 times its standard deviation from its mean, and were Winsorised (Hastings, Mosteller, & Tukey, 1947). This condition-blind preprocessing step minimizes the impact of rare spikes in MR intensity while preserving the temporal structure of the responses in each voxel. Only a small fraction of a percentage of data points in the current experiments were marked as outliers (0.26-0.41% across Experiments and ROIs). Further, voxels that left the volume at any time point due to head motion were excluded from that run's analysis. Additionally, we found that during one session in Experiment 1, a combination of slice placement and head motion caused some regions to intermittently clip the edge of the volume. For this subject, we excluded 2 runs from the analysis of bilateral V2 and right hV4, based on drop-out of visually selective responses that was restricted to these ROIs.

fMRI Analysis: Mean BOLD

To calculate the average mean BOLD response for each experimental condition, we first converted the MR time series of each voxel from signal intensity units to units of percent signal change. Next, we calculated the average BOLD amplitude for each block, after shifting the response period by 2 TRs to account for hemodynamic lag. To account for baseline differences preceding each condition's onset, the average of the 2 TRs immediately preceding each block was subtracted from its mean. For every stimulus block, we obtained a measure of the mean BOLD response in each ROI, and further calculated the overall BOLD response across all blocks for each experimental condition.

Conditions were defined relative to lateralized ROI responses such that, for example, a block in which the left target patch was salient was labeled 'salient' in the right hemisphere and 'non-salient' in the left hemisphere; data were pooled in this way across the left and right hemisphere ROIs. We also combined data across specific feature values (e.g. when the salient target was vertical vs. horizontal in Experiment 1), which were not pertinent to our hypotheses. In each bilateral region, we performed a 2 x 2 repeated measures ANOVA to characterize the effects of salience and spatial attention, as well as their interaction.

2.3 Results

fMRI: Experiment 1

In this experiment, orientation was used to define one salient target location in an array of gratings, while the participant spatially attended to the salient target or to a non-salient target in the contralateral hemifield (see Figure 2.1A). We compared fMRI responses in early visual areas evoked by the salient grating and by the contralateral non-salient grating across attended and unattended conditions, pooling the data across the left and right ROIs.

The time course of mean BOLD responses in the LGN and areas V1-hV4 are shown in the top panel of Figure 2.2B, while the mean response amplitudes observed in each block are plotted in the bottom panel. Figure 2.2A shows the modulatory effects of both attention and salience as average difference in BOLD response. Spatial attention led to positive increases in the BOLD response throughout the visual pathway, in both subcortical and cortical regions of interest. We observed reliable effects of top-down attentional modulation in all individual cortical visual areas (V1: $F(1,5) = 7.59$, $p < 0.041$; V2: $F(1,5) = 8.08$, $p < 0.037$; V3: $F(1,5) = 23.7$, $p < 0.005$; hV4: $F(1,5) = 19.2$, $p < 0.008$). This effect was likewise present in the LGN ($F(1,5) = 9.11$, $p < 0.030$), consistent with previous fMRI reports that the human LGN can be reliably modulated by spatial attention (Ling et al., 2015; O'Connor et al., 2002; Schneider & Kastner, 2009).

In contrast, we found that the effect of salience was evident only in the visual cortex. No reliable difference between salient and non-salient items was observed in the LGN ($F(1,5) = 0.028$, $p = 0.87$). Early visual cortical areas, however, exhibited clear and reliable effects of salience, with higher mean BOLD responses to the salient grating than to the non-salient grating that matched the orientation of the surrounding context (V1: $F(1,5) = 9.47$, $p < 0.003$; V2: $F(1,5) = 8.53$, $p < 0.034$; V3: $F(1,5) = 42.8$, $p < 0.002$; hV4: $F(1,5) = 7.35$, $p < 0.043$).

Interestingly, the main effects of salience and spatial attention did not significantly interact in

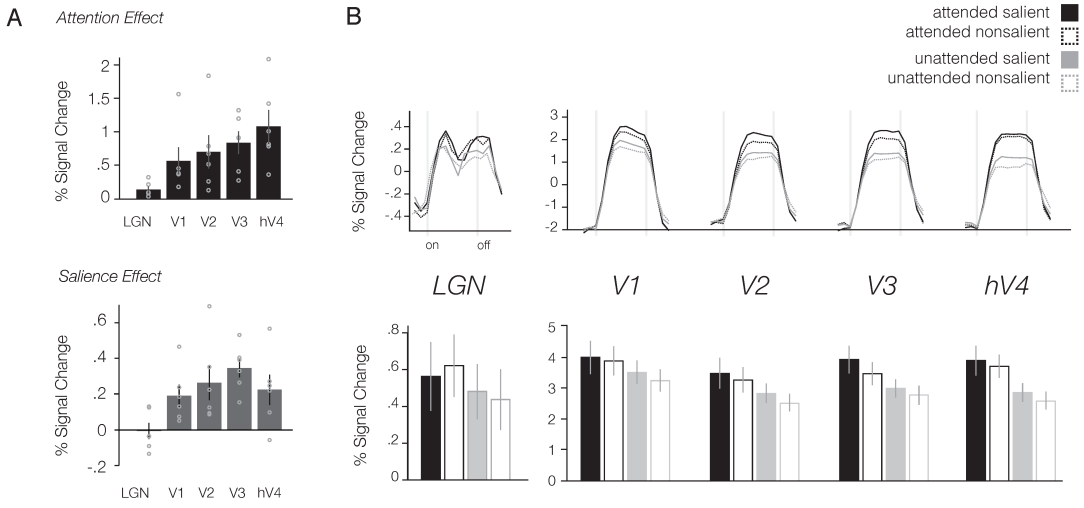

Figure 2.2. Results of Experiment 1 for the four salience/attention conditions. A) Magnitude of the attention and salience effects across ROIs, computed as difference in percent signal change. Dots show the effect in individual subjects, and error bars depict one standard error (across subjects). In the LGN, only attention significantly modulated BOLD responses; neither the main effect of salience nor the interaction effect were significant. A significant effect of salience first emerges in V1 and is evident in each cortical ROI. Attention also modulated BOLD responses in V1 through hV4, but the two effects do not reliably interact in any region. B) The top panel shows mean ROI time-courses time-locked to the beginning of each experimental block, which lasted 8 TRs and is demarcated by dotted lines. The bottom panel shows the same data averaged across the block, offset to account for hemodynamic lag, and normalized by subtracting the mean response of the 2 TRs immediately preceding the block. Error bars on block averages depict one standard error in the mean BOLD response across subjects.

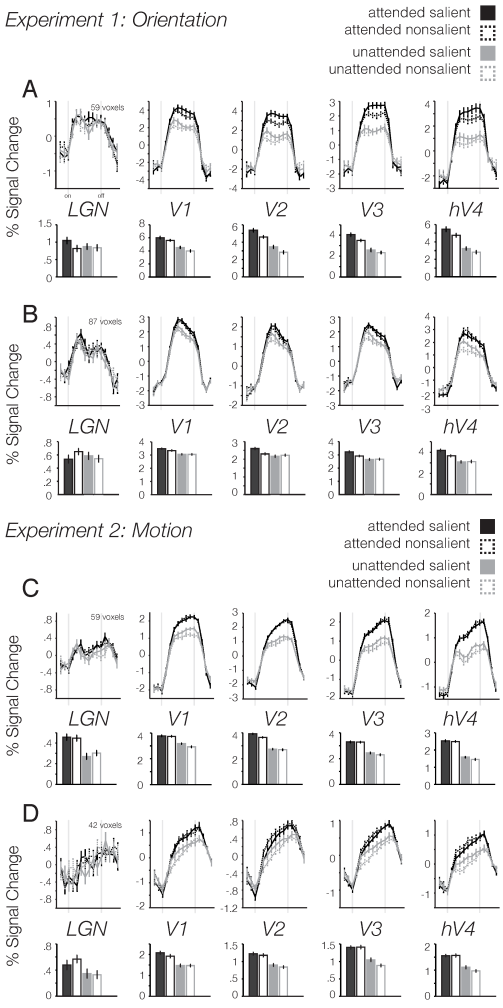

Figure 2.3. Examples of average fMRI time courses from representative subjects in Experiment 1 (A, B) and Experiment 2 (C, D). Error bars depict +/- 1 SEM for each experimental condition, and each of the four panels shows data from a different individual. As expected, there is some variability between subjects; fMRI responses in the LGN are also more variable than those in early visual areas, partly because of the LGN's smaller size and the presence of greater physiological noise in midbrain structures.

any of our measured ROIs (V1: $F(1,5) = 1.83$, $p = 0.24$; V2: $F(1,5) = 0.14$, $p = 0.73$; V3: $F(1,5) = 2.34$, $p = 0.19$; hV4: $F(1,5) = 0.12$, $p = 0.74$). That is, the enhancement of salient items is not contingent on the observer's attentional state or goals. The LGN, which did not show a significant main effect of salience but was reliably modulated by attention, likewise showed no interaction effect ($F(1,5) = 1.07$, $p = 0.35$).

As shown in Figure 2.2A, comparison of the magnitude of these effects across cortical visual areas reveals a significant difference between regions (ANOVA, $F(3,15) = 5.74$; $p < .009$), suggesting a trend of increasing attention modulation along the visual hierarchy (F-test, $t(2) = 9.00$, $p < .013$). There appeared to be no difference in the magnitude of salience modulation across visual areas V1 through V4 ($F(3,15) = 1.68$, $p = 0.21$). Including the LGN in this comparison, however, did yield significant differences across regions of interest ($F(4,29) = 7.8$, $p < 8.0 \times 10\text{-}4$). Specifically, the salience modulation of the LGN was not significantly different from zero ($t(5) = 0.16$, $p = 0.87$), and significantly weaker than the salience effect observed in V1 ($t(5) = 4.7$, $p < 0.0054$).

We performed a Bayes factor analysis to estimate the likelihood that the results from the LGN could have arisen from a null effect of salience. We calculated the JZS Bayes Factor (Rouder, Speckman, Sun, Morey, & Iverson, 2009) applying a scale factor of 1 for the prior on effect size for the alternative hypothesis, and obtained a value of 3.44 in favor of the null hypothesis. It has been suggested that odds factors greater than 3 should be considered as evidence in favor of a hypothesis (Jeffreys, 1961); thus, the LGN data are not strongly conclusive but do tend to favor the null hypothesis.

To further test for the possible presence of an interaction effect between salience and spatial attention, we performed an ANOVA on the data of individual subjects, using the fMRI response amplitudes observed on individual fMRI blocks for each participant (Figure 2.6). Consistent with our group-analysis results, we observed statistically significant main effects of salience and top-down attention in early visual areas for the majority of individual subjects. However, the interaction between salience and attention did not reach statistical significance in any ROI for any subject. To illustrate the quality of data collected, as well as the differences between measurements of the LGN and of cortex, representative single-subject data from this Experiment is shown in Figure 2.3A/B.

All preceding analyses were performed on the mean activations of the functional ROIs; in the cortex, these were selected from individual retinotopic regions, and defined as the 100 most selective voxels in each hemisphere based on independent localizer runs. We assessed whether this ROI size criteria had a meaningful effect on the pattern of our results by calculating the

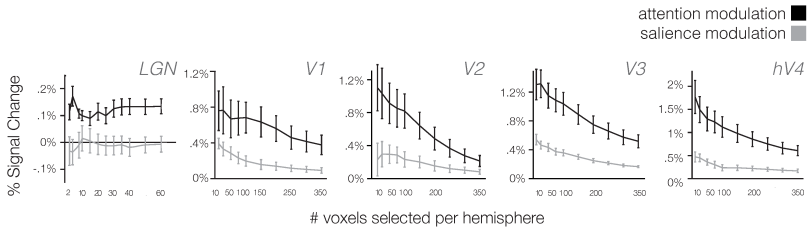

Figure 2.4. Magnitude of attention and salience effects plotted as a function of ROI size. We calculated the difference in BOLD response for attended minus unattended conditions and salient minus non-salient conditions, for a wide range of ROI sizes, ranging from 2-60 maximum voxels per hemisphere in the LGN and from 10-350 voxels per hemisphere in individual cortical visual areas. Error bars indicate one standard error across subjects. Our findings of positive salience and attention effects in each cortical ROI are highly consistent across a wide range of ROI sizes.

magnitude of attentional and salience-based modulation for a wide range of ROI sizes. Figure 2.4 plots the magnitude of attention and salience-based enhancement as a function of the number of voxels selected from each ROI. As can be seen, these effects are broadly consistent across a wide range of ROI sizes.

The results of Experiment 1 suggest that effects of orientation-defined salience emerge in the primary visual cortex, with a similar degree of enhancement observed in higher extrastriate visual areas. Although orientation-specific responses have recently been demonstrated in the human LGN (Ling et al., 2015), here we find no evidence of enhanced responses to orientation-defined salience at this subcortical site. However, we do find that spatial attention reliably modulates responses in the LGN, in agreement with previous fMRI studies (O'Conner et al., 2002; Ling et al., 2015). The finding that modulatory effects of visual salience and top-down attention emerged at different levels of the visual pathway provides support for the proposal that these mechanisms operate independently and are functionally distinct. Consistent with this proposal, the modulatory influences of spatial attention and orientation-defined salience appeared to be separable and additive in early visual areas of interest.

fMRI: Experiment 2

Do the effects of orientation-based salience generalize to other feature domains? Most studies that find salience enhancement in early visual areas have focused on orientation as the defining feature (Joo et al., 2012; Z. Li, 2002; Schallmo et al., 2016; Zhang et al., 2012); however, direction-selective contextual interactions in V1 have been reported for spatially separated moving stimuli (Harrison et al., 2007). In Experiment 2, we were motivated to test whether motion-defined salience would also yield similar effects of top-down attention and stimulus-driven salience across the visual hierarchy. In this experiment, one Gabor grating drifted in a direction opposite to that of all other gratings in the display and was deemed salient. Meanwhile, observers performed an attentionally demanding task on either the salient item or on a non-salient item, as illustrated in Figure 2.1B.

As can be seen in Figure 2.5, mean BOLD responses were consistently greater for attended than unattended items. A repeated-measures ANOVA indicated that this modulatory effect of attention was statistically significant in all regions of interest, including the lateral geniculate nucleus (LGN: $F(1,5) = 27.5$, $p < .004$; V1: $F(1,5) = 47.9$, $p < 9.7 \times 10\text{-}4$; V2: $F(1,5) = 44.0$, $p < 0.0013$; V3: $F(1,5) = 37.6$, $p < 0.002$; hV4: $F(1,5) = 96.6$, $p < 1.9 \times 10\text{-}4$). While absolute BOLD amplitudes in the LGN were variable across subjects, as evidenced by the size of the error bars in Figure 2.5B (which correspond to +/- 1 S.E.M. across the six subjects), the within-subject effect of attention was statistically reliable in the LGN. Moreover, the effect of attention appeared highly consistent in both LGN and visual cortex, with every participant showing attentional modulations in

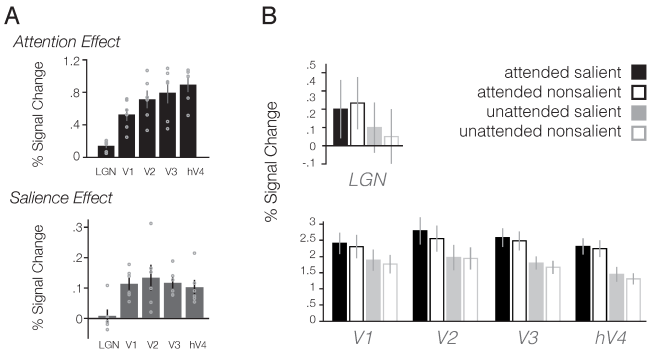

Figure 2.5. Mean BOLD amplitudes in Experiment 2, in which salience was defined by the direction of drifting motion of Gabor patches. A) Attention and salience effects (percent signal change difference) across ROIs, with dots plotting individual subjects' results and error bars showing one standard error across subjects. B) Results in each condition are averaged across the stimulus block and normalized by a pre-stimulus window for each condition. The pattern of results follows that of the first experiment: attentional enhancement is evident in each ROI, including the LGN, while salience modulated only cortical mean BOLD responses. In every region studied, the effects **of** attention and salience did not significantly interact.

the predicted direction. This is illustrated in Figure 2.5A, which depicts individual subjects' effect amplitudes as grey dots overlaid on the mean effect across ROIs; representative individual subject data for Experiment 2 is shown in Figure 2.3C/D.

In the visual cortex, we again observed significant modulation by salience: non-salient gratings that drifted in the same direction as the contextual gratings elicited weaker mean BOLD responses than did items that drifted in the opposite direction. Significant enhancement of salient items was observed in each of the cortical regions of interest, beginning in V1 ($F(1,5) = 33.7$, $p < .003$) and persisting through V2-hV4 (V2: $F(1,5) = 10.0$, $p < 0.026$; V3: $F(1,5) = 42.5$, $p < 0.002$; hV4: $F(1,5) = 18.1$, $p < .009$). However, this motion-defined salience did not significantly modulate activity in the LGN ($F(1,5) = 0.14$, $p = 0.72$).

Critically, the enhancements of BOLD activity elicited by attention and salience did not appear to interact in this experiment. Analysis of variance indicated that salience and attention conditions did not lead to a significant interaction effect in any of the cortical regions of interest (V1: $F(1,5) = .015$, $p = 0.91$; V2: $F(1,5) = 4.4$, $p = 0.09$; V3: $F(1,5) = 0.083$, $p = 0.78$; hV4: $F(1,5) = .62$, $p = 0.47$), nor in the LGN $F(1,5) = 1.90$, $p = 0.23$). Thus, the degree of enhancement observed for salient items appeared comparable under conditions of attention and inattention.

Figure 2.5A shows the modulatory effects of spatial attention and of salience in Experiment 2 for each region of interest. Similar to the results of the first experiment, we see a trend toward increasing attentional modulation as one ascends from V1 to V4 (ANOVA $F(3,15) = 4.26$, $p = 0.023$; F-test $t(2) = 7.19$, $p < .019$). Salience modulations appear similar in magnitude across areas V1-hV4, and no reliable difference was found among these cortical ROIs ($F(3,15) = .38$, $p = 0.77$). A difference emerged across brain areas when data from the LGN was included in the analysis of variance ($F(4,29) = 7.46$, $p < 9.0 \times 10\text{-}4$) and paired comparisons indicated that salience modulation was significantly weaker in the LGN than in V1 ($t(5) = 4.88$, $p < .0046$). As in the previous experiment, salience modulation in the LGN did not significantly differ from zero ($t(5) = .38$, $p = 0.72$). In this case, a Bayes factor analysis indicated a value of 3.25, in moderate favor of the null hypothesis. We again note the difficulty in interpreting a null effect in the LGN: small effect sizes, when present, are difficult to detect, and the increased physiological noise when measuring subcortical activity can impede the reliability of fMRI measures. However, the magnitude of salience modulation appears clustered around zero for the majority (5/6) of our individual subjects (Figure 2.5A), in contrast with individual effects of attention in the LGN or of salience in cortex.

Overall, the two experiments yielded very similar patterns of results, demonstrating that

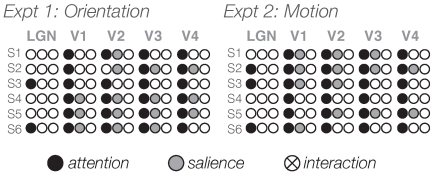

Figure 2.6. Individual subject ANOVA results for Experiments 1 (left) and 2 (right). Here, each experimental session was analyzed independently, with experimental run as the repeated measure. Filled circles indicate significant effects of attention (black), salience (grey), or their interaction (cross) at $p < 0.05$ level. Subject labels (S1-S6) are arbitrary and unmatched between the two experiments. While many subjects exhibit significant main effects of attention and salience, the interaction of these two factors is not significant in any ROI in any subject.

both orientation- and motion-defined salience lead to common effects of enhancement in the early visual cortex. These effects did not significantly interact in any region of interest, neither in the group analysis nor in analysis results for any individual subject, suggesting that stimulus-driven salience and top-down attention provide independent sources of modulatory influence at early stages of the visual system.

2.4 Discussion

In two experiments, we used high-field fMRI to characterize the effects of top-down attention and feature-defined salience at multiple levels of the visual hierarchy. We found that orientation- and motion-defined salience consistently enhanced responses to stimuli in areas V1 through hV4, implying a common mechanism for detecting local feature differences across spatially disparate stimuli. These feature-tuned contextual interactions are critical for the computation of the local salience of a region, as predicted by modeling work (Itti & Koch, 2001; Z. Li, 2002). Additionally, we found that directed spatial attention enhanced responses in all regions of interest, including the LGN; our findings add to a growing number of studies indicating that activity in the LGN can be altered by top-down attentional goals (Ling et al., 2015; McAlonan, Cavanaugh, & Wurtz, 2008; O'Connor et al., 2002; Schneider & Kastner, 2009). In contrast, we did not find evidence that LGN activity is reliably modulated by feature-defined salience. Although one must be cautious about the interpretation of null effects, the absence of a statistically reliable effect of salience in the LGN is consistent with the proposal that long-range feature-selective interactions depend on cortical mechanisms (Bair et al., 2003; Nassi et al., 2013; Shushruth et al., 2013; Shushruth, Ichida, Levitt, & Angelucci, 2009). Our results support the proposal that at the earliest stages of visual processing, salience is distinct from mechanisms of top-down attention.

Consistent with this view, we found that the effects of salience did not lead to significant interactions with the effects of attention in any of our regions of interest. Our findings suggest that the effects of top-down attention and bottom-up salience are summed in a simple additive manner, such that both mechanisms distinctly inform the prioritization of items within a visual scene. These results concur with a recent fMRI study by Schallmo et al. (2016) though it should be noted that more complex interactions between bottom-up mechanisms and top-down attention may occur if attention is allowed to spread to other neighboring items. For example, feature-based attention may counteract the influence of feature-tuned suppression with certain stimulus configurations and tasks. Flevaris and Murray (2015) found that an attended target grating evoked greater responses when flanked by orthogonally oriented gratings than when the target and flankers were iso-oriented. When participants attended to one of the flanking gratings, however, attention appeared to spread to the target if the gratings shared a common orientation. This reported interaction presumably emerged due to the spatial spreading of feature-based attention (Saenz, Buracas, & Boynton, 2002). In the present study, attention was cued to shift between target locations in separate hemifields, thereby minimizing the spread of attention between the lateralized target regions. Under these conditions, we observed separate additive effects of top-down attention and salience at each of the target locations. We should note, however, that the limited range of contrast levels in the current experiments may bias us toward finding additive effects if the neural response function is saturated (Carrasco, 2006).

Our results support the view that processing of feature-defined salience in V1 relies on horizontal interactions within the region (Adesnik et al., 2012; Stettler, Das, Bennett, & Gilbert, 2002) as well as more long-range effects of feedback from higher extrastriate areas to V1 (Angelucci & Bressloff, 2006; Bair et al., 2003; Nassi et al., 2013). Such feature-selective interactions are believed to underlie a variety of early contextual effects (Stettler et al., 2002)

including tilt repulsion (Dragoi, Rivadulla, & Sur, 2001; Jin, Dragoi, Sur, & Seung, 2005), collinear enhancement (Kapadia et al., 1995), and figure-ground processing (Lamme, 1995); the current experiments considered how these early visual interactions may function over spatially disparate elements in a scene, which is critical toward building a behaviorally relevant representation of salience. Of course, further processing of salience occurs in higher-level extrastriate and parietal regions as the observer orients his or her attention during visual search (Melloni, van Leeuwen, Alink, & Müller, 2012), integrates many features in a naturalistic environment (Bogler, Bode, & Haynes, 2011), shifts attention (Kincade, Abrams, Astafiev, Shulman, & Corbetta, 2005), or plans eye movements (Fecteau & Munoz, 2006; Gottlieb, Kusunoki, & Goldberg, 1998; Mazer & Gallant, 2003). The current work informs the wider study of salience by investigating the early cortical stages involved in these processes, and by demonstrating that the representation of this information is unaffected by the observer's attentional focus.

Our results in the LGN suggest that feature-selective surround effects in this region are either negligible or too weak to be reliably detected in the current paradigm, which relied on spatially separated gratings to manipulate salience. Existing neurophysiological work is consistent with these views. Studies of the LGN in cats and monkeys have demonstrated size-tuned responses to gratings due to strong suppression from the immediate surround, which likely arises from mechanisms originating from the retina (Alitto & Usrey, 2008) as well as effects of cortical feedback (Jones et al., 2012). While some orientation bias has been reported in LGN neurons in several species (e.g. cat: Suematsu et al., 2012; owl monkey: Cheong et al., 2013; Xu et al., 2002; and macaque: Smith et al., 1990), positive effects of feature-tuned suppression in this structure have been predominantly reported in the cat (Cudeiro & Sillito, 1996; Jones et al., 2000; Naito et al., 2007). Given the known anatomical differences between cat and primate LGN, our null effect could reflect a lack of feature-tuned surround suppression in the region. It is also possible that feature-tuned suppression in the LGN may operate at a more local spatial scale than our Gabor array displays were designed to probe (Angelucci & Bressloff, 2006), as suggested by our group's positive report of modulatory effects of orientation masking in this structure (Ling et al., 2015). Alternatively, it may be that particular elements of our experimental design precluded us from detecting modulatory effects: we were limited in testing only one contrast level in each experiment, and using a constant spatial frequency (1.5cpd) throughout the study.

These findings contribute to a broader understanding of how surround suppression functions toward a behaviorally relevant representation of the visual scene. The effects of surround suppression have been shown to fall off with distance; however, modulations in apparent contrast of a central stimulus can be detected even with spatial separation of several degrees between the center and surround (Cannon & Fullenkamp, 1991; Petrov & McKee, 2006). While the generalization of surround suppression to natural image or movie inputs has posed a challenge in vision research, recent work has suggested that the suppression may be gated by the degree of redundancy in a natural image (Coen-Cagli et al., 2015). Homogeneity in a natural image was found to elicit stronger suppression, which may act to reduce redundancy in visual input and code natural input more efficiently (Vinje & Gallant, 2000). fMRI studies of perceptual grouping report effects of global configuration, consistent with this view of redundancy reduction (Joo et al., 2012). Given the visual system's adaptive nature, sensitivity to the statistics of natural input may provide a bridge from fundamental mechanisms of suppression and contextual interactions to behaviorally relevant representations of salience in the environment (Coen-Cagli, Dayan, & Schwartz, 2012; Kayser, Körding, & König, 2004), especially as these mechanisms function across different spatial scales (Nurminen & Angelucci, 2014). By representing the influences of visual salience and top-down attention at the earliest stages of cortical visual processing, the visual system is able to

achieve a balance between automatic prioritization of local regions throughout the visual field and voluntary guidance based on current goals and tasks.

Acknowledgements

This work appeared in the *Journal of Neurophysiology* in July 2017. Sam Ling, Devin McCormack, and Frank Tong are co-authors on this project.

CHAPTER 3

Distinct effects of boundary detection and figure enhancement in human V1

3.1 Introduction

The visual system is sensitive to differences in visual features; this seemingly simple property of neural responses allows us to make complex visual inferences about the locations and boundaries of meaningful figures in the visual world. Differences in luminance or color lead to distinctly visible edges, which serve as a powerful cue for segmentation and detection. Segmentation becomes much more difficult, however, when the figure resembles the color, luminance, and spatial frequency contents of the surround (Figure 3.1A). How might the processing of feature differences at local and larger scales contribute to the visual detection of figures?

Neurophysiological studies of figure-ground processing have focused heavily on characterizing V1 responses in awake-behaving monkeys (Nothdurft, 1993; Treisman & Gelade, 1980; Wolfe, 1994). These studies find not only robust border detection responses near the boundary between figure and ground, but also a subsequent effect of figure enhancement that emerges during the sustained period of the V1 response. Curiously, figure enhancement is greatly attenuated if the animal must attend elsewhere to perform a visual task, while border responses remain intact. This has been taken to suggest that covert spatial attention may be a critical mediating factor for figure enhancement to occur (Poort et al., 2012; Schira, Fahle, Donner, Kraft, & Brandt, 2004). Notably, however, studies that have examined figure-ground modulations have commonly relied on reinforcement-learning paradigms to train the animal to search for and

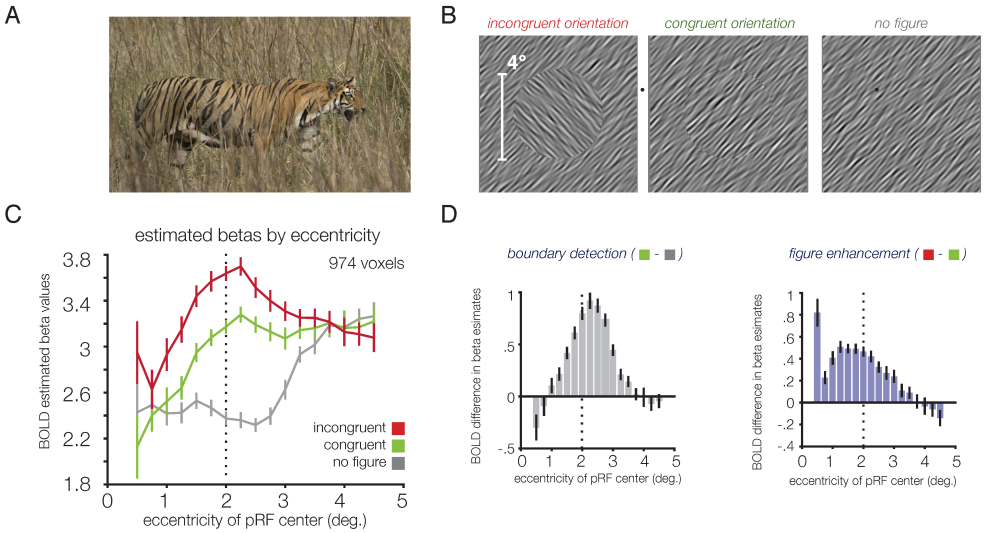

Figure 3.1. Sample stimuli and fMRI BOLD results as a function of pRF eccentricity. (A) Example of animal camouflage, which is pervasive in the natural visual world. Detection of figures is much more difficult when they resemble the color, luminance, and spatial frequency content of the surround. (B) Examples of spatially filtered oriented-noise stimulus displays used in the experiments. Three main conditions are depicted: (left) no figure; (middle) congruent orientation, in which the figure and surround are iso-oriented but sampled from distinct patches of noise, creating a visible phase-defined edge around the figure; (right) incongruent orientation, in which the figure and surround were orthogonally oriented. The spatial frequency depicted here is much lower than the experimental stimuli (0.5-8 cycles per degree) to improve visibility. (C) V1 BOLD responses plotted using mean beta weights for each experimental condition as a function of the eccentricity of each voxel's pRF center. Dotted line at 2° eccentricity indicates the figure-surround boundary. (D) V1 BOLD effects of boundary detection based on the difference in estimated beta weights between the congruent-orientation and no-figure conditions. (E) V1 BOLD effects of figure enhancement, defined as the difference between incongruent-orientation and congruent-orientation conditions. The spatial profile of this effect as a function of eccentricity is clearly distinct from that of panel D, and appears to extend through the full eccentricity range of the figure but to decline beyond 2°.

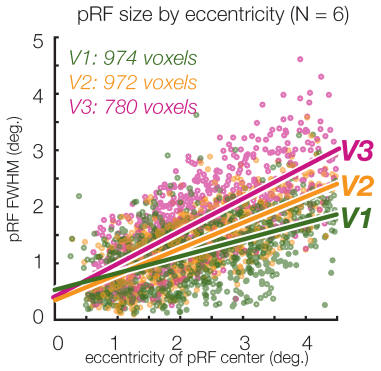

Figure 3.2. pRF estimates across subjects (N=6) and regions of interest. Consistent with prior reports, pRF size increases as a function of eccentricity, and as one ascend the visual heirarchy.

saccade to orientation- or motion-defined figures (cf. Itti & Koch, 2001; Koch & Ullman, 1985). It is now known that stimulus-contingent reinforcement can greatly enhance the visual salience of an item, and that such associative learning has long-term effects on attentional performance (Itti & Koch, 2001; Parkhurst et al., 2002). Thus, it remains unclear as to whether figure enhancement should be considered an early visual process that supports scene segmentation or a byproduct of covert spatial attention.

In this study, we used high-field fMRI at 7 Tesla to determine whether orientation-defined figures are selectively enhanced in an automatic non-attentive manner at early stages of visual processing. We relied on population receptive field (pRF) mapping to characterize the response fields of individual voxels (Henderson, 2003; Tatler et al., 2011) in order to distinguish figure enhancement from local responses elicited by the boundary between figure and surround. Observers performed an engaging visual task at central fixation while task-irrelevant figure/ground stimuli were presented. The experiment consisted of 3 main conditions designed to isolate the effects of figure enhancement and boundary detection. The orientation of the figure could appear either orthogonal to that of the surround (Figure 3.1B, incongruent orientation) or iso-oriented with the surround but sampled from a distinct patch of spatially filtered oriented noise, such than a phase-defined boundary is visible (Figure 3.1B, congruent orientation). In the latter condition, spatially misaligned phase information at the boundary should cue early-stage border detection mechanisms; however, the figure itself does not appear to be especially prominent. In the third, no figure condition, the entire display was filled with a single oriented texture.

We predicted that both orientation- and phase-defined boundaries between figure and surround (Figure 3.1B middle, left) should yield spatially local enhancement of BOLD responses. Of greater interest, we predicted that the entire figure region should be enhanced when the orientation of the figure differs from that of the surround, as the central figure appears much more salient in this condition. As a consequence, a comparison of V1 responses for the incongruent orientation condition minus the congruent orientation condition should reveal evidence of a distinct spatial profile consistent with figure enhancement, if indeed such processes take place in an automatic manner at early stages of visual processing.

3.2 Results

Distinct effects of boundary detection and figure enhancement in V1

Figure 3.1C shows V1 BOLD responses in each of the experimental conditions binned by the eccentricity of the estimated pRF center for each voxel. fMRI responses were generally stronger for the congruent-orientation condition than for the no-figure control, especially around 2° eccentricity, which corresponded to the radial location of the boundary. However, fMRI responses were greatest in the incongruent orientation condition, and this appeared to be true even for voxels with pRFs near the center of the figure. In the far surround (~4° eccentricity), the responses across all three conditions appeared much more comparable.

We tested for the predicted consequences of boundary detection by calculating responses to the congruent figure minus the no-figure control condition. This revealed enhanced V1 responses centered at the boundary between the figure and surround (Figure 3.1D). A modest degree of skewness towards farther eccentricities was also observed, consistent with the fact that

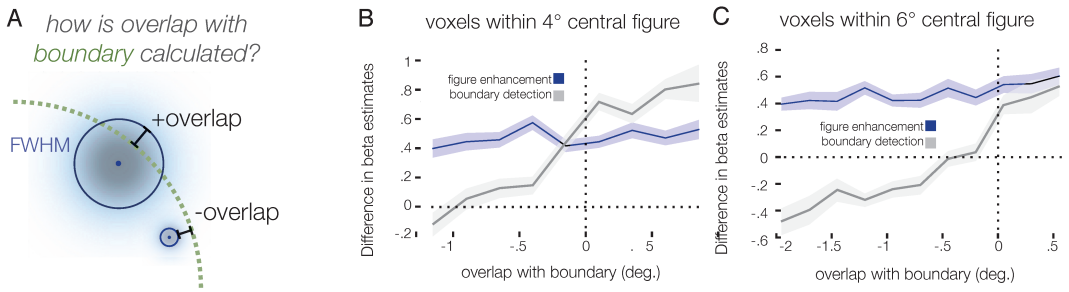

Figure 3.3. Boundary detection and figure enhancement plotted as a function of pRF extension beyond boundary. As illustrated in (A), two pRFs with centers at the same eccentricity within the figure might lead to differential boundary responses. We calculated whether the central FWHM-region of each pRF fell short of the boundary (negative 'overlap' values) or extended past the boundary (positive 'overlap' values). (B) V1 boundary detection and figure enhancement as a function of distance; while boundary detection (grey) is evident primarily in voxels whose pRFs overlap with the center/surround boundary (overlap > 0), figure enhancement (blue) occurs along the full extent of measured distances. Bins are 0.25° wide, and those containing fewer than 10 data points were trimmed. (C) Results from a control experiment using 6°-diameter figures. While boundary detection again primarily occurs in voxels whose pRFs overlap with the boundary (overlap < 0), figure enhancement persists even in voxels with up to 2° spatial separation with the center/surround boundary.

pRF sizes generally increase with eccentricity (Figure 3.2). These effects of boundary detection appeared distinct from the predicted effects of figure enhancement.

We calculated the spatial profile of responses for the incongruent orientation condition minus the congruent condition; this revealed enhanced activity that appeared to extend through the full 2° radius of the figure (Figure 3.1E). However, as some variability in pRF size is present at each eccentricity (Figure 3.2), one must ask whether these effects are primarily driven by voxels that, while centered within the figure, nevertheless receive input from the location of the boundary. To address this question, we performed another analysis that took into consideration both the center position and spatial extent of voxel pRFs. Specifically, we evaluated whether figure enhancement might be driven by voxels whose pRFs are centered within the figure but also extend over the boundary, in which case local processing of orientation differences might underlie the enhancement. We binned the response of voxels according to whether the most sensitive central region of their pRF, defined by the full-width-half-maximum (FHWM) envelope, was fully contained within the figure (negative values of 'overlap' with the boundary) or extended past the boundary (positive values of 'overlap'), as illustrated in Figure 3.3A. While the boundary detection effect (grey) occurs primarily in voxels whose pRFs overlap with the boundary, the magnitude of figure enhancement (blue) remains stable and positive for voxels whose response region (FWHM) falls within the figure and does not overlap with the boundary ($F(1,3) = .50$, $p = .53$).

We observed this same pattern of results in a separate control experiment using larger 6°-diameter figures (Supplementary Information). Figure enhancement persisted in the absence of directed attention even in voxels whose pRF response fields (FWHM) were more than 2° away from the center/surround boundary (Figure 3.3; $F(1,6) = 1.66$, $p = .25$). Together, these results imply that figure enhancement does not arise from local processing of feature contrast, as such effects would be expected to decline as a function of distance.

Finally, we used pRF measurements to identify voxels within V1 whose FWHM envelopes were fully contained within the figure or the surround (Figure 3.4A), as well as those whose central envelope overlapped with the boundary. Mean BOLD time courses and beta estimates for these regions of interest (ROIs) are shown in Figure 3.4B. These distinct ROIs reveal the predicted pattern of results: the figure ROI shows a significant effect of figure enhancement ($t(5) = 7.86$, $p = 5.4 \times 10^{-4}$) but not of boundary detection ($t(5) = 0.88$, $p = 0.42$); the boundary ROI exhibits both figure enhancement ($t(5) = 5.46$, $p = 2.8 \times 10^{-3}$) and boundary detection ($t(5) = 7.7$, $p = 7.4 \times 10^{-4}$), while the surround ROI does not yield significant differences between the experimental conditions. Results in areas V2 and V3 follow this same pattern (Figure 3.5). In sum, the presence of an

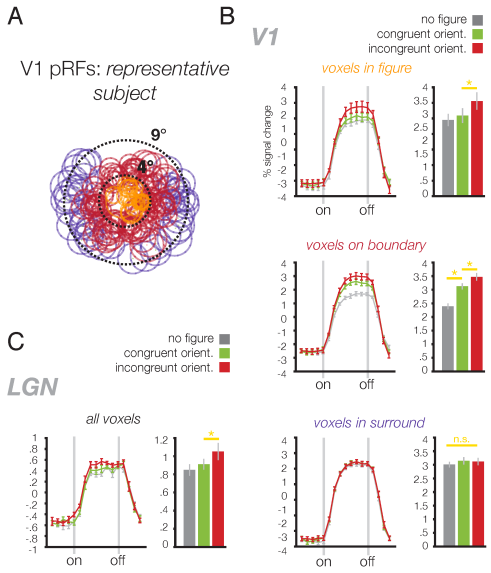

Figure 3.4. Average BOLD responses across figure, boundary, and surround-selective voxels. (A) V1 pRF locations of one representative subject (circles denote FWHM), illustrating how pRFs were sorted to create ROIs with voxels primarily responsive to the figure (orange, 205 voxels across participants), surround (purple, 293 voxels), or on the boundary (red, 476 voxels). Dotted lines denote the spatial extent of our mapping stimulus (9° diameter) and the figure region (4° diameter). (B) V1 results for these ROIs; left panels depict average BOLD time courses in each condition, with dotted lines marking the onset and offset of 16s stimulus blocks. Right panels depict averaged estimated GLM beta weights in each condition. Error bars indicate ±1 SEM between subjects. (C) Results from voxels in the LGN that overlapped with the center, surround, and boundary.

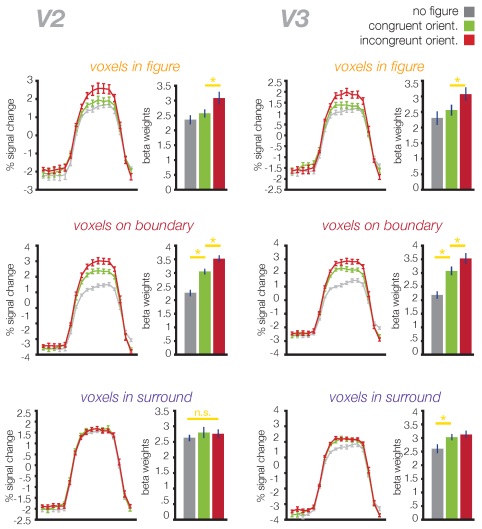

Figure 3.5. Average BOLD responses across figure, boundary, and surround-selective voxels in regions V2 and V3, plotted following the conventions of Figure 3.4.

orientation difference between the center and surround leads to enhancement of the entire figure region, whereas the presence of an iso-oriented phase-defined boundary reveals local responses to the border but no such global enhancement.

Figure-ground modulation in the LGN

While our primary aim was to resolve the spatial extent of figure-ground modulation in early visual cortex, we also measured mean BOLD activity in the lateral geniculate nucleus (LGN), which has recently been shown to exhibit enhancement of motion-defined figures in single neuron responses (Blakemore & Tobin, 1972; Cavanaugh et al., 2002; Nelson & Frost, 1978). Although the receptive fields of individual LGN neurons are small, each fMRI voxel covers a proportionally greater segment of this small subcortical structure than a similar voxel sampled from cortex. As such, the pRFs that we measured in the LGN were too large to clearly distinguish responses to the central figure region and those to the surround (eccentricity 2-4.5°); 96.2% of recorded voxels overlapped with the figure/surround boundary. However, we do see evidence of figure-ground modulation (incongruent orientation > congruent orientation, $t(5) = 2.89$, $p = 0.034$) across the full set of LGN voxels as shown in Figure 3.4C; there is additionally a non-significant trend of boundary detection ($t(5) = 2.27$, $p = 0.072$). The results suggest that activity in the visual thalamus reflects the segmentation and enhancement of figural regions, presumably via top-down feedback from visual cortex (Allman et al., 1985; Blakemore & Tobin, 1972; Chao-Yi & Wu, 1994; DeAngelis et al., 1994; Jones et al., 2001; Zipser et al., 1996).

Reconstructing responses in stimulus space

pRF modeling also provides an opportunity to visualize the pattern of voxel responses for a given visual area in stimulus space (cf. Adesnik et al., 2012; Cavanaugh, 2002; Gilbert & Wiesel, 1979; Lamme, Super, & Spekreijse, 1998a). Here, we used regularized linear regression to infer the effect in stimulus space (Figure 3.6) most likely to have yielded the pattern of response amplitudes among voxels given their pRFs. We used an L2 (ridge) penalty term to minimize overfitting, and ten-fold cross-validation to select the penalty that produced a

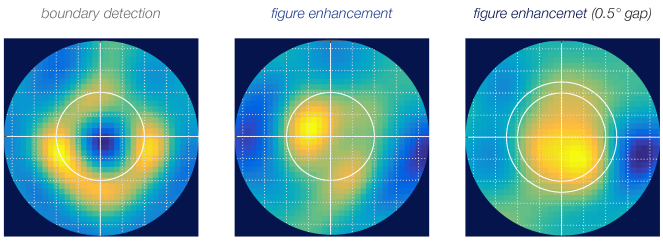

Figure 3.6. pRF-based reconstruction of voxelwise BOLD responses in stimulus space, using a multivariate regression-based method of estimation. Panels indicate the reconstructed effects of boundary detection (left) and figure enhancement for V1 data in the main experiment (middle) and in a control experiment which introduced a 0.5° gap between figure and surround (right).

minimal RMSE. This regression method can yield better spatial resolution in the reconstruction than simply scaling each voxel's pRF by its response and averaging across the visual field, as the latter approach is sure to introduce some blurring given the spatial spread of the Gaussian pRFs themselves. We performed this analysis on the differential patterns of V1 activity associated with boundary detection and figure enhancement for each subject, normalized the resulting images, and created averaged reconstructions across subjects as shown in Figure 3.6. The reconstruction results clearly support the distinct effects of boundary detection and figure enhancement, and go on to show that the inferred effect of figure enhancement is well described by an enhanced representation specific to the figural region.

Obscuring the local boundary does not eliminate figure enhancement in V1

We have demonstrated that orientation-defined figures produce enhanced BOLD responses across the figure region in early visual areas, and that this enhancement is spatially distinct from the local detection of a phase-defined boundary. We ran an additional control experiment to ascertain whether this enhancement depends on the figure directly abutting the surround, as such proximity may be needed to enhance competitive interactions between the neural populations representing the two stimuli. This experiment relied on the same stimulus configuration and task, but introduced a 0.5° gap (2.0-2.5° eccentricity) between the figure and the surround. The reconstruction of V1 activity again shows that figure enhancement (incongruent > congruent figure) occurred throughout the extent of the center region (Figure 3.6, right). Thus, we can conclude that figure enhancement cannot be explained simply in terms of an inward spread of a local boundary detection process. Orientation-defined figures lead to highly consistent effects in V1.

3.3 Discussion

Our study provides compelling evidence that distinct mechanisms support boundary detection and figure enhancement in the human primary visual cortex, and critically, that figural enhancement is not simply a byproduct of covert spatial attention. Unlike boundary detection, which serves to enhance local feature differences, figure enhancement performs a more integrative function of grouping regions that *share* one or more features that distinguish it from the surround. Figure enhancement is thought to be instantiated by feedback from higher-order visual areas (Nurminen & Angelucci, 2014). In support of this dual-stage segmentation process, human performance at discriminating figures is best matched by models that instantiate both a border detection process and a subsequent region-grouping or filling-in process (Vinje & Gallant, 2000). Our findings are consistent with the idea that feedback-driven figure enhancement is an automatic process of the early visual system (Coen-Cagli et al., 2015), rather than a consequence of directed spatial attention.

In the current study, we were able to reconstruct the spatial profile of figure-ground modulation in the human visual system, using a computational pRF-based approach that estimates the spatial response fields of individual voxels. This method allowed us to directly quantify and

visualize what aspects of figure-ground modulation in the human visual cortex can be attributed to border detection mechanisms and what aspects can be attributed to figure enhancement. Although some fMRI work has suggested that automatic effects of figure-ground processing are evident in V1 (Shushruth et al., 2013), null effects of figure enhancement have been more commonly reported for human V1 while positive effects have been consistently observed in higher visual areas (e.g., Nassi, Gómez-Laberge, Kreiman, & Born, 2014). Moreover, previous studies were not able to characterize the spatial profile of these distinct mechanisms, as we were able to carry out here using high-field imaging, pRF modeling, and computational methods for reconstruction.

We demonstrate not only that figure enhancement occurs in V1 (and likely LGN) in the absence of directed attention, but also that this enhancement is spatially distinct from local boundary detection. While a phase-defined boundary was present in the congruent condition, it did not lead to figure enhancement, consistent with the weaker perceptual salience of the figure region in this condition. It was also the case that figure enhancement occurred when the local orientation-defined boundary was obscured by introducing a 0.5° gap between the figure and surround (Figure 3.6C). Our results suggest that the presence of a local feature-defined boundary is neither necessary nor sufficient to observe V1 enhancement to figures with distinct orientation content.

This work builds upon several recent findings of V1 contributions to mechanisms of visual segmentation, including salience (Cannon & Fullenkamp, 1991; Petrov & McKee, 2006), visual grouping (Joo et al., 2012; Schallmo et al., 2016), perceptual filling-in (Bogler et al., 2013), and recent applications of pRF mapping to investigate these perceptual phenomena (Kok & de Lange, 2014). Together, these studies point to an important role of the early visual system in figure perception through a combination of both local mechanisms and automatic feedback mechanisms.

This study also highlights a useful and intuitive method for utilizing population receptive field modeling to reconstruct visual responses (Zhang et al., 2012), which can move toward bridging the gap between the resolution of spatial effects measured by neurophysiological recordings and human neuroimaging. The regression method allows for a finer spatial resolution of reconstruction than a simple summation of pRF responses weighted by their response amplitudes (Figure 3.7). The latter approach is necessarily constrained by the resolution of pRFs themselves – even perfectly noise-free data will yield blurry reconstructions if the Gaussian pRFs are large, relative to the stimulus, and greater effects of pRF blurring will occur when ascending the visual hierarchy. This regression-based reconstruction approach could be adopted more widely to estimate the spatial profile of fMRI BOLD effects at spatial scales that would prove challenging with standard approaches.

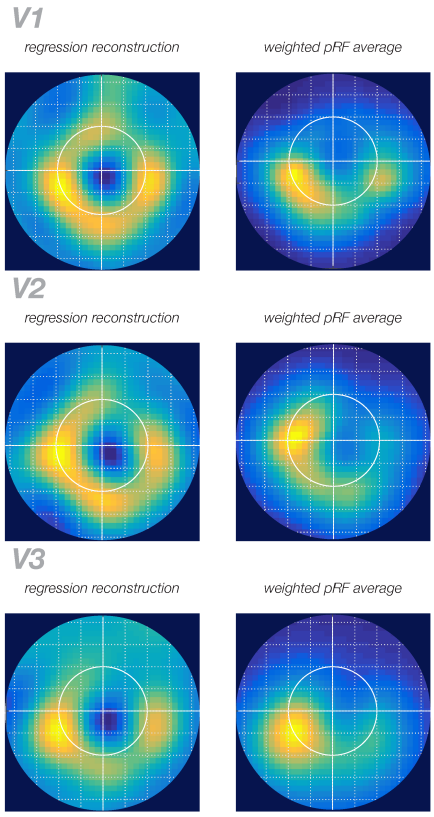

Figure 3.7. A comparison of the visualization of boundary detection using our pRF-based reconstruction method (left) and a standard weighting of pRFs as used by Kok et al. (2015) and others. The regression reconstruction yields clearer reconstructions, particularly as one ascends the visual hierarchy and pRF sizes become larger.

3.4 Supplemental Experimental Procedures

Scanning procedure

Six participants (four women) ages 23-31 were scanned in the main experiment; three of these participants (all women) also completed the control experiment, in which figures were 6° in diameter. All experiments were performed on the high-field 7-Tesla Philips Achieva MRI scanner at the Vanderbilt University Institute for Imaging Science, adhering to the guidelines of the Vanderbilt IRB. Each MRI session lasted 2 hours, during which we acquired the following images: i) 1-2 functional localizer runs using a flickering checkerboard to identify retinotopic regions in visual cortex and LGN that corresponded to the stimulus location ii) 7-8 fMRI runs to measure BOLD activity during the experiment, and iii) 5-8 fMRI runs to map population receptive fields in these voxels. Each of the run types lasted 4-6 minutes. BOLD activity was measured using single-shot, gradient-echo echoplanar T2*-weighted imaging, at a 2 mm isotropic voxel resolution (40 slices, TR 2000 ms, TE 35 ms; flip angle 63°; FOV 224 x 224; SENSE acceleration factor of 2.9).

fMRI preprocessing

Data were preprocessed using FSL and Freesurfer tools (documented and freely available for download at http://surfer.nmr.mgh.harvard.edu), beginning with 3D motion correction and linear trend removal, followed by slice-timing correction for pRF runs and a high-pass filter cutoff of 60s. Functional images were registered to a reconstructed anatomical space for each subject; this registration was first automated in FSL and then checked and corrected by hand. This allowed the alignment of the current fMRI data to the retinotopy data, which was collected in a separate session. The functional localizer was spatially smoothed using a 1-mm Gaussian kernel; no spatial smoothing was done for the experimental or pRF mapping runs. Further analyses were conducted using a custom Matlab processing stream. Each voxel's intensities were normalized by the mean of the time series, converting to mean percent signal change within each run. Outliers were defined as time points for which the voxel's response measured more than 3 times its standard deviation from its mean, and were Winsorised (Hastings et al., 1947). This condition-blind preprocessing step minimizes the impact of rare spikes in MR intensity while preserving the temporal structure of the responses in each voxel.

Population receptive field mapping

Population receptive fields (pRFs) reflect the location in space that best drives activity in the population of neurons in each voxel (Dumoulin & Wandell, 2008; Wandell & Winawer, 2015). Unlike standard retinotopy, pRF mapping resolves not only the location that best drives responses in each voxel, but also the spatial extent of this response field. In each 7T fMRI experimental session, we mapped population receptive fields in retinotopic areas V1-V3, using a 2D circular Gaussian model of pRF structure. In accord with prior work (Dumoulin & Wandell, 2008; Wandell & Winawer, 2015), population receptive field size in these regions increased with eccentricity within visual regions, and increased along the visual hierarchy from V1 to V3 (Figure 3.2). PRF modeling involves presenting a mapping stimulus that spans the visual field across time, and estimating the pRF parameters that most likely produced the measured BOLD response in a particular voxel. These parameters define the pRF's location, size, and gain. pRF properties for each voxel are assumed to reflect the combined RFs of the neural population in a voxel, and appear well-aligned with single-neuron receptive field properties, such as contralateral preference, increasing size at greater eccentricities, and increasing size as one ascends the visual hierarchy (Dumoulin & Wandell, 2008; Wandell & Winawer, 2015). Here, we mapped pRFs using a traveling-bar stimulus (developed by Kendrick Kay; described in Kay, Winawer, Mezer, & Wandell, 2013) that swept through a circular region with 4.5˚ radius, the maximal visible field of view at our 7T scanner; each mapping run lasted 5 minutes. Voxel-wise responses for each visual area were fitted with a 2D Gaussian pRF model using a custom Matlab pipeline. We used a dual-stage multidimensional

nonlinear minimization (Nelder-Mead) fitting procedure: each voxel was initially fitted in a downsampled stimulus space with a fixed Gaussian σ, and then these parameters were used to initialize a full model fitting in native stimulus space. Estimated parameters described each voxel's pRF position (X, Y), size (σ), and response amplitude; we can additionally convert these parameters to measures of polar angle, eccentricity, or full-width half-maximum (FWHM) to convey pRF size.

ROI localization and voxel selection

To define retinotopic visual areas V1-V3, each subject participated in a separate retinotopic mapping scan. We used a typical phase-encoded design (Engel et al., 1997; Wandell et al., 2007) in which subjects fixated while they viewed flickering checkerboards consisting of rotating wedges to map polar angle and expanding rings to map eccentricity. Retinotopy data was acquired using a Philips 3T Intera Achieva MRI scanner equipped with an 8-channel coil. Subjects were scanned using 3 mm isotropic resolution voxels (TR 2 s, TE 35 ms, flip angle 80˚, 28 slices, 192 x 192 FOV). Boundaries between retinotopic areas V1-V3 were delineated by hand, by identifying reversals in the phase of the polar angle map measurements; the resulting ROIs were aligned to the functional space of the current experiment using FSL and Freesurfer software. Additionally, 2-3 runs of functional localizers were collected in the main experimental sessions to identify the LGN in each subject. The localizer consisted of blocks of a flickering checkerboard stimulus spanning the full 9˚ field of view, and was designed to yield a large ROI that could be then refined by pRF model fitting.

For all analyses, we used these functional labels in conjunction with the pRF fitting results to define regions of interest. For each subject, all voxels in each visual area were fitted with the pRF model as described above. For further analyses, we used voxels whose pRF centers were within the range of the mapping stimulus (0.25°-4.5° eccentricity) and were larger than 0.1°; this limit trimmed instances in which the model predicted nearly no visual response to the mapping stimulus. Following this trimming procedure, we selected the top 33% of best-fitted voxels for each subject in each ROI, as indexed by the R^2 between observed and predicted data. In V1, this yielded fits with R^2 cutoffs that ranged from 0.62 to 0.81 in individual subjects (mean = 0.71); corresponding V1 ROIs for each subject ranged from 131 to 187 voxels bilaterally (mean = 162). In the LGN, R^2 cutoffs ranged from 0.16 to 0.27 (mean = 0.22), yielding ROIs that were 17-36 voxels in size (mean = 26.5).

Experimental displays and paradigm

Visual displays were dynamically generated and updated every 200ms during each 16s stimulus block, while observers performed a color-change detection task at central fixation. Sample displays for each of the three conditions in this experiment are illustrated in Figure 1. To minimize the prevalence of high spatial frequency energy artifacts between center and surround, we used oriented bandpass-filtered noise (100% contrast, bandpass filtered from 0.5-8 cycles per degree with a center orientation of 45° or 135°, ~20° FWHM in the orientation domain). Each orientation stimulus was dynamically generated from random white noise, bandpass filtered in the Fourier domain, convolved with a small Gaussian to minimize Gibbs ringing artifacts, and then converted back to the image domain. For the congruent-orientation figure condition, figure and surround images were generated from two different noise patterns to create a phase-misaligned figure-surround display.

In both the main experiment and control, participants passively viewed the figure-surround displays while they performed a color-change detection task at fixation. The fixation changed from black to red for increments of 200ms at random time intervals, occurring on average of 4 times per

16s experimental block. Participants reported these events by pressing a key on an MR-compatible button box; average percent correct across participants was 94.5% (SD = 5.8%).

We relied on an fMRI block paradigm, and presented 8 blocks of visual stimulation (16s per block) interleaved among 16-s fixation rest periods. The amplitude of the BOLD response during each stimulus block was then estimated using the standard general linear model for each voxel. These estimates were averaged across blocks to yield voxelwise mean betas in each experimental condition.

Reconstruction in stimulus space

pRF mapping was used to reconstruct the pattern of activity evoked by the figure-ground stimuli. Previous studies (e.g., Kok & de Lange, 2014) have adopted an approach of projecting the weighted activity of each voxel into image space by scaling its Gaussian pRF by the voxel's response to the display, and then calculating the linearly summed response of all voxels. In our work, we have found that a multivariate regression-based approach leads to sharper reconstructions. Specifically, ridge regression was used to estimate the contrast strength of every pixel in the display, which was then convolved with the population receptive field of every voxel to predict the measured pattern of voxel responses. The number of predictors was reduced by downsampling the visual space to 4 pixels per degree; equivalently, each predictor/pixel corresponded to 0.25° of the original image. This yielded an estimate of predictors (β) in pixel space that are most likely to have produced the measured BOLD responses across voxels, given the voxels' known pRFs. This can be described as a standard linear model: $y = X\beta + \varepsilon$, where $y \in \mathbb{R}^N$ represents the vector of voxel responses, $X \in \mathbb{R}^{N \times p}$ consists of a 2D matric of each voxel's pre-measured Gaussian pRF map, $\beta \in \mathbb{R}^p$ are the predictors to be estimated, ε is the error term, and p and N are pixels and voxels, respectively. Since there are many more predictors than voxel responses to be predicted, this leads to an inverse mapping problem, for which standard multivariate regression cannot find a unique solution. To constrain the estimation procedure and to minimize overfitting, ridge regression can be used to estimate the predictors for each pixel's contrast. Whereas linear regression seeks to minimize the sum of squared residuals, ridge regression applies an additional penalty term based on the sum of squared weights (β) to be estimated, thereby giving preference to solutions with smaller β values (Hoerl & Kennard, 1970), as follows: $\min_\beta \frac{1}{2N}\sum_{i=1}^{N}\left(y_i - x_i^T\beta\right)^2 + \lambda\|\beta\|_2^2$, where λ is the scaling factor of the penalty term. This is often referred to as L2 regularization. Ridge regression was performed for each individual subject to produce reconstruction images. These images were normalized within subjects, averaged, and smoothed with a small Gaussian kernel (σ = 0.75°) to create the reconstructions in Figure 3.6.

Supplementary results: V2 and V3

pRF sizes were somewhat larger in V2 and V3, as has been previously reported (Dumoulin & Wandell, 2008; Wandell & Winawer, 2015). Nevertheless, the pattern of results in V2 and V3 closely echoed that of V1. This is evident in Figure 3.5, which illustrates BOLD responses in each of the experimental conditions binned by voxels' pRF center eccentricity; larger peripheral pRFs in these regions are more likely to encroach on the 2° figure/surround boundary, and thus boundary detection effects are evident in a larger proportion of voxels than in V1. We averaged responses across voxels whose pRFs fell primarily in the figure, in the surround, or included the boundary, assuming a FWHM central region. Voxels with pRFs confined well within the figure showed significant effects of figure enhancement (V2: $t(5) = 5.37$, $p = 0.0030$; V3: $t(5) = 4.81$, $p = 0.0048$) but not boundary detection (V2: $t(5) = 1.38$, $p = 0.23$; V3: $t(5) = 1.12$, $p = 0.31$). Voxels whose pRFs fell on the boundary exhibited both effects (V2: t's(5) > 5.16, p's < 0.0036; V3: t's(5) > 4.27, p's < 0.0079). In V2, voxels responding only to the surround did not show evidence of either type

of enhancement (V2: t's(5) < 1.64, *p's* > 0.16). In V3, voxels responding to the surround also showed no figure enhancement (t(5) = 0.46, p = 0.80), though they did show a significant effect of boundary detection (t(5) = 2.91, p = .034). Of note, however, was that pRFs in this peripheral region of V3 were quite large: only 96 voxels across six participants were included in the surround-only subset of voxels.

Control experiment: larger figures

We performed a control experiment to evaluate whether larger figures would still lead to figure enhancement near the center of the figure, in a manner that could not be readily explained by local effects of surround suppression or boundary detection. This experiment followed the same procedures as the main study, but used figures that were 6° in diameter. Three subjects from the original experiment participated in this version of the study. In ROI-based analyses (similar to those in Figure 3.4), we observed a very similar pattern of results in V1: voxels responding to only the figure region exhibited figure enhancement (t(2) = 3.26, p = 0.074) but not boundary detection (t(2) = 1.05, p = 0.40), while voxels whose pRFs were situated on the boundary exhibited both effects (figure: t(2) = 6.72, p = 0.021; boundary: t(2) = 8.21, p = 0.015); across these 3 subjects, 23 voxels had pRFs responding only to the surround, and exhibited some effect of boundary detection (t(2) = 5.06, p = 0.036) but not figure enhancement (t(2) = 2.15, p = 0.16). As in the main experiment, nearly all LGN voxels overlapped with the boundary, and we again saw a significant effect of figure enhancement (t(2) = 8.19, p = 0.015).

Acknowledgements

This work was recently submitted for peer review in co-authorship with Frank Tong.

Chapter 4

Automatic cortical feedback mediates figure-ground modulation in the human lateral geniculate nucleus

4.1 Introduction

We perceive the visual world not as a single sheet of information, but as a parsed landscape of meaningful figures, objects, and surfaces. A core function of the visual system is to thus segregate two-dimensional pattern of light on the retina into figures and background regions. A difference in visual features, like contrast, color, spatial frequency and orientation, between the figure and its surroundings provides a powerful cue for segmentation. This property of vision is exploited by animals whose coats or skins match their surroundings, making them more difficult for predators to detect (e.g. Figure 3.1).

These dual mechanisms of *boundary detection* and *figure enhancement* have been characterized in the primary visual cortex (V1) of macaque monkeys (Lamme, 1995; Zipser et al., 1996) and humans (Chapter 3). When a feature-defined figure (e.g. a patch of horizontal oriented lines in a field of vertical oriented lines) is viewed, neurons responding to the edge of the figure will show a heightened response; this is thought to occur on the basis of local interactions between feature-tuned neurons in the early visual cortex. Subsequently, the figure region itself is enhanced, such that a neuron whose receptive field falls within the figure region will fire more than those responding to a matched region in the surround. This result has been recently confirmed in human V1 using fMRI population receptive field (pRF) mapping of individual voxels (Chapter 3). Unlike boundary enhancement, which detects local feature differences, figure enhancement acts to group regions that share one or more features. It is thought to be instantiated by feedback from higher-order visual areas, as indexed by the temporal profile (Poort et al., 2012) and layer-specificity of figure enhancement, relative to boundary detection (Self, van Kerkoerle, Supèr, & Roelfsema, 2013). In support of this dual-stage segmentation process, human performance at discriminating figures is best matched by models that instantiate both border detection and a subsequent filling-in process (Mumford et al., 1987; Wolfson & Landy, 1998).

Recently, a intriguing study by Jones and colleagues (Jones et al., 2015) found elevated responses to motion direction-defined figures in the lateral geniculate nucleus (LGN) of the thalamus in alert non-human primates. This work suggested that the heightened salience of figures may be fed back to the earliest anatomical processing stage possible in the visual system – an intriguing account of how higher-order processing may recurrently shape visual input toward the seamless and rapid segmentation of meaningful objects. Jones et al. (2015) rule out a pre-motor explanation of the observed neural enhancement by recording from some neurons with foveal RFs. In this subset of trials, the monkey did not have to make an eye movement to the figure and figure enhancement was still observed. However, this experimental paradigm does not allow the dissociation of figure enhancement from spatial attention. Directed spatial attention consistently modulates responses in the LGN (Ling et al., 2015; McAlonan et al., 2008; Poltoratski, Ling, McCormack, & Tong, 2017; Schneider & Kastner, 2009). Spatial attention has at times been proposed to underlie all (Lamme, Zipser, & Spekreijse, 1998b; Rossi et al., 2001) or most (Poort et al., 2012) of the observed figure-ground modulation in primate V1, although recent work in our lab (Chapter 3) found robust figure enhancement in human V1 in the absence of attention. This discrepancy between primate and human results may be the product of the extensive reward-based training necessary to perform visual experiments in non-human primates. This type of training has been shown to yield long-term effects in attentional performance (Anderson & Laurent,

2011; Anderson & Yantis, 2013), and makes it difficult to distinguish whether figure-ground modulation in the LGN is primarily the consequence of directed attention, or whether it occurs automatically.

The current study resolves these issues by measuring fMRI BOLD responses in humans, whose attention can be manipulated flexibly and easily, without the need for reward-based training. We used high-resolution fMRI at 7 Tesla to simultaneously measure activity in the visual cortex and in the LGN while presenting participants with two orientation-defined figures. In the first experiment, illustrated in Figure 4.1A, we manipulated the locus of spatial attention by asking participants to fixate on the center of the screen but to perform an engaging detection task on one of two lateralized figures, such that on each trial, one figure was attended and the other unattended. Doing so allowed us to quantify the effect of spatial attention on figure-ground modulation in the LGN, and to answer whether attention accounted for the enhancement of figures recently observed in this structure (Jones et al., 2015). We then performed a second experiment to determine the degree to which cortical feedback mediates this effect, presenting the figure and surround stimuli either to the same eye or to different eyes on a given trial (Figure 4.2A). The design of this experiment leverages the known properties of binocular processing in the early visual system, as V1 is considered the first stage of the visual hierarchy in which information from the two eyes is combined and processed binocularly. Thus, if figure-ground modulation occurs in the LGN when the figure and surround are presented to two different eyes, corticothalamic feedback must presumably be responsible for the effect. To our knowledge, these experiments provide novel evidence of figure-ground modulation in the human LGN; additionally, they implicate an underlying mechanism of automatic feedback.

4.2 Materials and Methods

In both experiments, we presented human subjects (N = 8) with two lateralized, orientation-defined figures, as illustrated in Figure 4.1, in a 16s block design interspersed with 16s blank periods. Stimuli were generated from bandpass-filtered oriented noise centered around 45° and 135° (see Supplementary Materials and Methods). Center and surround orientation could either be incongruent, yielding the percept of a figure, or congruent, and figure-ground modulation was defined as the difference in response to these two conditions. The figures were 4° in diameter and located 3° from the fixation point (center to center). A narrow equiluminant (0.15° wide) grey gap separated the figures from the surround, minimizing differences in local orientation-contrast information at the center/surround boundary between the two orientation conditions.

In Experiment 1, both the figures and surround were presented at 100% contrast. Participants were cued to detect a spatial frequency change in one of the two figures while maintaining central fixation; the change occurred with equal frequency at random independent intervals in each figure, but participants were cued to respond to only one of the lateralized figures. In each block, one of the figure regions was incongruent in orientation with the surround and the other was congruent, so that overall salience was matched between blocks (Poltoratski et al., 2017).

In Experiment 2, participants wore red/green analglyph filter glasses for the duration of the scan session to allow monocular presentation of red and green stimuli. In each 16s stimulus block, the two figures could appear in the same color (and thus, eye) as the surround, or in the other eye. In each block, one figure was congruent in orientation with the surround and the other incongruent; this allowed us to cleanly compare the effects of orientation in each eye condition. Participants performed an engaging detection task on the fixation point throughout the experimental run.

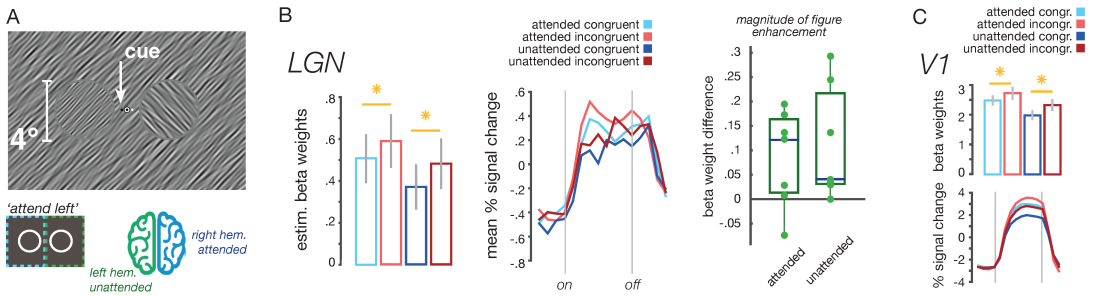

Figure 4.1. Design and results of Experiment 1, in which the locus of spatial attention was manipulated. Figure enhancement persists in both LGN and V1 even when the figure is unattended. A) Sample display from Experiment 1, in which a small cue directed participants to detect a brief spatial frequency increment (as visible in the left circular figure) on one of two orientation-defined figure. As illustrated in the bottom panel, this yielded concurrent measurements of responses to an attended and unattended figure in the two hemispheres. B) Estimated beta weights (left) across the experimental conditions in the LGN. Error bars depict ± SEM across participants (N = 7), and asterisks denote significant t values at a threshold of $p < .05$. Mean percent signal change time courses are also presented, averaged relative to the 16s on-stimulus blocks for each condition (grey vertical lines). C) The magnitude of figure enhancement (incongruent minus congruent orientation) in the LGN in both the attended and unattended condition. Circles indicate individual participant effect sizes. C) Results for early visual cortical area V1, which echo the pattern of results observed in the LGN.

We recorded neural activity using fMRI at 7 Tesla with 2mm isotropic voxel resolution (see Supplemental Materials and Methods). Voxels in both the LGN and early visual cortex responding to the position of the two figure regions were localized with a flickering checkerboard display; this data was used in conjunction with a separate session of retinotopy mapping (Engel et al., 1997; Wandell et al., 2007) to identify cortical regions of interest in V1-hV4. The LGN was identified on the basis of the functional localizer as a contiguous cluster of voxels in the medial subcortex using a t-statistic threshold of no less than 2.8; individual LGN regions of interest (ROIs) were selected to be maximally lateralized to avoid including other regions of the thalamus. The LGN is more readily activated by visual stimulation than other subcortical regions, and there is evidence to suggest that functional localizers that rely on passive viewing, as ours did, do not activate pulvinar activity as strongly as the LGN (Kastner et al., 2004). All aspects of this study followed the guidelines of the Vanderbilt University Institutional Review Board.

4.3 Results

Directed attention is not necessary for figure-ground modulation in the LGN

If figure enhancement occurs in the human LGN, it is critical to determine whether it is primarily the result of directed spatial attention, which has been shown to robustly modulate responses in this structure (Ling et al., 2015; O'Connor et al., 2002; Schneider & Kastner, 2009), or if it is the signature of an automatic visual process. In the first experiment, we manipulated the locus of participants' spatial attention such that on each trial, one lateralized figure was attended and one was unattended, as illustrated in Figure 4.1A. Because each hemisphere in the early visual system represents the contralateral visual hemifield, this allowed us to define experimental conditions relative to lateralized ROI responses such that, for example, a block in which the participant performed the task on the right figure was 'attended' in the left hemisphere and 'unattended' in the right hemisphere; data were then pooled according to condition across the left and right hemisphere ROIs. After standard pre-processing of the MR BOLD signal (see Supplemental Materials and Methods), neural responses were normalized to mean percent signal change, and responses to individual stimulus blocks were fitted with a general linear model (GLM). Estimated beta values were then sorted by experimental condition. In each bilateral region, we performed a 2 x 2 repeated measures ANOVA to characterize the effects of orientation

(incongruent or congruent relative to the orientation of the surround) and spatial attention, as well as their interaction.

In our design, figure enhancement should lead to an increased BOLD response in the incongruent-orientation condition relative to the congruent-orientation condition. Results from this experiment are shown in Figure 4.1. A main effect of figure enhancement was observed in the LGN ($F(1,6) = 21.2$, $p = .0037$), consistent with recent single- and multi-unit recordings in the macaque (Jones et al., 2015). We also observed heightened LGN responses for attended figure regions ($F(1,6) = 19.6$, $p = .0044$), as predicted by prior work in our lab and others (Ling et al., 2015; O'Connor et al., 2002; Schneider & Kastner, 2009). More importantly, figure enhancement was observed both when the observer attended to the figure (one-tailed $t(6) = 2.26$, $p = .032$) and when spatial attention was directed to the opposite hemifield ($t(6) = 2.52$, $p = .023$); the main effects of attention and orientation did not interact in the LGN ($F(1,6) = .15$, $p = .71$). This pattern of results was evident not only in the group average, but also in individual subjects (Figure 4.1C). It appears that figure-ground modulation in this subcortical structure is the result of automatic visual processing, rather than a consequence of directed spatial attention or training.

In V1, attention led to a similar enhancement of responses at the attended location ($F(1,6) = 48.5$, $p = 4.4 \times 10^{-4}$), and figure enhancement occurred both when the figure was attended ($t(6) = 2.73$, $p = .017$) and when it was not ($t(6) = 5.0$, $p = .0013$; main effect of orientation ($F(1,6) = 23.3$, $p = .0029$). Results in areas V2-hV4 closely followed this pattern (Figure 4.3), with one exception: V1 attention and orientation did not interact ($F(1,6) = .87$, $p = .39$), while in higher areas this interaction met or reached significance ($F's(1,6) = 8.9\text{-}5.7$, $p's = .024\text{-}.054$), with a trend of larger magnitude of figure enhancement when the stimulus was unattended.

Dichoptic presentation still yields figure enhancement in the LGN, implicating feedback

Inputs from the two eyes are thought to first converge in the early visual cortex; in the LGN, input from the contralateral eye is processed in layers 6, 4, and 1, while information from the ipsalateral eye is input to layers 2, 3, and 5. We leveraged this property of the visual system to deduce the role of cortical feedback in figure-ground modulation in LGN, by presenting the figure and surround stimuli either the same eye or to different eyes, as illustrated in Figure 4.2A. As predicted by Experiment 1, we found consistent effects of figure enhancement in the LGN when the figure and surround were presented to the same eye ($t(7) = 3.13$, $p = .0084$), and a main effect of orientation ($F(1,7) = 15.6$, $p = .0056$). Moreover, figure enhancement persisted even when the figure and surround were presented to two different eyes ($t(7) = 2.28$, $p = .028$). In the LGN, the effect of eye condition (that is, whether stimuli were presented monoptically or dichoptically) was not significant ($F(1,7) = 0.33$, $p = .58$), nor did it interact with the effect of orientation ($F(1,7) = 0.22$, $p = .65$). These findings support the hypothesis that feedback from visual

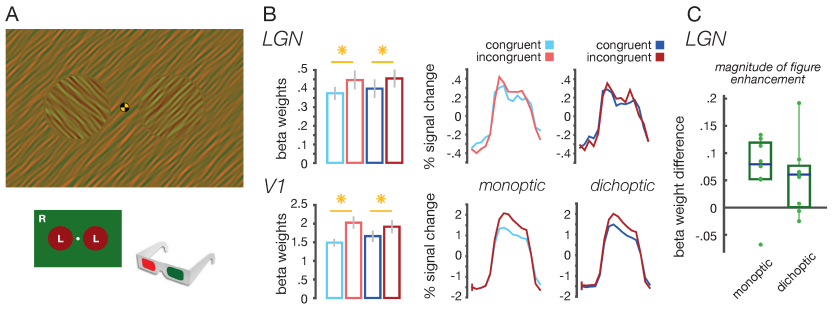

Figure 4.2. Design and results of Experiment 2, in which the figure and surround stimuli were presented either to the same eye (middle panel) or to different eyes (right panel). Plotting follows the conventions of Figure 2. A) Sample display from Experiment 2, designed to be viewed through red/green anaglyph glasses (the reader is encouraged to do so). As illustrated in the bottom panel, this example presents the two figure regions (one incongruent-oriented and one congruent-oriented to the surround) to the two different eyes. B) Estimated beta weights and time courses across conditions in Experiment 2. C) Magnitude of figure enhancement in the monoptic and dichoptic presentation conditions in the LGN. Figure enhancement was evident in the LGN even when the figure and the surround were presented dichoptically, despite the predominantly monocular processing of visual input by LGN neurons. This implicates a mechanism of corticothalamic feedback for figure-ground modulation in the LGN.

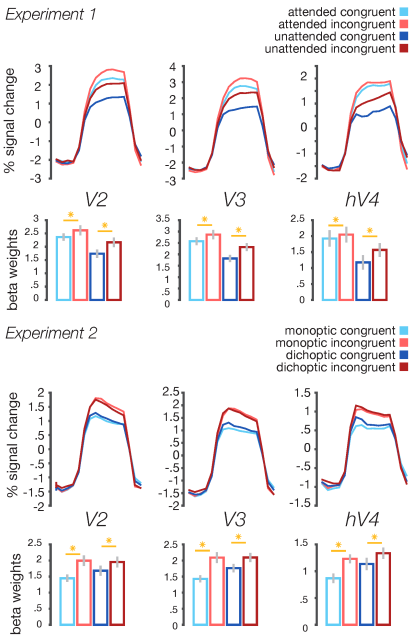

Figure 4.3. Results of Experiments 1 and 2 in regions V2-hV4, which follow those of V1 in both studies.

cortex to LGN mediates figure-ground modulation, in a manner that is independent of whether this percept requires combining visual information from the two eyes.

In early visual cortex, we found again strong effects of figure enhancement in V1 ($F(1,7) = 27.0$, $p = .0013$), both when the stimuli were presented monoptically ($t(7) = 5.68$, $p = .00038$) and dichoptically ($t(7) = 4.12$, $p = .0022$). As expected from the binocular integration that occurs in V1, eye condition and orientation interact in this region ($F(1,7) = 34.2$, $p = .0006$), with stronger effects of surround suppression observed under monocular presentation. This is consistent with the notion that binocular information undergoes additional processing in the cortex that does not occur in the LGN, where we did not find an effect of eye condition. Regions V2-hV4 each follow this pattern of results, as illustrated in Figure 4.3.

4.4 Discussion

In two experiments, we have demonstrated that figure-ground modulation in the human LGN occurs (1) in the absence of directed attention to the figure and (2) when the figure and surround are presented to different eyes, implying a mechanism of feedback. These results urge us to revisit the idea that the visual thalamus is a passive way-station, and instead consider its active and recurrent role in visual perception. Responses in this region reflect not only visual input, but also complex processing in higher-order regions, which likely act to build predictions about the environment and resulting visual input (Rao & Ballard, 1999).

The detection and segmentation of meaning figures in the environment is a core function of the visual system. It has been demonstrated that figures defined by one or more visual features, like orientation color, luminance, or binocular disparity (Lamme, 1995; Zipser et al., 1996), yield heightened responses in the early visual cortex through dual mechanisms of local boundary detection and a subsequent, feedback-driven figure enhancement (Self et al., 2013). While boundary detection relies on known mechanisms that detect local feature differences, like lateral inhibitory interactions (Blakemore, Carpenter, & Georgeson, 1970; Gilbert, 1992; Stettler et al., 2002), figure enhancement serves to group regions of visual input that share features. Here, we show that this enhancement is present in the LGN, a subcortical structure that lies between the retina and cortex, and is anatomically the earliest possible stage of visual processing than can be modulated by cortical feedback.

While the LGN is traditionally thought to have little feature selectivity, this and other recent work points to more complex processing of visual information in the LGN. Recent studies in rodents and monkeys have demonstrated some degree of orientation selectivity in the activity of LGN neurons (Cheong et al., 2013; Lien & Scanziani, 2013; Ling et al., 2015; W. Sun, Tan, Mensh, & Ji, 2016; Vidyasagar & Eysel, 2015). While a component of this selectivity arises from an elongation of retinal receptive fields (Leventhal & Schall, 1983; Smith et al., 1990), neuroimaging work in our lab (Ling et al., 2015) has revealed another component is independent of radial preference, which reflects the influence of directed attention likely instantiated by top-down cortical feedback. Orientation-tuned V1 neurons appear to send feedback signals to the LGN in a retinotopically organized manner that coincides with the spatial structure of their oriented receptive

field (Wang, Jones, Andolina, Salt, & Sillito, 2006a). In this respect, even if the individual LGN neurons are weakly tuned or untuned for orientation, the pattern of retinotopic modulation in the LGN via feedback can be informed by the nature of orientation processing in the early visual cortex (Briggs & Usrey, 2008). The current work finds that feedback can likewise carry second-order orientation information from the cortex to the LGN, as responses to the figure region are modulated by the relative orientation of the figure and the surround.

Importantly, this figure enhancement occurs even when spatial attention is directed away from the figure, contrary to accounts of figure-ground modulation as primarily a consequence of directed attention (Lamme, Zipser, & Spekreijse, 1998b; Poort et al., 2012; Rossi et al., 2001)Given the consistent modulation of LGN activity by both spatial (Ling et al., 2015; O'Connor et al., 2002; Poltoratski et al., 2017; Schneider & Kastner, 2009) and feature-based attention (Schneider, 2011), it is important to rule out an attention-driven explanation for the findings of figure-ground modulation in this structure (Jones et al., 2015). This question can be more readily tested in human subjects, who can perform our relatively complex attentional task without extensive training. Our results strongly suggest that figure-ground modulation is an automatic, pre-attentive function of the early visual system; while orientation-defined figures may be more visible to the observer, this bottom-up salience is separable from top-down attention in the early visual system (Poltoratski et al., 2017). In cases whether neither attention nor awareness are possible, such as following anesthesia, figure-ground modulation appears to be extinguished (Lamme, Zipser, & Spekreijse, 1998b).

There is growing evidence that vision is not a uni-directional, hierarchical process; instead, this work and others highlight the important functional role of feedback connectivity even in automatic visual processes. In addition to modulation by attention (Gandhi et al., 1999; O'Connor et al., 2002; Schneider & Kastner, 2009; Somers et al., 1999), responses at the earliest stages of the visual system are altered by complex shape and figure processing typically attributed to higher-order areas (Kok & de Lange, 2014; Lamme, Super, & Spekreijse, 1998a; Poort et al., 2012). In this way, even unattended stimuli are processed in a combination of 'top-down' and 'bottom-up' mechanisms, such that input at lower stages is constantly influenced by feedback from higher-order processing. Our demonstration of attention-independent figure modulation in the LGN is particularly striking, as the detection and segmentation of figures is often attributed to regions beyond the primary visual cortex, including V4 (Kastner, De Weerd, & Ungerleider, 2000; Kourtzi, Tolias, Altmann, Augath, & Logothetis, 2003) and object-selective regions in the lateral occipital cortex (Grill-Spector, Kushnir, Edelman, & Avidan, 1999; Vinberg & Grill-Spector, 2008). Here, we see that this figure-ground modulation alters responses in the LGN, which is anatomically the earliest stage of visual processing that can be modified by feedback, and which gates the majority of visual input between the eyes and the brain. Importantly, neurons in the LGN have at most modest orientation selectivity, but via this recurrent processing are able to represent complex context-dependent feature information. This work urges us to consider the active role of feedback at all stages of the visual system, and how this feedback can mitigate core aspects of perception across multiple stages of the visual hierarchy.

Acknowledgements

We hope to imminently submit this work for peer review; Devin McCormack, Allen Newton, Alexander Maier, and Frank Tong have contributed to this project.

Supplemental Materials and Methods

Participants

Eight experienced subjects (five females) ages 23-31participated in the first experiment, and eight (six females, 24-31) participated in the second; all received monetary compensation. One participant's data was excluded from Experiment 1 after an imaging artifact located over his occipital pole was evident across functional images. The remaining seven subjects each also participated in Experiment 2, including author SP.

Stimuli & Design

In both studies, stimulus displays were created using randomly generated bandpass-filtered oriented noise (0.5-4cpd, 100% contrast, 45° and 135° with a filter width of ±10°). Within each 16s stimulus block in the functional experiments, noise patterns were dynamically regenerated every 200ms; these stimulus blocks were interspersed with 16s grey rest blocks. In each stimulus, two circular figures 4° in diameter appeared to the left and the right of central fixation (3° center-to-center); one of these was always congruent in orientation with the surround, and the other was incongruent. This minimized variability in overall display salience across blocks. Similarly, to match local orientation information at the boundary of each figure region with the surround, a 0.15° greyscale gap circled both figure regions.

In Experiment 1, the figure and surround stimuli were presented binocularly in greyscale. While maintaining central fixation, subjects were directed to respond when they saw a brief (200ms) spatial frequency decrement, which shifted the range of the noise in one of the figures to 1.5-12cpd. These increments occurred, on average, 4 times in each figure in each 16s block; the timing of their occurrence was independently determined for each figure, although participants were instructed to attend and respond to only one figure in each block. The cue, which appeared 1s before the onset of each block, consisted of a pair of dots (0.1°), one white and one black, to the left and the right of fixation. Each participant was told to attend to the side indicated by one of these two colors throughout the experiment, and cue color was counterbalanced across participants. Experimental conditions of orientation (that is, whether each the orientation in each figure region was congruent or incongruent with the surround) and attention were randomly ordered and counterbalanced within each experimental run, which lasted approximately 4.5 minutes. Each participant completed 11-16 functional runs of this task; behavioral performance was at 74% hit rate (stdev. = 11.8%).

In Experiment 2, we used red/green anaglyph glasses to present the oriented noise comprising the figure regions and the surround to either the same eye or to different eyes. Participants wore the glasses throughout the session, with the red lens always placed over the right eye; thus, red oriented noise would be visible to the left eye but filtered by the lens on the right eye and vice versa. The figures and surround were never shown binocularly, and eye of presentation was counterbalanced across blocks. To create the colored stimuli, presented oriented noise using only the red and green color channels of the projector; the green color channel was reduced from a range of 0-256 to a range of 0-200 to more closely match the apparent luminance of the red, but our experimental design did not necessitate luminance matching the two colors. Throughout the experimental run, participants detected a brief (200ms) contrast decrement on the central fixation, which was a circular 2 x 2 checkerboard (as illustrated in Figure 2A) to facilitate fusion of dichoptic stimuli. These decrements were 50% contrast in magnitude, and occurred on average 4 times per block. Participants detected these events with 84% accuracy (stdev. = 10.3%) and each performed between 12 and 16 functional runs.

fMRI Scanning Parameters

All functional data was collected at the Vanderbilt University Institute for Imaging Science's research-dedicated 7 Tesla magnet using 2mm isotropic voxel resolution and 2s TR. BOLD activity

was measured using single-shot, gradient-echo echoplanar T2*-weighted imaging, at a 2 mm isotropic voxel resolution (40 slices, TR 2000 ms, TE 35 ms; flip angle 63°; FOV 224 x 224; SENSE acceleration factor of 2.9). Additionally, each subject underwent a separate session of retinotopic mapping, which used a standard phase-encoded design (CITE) and a stimulus of flashing checkerboard wedges and expanding rings. Retinotopy data was acquired using a Philips 3T Intera Achieva MRI scanner equipped with an 8-channel coil. Subjects were scanned using 3 mm isotropic resolution (TR 2 s, TE 35 ms, flip angle 80˚, 28 slices, 192 x 192 FOV).

ROI localization

Boundaries between retinotopic areas V1-hV4 were delineated in each participant by hand, identifying reversals in the phase of the polar angle map measurements. These ROIs were aligned to the functional space of the current experiment using FSL and Freesurfer software, and used in addition to a functional localizer to identify regions within each retinotopic area that responded to the two bilateral figure regions of our experiment displays. In the experimental scan session, we ran 2-3 runs of a visual localizer consisting of full contrast flickering checkerboards at each of the two figure locations; these were presented in 16s alternating blocks. Subsequently, ROIs were selected from the conjunction of retinotopy and a statistical map of the left vs. right contrast of our functional localizer. We report results from the 100 most functionally selective voxels as defined by the t-statistic map in each lateralized ROI in early visual areas.

LGN regions of interest were identified in each participant by functional localization as described in Materials and Methods. ROIs in subjects spanned 23-48 voxels bilaterally in Experiment 1 (mean = 37.0, stdev. = 8.3), and 19-57 voxels bilaterally in Experiment 2 (mean = 37.3, stdev. = 13.1).

Data processing

Data were preprocessed using FSL and Freesurfer tools (documented and freely available for download at http://surfer.nmr.mgh.harvard.edu), beginning with 3D motion correction and linear trend removal, followed by slice-timing correction for pRF runs and a high-pass filter cutoff of 60s. Functional images were registered to a reconstructed anatomical space for each subject; this registration was first automated in FSL and then checked and corrected by hand. This allowed the alignment of the current fMRI data to the retinotopy data, which was collected in a separate session. The functional localizer was spatially smoothed using a 1-mm Gaussian kernel; no spatial smoothing was done for the experimental or pRF mapping runs. Further analyses were conducted using a custom Matlab processing stream. Each voxel's intensities were normalized by the mean of the time series, converting to mean percent signal change within each run. Outliers were defined as time points for which the voxel's response measured more than 3 times its standard deviation from its mean, and were Winsorised (Hastings et al., 1947). Finally, a general linear model was fitted to the time course of each run to generate an estimated beta weight for each stimulus blocks; these beta weights were then averaged by experimental condition for the reported analyses.

CHAPTER 5

Conclusions

5.1 Summary of findings

In this dissertation, I sought to characterize the neural mechanisms of feature-tuned suppressive effects across multiple stages of the visual hierarchy, and to evaluate how these effects yield complex perceptual outcomes, like the computation of salience or the detection and segmentation of orientation-defined figures. To do so, I have also evaluated the role of top-down directed attention in these perceptual mechanisms of contextual processing.

I have presented evidence for figure-ground modulation in the human LGN, suggesting that feature-tuned figure enhancement can be instantiated in this structure via corticothalamic feedback. As figure-ground modulation is thought to rely on form- and object-selective regions beyond the primary visual cortex (Grill-Spector, Kushnir, Edelman, Itzchak, & Malach, 1998; Kastner et al., 2000; Kourtzi et al., 2003; Vinberg & Grill-Spector, 2008), this implies a robust transfer of information from higher-order visual regions to the earliest possible anatomical stage of the visual system that can be modulated by neural feedback. This is particularly interesting because most individual LGN neurons themselves exhibit only modest orientation tuning (Smith et al., 1990; Suematsu et al., 2012; Xu et al., 2002), and untuned extraclassical receptive-field suppression (Alitto & Usrey, 2008; Sceniak et al., 2006). My work shows that, via feedback, LGN responses can be sensitive not only to orientation information, but also to contextual effects of orientation. As illustrated by our failure to find feature-defined salience in the LGN, the feedback mechanisms that carry this information likely involve more restricted interactions in retinotopic space, allowing for contextual influences from nearby surrounding regions but not across a broad portion of the visual field.

In the visual cortex, I characterized how feature contrast can function to segment a particular item or figure from a surround. I explored the neural mechanisms of figure-ground modulation, the study of which had been largely limited to non-human primates (e.g., Lamme, 1995; Poort et al., 2012; Zipser et al., 1996). To investigate figure-ground processing in the human visual system, I developed a pRF-based reconstruction method that allowed me to project measured fMRI effects into stimulus space. I found that figure enhancement in V1 can occur in the absence of directed attention to the figure, and in a way that is spatially distinct from local boundary detection. Thus, while prior work in the macaque has suggested that figure-enhancement may be largely the reflection of spatial attention (Poort et al., 2012; Rossi et al., 2001), our studies posit that this is indeed an automatic mechanism of early visual processing in the human. I also demonstrated that across multiple items in a display, feature contrast can yield a computation of visual salience. This process likewise occurs independently of directed spatial attention: an item that is salient is enhanced in V1-hV4 both when the observer attends to it, and when attention is directed away. Like figure enhancement, the computation of salience likely relies on both horizontal connections within V1 and feedback from higher visual areas (Bair et al., 2003; Stettler et al., 2002), which has been shown to increase the spatial extent of suppressive effects (Nassi et al., 2014).

5.2 The neural basis of figure processing

A primary goal of this proposal was to characterize how figures are enhanced at the earliest stages of the visual system. In V1, figure-ground modulation is thought to be comprised of two

distinct stages: an initial *boundary detection*, in which responses to local feature differences are heightened, and a subsequent *figure enhancement*, by which regions within the figure that share features are grouped (Poort et al., 2012). A prominent account suggests that this figure enhancement in early visual cortex relies on feedback from higher-order regions; this is evidenced by measurements of the relative latency of figure enhancement and boundary detection at different levels of the visual cortical hierarchy (Poort et al., 2012) and the laminar profile of these effects (Self et al., 2013). Most recently, researchers demonstrated that microstimulation in V4 selectively affected responses to figures in V1 (Klink, Dagnino, Gariel-Mathis, & Roelfsema, 2017). It is likely that in both humans and non-human primates, regions beyond V1, including those selective for complex objects (LOC; Grill-Spector et al., 1998; Kourtzi & Kanwisher, 2001), shape (V4; Kourtzi et al., 2003; Pasupathy & Connor, 2002) and border ownership (V2, predominantly; Qiu & Heydt, 2005; H. Zhou, Friedman, & Heydt, 2000), influence V1 responses to figures via feedback.

It is interesting to note, however, that few studies in humans have examined the early visual mechanisms of figure processing that have been reported in the non-human primate research. Instead, most fMRI studies have focused on form and object processing in higher-level visual areas like hV4 (Grill-Spector et al., 1999; Kastner et al., 2000; Kourtzi et al., 2003; Kourtzi & Kanwisher, 2001; Vinberg & Grill-Spector, 2008; but see Scholte, Jolij, Fahrenfort, & Lamme, 2008). An early attempt to measure BOLD responses to texture-defined boundaries did not find effects in V1 (Kastner et al., 2000); even as neuroimaging methods have become more sensitive to the influences of form perception in early visual cortex (e.g., Murray, Kersten, Olshausen, Schrater, & Woods, 2002), the spatial resolution of typical fMRI does not allow researchers to resolve the spatial extent of figure-ground modulation.

In Chapter 3, I addressed this gap in knowledge, using a novel, pRF-based analysis to resolve the spatial extent of figure enhancement in human early visual cortex. I report that this figure enhancement appears to occur over the full extent of the figure region in V1; further, the spatial profile of this figure enhancement is distinct from that of a mechanism of local boundary detection. While influential models of computer vision (Grossberg & Mingolla, 1985) and of human behavior (Mumford et al., 1987; Wolfson & Landy, 1998) have proposed this dual-stage figure-ground processing, other modeling work has argued that figure enhancement can be instantiated solely through local, horizontal interactions at the boundary (Z. Li, 2009; Zhaoping, 2003). Importantly, if the latter account were true, figure enhancement would be expected to diminish as a function of distance from the boundary for larger figures (e.g. Rossi et al., 2001). This is not what I observed in our fMRI studies; as illustrated in Figure 3.3, figure enhancement did not decrease as a function of distance from the boundary, even for larger figures of 6° diameter. Instead, our results support the notion of separable mechanisms of boundary detection and figure enhancement, the latter likely instantiated by feedback.

Importantly, I found robust figure enhancement in V1 in the absence of directed attention, both when observers performed a fixation task located at the center of the stimulus (Chapter 3) and when the observer attended to a spatially distinct portion of the visual field (Chapter 4). These findings address an open question in the literature about the role of spatial attention in figure enhancement: does figure enhancement occur in an automatic, stimulus-driven way, or is it primarily the consequence of directed attention? In the non-human primate literature, effects of figure enhancement have not always been replicated when the animal does not saccade to the figure (Rossi et al., 2001), nor is enhancement seen under conditions of anesthesia (Lamme, Zipser, & Spekreijse, 1998b). To address this, Poort and colleagues (2012), measured figure enhancement when trained monkeys performed a task on the figure and when they attended away from it; they found that while attention minimally affected the magnitude of boundary detection

responses, figure enhancement was greatly reduced (by ~50% magnitude) when the animal attended away from the figure.

While these authors concluded that figure enhancement is driven strongly by spatial attention, the current results speak to the enhancement that remains, even in the absence of directed attention. I found consistent figure-ground modulation when participants attended away from the figure in Chapter 4, as well as enhancement over the entire figure region when participants performed an attentionally demanding task in a small (0.5°) fixation region at the center of the figure. This spatially extensive enhancement was observed even when the figure was 6° in diameter, making it highly unlikely that this figure enhancement resulted from an outward spread of spatial attention.

The discrepancies between findings of figure enhancement in the absence of attention in non-human primates (Poort et al., 2012; Rossi et al., 2001) perhaps speak to the persistent effects of the training used in animal research. For training on any experimental paradigm or task, animals typically undergo many months of reward-based training; in many of the studies discussed here, animals were trained to saccade to the orientation-defined figure (Jones et al., 2015; Lamme, 1995; Poort et al., 2012; Zipser et al., 1996). Even if recordings are subsequently done without this saccade task, reward-based training may alter the animal's attention to or representations of the trained stimuli. In humans, such training has been shown to yield long-lasting changes in the salience of stimuli associated with reward (Anderson & Laurent, 2011), which can persist for several months following a modest number of training trials (Anderson & Yantis, 2013). This underscores the benefit of researching the effects of attention in humans, who require minimal training with and exposure to experimental stimuli.

5.3 Figure enhancement and orientation processing in the LGN

An important finding of this dissertation was the consistent modulation of LGN responses by figure-ground modulation (Chapter 4). While recurrent corticothalamic connections to the LGN have been hypothesized to have an important functional role in models of figure-ground modulation (Grossberg, Mingolla, & Ross, 1997), evidence of such effects in LGN neurons has been reported in only one recent study to date (Jones et al., 2015). However, this study in non-human primates could not distinguish figure-ground modulation from the effects of spatial attention, which have been shown to robustly enhance neural responses in the LGN (Ling et al., 2015; O'Connor et al., 2002; Schneider & Kastner, 2009). As discussed above, Jones et al. (2015) required animals to make a saccade to the orientation-defined figure; although in a subsequent experiment the figures were presented while the animal fixated passively, it is likely that reward-based training with the figures had some effect on how the animals attended to them.

In Chapter 4, I thus asked whether figure-ground modulation in the LGN is primarily a consequence of directed attention to the figure, or if it can occur in an automatic, stimulus-driven way. Somewhat surprisingly, I found that enhancement of orientation-defined figures in the LGN occurred both when participants attended to the figure and when they attended away. This finding suggests that while LGN neurons themselves have only modest preferences for orientation (Cheong et al., 2013; Smith et al., 1990; Suematsu et al., 2012; Xu et al., 2002) which may be inherited from the retina (Leventhal & Schall, 1983), their responses can be modulated by *second-order orientation* – that is, the orientation of the figure relative to that of the surround.

Do individual LGN neurons show orientation-tuned suppression? While feature-tuned surround suppression has been demonstrated in LGN neurons of the cat (Cudeiro & Sillito, 1996; Jones et al., 2000; Naito et al., 2007; Sillito et al., 1993; C. Sun et al., 2004), only untuned suppression has been reported in studies of non-human primates (Alitto & Usrey, 2008; Sceniak et

al., 2006). Previous work from our lab has demonstrated the presence of orientation information in the human LGN (Ling et al., 2015), which agrees with the reported orientation biases and preferences in individual neurons (Smith et al., 1990; Suematsu et al., 2012; Xu et al., 2002). However, fMRI is unable to resolve the tuning profile of individual neurons: the responses that we measure are from a population of cells, and the temporal lag of BOLD signal typically precludes analyses that can distinguish feed-forward from feedback properties of responses.

The second study of Chapter 4 leads us to believe that feedback from visual cortex, rather than feedforward processes inherited from the retina or intrinsic to the LGN, are responsible for figure-ground modulation in the structure. Evidence for this comes from our measures of figure-ground modulation while presenting the figure and surround to either the same eye, or to different eyes. In the latter condition, information from the two eyes remains largely segregated during the initial feedforward sweep of processing until it arrives at visual area V1, the first stage of the visual system at which neurons are typically sensitive to input from the two eyes. Thus, our finding that figure-enhancement occurs in the LGN even in this dichoptic presentation condition suggests that this processing is carried out in cortex and then fed back to the LGN. Our findings are generally consistent with a recent proposal that corticothalamic feedback can serve to enhance the responsiveness or gain of neurons in the LGN (Briggs & Usrey, 2008), although it is perhaps surprising that this feedback occurs in the absence of directed attention, in an automatic, stimulus-driven way.

Orientation-tuned V1 neurons appear to send feedback signals to the LGN in a retinotopically organized manner that coincides with the spatial structure of their oriented receptive field (Wang, Jones, Andolina, Salt, & Sillito, 2006a). This spatial specificity may explain why we do not see enhancement of salient items in the LGN (Chapter 2). On one hand, the displays and task that I used to probe salience in Chapter 2 and figure-ground modulation in Chapter 3 are quite similar: in both, the lateralized salient item/figure region is defined by a 90° difference in orientation from the surround items/surround region. However, the spatial separation between items in the salience experiments (effectively ~1.8°) was much larger than in the gap separation (0.5°) between the figure and surround in Chapter 3. Together, these experiments may suggest that orientation-tuned contextual interactions in the LGN are instantiated by mechanisms of automatic feedback with a limited retinotopic spread.

5.4 Recurrent processing of contextual information in the early visual system

We often characterize and contrast 'bottom-up,' stimulus-driven processes with 'top-down' influences of the observer's goals and cognitive knowledge; however, this simplified view does not account for the pervasive recurrent processing documented in many studies. The complex, feedback-driven processing explored in this dissertation can result in the representation of information by neurons that are themselves untuned for this information (as in orientation-contrast in the LGN), as well as information beyond the spatial extent of the neuron's response field. Considering the active role of stimulus-driven feedback in the early visual system may challenge the way that we think about these regions' tuning for visual features: for example, a 'V1-like' Gabor bank model of feature detection may be quite poor in predicting V1 activity after the initial feed-forward sweep, even in the absence of typically 'top-down' modulations of directed attention or semantic knowledge.

Corticothalamic feedback from V1 to the LGN in cortex is pervasive: it is estimated that only ~10% of synaptic input to LGN neurons arrives from the retina, while ~30% of the input is fed back from layer 6 of primary visual cortex (for review: Sillito, Cudeiro, & Jones, 2006). This feedback is topographically organized and functionally aligned across receptive field properties

(Wang, Jones, Andolina, Salt, & Sillito, 2006b), and can influence LGN neurons by modulating their contrast gain (Ahlsén, Lindström, & Lo, 1985; Jones et al., 2012; Przybyszewski, Gaska, Foote, & Pollen, 2000; Webb et al., 2002), spatial response properties (e.g., Jones et al., 2012; Webb et al., 2002) or temporal response properties (reviewed in Ghodrati, Khaligh-Razavi, & Lehky, 2017). Broadly, it has been theorized that this feedback can function to enhance the transmission of relevant information from the thalamus to the cortex, either by increasing the gain or responsiveness of neurons or by increasing the reliability of thalamic responses (Briggs & Usrey, 2008). In this way, the role of the LGN as a 'relay-station' is far from passive, as feedback from V1 can profoundly impact quality and signal strength of the information that passes through this structure and onward to V1.

Why would contextual information, and in particular, 'high-level' processing of figural information, be fed back to the earliest possible anatomical stage of visual processing? Why is it the case that neural populations in the LGN, which provide the majority of input to the visual cortex, are modulated in this way? We may consider the notion of predictive coding of information (Rao & Ballard, 1999), in which each stage of processing reflects the complexities of higher-order stages and represents deviations from the predicted input, rather than a passive detection of the absolute values of the features of that input. According to this theory, which has been used to computationally predict enhancement of orientation-tuned figures and other contextual effects (Rao & Ballard, 1999), feedback connections from higher visual areas carry predictions about visual input. Subsequent feed-forward connections convey the residual error of these predictions. Critically, higher visual areas integrate information over a larger portion of the visual field than the regions to which they feed back; this allows feature differences to be amplified, as predictions carry information about the surrounding context, and the resulting 'residual error' reflects cases in which the broader context cannot predict the feature value observed.

The findings of this dissertation underscore this notion that recurrent interactions between the LGN and early visual cortex are involved in the processing of visual-contextual information. This stimulus-driven processing consists not only of simple feature detection, but yields contextual effects that lead to complex perceptual phenomena, including the detection and segmentation of figures. It is possible that the LGN is sensitive to additional stimulus-driven contextual effects, like contour integration (Field, Hayes, & Hess, 1993; W. Li, Piëch, & Gilbert, 2006; Z. Li, 1998), visual crowding (J. Chen et al., 2014; Whitney & Levi, 2011), or the perception of illusory surfaces/figures (Kok & de Lange, 2014; Mendola, Dale, Fischl, Liu, & Tootell, 1999; Stanley & Rubin, 2003) – although, as I have demonstrated, the prevalence of these effects in the LGN may constrained by the spatial specificity of corticothalamic feedback. It will be of considerable interest for future studies to examine whether such effects can indeed be measured in the LGN.

References

Adesnik, H., Bruns, W., Taniguchi, H., Huang, Z. J., & Scanziani, M. (2012). A neural circuit for spatial summation in visual cortex. *Nature*, *490*, 226–231.

Ahlsén, G., Lindström, S., & Lo, F. S. (1985). Interaction between inhibitory pathways to principal cells in the lateral geniculate nucleus of the cat. *Experimental Brain Research*, *58*, 134–143.

Alitto, H. J., & Usrey, W. M. (2008). Origin and dynamics of extraclassical suppression in the lateral geniculate nucleus of the macaque monkey. *Neuron*, *57*, 135–146.

Allman, J., Miezin, F., & McGuinness, E. (1985). Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. *Annual Review of Neuroscience*, *8*, 407–430.

Anderson, B. A., & Laurent, P. A. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences of the United States of America*, *108*, 10367–10371.

Anderson, B. A., & Yantis, S. (2013). Persistence of value-driven attentional capture. *Journal of Experimental Psychology*, *39*, 6–9.

Angelucci, A., & Bressloff, P. C. (2006). Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate V1 neurons. *Progress in Brain Research*, *154*, 93–120.

Bair, W., Cavanaugh, J. R., & Movshon, J. A. (2003). Time course and time-distance relationships for surround suppression in macaque V1 neurons. *Journal of Neuroscience*, *23*, 7690–7701.

Barlow, H. (1961). Possible principles underlying the transformation of sensory messages. *Sensory Communication*, 217–234.

Beck, D., & Kastner, S. (2005). Stimulus context modulates competition in human extrastriate cortex. *Nature Neuroscience*, *8*, 1110–1116.

Blakemore, C., & Tobin, E. A. (1972). Lateral inhibition between orientation detectors in the cat's visual cortex. *Experimental Brain Research*, *15*, 439–440.

Blakemore, C., Carpenter, R. H., & Georgeson, M. A. (1970). Lateral inhibition between orientation detectors in the human visual system. *Nature*, *228*, 37–39.

Bogler, C., Bode, S., & Haynes, J.-D. (2011). Decoding Successive Computational Stages of Saliency Processing. *Current Biology*, *21*, 1667–1671.

Bogler, C., Bode, S., & Haynes, J.-D. (2013). Orientation pop-out processing in human visual cortex. *Neuroimage*, *81*, 73–80.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.

Brefczynski, J. A., & DeYoe, E. A. (1999). A physiological correlate of the “spotlight” of visual attention. *Nature Neuroscience*, *2*, 370–374.

Briggs, F., & Usrey, W. M. (2008). Emerging views of corticothalamic function. *Current Opinion in Neurobiology*, *18*, 403–407.

Cannon, M. W., & Fullenkamp, S. C. (1991). Spatial interactions in apparent contrast: inhibitory effects among grating patterns of different spatial frequencies, spatial positions and orientations. *Vision Research*, *31*, 1985–1998.

Carandini, Matteo, and David J. Heeger. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience,* 13, 51-62.

Carrasco, M. (2006). Covert attention increases contrast sensitivity: psychophysical, neurophysiological and neuroimaging studies. In *Visual Perception - Fundamentals of Vision: Low and Mid-Level Processes in Perception* (Vol. 154, pp. 33–70). Elsevier.

Cavanaugh, J. R. (2002). Nature and Interaction of Signals From the Receptive Field Center and Surround in Macaque V1 Neurons. *Journal of Neurophysiology*, *88*, 2530–2546.

Cavanaugh, J. R., Bair, W., & Movshon, J. A. (2002). Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *Journal of Neurophysiology*, *88*, 2547–2556.
Chao-Yi, L., & Wu, L. (1994). Extensive integration field beyond the classical receptive field of cat's striate cortical neurons—classification and tuning properties. *Vision Research*, *34*, 2337–2355.
Chen, J., He, Y., Zhu, Z., Zhou, T., Peng, Y., Zhang, X., & Fang, F. (2014). Attention-dependent early cortical suppression contributes to crowding. *Journal of Neuroscience*, *34*, 10465–10474.
Cheong, S. K., Tailby, C., Solomon, S. G., & Martin, P. R. (2013). Cortical-like receptive fields in the lateral geniculate nucleus of marmoset monkeys. *Journal of Neuroscience*, *33*, 6864–6876.
Coen-Cagli, R., Dayan, P., & Schwartz, O. (2012). Cortical Surround Interactions and Perceptual Salience via Natural Scene Statistics. *PLoS Computational Biology*, *8*, e1002405.
Coen-Cagli, R., Kohn, A., & Schwartz, O. (2015). Flexible gating of contextual influences in natural vision. *Nature Neuroscience*, *18*, 1648–1655.
Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215.
Cudeiro, J., & Sillito, A. M. (1996). Spatial frequency tuning of orientation-discontinuity-sensitive corticofugal feedback to the cat lateral geniculate nucleus. *The Journal of Physiology*, *490*, 481–492.
DeAngelis, G. C., Freeman, R. D., & Ohzawa, I. (1994). Length and width tuning of neurons in the cat's primary visual cortex. *Journal of Neurophysiology*, *71*, 347–374.
Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.
Dragoi, V., Rivadulla, C., & Sur, M. (2001). Foci of orientation plasticity in visual cortex. *Nature*, *411*, 80–86.
Dumoulin, S. O., & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *Neuroimage*, *39*, 647–660.
Engel, S. A., Glover, G. H., & Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex*, *7*, 181–192.
Fecteau, J. H., & Munoz, D. P. (2006). Salience, relevance, and firing: a priority map for target selection. *Trends in Cognitive Sciences*, *10*, 382–390.
Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local "association field." *Vision Research*, *33*, 173–193.
Flevaris, A. V., & Murray, S. O. (2015). Attention Determines Contextual Enhancement versus Suppression in Human Primary Visual Cortex. *The Journal of Neuroscience*, *35*, 12273–12280.
Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *96*, 3314–3319.
Ghodrati, M., Khaligh-Razavi, S.-M., & Lehky, S. R. (2017). Towards building a more complex view of the lateral geniculate nucleus: Recent advances in understanding its role. *Progress in Neurobiology*. doi:10.1016/j.pneurobio.2017.06.002
Gilbert, C. D. (1992). Horizontal integration and cortical dynamics. *Neuron*, *9*, 1–13.
Gilbert, C. D., & Wiesel, T. N. (1979). Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature*, *280*, 120–125.
Gottlieb, J. P., Kusunoki, M., & Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, *391*, 481–484.
Grill-Spector, K., Kushnir, T., Edelman, S., & Avidan, G. (1999). Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron*, *24*, 187–203.

Grill-Spector, K., Kushnir, T., Edelman, S., Itzchak, Y., & Malach, R. (1998). Cue-invariant activation in object-related areas of the human occipital lobe. *Neuron-Cambridge Ma-*, *21*, 191–202.
Grossberg, S., & Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception & Psychophysics*, *38*, 141–171.
Grossberg, S., Mingolla, E., & Ross, W. D. (1997). Visual brain and visual perception: How does the cortex do perceptual grouping? *Trends in Neurosciences*. doi:10.1016/S0166-2236(96)01002-8
Harrison, L. M., Stephan, K. E., Rees, G., & Friston, K. J. (2007). Extra-classical receptive field effects measured in striate cortex with fMRI. *Neuroimage*, *34*, 1199–1208.
Hastings, C., Jr, Mosteller, F., & Tukey, J. W. (1947). Low moments for small samples: a comparative study of order statistics. *The Annals of Mathematical Statistics*, *8*, 413–426.
Henderson, J. M. (2003). Human gaze control during real-world scene perception. *Trends in Cognitive Sciences*, *7*, 498–504.
Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, *12*, 55–67.
Hopf, J.-M., Noesselt, T., Tempelmann, C., Braun, J., Schoenfeld, M. A., & Heinze, H.-J. (2004). Popout modulates focal attention in the primary visual cortex. *Neuroimage*, *22*, 574–582.
Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*, 215–243.
Itti, L., & Koch, C. (2001). Computational modelling of visual attention : Abstract : Nature Reviews Neuroscience. *Nature Reviews Neuroscience*, *2*, 194–203.
Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *20*, 1254–1259.
Jeffreys, S. H. (1961). The Theory of Probability, 3rd: Oxford University Press.
Jehee, J. F. M., Brady, D. K., & Tong, F. (2011). Attention improves encoding of task-relevant features in the human visual cortex. *Journal of Neuroscience*, *31*, 8210–8219.
Jin, D. Z., Dragoi, V., Sur, M., & Seung, H. S. (2005). Tilt Aftereffect and Adaptation-Induced Changes in Orientation Tuning in Visual Cortex. *Journal of Neurophysiology*, *94*, 4038–4050.
Jones, H. E., Andolina, I. M., Ahmed, B., Shipp, S. D., Clements, J. T. C., Grieve, K. L., et al. (2012). Differential Feedback Modulation of Center and Surround Mechanisms in Parvocellular Cells in the Visual Thalamus. *The Journal of Neuroscience*, *32*, 15946–15951.
Jones, H. E., Andolina, I. M., Oakely, N. M., Murphy, P. C., & Sillito, A. M. (2000). Spatial summation in lateral geniculate nucleus and visual cortex. *Experimental Brain Research*, *135*, 279–284.
Jones, H. E., Andolina, I. M., Shipp, S. D., Adams, D. L., Cudeiro, J., Salt, T. E., & Sillito, A. M. (2015). Figure-ground modulation in awake primate thalamus. *Proceedings of the National Academy of Sciences of the United States of America*, *112*, 7085–7090.
Jones, H. E., Grieve, K. L., & Wang, W. (2001). Surround suppression in primate V1. *Journal of Neurophysiology*, *86*, 2011–2028.
Joo, S. J., Boynton, G. M., & Murray, S. O. (2012). Long-range, pattern-dependent contextual effects in early human visual cortex. *Current Biology*, *22*, 781–786.
Kapadia, M. K., Ito, M., Gilbert, C. D., & Westheimer, G. (1995). Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron*, *15*, 843–856.
Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, *23*, 315–341.
Kastner, S., De Weerd, P., & Ungerleider, L. G. (2000). Texture segregation in the human visual cortex: A functional MRI study. *Journal of Neurophysiology*, *83*, 2453–2457.

Kastner, S., O'Connor, D. H., Fukui, M. M., Fehd, H. M., Herwig, U., & Pinsk, M. A. (2004). Functional Imaging of the Human Lateral Geniculate Nucleus and Pulvinar. *Journal of Neurophysiology*, *91*, 438–448.
Kay, K., Winawer, J., Mezer, A., & Wandell, B. (2013). Compressive spatial summation in human visual cortex. *Journal of Neurophysiology*, *110*, 481–494.
Kayser, C., Körding, K., & König, P. (2004). Processing of complex stimuli and natural scenes in the visual cortex. *Current Opinion in Neurobiology*, *14*, 468.
Kincade, J. M., Abrams, R. A., Astafiev, S. V., Shulman, G. L., & Corbetta, M. (2005). An Event-Related Functional Magnetic Resonance Imaging Study of Voluntary and Stimulus-Driven Orienting of Attention. *The Journal of Neuroscience*, *25*, 4593–4604.
Klink, P. C., Dagnino, B., Gariel-Mathis, M.-A., & Roelfsema, P. R. (2017). Distinct Feedforward and Feedback Effects of Microstimulation in Visual Cortex Reveal Neural Mechanisms of Texture Segregation. *Neuron*, *95*, 209–220.e3.
Knierim, J. J., & van Essen, D. C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *Journal of Neurophysiology*, *67*, 961–980.
Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, *4*, 219–227.
Koene, A. R., & Zhaoping, L. (2007). Feature-specific interactions in salience from combined feature contrasts: evidence for a bottom-up saliency map in V1. *Journal of Vision*, *7*, 6.1–14.
Kok, P., & de Lange, F. P. (2014). Shape Perception Simultaneously Up- and Downregulates Neural Activity in the Primary Visual Cortex. *Current Biology*, *24*, 1531–1535.
Kourtzi, Z., & Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science*, *293*, 1506–1509.
Kourtzi, Z., Tolias, A. S., Altmann, C. F., Augath, M., & Logothetis, N. K. (2003). Integration of Local Features into Global Shapes. *Neuron*, *37*, 333–346.
Lamme, V. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *The Journal of Neuroscience*, *15*, 1605–1615.
Lamme, V. A., Super, H., & Spekreijse, H. (1998a). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology*, *8*, 529–535.
Lamme, V. A., Zipser, K., & Spekreijse, H. (1998b). Figure-ground activity in primary visual cortex is suppressed by anesthesia. *Proceedings of the National Academy of Sciences of the United States of America*, *95*, 3263–3268.
Leventhal, A. G., & Schall, J. D. (1983). Structural basis of orientation sensitivity of cat retinal ganglion cells. *Journal of Comparative Neurology*, *220*, 465–475.
Levitt, J. B., & Lund, J. S. (1997). Contrast dependence of contextual effects in primate visual cortex. *Nature*, *387*, 73–76.
Li, W., Piëch, V., & Gilbert, C. D. (2006). Contour saliency in primary visual cortex. *Neuron*, *50*, 951–962.
Li, Z. (1998). A neural model of contour integration in the primary visual cortex. *Neural Computation*, *10*, 903–940.
Li, Z. (1999). Contextual influences in V1 as a basis for pop out and asymmetry in visual search. *Proceedings of the National Academy of Sciences of the United States of America*, *96*, 10530–10535.
Li, Z. (2002). A saliency map in primary visual cortex. *Trends in Cognitive Sciences*, *6*, 9–16.
Li, Z. (2009). Visual segmentation by contextual influences via intra-cortical interactions in the primary visual cortex. *Network: Computation in Neural Systems*. doi:10.1088/0954-898X_10_2_305
Lien, A. D., & Scanziani, M. (2013). Tuned thalamic excitation is amplified by visual cortical circuits.

Nature Neuroscience, *16*, 1315–1323.
Ling, S., Pratte, M. S., & Tong, F. (2015). Attention alters orientation processing in the human lateral geniculate nucleus. *Nature Neuroscience*, *18*, 496–498.
Marcus, D. S., & van Essen, D. C. (2002). Scene segmentation and attention in primate cortical areas V1 and V2. *Journal of Neurophysiology*, *88*, 2648–2658.
Mazer, J. A., & Gallant, J. L. (2003). Goal-Related Activity in V4 during Free Viewing Visual Search. *Neuron*, *40*, 1241–1250.
McAlonan, K., Cavanaugh, J., & Wurtz, R. H. (2008). Guarding the gateway to cortex with attention in visual thalamus. *Nature*, *456*, 391–394.
McDonald, J. S., Seymour, K. J., Schira, M. M., Spehar, B., & Clifford, C. W. G. (2009). Orientation-specific contextual modulation of the fMRI BOLD response to luminance and chromatic gratings in human visual cortex. *Vision Research*, *49*, 1397–1405.
McMains, S., & Kastner, S. (2011). Interactions of top-down and bottom-up mechanisms in human visual cortex. *Journal of Neuroscience*, *31*, 587–597.
Melloni, L., van Leeuwen, S., Alink, A., & Müller, N. G. (2012). Interaction between bottom-up saliency and top-down control: how saliency maps are created in the human brain. *Cerebral Cortex*, *22*, 2943–2952.
Mendola, J. D., Dale, A. M., Fischl, B., Liu, A. K., & Tootell, R. B. (1999). The representation of illusory and real contours in human cortical visual areas revealed by functional magnetic resonance imaging. *Journal of Neuroscience*, *19*, 8560–8572.
Mumford, D., Kosslyn, S. M., Hillger, L. A., & Herrnstein, R. J. (1987). Discriminating figure from ground: the role of edge detection and region growing. *Proceedings of the National Academy of Sciences of the United States of America*, *84*, 7354–7358.
Murray, S. O., Kersten, D., Olshausen, B. A., Schrater, P., & Woods, D. L. (2002). Shape perception reduces activity in human primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *99*, 15164–15169.
Naito, T., Sadakane, O., Okamoto, M., & Sato, H. (2007). Orientation tuning of surround suppression in lateral geniculate nucleus and primary visual cortex of cat. *Neuroscience*, *149*, 962–975.
Nassi, J. J., Gómez-Laberge, C., Kreiman, G., & Born, R. T. (2014). Corticocortical feedback increases the spatial extent of normalization. *Frontiers in Systems Neuroscience*, *8*, 226.
Nassi, J. J., Lomber, S. G., & Born, R. T. (2013). Corticocortical feedback contributes to surround suppression in V1 of the alert primate. *Journal of Neuroscience*, *33*, 8504–8517.
Nelson, J. I., & Frost, B. J. (1978). Orientation-selective inhibition from beyond the classic visual receptive field. *Brain Research*, *139*, 359–365.
Nothdurft, H. C. (1993). The role of features in preattentive vision: comparison of orientation, motion and color cues. *Vision Research*, *33*, 1937–1958.
Nothdurft, H. C., Gallant, J. L., & van Essen, D. C. (1999). Response modulation by texture surround in primate area V1: correlates of “popout” under anesthesia. *Visual Neuroscience*, *16*, 15–34.
Nurminen, L., & Angelucci, A. (2014). Multiple components of surround modulation in primary visual cortex: Multiple neural circuits with multiple functions? *Vision Research*, *104*, 47–56.
O'Connor, D. H., Fukui, M. M., Pinsk, M. A., & Kastner, S. (2002). Attention modulates responses in the human lateral geniculate nucleus. *Nature Neuroscience*, *5*, 1203–1209.
Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research*, *42*, 107–123.
Pasupathy, A., & Connor, C. E. (2002). Population coding of shape in area V4 - ProQuest. *Nature Neuroscience*.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
Petrov, Y., & McKee, S. P. (2006). The effect of spatial configuration on surround suppression of contrast sensitivity. *Journal of Vision*, *6*, 224–238.
Poltoratski, S., Ling, S., McCormack, D., & Tong, F. (2017). Characterizing the effects of feature salience and top-down attention in the early visual system. *Journal of Neurophysiology*, jn.00924.2016.
Poort, J., Raudies, F., Wannig, A., Lamme, V. A. F., Neumann, H., & Roelfsema, P. R. (2012). The Role of Attention in Figure-Ground Segregation in Areas V1 and V4 of the Visual Cortex. *Neuron*, *75*, 143–156.
Przybyszewski, A. W., Gaska, J. P., Foote, W., & Pollen, D. A. (2000). Striate cortex increases contrast gain of macaque LGN neurons. *Visual Neuroscience*, *17*, 485–494.
Qiu, F. T., & Heydt, von der, R. (2005). Figure and Ground in the Visual Cortex: V2 Combines Stereoscopic Cues with Gestalt Rules. *Neuron*, *47*, 155–166.
Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*, 79–87.
Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, *61*, 168–185.
Roelfsema, P. R., Tolboom, M., & Khayat, P. S. (2007). Different Processing Phases for Features, Figures, and Selective Attention in the Primary Visual Cortex. *Neuron*, *56*, 785–792.
Rossi, A., Desimone, R., & Ungerleider, L. (2001). Contextual modulation in primary visual cortex of macaques. *Journal of Neuroscience*, *21*, 1698–1709.
Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237.
Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, *5*, 631–632.
Sceniak, M. P., Chatterjee, S., & Callaway, E. M. (2006). Visual Spatial Summation in Macaque Geniculocortical Afferents. *Journal of Neurophysiology*, *96*, 3474–3484.
Schallmo, M.-P., Grant, A. N., Burton, P. C., & Olman, C. A. (2016). The effects of orientation and attention during surround suppression of small image features: A 7 Tesla fMRI study, *16*, 19–21.
Schira, M. M., Fahle, M., Donner, T. H., Kraft, A., & Brandt, S. A. (2004). Differential contribution of early visual areas to the perceptual process of contour processing. *Journal of Neurophysiology*, *91*, 1716–1721.
Schneider, K. A. (2011). Subcortical Mechanisms of Feature-Based Attention. *The Journal of Neuroscience*, *31*, 8643–8653.
Schneider, K. A., & Kastner, S. (2009). Effects of sustained spatial attention in the human lateral geniculate nucleus and superior colliculus. *Journal of Neuroscience*, *29*, 1784–1795.
Scholte, H. S., Jolij, J., Fahrenfort, J. J., & Lamme, V. A. F. (2008). Feedforward and recurrent processing in scene segmentation: electroencephalography and functional magnetic resonance imaging. *Journal of Cognitive Neuroscience*, *20*, 2097–2109.
Self, M. W., Peters, J. C., Possel, J. K., Reithler, J., Goebel, R., Ris, P., et al. (2016). The Effects of Context and Attention on Spiking Activity in Human Early Visual Cortex. *PLoS Biology*, *14*, e1002420.
Self, M. W., van Kerkoerle, T., Supèr, H., & Roelfsema, P. R. (2013). Distinct Roles of the Cortical Layers of Area V1 in Figure-Ground Segregation. *Current Biology*, *23*, 2121–2129.
Shushruth, S., Ichida, J. M., Levitt, J. B., & Angelucci, A. (2009). Comparison of spatial summation properties of neurons in macaque V1 and V2. *Journal of Neurophysiology*, *102*, 2069–2083.

Shushruth, S., Nurminen, L., Bijanzadeh, M., Ichida, J. M., Vanni, S., & Angelucci, A. (2013). Different orientation tuning of near- and far-surround suppression in macaque primary visual cortex mirrors their tuning in human perception. *The Journal of Neuroscience*, *33*, 106–119.
Sillito, A. M., Cudeiro, J., & Jones, H. E. (2006). Always returning: feedback and sensory processing in visual cortex and thalamus. *Trends in Neurosciences*, *29*, 307–316.
Sillito, A. M., Cudeiro, J., & Murphy, P. C. (1993). Orientation sensitive elements in the corticofugal influence on centre-surround interactions in the dorsal lateral geniculate nucleus. *Experimental Brain Research*, *93*, 6–16.
Sillito, A. M., Grieve, K. L., Jones, H. E., Cudeiro, J., & Davis, J. (1995). Visual cortical mechanisms detecting focal orientation discontinuities. *Nature*, *378*, 492–496.
Silver, M. A., & Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences*, *13*, 488–495.
Smith, E. L., Chino, Y. M., Ridder, W. H., Kitagawa, K., & Langston, A. (1990). Orientation bias of neurons in the lateral geniculate nucleus of macaque monkeys. *Visual Neuroscience*, *5*, 525–545.
Somers, D. C., Dale, A. M., Seiffert, A. E., & Tootell, R. B. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *96*, 1663–1668.
Squire, R. F., Noudoost, B., Schafer, R. J., & Moore, T. (2013). Prefrontal contributions to visual selective attention. *Annual Review of Neuroscience*, *36*, 451–466.
Stanley, D. A., & Rubin, N. (2003). fMRI activation in response to illusory contours and salient regions in the human lateral occipital complex. *Neuron*, *37*, 323–331.
Stettler, D. D., Das, A., Bennett, J., & Gilbert, C. D. (2002). Lateral Connectivity and Contextual Interactions in Macaque Primary Visual Cortex. *Neuron*, *36*, 739–750.
Suematsu, N., Naito, T., & Sato, H. (2012). Relationship between orientation sensitivity and spatiotemporal receptive field structures of neurons in the cat lateral geniculate nucleus. *Neural Networks*, *35*, 10–20.
Sun, C., Chen, X., Huang, L., & Shou, T. (2004). Orientation bias of the extraclassical receptive field of the relay cells in the cat's dorsal lateral geniculate nucleus. *Neuroscience*, *125*, 495–505.
Sun, W., Tan, Z., Mensh, B. D., & Ji, N. (2016). Thalamus provides layer 4 of primary visual cortex with orientation- and direction-tuned inputs. *Nature Neuroscience*, *19*, 308–315.
Tatler, B. W., Hayhoe, M. M., Land, M. F., & Ballard, D. H. (2011). Eye guidance in natural vision: reinterpreting salience. *Journal of Vision*, *11*, 5–5.
Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*, 97–136.
Vidyasagar, T. R., & Eysel, U. T. (2015). Origins of feature selectivities and maps in the mammalian primary visual cortex. *Trends in Neurosciences*, *38*, 475–485.
Vinberg, J., & Grill-Spector, K. (2008). Representation of Shapes, Edges, and Surfaces Across Multiple Cues in the Human Visual Cortex. *Journal of Neurophysiology*, *99*, 1380–1393.
Vinje, W., & Gallant, J. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, *287*, 1273–1276.
Wandell, B. A., & Winawer, J. (2015). Computational neuroimaging and population receptive fields. *Trends in Cognitive Sciences*, *19*, 1–9.
Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, *56*, 366–383.
Wang, W., Jones, H. E., Andolina, I. M., Salt, T. E., & Sillito, A. M. (2006a). Functional alignment of feedback effects from visual cortex to thalamus. *Nature Neuroscience*, *9*, 1330–1336.

Wang, W., Jones, H. E., Andolina, I. M., Salt, T. E., & Sillito, A. M. (2006b). Functional alignment of feedback effects from visual cortex to thalamus. *Nature Neuroscience*, *9*, 1330–1336.
Webb, B. S., Tinsley, C. J., Barraclough, N. E., Easton, A., Parker, A., & Derrington, A. M. (2002). Feedback from V1 and inhibition from beyond the classical receptive field modulates the responses of neurons in the primate lateral geniculate nucleus. *Visual Neuroscience*, *19*, 583–592.
Whitney, D., & Levi, D. M. (2011). Visual crowding: a fundamental limit on conscious perception and object recognition. *Trends in Cognitive Sciences*, *15*, 160–168.
Wolfe, J. M. (1994). Guided Search 2.0 A revised model of visual search. *Psychonomic Bulletin & Review*, *1*, 202–238.
Wolfson, S. S., & Landy, M. S. (1998). Examining edge- and region-based texture analysis mechanisms. *Vision Research*, *38*, 439–446.
Xu, X., Ichida, J., Shostak, Y., Bonds, A. B., & Casagrande, V. A. (2002). Are primate lateral geniculate nucleus (LGN) cells really sensitive to orientation or direction? *Visual Neuroscience*, *19*, 97–108.
Zenger-Landolt, B., & Heeger, D. (2003). Response suppression in V1 agrees with psychophysics of surround masking. *The Journal of Neuroscience*, *23*, 6884–6893.
Zhang, X., Zhaoping, L., Zhou, T., & Fang, F. (2012). Neural Activities in V1 Create a Bottom-Up Saliency Map. *Neuron*, *73*, 183–192.
Zhaoping, L. (2003). V1 mechanisms and some figure–ground and border effects. *Journal of Physiology-Paris*, *97*, 503–515.
Zhou, H., Friedman, H. S., & Heydt, von der, R. (2000). Coding of Border Ownership in Monkey Visual Cortex. *The Journal of Neuroscience*, *20*, 6594–6611.
Zipser, K., Lamme, V., & Schiller, P. (1996). Contextual modulation in primary visual cortex. *The Journal of Neuroscience*, *16*, 7376–7389.