

2010

Dynamic and Integrative Properties of the Primary Visual Cortex

Justin N. J. McManus

Follow this and additional works at: http://digitalcommons.rockefeller.edu/student_theses_and_dissertations



Part of the [Life Sciences Commons](#)

Recommended Citation

McManus, Justin N. J., "Dynamic and Integrative Properties of the Primary Visual Cortex" (2010). *Student Theses and Dissertations*. Paper 80.

This Thesis is brought to you for free and open access by Digital Commons @ RU. It has been accepted for inclusion in Student Theses and Dissertations by an authorized administrator of Digital Commons @ RU. For more information, please contact mcsweej@mail.rockefeller.edu.



DYNAMIC AND INTEGRATIVE PROPERTIES
OF THE PRIMARY VISUAL CORTEX

A Thesis Presented to the Faculty of
The Rockefeller University
in Partial Fulfillment of the Requirements for
the degree of Doctor of Philosophy

by

Justin N. J. McManus

June 2010

DYNAMIC AND INTEGRATIVE PROPERTIES OF THE PRIMARY VISUAL CORTEX

Justin N. J. McManus, Ph.D.

The Rockefeller University 2010

The ability to derive meaning from complex, ambiguous sensory input requires the integration of information over both space and time, as well as cognitive mechanisms to dynamically shape that integration. We have studied these processes in the primary visual cortex (V1), where neurons have been proposed to integrate visual inputs along a geometric pattern known as the association field (AF). We first used cortical reorganization as a model to investigate the role that a specific network of V1 connections, the long-range horizontal connections, might play in temporal and spatial integration across the AF. When retinal lesions ablate sensory information from portions of the visual field, V1 undergoes a process of reorganization mediated by compensatory changes in the network of horizontal collaterals. The reorganization accompanies the brain's amazing ability to perceptually "fill-in", or "see", the lost visual input. We developed a computational model to simulate cortical reorganization and perceptual fill-in mediated by a plexus of horizontal connections that encode the AF. The model

reproduces the major features of the perceptual fill-in reported by human subjects with retinal lesions, and it suggests that V1 neurons, empowered by their horizontal connections, underlie both perceptual fill-in and normal integrative mechanisms that are crucial to our visual perception. These results motivated the second prong of our work, which was to experimentally study the normal integration of information in V1. Since psychophysical and physiological studies suggest that spatial interactions in V1 may be under cognitive control, we investigated the integrative properties of V1 neurons under different cognitive states. We performed extracellular recordings from single V1 neurons in macaques that were trained to perform a delayed-match-to-sample contour detection task. We found that the ability of V1 neurons to summate visual inputs from beyond the classical receptive field (cRF) imbues them with selectivity for complex contour shapes, and that neuronal shape selectivity in V1 changed dynamically according to the shapes monkeys were cued to detect. Over the population, V1 encoded subsets of the AF, predicted by the computational model, that shifted as a function of the monkeys' expectations. These results support the major conclusions of the theoretical work; even more, they reveal a sophisticated mode of form processing, whereby the selectivity of the whole network in V1 is reshaped by cognitive state.

For my father, Thomas J. McManus

ACKNOWLEDGMENTS

I am grateful to Drs. Wu Li and Charles Gilbert for their mentorship and support. I also thank Drs. Daniel Stettler and Valentin Piëch for stimulating discussions, on subjects from neuroscience in its broadest strokes to mathematical data analyses in their finest detail. Finally, I thank Drs. Homare Yamahachi and Sally Marik for their camaraderie and for fruitful collaborations.

TABLE OF CONTENTS

1	Introduction.....	1
1.1	The classical receptive field (cRF).....	1
1.2	Contextual interactions.....	4
1.3	Horizontal connections and their role in sensory integration.....	14
1.4	Cortical plasticity and reorganization.....	18
1.5	Perceptual fill-in.....	23
1.6	Cognitive (top-down) influences on sensory processing.....	24
1.7	Cortical sources of top-down control.....	32
1.8	Current work.....	43
2	A computational model of perceptual fill-in following retinal degeneration.....	45
2.1	Summary.....	45
2.2	Introduction.....	46
2.3	Methods.....	51
2.3.1	Qualitative description and general overview.....	51
2.3.2	Stage 1: Simulation of cortical activity.....	55
2.3.3	Stage 2: Illustration of perceptual fill-in.....	62
2.3.4	Technical details of the model.....	63
2.4	Results.....	86

2.5	Discussion.....	99
2.6	Acknowledgments.....	106
2.7	Footnotes.....	106
3	Shape selectivity and its top-down modulation in primary visual cortex.....	110
3.1	Summary.....	110
3.2	Introduction.....	111
3.3	Results.....	119
3.3.1	Contour detection task.....	119
3.3.2	Stimulus generation algorithm.....	124
3.3.3	Geometric tuning surfaces.....	126
3.3.4	Shape selectivity and task dependence over the population.....	144
3.3.5	Time course of geometric selectivity.....	152
3.4	Discussion.....	160
3.5	Experimental Procedures.....	171
3.5.1	Electrophysiology.....	171
3.5.2	Data acquisition and stimulus design.....	172
3.5.3	Task design.....	173
3.5.4	Automated stimulus generation.....	176

3.5.5	<i>A priori</i> stimulus additions.....	180
3.5.6	Three-dimensional tuning surfaces.....	181
3.5.7	5-bar population analysis with Sammon's mapping.....	183
3.5.8	Permutation statistics.....	184
3.5.9	Two-way ANOVA.....	186
3.6	Supplemental Figures.....	188
3.7	Supplemental Experimental Procedures.....	194
3.7.1	Eye tracking analysis.....	194
3.8	Acknowledgments.....	195
4	Future Directions.....	196
4.1	Introduction.....	196
4.2	Relating AMD Psychophysics to Computational Simulation.....	196
4.3	Improving the Computational Model.....	199
4.4	Testing the Theoretical Model with Multielectrode Arrays.....	200
4.5	A Direct Test of the Adaptive Processor Theory.....	203
4.6	Localizing the Cortical Origin of the Task-Dependent AF.....	205
4.7	Resolving the Speed of AF Plasticity and its Role in Object Perception.....	208
4.8	Studying the Mechanisms of Top-Down Control.....	213
4.9	Conclusion.....	218

5	References.....	220
---	-----------------	-----

LIST OF FIGURES

Figure 1. Perceptual grouping rules for contour integration and saliency.....	11
Figure 2. The association field (AF).....	13
Figure 3. The macular degeneration simulated in the model.....	53
Figure 4. The patterns of lateral connectivity between our V1 model neurons.....	60
Figure 5. A demonstration of perceptual fill-in, using the circle.....	85
Figure 6. A full set of simulations on an aerial view of Rockefeller University.....	88
Figure 7. The Dayan caricature.....	90
Figure 8. The Einstein photograph.....	91
Figure 9. The dining room.....	92
Figure 10. Non-directional image reconstructions.....	95
Figure 11. Trial and task design.....	118
Figure 12. Optimization algorithm for stimulus generation.....	123
Figure 13. Representation of neural shape selectivity.....	128
Figure 14. Shape selectivity for two neurons recorded during the line detection task.....	136
Figure 15. Circular optima under the circle detection task.....	139

Figure 16. Heat maps with wave-like optima under the sinusoid detection task.....	142
Figure 17. Neurons preferred different 5-bar contours under different task conditions.....	147
Figure 18. Neuronal heat maps are reshaped by task.....	151
Figure 19. Temporal evolution of geometric selectivity.....	156
Figure 20. Time course of 5-bar shape selectivity.....	159
Figure 21. Facilitation by co-circular stimuli with a small radius of curvature.....	189
Figure 22. The eye position distributions during the presentation of the strongest and weakest 5-bar stimuli, for all three monkeys.....	191
Figure 23. Pearson correlation coefficients between neural responses to geometric stimuli and three features of the monkeys' eye traces.....	193

1 INTRODUCTION

1.1 THE CLASSICAL RECEPTIVE FIELD (cRF)

The fundamental functional properties of neurons in the primary visual cortex (V1) were first described by Hubel and Wiesel (Hubel and Wiesel, 1959, 1962, 1968). These findings have been expanded over the ensuing decades with the help of improved electrical recording techniques and pharmacological manipulations (Reid and Alonso, 1995; Das, 1996; Lampl et al., 2001; Martinez and Alonso, 2001; Rust et al., 2005; Touryan et al., 2005), experimental preparations using alert rather than anesthetized primates (Motter, 1993), the advent of optical recoding (Grinvald et al., 1986; T'so et al., 1990), advanced methods of anatomical tracing (Stettler et al., 2002), and the concurrent development of computing power and mathematical models (Martinez and Alonso, 2003; Simoncelli et al., 2004; Rust et al., 2005; Touryan et al., 2005). Taken together, these studies constitute an enormous effort to elucidate the properties of V1 neurons within their classical RFs (cRFs); to map the spatial distribution of these response properties across the cortical surface; and to describe the anatomical connections underlying the earliest sensory transformations in the visual cortex. A large number of studies have investigated various

properties of V1 neurons in the cRF, including: their cRF dimensions (Hubel and Wiesel, 1974; Dow et al., 1981; De Valois et al., 1982a; Webster and De Valois, 1985), spatial frequency (De Valois et al., 1982a; Foster et al., 1985; Hawken and Parker, 1987), orientation tuning bandwidth (De Valois et al., 1982b), chromatic selectivity (Hubel and Wiesel, 1959, 1968; Livingstone and Hubel, 1984; Thorell et al., 1984; T'so and Gilbert, 1988; Lennie et al., 1990; Johnson et al., 2001; Conway et al., 2002; Landisman and Ts'o, 2002), contrast sensitivity (Albrecht and Hamilton, 1982; Sclar and Freeman, 1982), direction selectivity (Hubel and Wiesel, 1959, 1974; Reid et al., 1991; Jagadeesh et al., 1993; Priebe and Ferster, 2005), disparity tuning (Poggio and Fischer, 1977; Fischer and Kruger, 1979; Poggio et al., 1985), and ocular dominance (Hubel and Wiesel, 1959, 1974). Optical imaging studies have revealed the topography of many of these properties across the surface of V1 (Grinvald et al., 1986; T'so et al., 1990; Weliky et al., 1996; Hubener et al., 1997; White et al., 2001; Yang et al., 2007), and two-photon imaging has been applied to study the microstructure of this organization (Ohki et al., 2005). Proportionate effort has been invested in settling controversies over the anatomical mechanism of orientation selectivity (Das, 1996; Lampl et al., 2001; Mooser et al., 2004) and the feed-forward convergence of inputs that generates complex cells (Martinez and Alonso, 2001, 2003).

What has emerged is a lucid picture of the very earliest cortical processing done in the visual system. The detail that has crystallized from this work, however, is in inverse proportion to its breadth. Only very specific aspects of the simplest cortical processing are well understood. For instance, comparatively little is known about the computations performed by interactions between the superficial and deep cortical layers within a column. Moreover, it has been known for decades that pyramidal cells in the superficial layers of V1 send out horizontal axon collaterals that synapse on other V1 superficial cells with similar functional properties (Gilbert and Wiesel, 1979; Rockland and Lund, 1982; Gilbert and Wiesel, 1983, 1989; McGuire et al., 1991; Malach et al., 1993; Yoshioka et al., 1996; Bosking et al., 1997; Schmidt et al., 1997; Sincich and Blasdel, 2001; Angelucci et al., 2002; Stettler et al., 2002). But the specificity of these connections is understood in broad strokes (Stettler et al., 2002), and the computations they implement are highly controversial (Lamme et al., 1998; Angelucci and Bressloff, 2006; Gilbert and Sigman, 2007). It is now well known that neurons are powerfully modulated by stimuli outside the cRF, in a set of influences grouped under the rubric “contextual interactions” (Albright and Stoner, 2002). By definition, contextual stimuli positioned outside the cRF

cannot drive neural activity on their own, but they can profoundly alter responses to stimuli within the confines of a neuron's cRF. It is also well appreciated that attention can influence visual responses throughout the cortex (Reynolds and Chelazzi, 2004; Maunsell and Treue, 2006). However, major questions remain about how contextual interactions and attention might act in tandem to shape neural responses to complex scenes, under normal behavioral conditions.

1.2 CONTEXTUAL INTERACTIONS

Since the discovery of modulatory influences from the extra-classical surround (Maffei and Fiorentini, 1976; Fries et al., 1977; Allman et al., 1985), understanding the influence of contextual stimuli on the cRF has been a major goal of visual neuroscience. In fact, calling these effects “modulatory” underemphasizes their importance. Of the total number of inputs received by V1 neurons, the feed-forward afferents that compose the bottom-up input to the cRF are a small fraction; inter-areal projections and long-range intrinsic cortical connections dominate the synapse number (Levay and Gilbert, 1976; Peters and Payne, 1993; Ahmed et al., 1994; Stepanyants et al., 2009). While the synapses made by modulatory inputs to

superficial layer neurons are weaker than the bottom-up inputs from layer 4 (Hirsch and Gilbert, 1991; Yoshimura et al., 2000; Feldmeyer et al., 2002; Chisum et al., 2003), their sheer number indicates the complexity of the computations they subserve. When detailed computational models of V1 complex cells are used to predict neuronal responses to natural stimuli, the models fail to explain the neural activity, even qualitatively (Prenger et al., 2004; Olshausen and Field, 2006; Yen et al., 2007). These models fail because they are based solely on bottom-up characterizations of the RF, while neglecting contextual inputs.

The physiological importance of stimuli from beyond the cRF is paralleled by their perceptual preeminence. The appearance of local image regions—e.g., image segments falling within the V1 cRF—is often rigidly determined by the characteristics of the surrounding visual scene. This is a ubiquitous feature of visual perception: contextual influences shape our interpretation of stimulus color, brightness, size, position, saliency, and form (Albright and Stoner, 2002). Moreover, many of the most compelling visual illusions, which yield insights into how the visual cortex processes information, are based on contextual interactions (Eagleman, 2001).

The perceptual ubiquity of contextual effects is reflected in their widespread distribution throughout visual cortex: these influences shape the responses of neurons in all visual cortical areas with small to intermediate cRF sizes. In the middle temporal (MT) visual area, where neurons are responsive to the direction of moving stimuli in their cRF, the direction of motion in the surround can either facilitate or suppress the response to the cRF stimulus (Allman et al., 1985; Xiao et al., 1997). Importantly, the sign of the interaction depends upon the relative stimulus motions. Stimulus motion in the surround that matches the direction in the cRF will powerfully suppress the neural response, while surround motion in the opposite direction can enhance it (Allman et al., 1985). Similarly, an array of contextual line segments will suppress the response of a V1 cell to an optimally oriented bar in its cRF, but the suppression diminishes if all the background bars are made orthogonal to the bar in the cRF (Gilbert et al., 1990; Knierim and Van Essen, 1992; Kapadia et al., 1995; Li et al., 2000, 2001). The same result holds for stimuli comprised of oriented gratings (Maffei and Fiorentini, 1976; Fries et al., 1977; Nelson and Frost, 1978; Sillito et al., 1995) and textures (Lamme, 1995; Lee et al., 1998), as well as for stimuli in which the contrast is induced by binocular disparity (in MT) (Bradley and Andersen, 1998), rather than motion or orientation. These neurophysiological results closely

match the perceptual saliency of the stimuli within the cRF: local image regions delineated by some feature contrast with the surrounding visual scene are much more salient than image regions that are homogenous with the background (Treisman and Gelade, 1980; Julesz, 1981; Treisman and Gormican, 1988). Various features of the intra-extra RF interactions have been probed, including their strength, sign (excitatory versus inhibitory), spatial properties, time course, contrast dependence, symmetry, and selectivity (Allman et al., 1985; Knierim and Van Essen, 1992; Kapadia et al., 1995; Lamme, 1995; Zipser et al., 1996; Lee et al., 1998; Walker et al., 1999; Kapadia et al., 2000; Li et al., 2006).

Beyond perceptual saliency, the neural correlates of many other percepts have been found. “Illusory” contours—physically non-existent but perceptually visible lines—are encoded in V2 (von der Heydt et al., 1984). Neural responses that reflect the figure-ground segregation of texture patterns, and that assign “ownership” of image boundaries to the surfaces that create them, have also been found in early visual areas (Lamme, 1995; Zhou et al., 2000). “Brightness induction”, the dependence of a surface’s apparent brightness on surrounding surfaces, has a correlate in V1 (Rossi et al., 1996; Rossi and Paradiso, 1999; MacEvoy and Paradiso, 2001). The

mechanisms underlying “color constancy”, the perceptual constancy of surface color under different conditions of illumination, may be present in V4 (Zeki, 1983a, b). Lastly, the perception of three-dimensional shape from shading has matching response properties in V1 (Lee et al., 2002).

Interestingly, for all of the perceptual phenomena listed above, at least some researchers have found neural correlates in V1 (Grosf et al., 1993; Lee and Nguyen, 2001; Wachtler et al., 2003).

More germane to our own studies are the neural responses that mirror perceptual grouping phenomena, particularly contour integration and contour saliency. Early in the 20th century, Gestalt psychologists began describing the rules by which the visual system groups stimuli into perceptually unified wholes (Wertheimer, 1923). One of the most striking of these rules is “good continuation”, whereby the visual system perceptually integrates the segments of continuous contours into a coherent object. Figure 1 demonstrates this contour integration phenomenon: all the individual line segments in each image are similar to one another, but a select few of these are perceptually grouped into a smooth contour by virtue of their relative positions and orientations (Fig. 1B, bottom panel). As a result, these same line segments are also salient: along with their compatriots in the contour,

they pop-out from the background (this is known as contour saliency).

Recently, good continuation has been more precisely reformulated by the association field (AF) (Field et al., 1993), a theoretical pattern of pairwise interactions between locally oriented image regions. The AF specifies the position and orientation of line segments that are most readily bound by the visual system (Fig. 2). It has been shown by our laboratory that some of the perceptual linkages described by the AF (along its collinear axis, Fig. 2) have matching patterns of neural activity in the superficial layers of V1 (Kapadia et al., 1995; Kapadia et al., 2000). As noted earlier, complex cells are strongly inhibited by fields of line elements that perceptually obscure an optimally oriented bar in the cRF. Bringing some of these line elements into collinear alignment with the bar in the cRF relieves the response suppression (this is known as collinear facilitation), and the bars along the collinear contour become salient (Kapadia et al., 1995; Li et al., 2006). Up to a limit, lengthening the perceived contour by bringing more bars into collinear alignment progressively increases both the neuronal facilitation and the contour saliency (Kapadia et al., 1995; Li et al., 2006). In addition, the spatial distribution of the contextual interactions between iso-oriented line segments has been mapped (Kapadia et al., 2000). Iso-oriented lines tend to be facilitatory when they are aligned along the orientation axis of a neuron;

they tend to be inhibitory when aligned along the orthogonal axis (Kapadia et al., 2000). Psychophysically, the same stimuli influence perceptual judgments about the brightness (Polat and Sagi, 1994) and orientation of contour segments (Gilbert et al., 1990; Kapadia et al., 2000), in directions that are predicted by excitatory interactions between neurons with collinear RFs. Most recently, the connection between collinear facilitation in V1 and perceptual salience has been strengthened by extracellular recordings in macaques trained to perform a collinear contour detection task. Li et al. (2006) found that the magnitude of facilitation from increasingly long collinear contours correlated with the monkey's detection performance on the same contours. Information theoretic analyses determined that an ideal observer could perform as well as the monkey in the contour detection task, using only the responses from single V1 neurons (Li et al., 2006).

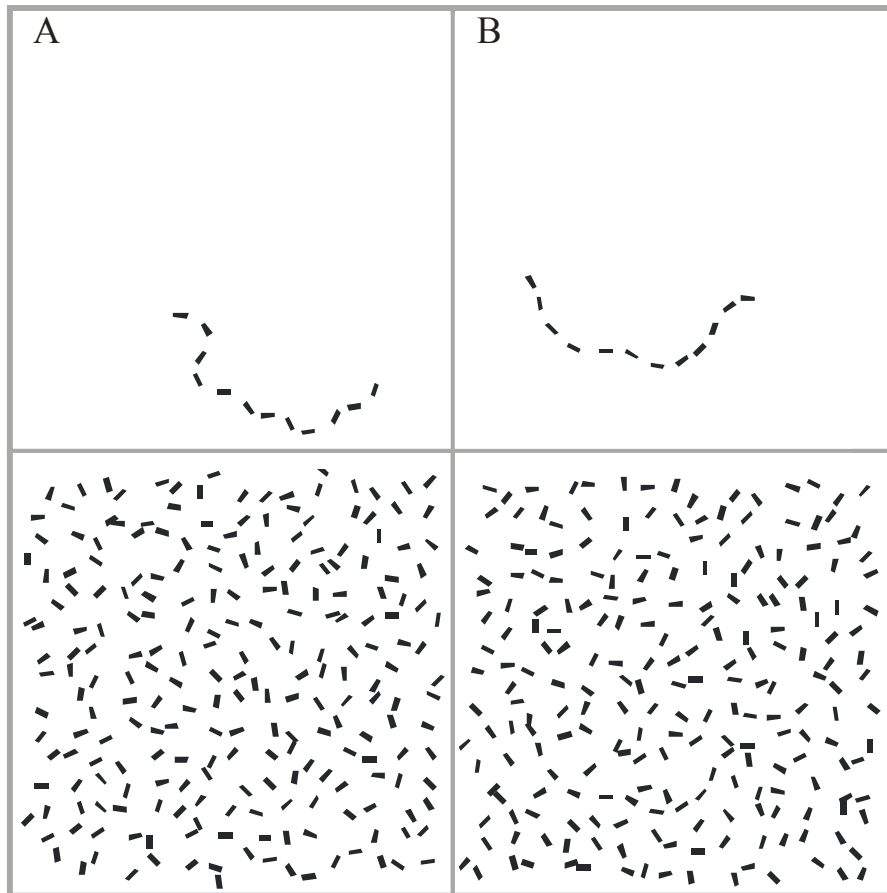



Figure 1. Perceptual grouping rules for contour integration and saliency. The top panels show the position of contours that are embedded in the bottom panels. The contour on the right, with relatively small differences in orientation between adjacent line elements, is salient even when embedded in a cluttered background of randomly oriented and positioned distracters. The difference in orientation between adjacent line elements is twice as large for the contour on the left, which does not pop-out from the background. The percepts of continuity and saliency are strongly dependent upon the relative orientations between local contour segments. Adapted from (Field et al., 1993), following (Gilbert and Sigman, 2007).

Figure 2. The association field (AF). An illustration of the theoretical AF, showing the optimal contextual interactions for the white line at the center of the figure. In the AF theory, the saliency of smooth contours derives from a broad set of local interactions between line segments. Any line segment can be perceptually linked with a range of other segments, and salient contours are characterized by local segments whose interactions sum up over the length of the contour. Given the relative positions of any two line segments, the strength of the interaction, or linkage, between them is a function of their relative orientations. Here, given the position and orientation of the white bar, we show the orientation of the line segment at each spatial position that elicits the strongest interaction with the white bar; the color of each segment specifies the strength of that linkage (warm colors: strong connection; cold colors: weak connection). In the standard AF model, the optimal interactions are between “co-circular” segments, which lie tangentially along the same circular arc. The strongest of these effects are between nearby segments that lie along a circle with an infinite radius (i.e., a straight line); as the radius of the circular arc connecting the two segments decreases, the linkage becomes weaker. Psychophysically, the strength of the interaction, or linkage, between line segments refers to the tendency of the visual system to bind them together into a perceptual whole. Physiologically, it refers to the connection strength between the neurons that are activated by the line segments. Gray scale bar: 0.5° .

Figure 2



1.3 HORIZONTAL CONNECTIONS AND THEIR ROLE IN SENSORY INTEGRATION

The anatomical connections that underlie contextual interactions are a matter of debate (Lamme et al., 1998; Angelucci and Bressloff, 2006; Gilbert and Sigman, 2007). One view is that the contextual influences in early visual areas reflect functional interactions between neurons in the same cortical area, mediated by horizontal connections between them. The alternative view is that feedback connections from areas later in the cortical hierarchy give rise to these properties, and that the responses in earlier areas simply echo computations that are performed elsewhere. A synthesis of these two ideas has emerged from some of the results presented in this thesis (see *Shape Selectivity and its Top-Down Modulation in Primary Visual Cortex*). Under this new view, also expressed elsewhere (Gilbert and Sigman, 2007), contextual influences depend upon dynamic interactions *between* intrinsic horizontal connections and feedback projections.

At least for the property of collinear facilitation, in the superficial layers of V1, there is a wealth of evidence that contextual effects depend upon the plexus of long-range horizontal connections running parallel to the cortical

surface. The surface of V1 is parceled into independent sets of vertical columns, running from the pia to the white matter, that contain neurons with response properties that are similar along a single dimension (Hubel and Wiesel, 1959, 1962; T'so et al., 1990). This property is not unique to V1, but has been observed elsewhere in the visual cortex (Albright et al., 1984; Hubel and Livingstone, 1987; Wang et al., 1996; Tsunoda et al., 2001) and even in other sensory modalities (Mountcastle, 1957; Imig and Adrian, 1977; Nelson et al., 1980; Reale and Imig, 1980). In V1, perhaps the best known are the columns in the orientation domain, which are characterized by neurons that respond to similar orientations. Anatomical tracing studies, using a variety of techniques and organisms, have demonstrated that the axon collaterals from pyramidal cells in one orientation column project preferentially to columns with a similar orientation preference (Gilbert and Wiesel, 1983, 1989; Bosking et al., 1997; Schmidt et al., 1997; Sincich and Blasdel, 2001; Angelucci et al., 2002; Stettler et al., 2002). In the tree shrew (Bosking et al., 1997), cat (Schmidt et al., 1997), and new world primate (Sincich and Blasdel, 2001), these projections also link columns with RF positions that are visuotopically aligned along the orientation axis of the RFs. These connections are ideally suited to mediate the collinear contour facilitation observed in V1, since they selectively link the same neurons that

respond to the individual line segments of the collinear contours (Chisum et al., 2003). This pattern of connections may provide the substrate for reciprocal facilitatory interactions that enhance the responses of interconnected neurons when optimal line segments are brought into their cRFs. Moreover, a psychophysical study by Li and Gilbert (2002) found a correspondence between the visuotopic extent of the horizontal connections and the maximum distance over which collinear contour segments can be perceptually grouped. In the macaque, the horizontal connections traverse about 3.5 mm of cortex from the position of the pyramidal cell body (Angelucci et al., 2002; Stettler et al., 2002). That anatomical distance corresponds to a visuotopic distance of 2° at parafoveal eccentricities. This is precisely the same distance at which the perceptual linkages between separated contour components breaks down (Li and Gilbert, 2002).

Theoretical considerations suggest V1 neurons may integrate spatial information over a much broader range of contours than the collinear stimuli used in previous experiments. Computational models of V1 used to explain perceptual grouping use interactions that include the full range geometric interactions in the AF (Grossberg and Mingolla, 1985; Guy and Medioni, 1996; Williams and Jacobs, 1997; Li, 1998; Raizada and Grossberg, 2003;

Ben-Shahar and Zucker, 2004; Ernst et al., 2004). While the horizontal connections preferentially link columns with similar orientation preferences, the selectivity of the connections is broadly tuned (Stettler et al., 2002), allowing for a wide range of connections. Moreover, statistical analyses of natural scenes have shown that co-circularity is the most likely geometric relationship between locally oriented image regions (Geisler et al., 2001; Sigman et al., 2001). That is, local image regions tend to lie tangent to circular arcs, with connecting arcs of lower curvature being most prominent. The same geometric relationships, with the same order of predominance, are described in the AF (Fig. 2). If the strength of horizontal connections between neurons is governed by Hebbian learning mechanisms, then neurons with co-circular RFs should exchange the strongest connections, since they experience coincident activation by natural stimuli. The notion that natural scenes might mold the horizontal connections into the shape of the AF is supported by the observation that juvenile monkeys (Kiorpes and Bassin, 2003) and young children (Kovacs et al., 1999) perform poorly on contour detection tasks. The ability of the visual system to link contour segments together depends on visual experience; it fully matures only after two years of postnatal experience in the monkey (Kiorpes and Bassin, 2003) and more than ten years in the human (Kovacs et al., 1999).

1.4 CORTICAL PLASTICITY AND REORGANIZATION

Cortical plasticity, both long- and short-term, has been used as a tool to study the circuitry underlying the integrative mechanisms in V1.

Cortical reorganization, resulting from permanent losses of input, provides a foothold for understanding normal mechanisms of sensory processing, by relating the resultant changes in neural activity with the anatomical changes that occur over the same time course. Likewise, short-term plasticity, induced by manipulations of the sensory input to the cRF, can yield insights into the dynamic mechanisms of cortical processing.

The irreversible loss of afferent input induces profound reorganizations of cortical function, in visual (Gilbert et al., 1990; Kaas et al., 1990; Heinen and Skavenski, 1991; Chino et al., 1992; Gilbert and Wiesel, 1992; Darian-Smith and Gilbert, 1994; Chino et al., 1995; Darian-Smith and Gilbert, 1995; Das and Gilbert, 1995a; Schmid et al., 1995; Calford et al., 2000; Calford et al., 2003; Giannikopoulos and Eysel, 2006; Keck et al., 2008), motor (Sanes et al., 1988; Sanes et al., 1990), somatosensory (Rasmusson, 1982; Merzenich et al., 1983a; Merzenich et al., 1983b; Merzenich et al., 1984;

Calford and Tweedale, 1988; Rasmusson, 1988; Calford and Tweedale, 1991), and auditory areas (Robertson and Irvine, 1989; Rajan et al., 1993). In V1, lesions that destroy corresponding regions of the retina in both eyes silence a circumscribed cortical region, the lesion projection zone (LPZ), where neurons have lost the bottom-up inputs that normally drive their function (Gilbert and Wiesel, 1992). However, within minutes of creating laser-induced lesions, neurons just inside the LPZ boundary expand their cRFs by about five-fold in area (Gilbert and Wiesel, 1992; Darian-Smith and Gilbert, 1995). The expansions bring large regions of the undamaged retinal surface into their purview, while maintaining tuning for orientation, movement direction, and binocularity. Over the ensuing weeks to months, neurons up to 3.5 mm inside the LPZ boundary progressively regain stimulus-driven activity, by shifting their RFs from inside the lesion-induced blind spot, or scotoma, to the surrounding area of visual space, the periscotoma. Temporally, the reorganization occurs like a traveling wave of recovery, starting at the LPZ boundary and moving inward over the span of about 4 weeks (Giannikopoulos and Eysel, 2006). Levels of stimulus-driven and spontaneous activity are highest along the vanguard of the wave, but they return to normal over the next several weeks (Giannikopoulos and Eysel, 2006). Similarly, the RFs of LPZ neurons are largest during their

initial phase of recovery, and while their dimensions decrease over the ensuing weeks, they remain larger than normal RFs at the same eccentricity even one year after the lesion (Giannikopoulos and Eysel, 2006). Moreover, while orientation tuning is qualitatively preserved, the mean width of the recovered tuning curves increases linearly with distance from the LPZ border. The orientation selectivity for neurons at the border of the LPZ is normal, but it broadens by more than a factor of two for neurons at the furthest reaches of the reorganization, about 3.5 mm into the LPZ (Giannikopoulos and Eysel, 2006). The same pattern is observed transiently for RF sizes during the initial phase of recovery: RF sizes are larger for neurons deeper in the LPZ, before the subsequent process of RF narrowing.

The substrate for this reorganization is cortical in origin, and is very likely the supragranular network of horizontal connections (Darian-Smith and Gilbert, 1995). Even after the neuronal recovery in the V1 LPZ has matured, the corresponding region of the LPZ in the lateral geniculate nucleus (LGN) of the thalamus remains silent (Eysel et al., 1981; Darian-Smith and Gilbert, 1995). Therefore, the major source of visual input that drives normal activity in V1 never recovers. Moreover, the lateral spread of geniculocortical afferents (connections from the LGN to V1) is about 1 mm in radius

(Darian-Smith and Gilbert, 1995). The limits of this envelope, which does not expand following retinal lesions, rules out thalamic inputs from healthy retinal regions as the source of the recovered activity deep in the LPZ. The mechanism of reorganization must therefore be cortical. The most likely anatomical substrate is the pattern of long-range horizontal connections in V1. The horizontal projections running from neurons outside the LPZ to neurons inside it sprout dramatically following retinal lesions (Darian-Smith and Gilbert, 1994). The anatomical spread of these connections, both in normal and reorganized cortices, closely corresponds to the lateral spread of reorganization into the LPZ. Moreover, the density of the horizontal projections into the LPZ changes as a function of time following the retinal lesions, along a trajectory that parallels the time course of the functional recovery (Yamahachi et al., 2009). Optical imaging studies demonstrate that the pattern of orientation columns in the LPZ is preserved after making retinal lesions (Das and Gilbert, 1995a). This result implicates a set of connections, like the horizontal collaterals, that respect the preexisting functional architecture of the cortex. Lastly, the increased RF size and the broadening of orientation tuning with distance into the LPZ offer precise mechanistic hints. Both observations are consistent with reorganization via a “bowtie-shaped” pattern of projections (like the one in Fig. 2) from neurons

just outside the LPZ to deprived cells inside it. Neurons deep in the LPZ will receive connections only from the ends of the bowtie pattern, which are broad along both spatial and orientation dimensions. Conversely, neurons closer to the LPZ border will receive the strongest connections from nearby neurons, close to the center of the pattern, where the spatial positions and orientation preferences are narrower in range. (See Fig. 2. Here, the center of the pattern represents an LPZ neuron gathering inputs from the surrounding field. For neurons deep in the LPZ, most of the field is inactive, because it corresponds to lesioned areas of visual space. For neurons nearer the border, more of the field toward the center is still available—i.e., undamaged.)

The cortical reorganization described above involves irreversible, long-term changes in anatomy and cortical function. Analogous observations, however, have been reported from studies using short-term, dynamic reorganization. “Artificial scotomata” (artificial retinal lesions) can be created by depriving V1 RFs of visual input while stimulating the extra-classical RF surround (Gerrits et al., 1966; Pettet and Gilbert, 1992). This mimics the loss of afferent input in circumscribed retinal areas, if the deprivation is maintained (in anesthetized animals) over the course of minutes. When the bottom-up input to the center of V1 RFs is temporarily “ablated” in this way, RFs

undergo a five-fold expansion in area, which can be reversed by visually stimulating the center the RF (Pettet and Gilbert, 1992). Like the permanent reorganization described earlier, the mechanism for this effect is also cortical in origin (Volchan and Gilbert, 1994), and it likely reflects changes in the strength of horizontal connections between superficial layer V1 neurons (Das and Gilbert, 1995b).

1.5 PERCEPTUAL FILL-IN

The processes of long- and short-term cortical reorganization, introduced in the previous section, are associated with profound alterations in visual perception. Most remarkably, human subjects who suffer from pathological (Gerrits and Timmerman, 1969; Schuchard, 1993, 1995; Zur and Ullman, 2003) or self-induced (Craik, 1966) retinal damage experience “perceptual filling-in”: the missing regions of the visual field, which fall on non-responsive regions of the photoreceptor mosaic, are perceptually filled-in by the brain. Artificial scotomata also cause similar filling-in processes (Gerrits et al., 1966; Ramachandran and Gregory, 1991) and can influence perceptual discriminations (Kapadia et al., 1994). Subjects with retinal damage do not simply ignore the ablated portion of the visual field, since their perception in

those regions contains consciously discernible features that are shaped by the visible portions of the surrounding scene (Zur and Ullman, 2003). Moreover, the perceptual fill-in is not caused by eye movements and a temporal integration of the image, since the effect occurs in fixating subjects. In the normal brain, with its full complement of sensory inputs, synthesizing information over the visual scene creates percepts like contour integration and illusory contours. Some of the same processes may underlie perceptual fill-in after retinal damage, including the perceptual completion of lines, gratings, and textures that span a scotoma (Zur and Ullman, 2003).

1.6 COGNITIVE (TOP-DOWN) INFLUENCES ON SENSORY PROCESSING

The processing of complex stimuli relies on attentional control, since the brain cannot adequately process all the information impinging upon it from multiple, concurrent sensory streams (Kastner and Ungerleider, 2000; Chun and Marois, 2002). Instead, attention dynamically tailors sensory systems to improve their processing of specific, behaviorally relevant stimuli. A host of attentional influences have been described in the visual (Corbetta and Shulman, 2002; Reynolds and Chelazzi, 2004) and auditory cortices (Fritz et

al., 2007). In the visual cortex, mechanisms of attentional control have been classified according to the domain (space, feature, or object) in which they act. Among these, spatial attention—focusing on specific locations in the visual scene—and feature-based attention—highlighting stimuli that share a particular feature throughout the visual field—are the best studied (Reynolds and Chelazzi, 2004; Maunsell and Treue, 2006). Behaviorally, spatial attention enhances perceptual discriminations within the envelope of the attended area (Posner et al., 1977; Bashinski and Bacharach, 1980; Handy et al., 1996; Lee et al., 1999). Likewise, feature-based attention improves the detection of stimuli that possess the attended feature (Maunsell and Treue, 2006). Correspondingly, attention modifies neural responses. For instance, when spatial attention is directed to the location of a neuron's cRF, the cell's responses to the stimuli within its cRF are enhanced (Mountcastle et al., 1987; Spitzer et al., 1988; Motter, 1993; Treue and Maunsell, 1996; McAdams and Maunsell, 1999a). More importantly, spatial attention also provides a selective role. In visual areas with large RFs that can accommodate multiple stimuli, neural responses are driven preferentially by the stimulus in the attended location; other stimuli in the RF are effectively ignored (Moran and Desimone, 1985). In addition to visual location, attentional selection can be based on other stimulus attributes. When a

monkey attends to a particular stimulus feature, neurons containing multiple stimuli in their RF will respond selectively to the stimulus that contains the attended feature. If the selected stimulus matches a neuron's tuning properties for shape, color, orientation, or motion direction, then the attention will enhance the neural response to the ensemble of stimuli in its RF; if there is a mismatch between the stimulus and the neural tuning properties, the attentional impact will be inhibitory (Chelazzi et al., 1993; Treue and Maunsell, 1996; Chelazzi et al., 1998; Reynolds et al., 1999; Treue and Trujillo, 1999; Recanzone and Wurtz, 2000; Chelazzi et al., 2001; Reynolds and Desimone, 2003). These attentional effects typically act like a gain modulation, by uniformly boosting or suppressing neural responses to stimuli in the RF, without altering the width of neural tuning curves (Reynolds and Chelazzi, 2004; Maunsell and Treue, 2006). Mathematically, this response enhancement can be comprehensively modeled by generalizing the equation used to describe surround inhibition (Reynolds and Heeger, 2009), the well-known contextual effect whereby neural responses are normalized by the activity in the surrounding cortical population.

In addition to spatial locations and stimulus features, attention can encompass entire objects. This object-oriented mode of attention may highlight the full collection of an object's attributes, including the object's color, shape, orientation, motion, and the spatial positions it spans (Duncan, 1984; Egly et al., 1994; Roelfsema et al., 1998; O'Craven et al., 1999; Blaser et al., 2000). Simultaneous perceptual discriminations can easily be made about different features of the same object, whereas making judgments about two separate objects is more difficult (Duncan, 1984). Similarly, attending to one part of an object can facilitate discriminations on other segments of the same object (Egly et al., 1994). These attentional boons are reflected in neural activity. In V1, attention can be distributed over the length of a task-relevant contour, so that neurons whose cRFs fall along the contour are collectively facilitated (Roelfsema et al., 1998). Additionally, sensory responses in MT are selective for the motion of attended objects, compared to non-attended but spatially overlapping distracters, even when the motion of the attended objects is not task-relevant (O'Craven et al., 1999).

Early experiments suggested that attention modulates neural responses in extrastriate cortical areas, but not in V1 (Moran and Desimone, 1985). When stimuli are presented in isolation to the RFs of V1 neurons, attention to those

stimuli modulates the neural response weakly, compared to the attention effects seen in V4 (Haenny and Schiller, 1988). Experiments by Motter, however, demonstrated that the magnitude of attentional modulation throughout the visual cortex depends on the complexity of the visual scene (Motter, 1993). He found that attention had the strongest impact on neural responses when multiple stimuli were present, rather than just a single stimulus in the cRF. This makes sense from ecological and information theoretic perspectives. If the purpose of attention is to harness limited computational resources for processing complicated stimuli, then attentional influences should be most prominent for complex sensory scenes and least relevant for unnaturally simple ones.

Motter's findings, which provided the first hint of an interaction between attention and contextual effects, have since been expanded (Ito and Gilbert, 1999; Crist et al., 2001; Li et al., 2004, 2006). When complex stimuli are used, attentional influences that equal those in extrastriate areas are observed in V1. Attention, and the influence of perceptual task, can amplify or otherwise shape contextual effects on V1 neurons (Ito and Gilbert, 1999; Crist et al., 2001; Li et al., 2004, 2006). In particular, the collinear facilitation described earlier in this Introduction can be enhanced by more

than 40% when trained monkeys perform a contour detection task (Li et al., 2006). The interactions between perceptual task and the contextual effects in V1 suggests a mode of attentional modulation that is more nuanced than the response normalization described by Reynolds and Heeger (2009). Recently, recordings in MT (Ghose and Beryl, 2010) and V4 (David et al., 2008) have revealed a similarly sophisticated mode of attentional enhancement, one that fundamentally reshapes neural tuning properties rather than simply adjusting the response gain.

Interestingly, the recent studies showing qualitative changes in neuronal tuning functions were observed under behavioral tasks that manipulated monkeys' expectations (David et al., 2008; Ghose and Beryl, 2010). While the effect of attention on visual processing has been extensively investigated, the corresponding influence of expectation has been neglected (Summerfield and Egner, 2009). Attention lifts the computational burden on sensory systems to a yoke they can support. Expectation, on the other hand, may help resolve the sensory ambiguities that arise from the projection of a four-dimensional world onto a two-dimensional receptor surface (Summerfield and Egner, 2009). Prior information about the statistical regularities in a known environment could inform the perceptual interpretation of noisy or

ambiguous stimuli. Indeed, observers are better able to detect sensory stimuli in environments where they are more likely to occur (Bar, 2004). Moreover, visual discrimination can be improved by priming subjects with stimuli that increase the statistical frequency of behaviorally relevant stimulus features (Maljkovic and Nakayama, 1994). Theories of orbitofrontal cortex (OFC) function, proposed by Bar and colleagues, postulate that OFC may send feedback signals to visual areas to bias sensory processing towards the recognition of objects that are likely to occur in the visual environment (Bar et al., 2001; Bar et al., 2006; Fenske et al., 2006). Under this view, OFC uses the sensory input from the current environment to predict which stimuli are probably present, based on prior joint probability distributions. OFC is well-suited for the task (Summerfield and Egner, 2009): it has visual responses to objects and visual scenes; it integrates sensory inputs from different modalities—allowing it to build probability distributions of stimulus co-occurrences; and it may be involved in long-term memory retrieval, to access stored statistical relationships. Ullman has proposed a very similar mechanism for object recognition, whereby visual information gleaned from a bottom-up processing stage is used to predict which local features are likely present in an image (Ullman, 2007; Epshtein et al., 2008). In this countercurrent stream model, the local

image predictions are fed back to earlier sensory cortices to bias their detection of sensory stimuli. Theoretical work suggests that this feedback step is important for recognizing the detailed features that compose larger objects, like the mouth, eyes, and ears of a face (Epshtein et al., 2008). Perceptual fill-in evidence from human subjects with retinal damage supports these ideas. When geometric stimuli are partially occluded by a retinal lesion, subjects can perceptually fill-in the stimuli and will report the presence of a corresponding geometric figure. However, when subjects are primed to expect a letter, their perception is completely different: rather than seeing a geometric figure, they perceive the letter whose shape is consistent with the visible portions of the stimulus outside the lesion (Schuchard, 1995). This result suggests that expectation can mold perceptual fill-in. If the horizontal connections in V1 participate in the perceptual completion, the result further suggests that expectation may dynamically reshape the connectivity pattern to match the most probable stimulus relationships.

1.7 CORTICAL SOURCES OF TOP-DOWN CONTROL

The deployment of cognitive top-down control, and its modulation of visual processing, is likely controlled by distributed networks of cortical and subcortical (superior colliculus and thalamic pulvinar nuclei) brain regions (Kastner and Ungerleider, 2000; Miller and Cohen, 2001; Corbetta and Shulman, 2002; Reynolds and Chelazzi, 2004). The cortical regions that shape the allocation of spatial attention are the best studied. In the macaque, strong evidence has accumulated that the frontal eye field (FEF) and the parietal cortex, specifically the lateral intraparietal area (LIP), are major determinants of spatial attention. Both regions exchange reciprocal connections with each other, with prefrontal cortex (PFC, a major control center for goal-directed behavior), and with extrastriate visual areas (Blatt et al., 1990; Felleman and Van Essen, 1991; Webster et al., 1994; Miller and Cohen, 2001). In primates, there is a strong behavioral association between spatial attention and eye movements toward the attended stimuli; correspondingly, both the FEF and LIP send projections to the oculomotor system that controls eye movements (Schiller and Tehovnik, 2005). Moreover, the neural activity in both cortical areas encodes multiple signals pertaining to behavioral task, visual salience, and saccadic motor plans. In

particular, LIP may contain a “priority map”, representing the behavioral importance of each point in space as a function of both sensory (bottom-up) inputs and cognitive (top-down) demands (Bisley and Goldberg, 2010). [The integration of bottom-up and top-down information is a key feature for an attentional control area, since both influences are known to effect our attention (Corbetta and Shulman, 2002).] The distribution of activity over the LIP map could therefore shape how attention is allotted over visual space (Bisley and Goldberg, 2010). The repertoire of response properties expressed by LIP neurons supports this view. First, LIP neurons respond more strongly to attention-grabbing, visually salient stimuli, regardless of the specific feature that imbues a stimulus with salience (Gottlieb et al., 1998; Kusunoki et al., 2000; Balan and Gottlieb, 2006). Second, LIP responses are enhanced when the stimuli in their receptive fields are related to behavioral demands, as when they represent the target of a search task (Ipata et al., 2006a; Buschman and Miller, 2007), and they may encode the relative value of task-related rewards (Dorris and Glimcher, 2004; Sugrue et al., 2004). Moreover, there is an interaction between visual salience and task demands. While LIP neurons are normally responsive to salient stimuli, they can be inhibited by pop-out stimuli that they are behaviorally irrelevant; the ability to ignore salient distracters is correlated with the strength of this

inhibition (Ipata et al., 2006b). More generally, it has been demonstrated that LIP neurons perform a linear combination of inputs relating to bottom-up, top-down, and saccade-related signals (Ipata et al., 2009). Furthermore, using the increase in perceptual contrast sensitivity engendered by attention, it has been directly shown that LIP neurons encode the locus of spatial attention (Bisley and Goldberg, 2003), whether the attention is elicited by visual (bottom-up) or cognitive signals.

Many of the response features that contribute to the construction of the priority map in LIP are also reflected in FEF (Thompson and Bichot, 2005). For instance, FEF neurons also respond preferentially to stimuli when they are task-relevant targets (rather than irrelevant distracters) of a visual search task (Schall et al., 1995; Thompson et al., 1996; Thompson et al., 1997; Thompson et al., 2005b), whether or not the task requires a saccade to the target stimulus. In both LIP and FEF, when a discrepancy is introduced between the spatial location of attention and the direction of a monkey's saccade—behavioral features that are normally tightly coupled—neurons tend to encode the location of the stimulus rather than the saccade direction (Gottlieb and Goldberg, 1999; Murthy et al., 2001). These results suggest that the response enhancements in LIP and FEF really reflect attention,

rather than overt or latent saccade planning. In FEF, neural activity is also correlated with performance on attention-demanding search tasks, and the magnitude of neural selectivity for target versus distracter shapes coincides with their perceptual similarity/dissimilarity (Bichot et al., 2001b; Thompson et al., 2005a). FEF neurons also respond to both bottom-up and top-down aspects of stimulus relevance, with a selective preference for overriding top-down signals (Bichot et al., 2001a), although they are not inhibited by salient distracters. Moreover, microstimulation studies have shown that FEF activity can induce improvements in perceptual discrimination (Moore and Fallah, 2001, 2004), and matching modulations of V4 responses (Moore and Armstrong, 2003), that mimic the response gain and stimulus selection facets of spatial attention.

LIP and FEF have closely correspondent patterns of anatomical connections and response properties, and their likely analogs in the human brain are both major components of a cortical network that mediates task- and stimulus-dependent attention (Corbetta and Shulman, 2002). Lesions of both parietal and frontal regions cause deficits in attention to the contralateral visual field (known as visuospatial neglect), in the human (Heilman and Valenstein, 1972; Damasio et al., 1980; Vallar and Perani, 1986) and monkey (Welch and

Stuteville, 1958; Latto and Cowey, 1971; Lynch and McLaren, 1989; Gaffan and Hornak, 1997). It is therefore unclear how these two regions interact, or which facets of attentional control originally arise from each area. Recent experiments, however, have begun to tease apart the relative contributions of both areas, by determining the time course of neural selectivity for attended stimuli in each area. Simultaneous recordings in LIP and FEF reveal that LIP neurons encode the salience of pop-out stimuli before FEF cells, whereas FEF cells begin preferentially responding to task-relevant target stimuli before LIP in a visual search task (Buschman and Miller, 2007). These results hint that bottom-up stimulus salience may be first represented in LIP, while FEF may be the first area to reflect top-down attention during difficult behavioral tasks. The latter idea is supported by simultaneous recordings of spiking activity and local field potentials (LFPs) in FEF. The selectivity for the location of a target stimulus in a search task emerges first in FEF spiking activity and not until 30 ms later in the LFP (Monosov et al., 2008). Since spiking activity signals the outputs of cortical computations, and the LFP reflects dendritic inputs, the earlier emergence of selectivity in neuronal spikes implies that the selectivity is computed in FEF.

While a large body of work has investigated the roles that LIP and FEF might play in allocating spatial attention, much less is known about the cortical areas that might control other cognitive influences, or even how those influences might be deployed. The control centers for feature- and object-based attention, and for expectation, remain more elusive. However, evidence from functional magnetic resonance imaging (fMRI) and electrophysiology suggests that the cortical regions involved in visuospatial attention may overlap with those involved in other forms of attention and top-down control (Corbetta and Shulman, 2002). For instance, one electrophysiology study showed a form of feature-based selectivity among LIP neurons, in which the relevant feature was the direction of stimulus motion (Freedman and Assad, 2009). After training monkeys to differentiate between two categories of stimulus motion, a subset of LIP neurons became selective for one of the two motion categories. When the monkey made a perceptual discrimination about the motion of a stimulus placed far outside the RF of a category-specific LIP neuron, the neuron responded preferentially when the direction of motion matched the cell's own preference. This result suggests that neurons with the same category selectivity, corresponding to the full range of task-relevant visual space, may have been collectively activated throughout LIP. Similar patterns of neural

activity could potentially be used to implement feature-based attention. (The authors did not examine the influence of attention on LIP category selectivity, however.) Another study found that LIP neurons may encode an expectation signal for moving objects that are temporarily occluded (Assad and Maunsell, 1995); yet another found that cells in LIP can encode color information after training on a task that involves associating different colors with specific saccade directions (Toth and Assad, 2002). Similar observations have been made in FEF (Bichot et al., 1996). Given the role of LIP and FEF in spatial attention, these experiments suggest that non-spatial forms of selectivity in LIP and FEF neurons may induce an integrated mode of attention that acts over a combination of spatial and non-spatial domains. Finally, fMRI studies have shown overlapping patterns of brain activation when human subjects perform tasks involving either spatial or feature-based attention (Shulman et al., 1999), further suggesting that some of the same cortical machinery may be used for more than one type of attentional control.

Another cortical area that features prominently in non-spatial top-down control is the prefrontal cortex (PFC). This region of the brain plays a very broad role in “executive function” and organizing complicated behaviors; in particular, it may contribute crucially to orchestrating the volley of

interacting sensory and cognitive processes in the brain (Miller and Cohen, 2001). In addition to LIP and FEF, the PFC may play a major part in managing the influence of cognitive control over sensory processing. The importance of the PFC in this respect has been highlighted by its interactions with inferotemporal cortex (IT) (Fuster et al., 1985; Tomita et al., 1999), during tasks involving attention, working memory, and memory recall. Neurons in IT exhibit an arguably feature-based form of attention, whereby neuronal responses are either enhanced or suppressed depending on whether a task-relevant cue shape matches the neuron's stimulus preferences (Chelazzi et al., 1993; Chelazzi et al., 1998). These attentional enhancements in IT responses are even present during the delay period in a delayed-match-to-sample task, a neuronal correlate of working memory. However, transient inactivation of the PFC reduces IT firing rates and abolishes stimulus selectivity during this delay period (Fuster et al., 1985), obliterating the memory trace. Further evidence suggests that, in IT, task-dependent response enhancements to stimuli that match a cue require PFC-IT interactions (Miller et al., 1996), because PFC maintains the working memory for the cue. It has also been shown that category-specific selectivity in IT, which emerges after monkeys memorize associations between categories of cue shapes and particular "choice" shapes, arises from

interactions with PFC (Tomita et al., 1999). In this study, PFC activity encoded the learned relationships between stimulus shapes, and it was sufficient to activate the resulting category specificity in IT, even in the absence of other inputs.

Moreover, the top-down influences from PFC to IT may extend to the realm of object recognition. As noted earlier, a recent theory suggests that a particular area in the PFC, the orbitofrontal cortex (OFC), may contribute to object recognition by manipulating extrastriate visual processing as a function of expectation (Bar et al., 2006; Fenske et al., 2006). The theory is based on evidence from fMRI (Bar et al., 2001) and magnetoencephalography (MEG) (Bar et al., 2006) that: 1.) OFC rapidly acquires low spatial frequency (LSF) information about the visual scene (possibly from the dorsal magnocellular pathway); and 2.) OFC uses the LSF information to bias responses in the temporal cortex for the detection of objects whose presence is conditionally probable, given the available information from the incoming image. Interestingly, this role for OFC in object recognition is also closely related to memory recall: it requires the retrieval, activated by the current image, of learned statistics about stimulus-stimulus co-occurrences. This coincides with proposals that an integral

function of the PFC is to maintain working memory and to retrieve long-term memories (Kastner and Ungerleider, 2000; Miller and Cohen, 2001), and with the evidence for PFC-IT interaction described in the previous paragraph.

In the context of the experimental work described in this thesis (Chapter 3: *Shape Selectivity and its Top-Down Modulation in Primary Visual Cortex*), the interactions between PFC (including OFC) and IT may be particularly important. The PFC, IT, and V1 form an interconnected network. The PFC, especially OFC, makes strong reciprocal connections with IT (Felleman and Van Essen, 1991; Cavada et al., 2000; Miller and Cohen, 2001). Moreover, both IT and PFC/OFC are connected with V1 through serial pathways of connections in the ventral stream (Felleman and Van Essen, 1991; Tanaka, 1996; Miller and Cohen, 2001). This network of anatomical connections suggests that PFC-IT activity may be reflected in the feedback inputs to V1 neurons. Furthermore, the behavioral task used in our electrophysiology study may have activated the object- and expectation-related activity in the PFC-IT network. Given the broad roles of PFC, and specifically OFC, in controlling goal-directed behavior, this region is almost certainly involved in the task performed by our monkeys. Since the task required monkeys to hold

the shape of a cue in working memory, the involvement of PFC may have been especially important. Moreover, IT has a well-established role in object recognition (Tanaka, 1996), one that was likely engaged by our contour detection/discrimination task. Even further, the delayed-match-to-sample structure of the behavioral task may have recruited expectation mechanisms in the brain. Expectation could have wielded an influence similar to the theoretical effect described above, whereby OFC modulates IT responses as a function of stimulus-derived expectations. If the OFC-IT interaction is important for object recognition, then it might be possible to activate it with task-related signals, in addition to the proposed bottom-up mode of activation—just as the saliency maps in LIP and FEF reflect both bottom-up and top-down signals. Given the dense interconnections between PFC-IT-V1, the top-down modulations deployed by the PFC/OFC would probably not be restricted to the temporal cortex, but could spread to the occipital pole, and might explain the dynamic task-related responses of V1 neurons reported here.

1.8 CURRENT WORK

The preceding Introduction sketches the connections between sensory integration in normal and reorganized cortices; between neural responses and perception; and between integrative mechanisms and cognitive influences like attention and expectation. The work presented below explores these connections in more detail. First, the computational modeling study establishes a lower bound for the perceptual recovery that might emerge during cortical reorganization, if the reorganization is mediated by long-range horizontal connections in V1. The model is based on the pattern of interactions described by the AF. This pattern is the basis for much of the current thinking about the integrative properties of V1, so we carried out extracellular recordings in the behaving macaque to examine the geometric interactions that are actually expressed in V1. We observed striking new patterns of stimulus selectivity in V1, patterns which were consistent with the AF over the neural population but which changed dynamically with perceptual task and expectation. Given the functional architecture of V1, and the time course with which V1 neurons express task-dependent geometric selectivity, our experimental data are consistent with a model in which top-down interactions dynamically reshape the horizontal connections between

V1 cells. Our findings suggest that the integrative properties of V1 underlie the perceptual grouping of aligned contour components. They further suggest that the same functional attributes underlie perceptual fill-in following retinal degeneration. The mechanistic connections between contour saliency and perceptual fill-in, and the similarity in their modulation by task and expectation, imply that the top-down modulation of intrinsic connectivity is a general processing mechanism in V1.

The following two chapters are drawn from work that is already published or is currently under publication. The coauthors on those manuscripts provided important mentorship roles, while I designed and conducted the experiments; conceived and performed the data analyses; and wrote the papers.

Throughout this thesis, I use the first person plural (e.g., “we”) to acknowledge that I did not singularly arrive upon the ideas or results presented here, as if in a vacuum, but through an evolving exchange of ideas between my collaborators and myself.

2 A COMPUTATIONAL MODEL OF PERCEPTUAL FILL-IN FOLLOWING RETINAL DEGENERATION

Justin N. J. McManus, Shimon Ullman, and Charles D. Gilbert

Journal of Neurophysiology 99, 2086-2100 (2008)

2.1 SUMMARY

The ablation of afferent input results in the reorganization of sensory and motor cortices. In the primary visual cortex (V1), binocular retinal lesions deprive a corresponding cortical region [the lesion projection zone (LPZ)] of visual input. Nevertheless, neurons in the LPZ regain responsiveness by shifting their receptive fields (RFs) outside the retinal lesions; this reemergence of neural activity is paralleled by the perceptual completion of disrupted visual input in human subjects with retinal damage. To determine whether V1 reorganization can account for perceptual fill-in, we developed a neural network model that simulates the cortical remapping in V1. The model shows that RF shifts mediated by the plexus of spatial- and orientation-dependent horizontal connections in V1 can engender filling-in that is both robust and consistent with psychophysical reports of perceptual completion. Our model suggests that V1 reorganization may underlie

perceptual fill-in, and it predicts spatial relationships between the original and remapped RFs that can be tested experimentally. More generally, it provides a general explanation for adaptive functional changes following CNS lesions, based on the recruitment of existing cortical connections that are involved in normal integrative mechanisms.

2.2 INTRODUCTION

The removal of afferent input to the motor, somatosensory, auditory, and visual cortices triggers reorganization in those areas, with concomitant changes in neural function and behavior. In particular, the primary visual cortex (V1) undergoes topographic reorganization after the removal of sensory input from the retina (Gilbert et al., 1990; Kaas et al., 1990; Heinen and Skavenski, 1991; Chino et al., 1995; Darian-Smith and Gilbert, 1995; Das and Gilbert, 1995a; Schmid et al., 1995; Calford et al., 2000; Baker et al., 2005). When focal binocular retinal lesions are made in the adult cat or macaque, the loss of sensory input silences a region of the lateral geniculate nucleus (LGN) while spurring reorganization in the corresponding cortical region, known as the lesion projection zone (LPZ) (Darian-Smith and Gilbert, 1995). Neurons ≤ 4 mm from the boundary of the LPZ regain

responsiveness by shifting their receptive fields (RFs) from inside the lesion, or scotoma, to the region of visual space immediately outside the scotoma perimeter (the periscotoma).

Interestingly, these profound changes in V1 topography and neural function are paralleled by equally impressive perceptual consequences. Human subjects with artificially induced retinal lesions or with macular degeneration do not receive visual input from the damaged portions of their retinae, yet their visual perception is continuous and complete (Craig, 1966; Gerrits and Timmerman, 1969; Schuchard, 1993, 1995; Burke, 1999; Zur and Ullman, 2003). Rather than perceiving holes or blind spots in their vision, these individuals can often “see” portions of the visual scene that are masked by their retinal lesions. Moreover, this phenomenon—known as perceptual fill-in or perceptual completion—is not the result of passively ignoring the missing regions of the retinal image, but rather it is an active filling-in of the lost regions of the visual field (Zur and Ullman, 2003). Individuals with retinal damage do not simply ignore the lesion-induced blind spots, because they really “see” objects at the locations corresponding to their lesions, and their percepts within the scotomata can even contain anomalies and distinguishing features.

Convincing evidence for the active nature of perceptual completion comes from psychophysical studies on subjects with age-related macular degeneration (AMD), a common disease among the elderly in which lesions progressively develop on the retinae (Zur and Ullman, 2003). When subjects with advanced AMD are shown regular arrays of dots that are partially occluded by a very large retinal lesion (spanning 5° - 7° in radius), the subjects can actually count the number of dots that fall within the lesion. When the same subjects are shown two-dimensional gratings, the grating lines appear to run continuously through the scotoma, but the perceptually reconstructed lines in the scotoma appear blurrier and at lower contrast than the lines outside the scotoma. Furthermore, this non-uniformity in their perception depends upon the attributes of the grating lines: subjects experience better perceptual fill-in as the spatial frequency of the gratings is increased. Under unfavorable conditions, advanced AMD subjects even report gaps in the perception of objects they can only partially fill-in. If AMD subjects have a continuous perception only because they ignore the holes in their vision, then they should never even notice stimuli within their scotomata, much less be able to report distinguishing characteristics about

them, and their perception should not depend on the attributes of the environment they ignore.

The evidence for an active mode of perceptual completion is clear, and the anatomical and functional changes during V1 reorganization are well described. What remains unclear, however, is how the physical changes in cortical structure and function engender perceptual completion. Nevertheless, it seems likely that the substrate for cortical reorganization (and thus perceptual fill-in) is intrinsic to the cortex. The degree of reorganization observed along the visual pathway antecedent to the cortex, most notably at the LGN, is much more limited than the cortical reorganization (Eysel, 1982; Darian-Smith and Gilbert, 1995). And since the lateral spread of geniculocortical afferents to V1 is insufficient to account for the extent of the observed reorganization, the next most likely substrate is the plexus of long-range horizontal connections in V1 (Darian-Smith and Gilbert, 1995; Calford et al., 2003). V1 reorganization is accompanied by a sprouting of these horizontal axon collaterals from outside to inside the LPZ, whereby the preexisting pattern of connections is intensified and reinforced (Darian-Smith and Gilbert, 1994). Moreover, intrinsic optical imaging in the cat demonstrates that reorganization preserves the orientation column

architecture in V1 (Das and Gilbert, 1995a). Since the horizontal projections preferentially link orientation columns with similar orientation preference (Gilbert and Wiesel, 1989; Bosking et al., 1997; Stettler et al., 2002; Chisum et al., 2003; Shmuel et al., 2005), this preservation of orientation column structure further suggests a role for the horizontal connections in the remapping (Das and Gilbert, 1995a).

If the horizontal connections play an important role in cortical reorganization—and if the reorganization respects their preexisting pattern of connectivity—then it is tempting to speculate that the reorganized cortex should retain its ability to perform the computations implemented by those connections. A prominent role that has been proposed for the horizontal connections in normal V1, based on evidence from electrophysiological, anatomical, and psychophysical studies, as well as from theoretical considerations, is contour integration (Kapadia et al., 1995; Li, 1998; Kapadia et al., 2000; Sigman et al., 2001; Li and Gilbert, 2002; Stettler et al., 2002; Li et al., 2006). Given the premise that the reorganized cortex should retain the functions implemented by the horizontal connections, and the idea that these connections mediate contour integration and contour saliency, the

implication is that V1 reorganization may preserve the perceptual integrity of contours occluded by retinal damage.

Here, we show how the functional changes that accompany V1 reorganization can be linked to the perceptual changes reported by human subjects with retinal lesions. We developed a neural network model to simulate the recovered neural activity in the LPZ, and the perceptual fill-in that might result, if the cortical reorganization in V1 is mediated by a network of long-range horizontal connections that normally underlies contour integration. Our goal was to determine if the lateral connections in V1 can account for the robust perceptual completion reported by subjects with retinal damage.

2.3 METHODS

2.3.1 QUALITATIVE DESCRIPTION AND GENERAL OVERVIEW

THE PARADIGM. To explore the relationship between V1 reorganization and perceptual fill-in, we modeled the neural and perceptual changes that might accompany retinal deterioration in human subjects. The retinal degeneration simulated in the model includes idealized features of the dry

form of AMD: large, continuous areas of photoreceptor loss, or scotomata, and intervening regions of diffuse, “salt-and-pepper” photoreceptor loss (Fig. 3). The input to the model is a stationary visual scene, projected onto the lesioned retinae of a hypothetical subject who fixates on the center of the image (e.g., Fig. 6B). The output from the model illustrates what this individual might perceive when looking at the input image, if the reorganization in V1 is mediated by geometrically specific horizontal connections (e.g., Fig. 4, D and E). The model overlays the input image with an array of black blotches (representing the “blind spots” caused by retinal lesions), and it computes an image that allows the reader to perceive what a subject might “see” underneath the lesion-induced blind spots.¹

Our simulations were carried out in two stages. In the first stage, we simulated an ensemble of complex cells, in a hypothetical retinal lesion subject, from the superficial layers of the V1 LPZ. Our model focuses on the properties of superficial-layer complex cells because the experimental evidence indicates that these are among the cells principally involved in the cortical reorganization. The axonal sprouting thought to underlie cortical reorganization occurs predominantly among the horizontal collaterals of




Figure 3. The macular degeneration simulated in the model. (A) an example of the binary mask $M(x,y)$ that simulates retinal deterioration in the macula, including representations of both the geographic and nongeographic atrophy encountered in AMD. Black pixels represent regions of photoreceptor loss [where $M(x,y) = 0$]; white pixels correspond to responsive regions of the photoreceptor mosaic [where $M(x,y) = 1$]. In this example, the black pixels cover a total of 76.6% of the simulated portion of the retina. (B) the nongeographic component of the retinal atrophy depicted in (A), modeling diffuse photoreceptor loss over 60% of the retina. (C) the geographic component of the atrophy $M(x,y)$. The lesions, which simulate scotomata from an early-to-intermediate stage of AMD, cover 41.4% of the retina in this example. Wherever the two components of retinal deterioration overlap in $M(x,y)$, the form of atrophy is considered to be geographic rather than diffuse. (Effectively, then, the geographic lesions cover 41.4% of the image, and the diffuse photoreceptor loss covers 60% of the remaining retina, or 35.2%.) Here, the simulated portion of the retina is taken to span the central 15 degrees of the visual field, each pixel spans 2.2 minutes of arc, and the average area of the geographic lesions is 1.2 square degrees of visual arc. There are 79 lesions that dot the central retina, distributed over an area of 225 square degrees of visual arc.

pyramidal cells in the superficial layers of V1 (Darian-Smith and Gilbert, 1994). These superficial-layer cells tend to have complex RFs (Hubel and Wiesel, 1968), they participate both pre- and postsynaptically in the plexus of horizontal connections, and their extrastriate projections constitute the output from V1 to the rest of the visual cortex.

In the second stage, we computed an image that demonstrates to the reader the visual perception we expect to emerge from the neural responses simulated during the first phase. The activity in the simulated LPZ signals the presence of continuous contours that pass through the retinal lesions, so the model generates its output image by filling-in the input image along these contours. This second stage of our model is intended to produce images that allow the reader to perceive the visual scene like a subject with the simulated retinal damage. (We refer to this process as “illustrating the perceptual fill-in.”) Since the visual cortex of the retinal lesion subjects undergoes a compensatory process of reorganization, the images computed by our model are not simply black, or even left blank, where the input has been ablated. Rather, they are actively filled-in by means of a neural network model of V1 reorganization and neuronal responses. These images of perceptual fill-in are designed to produce activity patterns in the visual

cortex of healthy readers that resemble the patterns we expect in retinal lesion subjects when they view the visual scene. To the extent that the reader's perception coincides with the reports from lesion subjects, our model explains the cortical reorganization and perceptual fill-in phenomena.

2.3.2 STAGE 1: SIMULATION OF NEURAL ACTIVITY

The first stage of the model, the simulation of neuronal activity, is itself composed of two successive computations, which simulate neural activity in the retina/LGN and then in V1. The functional recovery of neuronal activity begins, in a limited way, at the retina and LGN, but the truly extensive reorganization takes place in V1.

RETINAL/LGN COMPUTATIONS. In our model, the image underneath the diffuse photoreceptor loss—but not underneath the large retinal lesions—is recovered by retinal ganglion cells (RGCs) and LGN neurons. The RFs of these cells contain a mixture of responsive photoreceptors juxtaposed with destroyed receptors (Fig. 3B), and the visual signal transduced by the functional photoreceptors diffuses through the unresponsive region in each RF. The retinal image is relayed to the cortex,

but each spot of diffuse retinal damage is filled-in with a Gaussian average of the image on the adjacent spots of healthy retina (e.g., Fig. 6C). This process is modeled on a very small spatial scale. It simulates mechanisms that are potentially mediated by a re-weighting of local lateral interactions in the retina, although a re-weighting of converging inputs onto RGCs and LGN neurons may also play a role. Since the RGCs and the LGN never recover from large retinal lesions, the image relayed to the cortex still contains the large blind spots.

The signal transmitted to V1 is the partially recovered retinal image—rather than the spiking activity of explicitly modeled neurons from the retina and LGN—because the processing done by RGCs and LGN neurons is subsumed by the complex cells in our network.

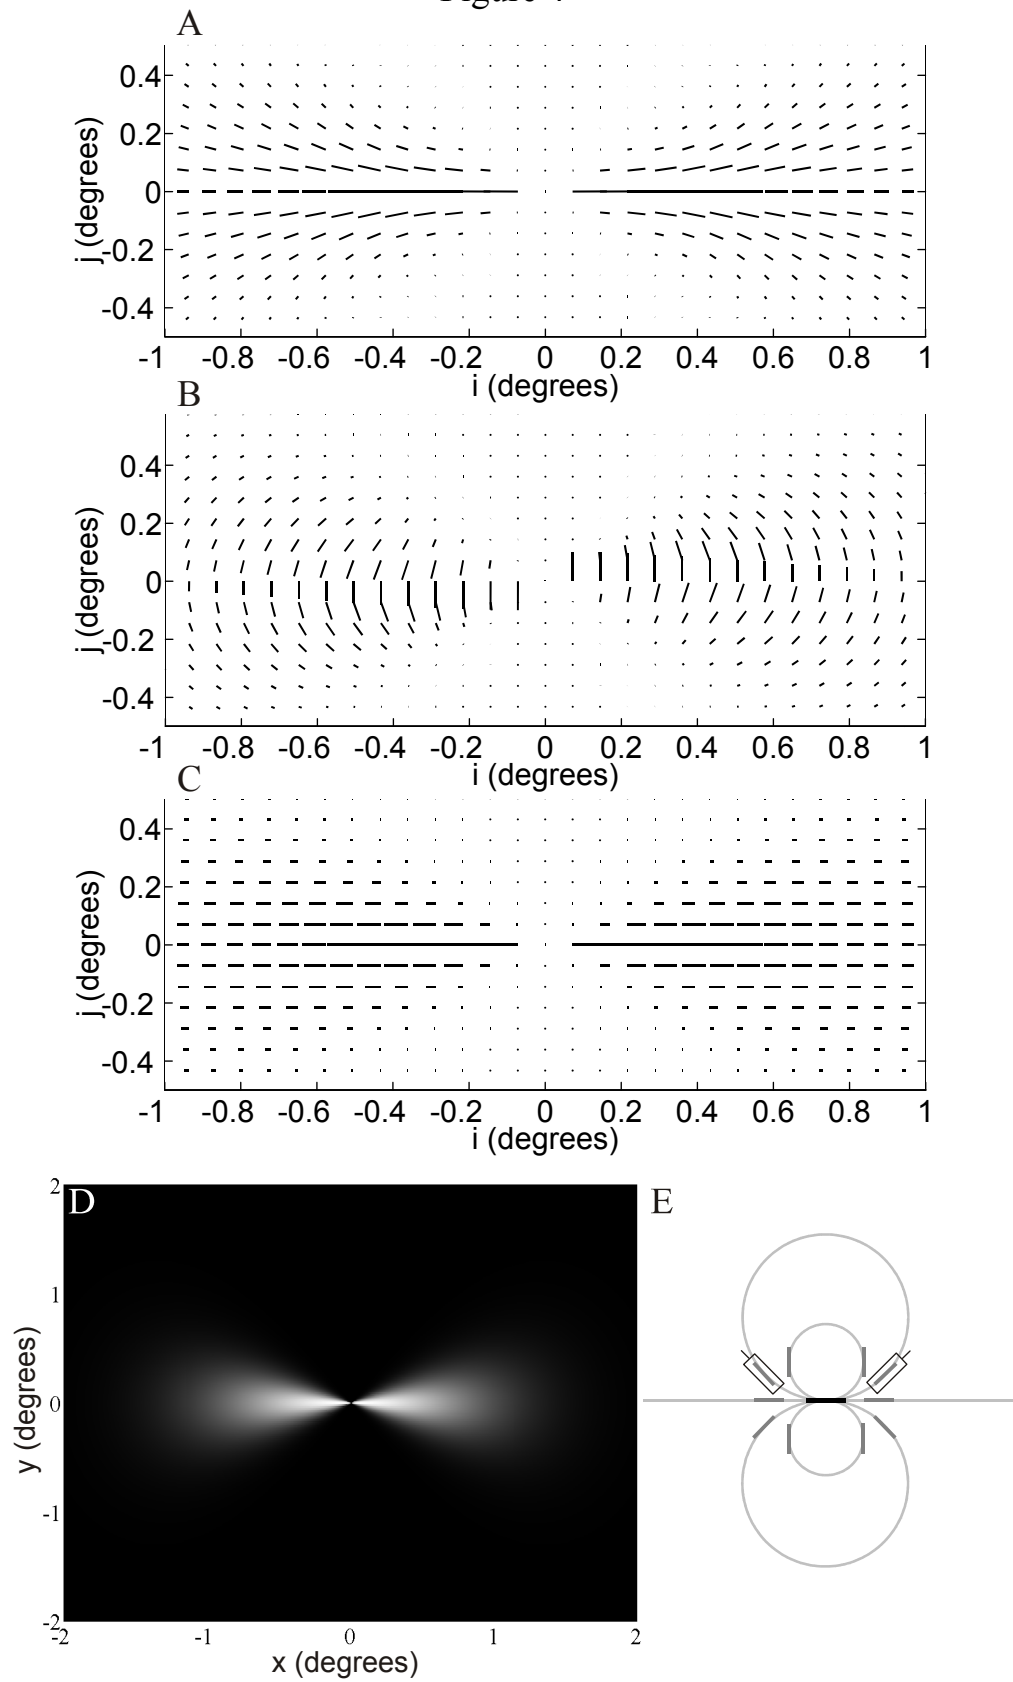
V1 COMPUTATIONS. In V1, we simulated a network of complex cells interconnected by a pattern of spatial- and orientation-selective horizontal connections. The connectivity pattern, called the association field (AF) (Field et al., 1993) and denoted by **K** throughout this work, models the neural interactions mediated by the long-range horizontal connections. Consistent with the anatomical and physiological data reviewed in the

Introduction, these connections underlie contour integration and contour saliency in the “healthy” regions of our model cortex, and they resuscitate neural activity in the simulated LPZ. All of the neurons in our network, both in the LPZ and in the surrounding cortical area (termed the peri-LPZ), receive lateral connections from presynaptic neurons, which send out their axon collaterals according to the geometric pattern specified by **K**. The strength of the lateral connections between each pair of model neurons depends on the relative spatial positions of their RFs and on their preferred orientations (POs). We tested two geometric patterns of connectivity, both of which are consistent with the known anatomical and physiological data. The first of these is called the collinear AF, in which the strongest network connections are between cells with collinear RFs (collinear RFs lie along a straight line whose orientation matches the POs of the neurons; Fig. 4E). The second is the co-circular AF, whereby co-circular RFs enjoy the strongest connection strengths. Co-circular RFs lie along a circular arc that passes through both RFs, where the PO of each RF is tangent to the circular arc (Fig. 4E). See the *Technical Details* below and Fig. 4 for the mathematical and graphical descriptions of these alternative connectivity patterns.

Figure 4. The patterns of lateral connectivity between our V1 model neurons. Each presynaptic cell in the peri-lesion projection zone (peri-LPZ) is associated with a vector field whose origin lies at the presynaptic receptive field (RF) and whose direction describes the preferred orientations (POs) of the postsynaptic neurons that receive the strongest connections from the presynaptic cell. Here, we depict the vector fields that describe co-circular (A and B) and collinear (C) patterns of connectivity, with each field shown in a rotated coordinate system whose x -axis is collinear with the PO of the presynaptic RF. (A) the co-circular vector field of peak excitatory connections made by the presynaptic neuron at the center of the field. At each position, the line orientation represents the direction of the field, while the line length represents the field magnitude. The line orientation at (i,j) , therefore, is the PO of the RF centered at (i,j) that receives the strongest connection from the presynaptic cell. The length of the line, in turn, is the connection strength between the 2 cells. Moreover, the connection strength \mathbf{K} between the presynaptic neuron and any arbitrary cell at (i,j) drops off exponentially as the PO of the postsynaptic cell deviates from the line direction at (i,j) (see *Eq. 14*). (B) the vector field of peak inhibitory connections made by a presynaptic cell in a co-circular network. The inhibitory field, obtained by rotating the vectors in the excitatory field by 90° , describes the POs of the postsynaptic cells that are most strongly inhibited by the presynaptic cell. (C) the alternative vector field of excitatory “collinear” interactions. (For the sake of brevity, the corresponding collinear inhibitory field obtained by rotating the vectors in the excitatory field is not shown.) The collinear field is unidirectional, and the presynaptic cell forms its most efficacious synapses onto those neurons that share its PO, regardless of the relative spatial positions of the pre- and postsynaptic RFs. Although the geometry of the collinear field differs from its co-circular counterpart, both fields have the same magnitude.

Figure 4, continued. The patterns of lateral connectivity between our V1 model neurons. (D) the vector field magnitude, represented here by intensity (bright white indicates large magnitude; dark gray indicates small magnitude) rather than by line length as above. The magnitude of the field was set to match electrophysiological and psychophysical measurements of the spatial extent of contextual interactions (Kapadia et al., 2000; Li and Gilbert, 2002; Li et al., 2006). (E) an illustration of co-circularity. The geometric relationships between the central black bar and the dark gray bars surrounding it demonstrate co-circularity, whereby any two oriented line segments that lie along the same circular arc are co-circular. Note that the black bar at the center of the drawing is the only line segment that is co-circular to every other bar. (Every gray bar is co-circular with exactly two other bars, including the black one.) The two bars resting along the line running through the black bar are said to be collinear with the central bar, but they also represent a special case of co-circularity, whereby the tangent circle that links the two bar elements has an infinite radius. Any two neurons and their RFs are said to be co-circular (collinear) if their preferred bar stimuli are co-circular (collinear). The thin rectangles represent the RFs of two co-circular neurons: the two gray bars, which illustrate the optimally oriented and positioned bar stimuli for those cells, are co-circular.

Figure 4



Given the connectivity pattern \mathbf{K} , in conjunction with the classical energy model of the complex cell (Adelson and Bergen, 1985; Dayan and Abbott, 2001), we simulated the activity of a network of V1 superficial-layer neurons. The model RFs were distributed over a grid of spatial positions at each pixel in the input image, with five RF sizes and eight POs at each location. (We evenly spaced the RF sizes between 0.2° and 1.0° and the POs between 0 and 157.5° .) All the cells in our model received a normalized combination of bottom-up and lateral inputs, where the relative weight of the two contributions was determined by the amount of retinal damage in each cell's RF (see the *Technical Details*). As in the real cortex, healthy neuronal responses were driven by bottom-up input from the retinal RF and modulated both by geometric lateral inputs and surround inhibition. On the other hand, neuronal responses in the LPZ were driven primarily or exclusively by lateral inputs from the adjacent region of normal cortex, where the synaptic weights between healthy presynaptic cells and the postsynaptic LPZ neurons were specified by the connection pattern \mathbf{K} . At stimulus onset, the stimulation of intact retinal loci seeded the cortical activity in the healthy cortex, which then coursed recurrently through the network \mathbf{K} and engendered new activity in the otherwise silent LPZ.

2.3.3 STAGE 2: ILLUSTRATION OF PERCEPTUAL FILL-IN

The simulated activity in the cortex constitutes a prediction of the image orientations underlying the retinal lesions, and the model uses this neuronal prediction to illustrate the perception of a human subject with a specific pattern of retinal deterioration (see the *Technical Details*). We postulated that the perception of the occluded visual scene arises from the joint activity in the peri-LPZ and the LPZ, and that it minimizes brightness discontinuities along the continuous features that are predicted to run through the lesions. Experimental evidence (Li et al., 2006) shows that a particular response property of V1 neurons—namely, contour facilitation—accounts for the perceptual phenomena of contour integration and contour saliency. Therefore, our model assumes that the facilitation of LPZ neurons by contours that pass through retinal lesions should lead to the integration and saliency of those contours within the scotomata—that is, to their perceptual completion.

In order to compute the output image illustrating perceptual completion, the model predicts the orientation around each point in the lesioned retinal image. The predicted orientation is derived from a principal components

analysis (PCA) of the neural responses corresponding to each point on the lesioned retinae. The PCA deciphers the underlying image orientation encoded by the activity in the LPZ, and it yields a measure of the saliency of the encoded orientation (which can also be interpreted as a measure of how reliably the neural responses predict the occluded image orientation). The output from our model is the image that selectively minimizes the luminance difference along the directions of the predicted image orientations (see the *Technical Details*).

2.3.4 TECHNICAL DETAILS OF THE MODEL

THE PARADIGM. Let the visual input to the model, defined over a domain in $\mathbb{N} \times \mathbb{N}$, be the gray scale image $I(x,y) : [1,A] \times [1,B] \rightarrow \mathbb{R}$;

let the simulated pattern of retinal deterioration be the mask

$M(x,y) : [1,A] \times [1,B] \rightarrow \{0,1\}$, which is zero (i.e., occluding) wherever the retinae have atrophied and unity (i.e., transparent) wherever both retinae are healthy. The simulated image on the retinae is then the input image overlaid with the retinal degeneration (e.g., Fig. 6B), and it is given by:

$$I^R(x,y) = I(x,y)M(x,y) \quad (1)$$

We define the regions of retinal atrophy by the set

$$\Omega \equiv \{(x, y) \mid M(x, y) = 0\} = S_0 \cup \bigcup_{i=1}^k S_i \quad (2)$$

where S_0 is the set of all pixels (x, y) that constitute the diffuse photoreceptor loss (Fig. 3B), k is the number of geographic lesions on the simulated retinae, and S_i is the set of all pixels within the i^{th} scotoma.

The retinal damage in our model, which is defined by a single mask (rather than one for each eye), may be interpreted as either monocular or binocular. Under the binocular interpretation, any point (x, y) in the mask refers to a pair of corresponding points on each retina. After a monocular lesion is made, there is a shift in ocular dominance in the LPZ so that the input from the healthy, normal eye either exclusively drives (Chino et al., 1992), or at least dominates (Calford et al., 2000), neuronal responses in the superficial layers of the V1 LPZ. Since the cells in our network model lie in these superficial layers, we can make the simplifying assumption that cells which receive masked input from one retina, but healthy input from corresponding points on the other retina, are driven by the healthy input, as if both retinae were undamaged. So, effectively, the only remaining retinal lesions in $M(x, y)$ are those that occur at matching, homonymous locations on both retinae. This interpretation is experimentally justified (Chino et al., 1992; Calford et al.,

2000), and it makes our model similar to an explicitly binocular model—one with partially mismatched patterns of retinal degeneration on two separate retinae—whereby the healthy subregions of each retina dominate the inputs to binocular cells. Alternatively, our model can be interpreted as an exclusively monocular one, if the mask is taken to define the retinal damage in only one eye and the hypothetical subject views the world through only the correspondingly lesioned eye. In either case, binocular interactions are ignored in our model.

RETINAL/LGN COMPUTATIONS. The visual signal, I^{V1} , transmitted from the retina to V1 is given by:

$$I^{V1}(x, y) = \begin{cases} I^R(x, y) & \text{for all } (x, y) \notin S_0 \\ \frac{\sum_{(i,j) \in R_{x,y}} G^{(x,y)}(i, j) I^R(i, j)}{\sum_{(i,j) \in R_{x,y}} G^{(x,y)}(i, j)} & \text{for all } (x, y) \in S_0 \end{cases} \quad (3)$$

where

$$G^{(x,y)}(i, j) = \frac{1}{2\pi\sigma^2} e^{-[(i-x)^2 + (j-y)^2]/2\sigma^2} \quad (4)$$

and

$$R_{x,y} = \{(a, b) | (a, b) \notin \Omega \text{ and } \sqrt{(x-a)^2 + (y-b)^2} \leq \delta\} \quad (5)$$

Here, $G^{(x,y)}(i, j)$ is a two-dimensional Gaussian centered at the masked point (x, y) , and $R_{x,y}$ is the set of all unmasked points surrounding (x, y) within a

radius of δ pixels. The radius over which the Gaussian averaging operates, δ , is set at the lower limit of the dimensions needed to fill-in the diffuse, “salt-and-pepper” atrophy. And the parameter σ is set so that

$$\sum_{(i,j) \in P_{x,y}} G^{(x,y)}(i,j) \approx 1$$

where $P_{x,y}$ is the set of all pixels surrounding (x,y) within a radius of δ pixels. This condition simply scales the width of the Gaussian to the size of the averaging disk $R_{x,y}$. (Our simulations were run with $\delta = 2$ pixels, corresponding to between 4.4 and 5.8 minutes of arc, and $\sigma = 0.833$.) We refer to the image I^{V1} as the partially filled-in image; this is the image that is filtered by the complex cells in our V1 network, as described below.

MODEL V1 NEURONS. Let $\tilde{n}_{r,\theta}^{x,y}$ be the V1 neuron with the following stimulus response characteristics: RF center (x,y) , RF width r , and PO θ .

Then we compute $n_{r,\theta}^{x,y}$, the depolarization elicited by the (partially filled-in) stimulus within that cell’s RF, via the relation:

$$n_{r,\theta}^{x,y} = (f_{r,\theta}^{x,y})^2 + (h_{r,\theta}^{x,y})^2 \quad (6)$$

Equation 6 is an expression of the classical energy model of the complex cell, whereby the neural response is simulated by convolving the visual stimulus (in our case, I^{V1}) with two quadrature RF filters, and then summing

the squared filter responses (Adelson and Bergen, 1985; Dayan and Abbott, 2001).² The classical RF of each neuron $\tilde{n}_{r,\theta}^{x,y}$ is modeled with two square Gabor filters centered at (x,y) with RF width r and orientation θ ; the variables $f_{r,\theta}^{x,y}$ and $h_{r,\theta}^{x,y}$ denote the point-by-point multiplication of each of these filters with the image falling within the RF. As a convention, we set the L2-norm of the (unmasked) Gabor filters to unity, and we map the pixel values in the image I^{V1} into the interval $[0, 1]$. The spatial substructure of the RF filters is set to match the characteristics of superficial-layer V1 RFs, with the additional constraint that wherever the RFs are partially or completely damaged by a retinal lesion, the masked subregions of the Gabor filters are made unresponsive (i.e., they are set to zero). The RFs of our model neurons densely cover the input image. Centered over each pixel in the input image, there are 40 RFs with five different sizes (ranging from 0.2 to 1.0° in width) and eight different POs (spaced evenly between 0 and 157.5°).

The Gabor filters simulate how neurons respond to vertical inputs from stimuli falling within the classical RF. They do not encompass the extra-classical RF surround, which arises from lateral activation through the horizontal connections, and so they do not mediate reorganization following retinal lesions. When the RF of a cell is partially ablated, the corresponding

regions of the Gabor filters are also permanently ablated, so that the classical RF becomes less responsive. Moreover, it often happens that for any partially ablated Gabor filter, the retinal damage destroys, for instance, a larger portion of the filter's inhibitory subregion than of its excitatory one. Wherever that occurs, the magnitude (or responsiveness) of the stronger subregion is scaled down so that the total responsiveness of the excitatory and inhibitory lobes is balanced. That is, filter regions that are directly masked are nullified, while unmasked regions of the same filter may be scaled down so that the total magnitude of the unmasked inhibitory lobes equals the unmasked excitatory lobes. (Otherwise, unbalanced RF filters would cause chronic neural activity in many cells at the border of the LPZ, a phenomenon which is not observed in the long-term reorganized cortex.)

Let $m_{r,\theta}^{x,y}$ be the fraction of the RF of neuron $\tilde{n}_{r,\theta}^{x,y}$ that has been ablated by retinal damage. Additionally, let $G_{1',r,\theta}$ and $G_{2',r,\theta}$ be the Gabor filters of the partially masked neuron $\tilde{n}_{r,\theta}^{x,y}$, and let $G_{1,r,\theta}$ and $G_{2,r,\theta}$ denote the unmasked versions of these same Gabor filters, which would be revealed if we were only to remove the retinal damage overlying the RF. Then

$\sum_{i,j} |G_{1,r,\theta}(i,j)|$ denotes the sum over the absolute value of the unmasked

Gabor filter $G_{1,r,\theta}$, and we define the masked fraction of the RF as follows:

$$m_{r,\theta}^{x,y} \equiv \frac{\sum_{i,j} (|G_{1,r,\theta}(i,j)| - |G_{1',r,\theta}(i,j)|) + \sum_{i,j} (|G_{2,r,\theta}(i,j)| - |G_{2',r,\theta}(i,j)|)}{\sum_{i,j} (|G_{1,r,\theta}(i,j)| + |G_{2,r,\theta}(i,j)|)} \quad (7)$$

MODEL V1 CONNECTIVITY. The V1 neurons in our model are interconnected by a plexus of long-range horizontal connections. The lateral connectivity is based on a vector field that represents the strongest excitatory interactions in the network. Each presynaptic cell is associated with a vector field that describes the neurons which receive the strongest excitatory connections from that cell. For instance, let $\vec{F}_{x,y,\theta}$ be the vector field of a presynaptic neuron $(\tilde{n}_{r,\theta}^{x,y})$ with RF center (x,y) and PO θ . The direction of the field $\vec{F}_{x,y,\theta}(i,j)$ at each point (i,j) specifies the PO of the neuron with RF center (i,j) that receives the strongest excitatory connection from the presynaptic cell. Of all the neurons at (i,j) , it is the one whose PO makes it co-circular with the presynaptic cell that receives the strongest connection. The field magnitude at (i,j) , in turn, is the connection strength between these two co-circular cells. We therefore call $\vec{F}_{x,y,\theta}$ the co-circular vector field, since it specifies that the strongest excitatory connections in the network are

between co-circular neurons (see Fig. 4, A and D, for a graphical depiction of this field). The mathematical description of the field, which was modified from Guy & Medioni (1996), is given by

$$\vec{F}_{x,y,\theta}(i,j) = \begin{cases} \exp(-A_1(\gamma x')^2) \begin{bmatrix} \frac{x'}{|x'|} \\ 0 \end{bmatrix} & \text{if } y' = 0 \\ \frac{1}{v_{\text{mag}}} \exp(-A_1(\gamma x')^2 - B_1(\arctan(|y'/x'|))^2) \begin{bmatrix} \frac{x'}{|x'|} (\sqrt{R^2 - x'^2} + z) \\ \frac{|x'|y'}{|y'|} \end{bmatrix} & \text{if } y' \neq 0 \end{cases} \quad (8)$$

where

$$(x', y') = (i - x, y - j) \quad (9)$$

$$R = \frac{x'^2 + y'^2}{2y'} \quad (10)$$

$$z = C \exp(-A_2|\gamma x'|^3 - B_2(\arctan(|y'/x'|))^3) \quad (11)$$

$$v_{\text{mag}} = \sqrt{(\sqrt{R^2 - x'^2} + z)^2 + x'^2} \quad (12)$$

Here, γ is the width of the input image in degrees of arc subtended at the eye, divided by the width of the image in pixels.

The parameters in the equations are as follows: $A_1 = 1.2$ degrees of arc subtended at the eye⁻²; $B_1 = 5.7$; $A_2 = 24.1$ degrees of arc subtended at the eye⁻³; $B_2 = 2.85$; and $C = 50$. For simplicity, we express $\vec{F}_{x,y,\theta}$ in a rotated coordinate system whose x -axis is collinear with the PO of the presynaptic neuron. The exponential terms to the left of the vectors in *Eq. 8* determine the magnitude of the field at each position (i,j) (see Fig. 4D). Beside the exponential terms are the unit vectors that determine the co-circular geometry of the field, with a collinear bias near the origin given by z . This correction term z favors iso-orientation facilitation over co-circularity at very small spatial scales, in agreement with psychophysical measurements of contour saliency and with the statistics of natural scenes (Field et al., 1993; Geisler et al., 2001; Sigman et al., 2001).

In some of our simulations, to test the effect of reshaping the AF geometry through top-down interactions, we substituted this co-circular vector field with a unidirectional one (see Fig. 4C):

$$\vec{F}_{x,y,\theta}^{\text{cl}}(i,j) = \exp(-A_1(\gamma x')^2 - B_1(\arctan(|y'/x'|))^2) [x'/|x'| \quad 0]^T \quad (13)$$

expressed here with the same notation as the co-circular field above. We call $\vec{F}_{x,y,\theta}^{\text{cl}}$ the “collinear field”; it specifies that the strongest excitatory connections are between neurons that share the same orientation preference.

Given the vector field of maximally excitatory connections made by each presynaptic cell, we can now specify all of the connections in the network as a function of their deviation from co-circularity (or collinearity, depending upon which vector field is used as the basis for the network connectivity).

Let $D(\vec{F}_{x,y,\theta}(i,j),\alpha)$ be the angular difference (in radians) between the vector field direction at (i,j) and the orientation preference α . Additionally, let $\vec{F}_{x,y,\theta}^{\perp}$ be the field obtained by rotating each vector in $\vec{F}_{x,y,\theta}$ by 90° (Fig. 4B). Then $\mathbf{K}_{x,y,\theta}(i,j,\alpha)$, which denotes the sign and magnitude of the connection between a presynaptic neuron with RF center (x,y) and PO θ and a postsynaptic neuron at (i,j) with PO α , is given by

$$\begin{aligned} & \mathbf{K}_{x,y,\theta}(i,j,\alpha) \\ &= \begin{cases} \|\vec{F}_{x,y,\theta}(i,j)\| \exp(-D(\vec{F}_{x,y,\theta}(i,j),\alpha)^2 / 0.15) & \text{if } D(\vec{F}_{x,y,\theta}(i,j),\alpha) \leq \pi/4 \\ -\|\vec{F}_{x,y,\theta}(i,j)\| \exp(-D(\vec{F}_{x,y,\theta}^{\perp}(i,j),\alpha)^2 / 0.15) & \text{otherwise} \end{cases} \quad (14) \end{aligned}$$

If the geometric relationship between the pre- and postsynaptic neurons is close to co-circular (i.e., if $D(\vec{F}_{x,y,\theta}(i,j),\alpha) \leq \pi/4$), then the presynaptic

neuron makes an excitatory connection onto the postsynaptic cell, with a strength that falls off exponentially with the angular deviation from co-circularity. On the other hand, if the relationship between the two neurons is very different from co-circularity ($D(\vec{F}_{x,y,\theta}(i,j), \alpha) > \pi/4$), then the connection between the two neurons is modeled as an (implicitly disynaptic) inhibitory connection whose strength rises exponentially with the angular deviation from co-circularity. The number 0.15 in *Eq. 14* was chosen so that the connection strength between cells was very weak ($\sim 2\%$ of the vector field strength) when the postsynaptic PO differs by 45° ($\pi/4$ radians) from the field direction.

SIMULATION OF CORTICAL ACTIVITY. We denote the total membrane depolarization of neuron $\tilde{n}_{r,\theta}^{x,y}$ at time t by $S_{r,\theta}^{x,y,t}$, and the instantaneous firing rate, obtained by passing the depolarization through a static nonlinearity, as $g(S_{r,\theta}^{x,y,t})$. The activation $S_{r,\theta}^{x,y,t}$ arises from the depolarization ($n_{r,\theta}^{x,y}$) induced by the RF stimulus, and it is modulated by the sum of excitatory and inhibitory inputs from the association field \mathbf{K} . Besides the vertical inputs to the RF, each cell $\tilde{n}_{r,\theta}^{x,y}$ receives postsynaptic potentials from the network of horizontal collaterals extending from the presynaptic

cells $\tilde{n}_{r,\alpha}^{i,j}$ in the peri-LPZ. The sign and strength of the synapse between each pre- and postsynaptic cell pair is given by the product of a scaling factor γ_1 and the connection $\mathbf{K}_{i,j,\alpha}(x,y,\theta)$. Therefore at time t , the synaptic potentials from the horizontal connections in postsynaptic cell $\tilde{n}_{r,\theta}^{x,y}$ are given by the firing rates $g(S_{r,\alpha}^{i,j,t-1})$ of the presynaptic cells, multiplied by the term $\gamma_1 \mathbf{K}_{i,j,\alpha}(x,y,\theta)$ (see *Eq. 15*). The effect of these potentials on the postsynaptic cell is gated by a toggle-like term (given by $\{\xi G_{(1-p)}(1 - m_{r,\theta}^{x,y}) + [(f_{r,\theta}^{x,y})^2 + (h_{r,\theta}^{x,y})^2] G_p(m_{r,\theta}^{x,y})\}$), which tends to multiply the potentials either by the responses of the (normalized) postsynaptic RF filters or by a constant factor ξ , depending upon the degree of RF degeneration, $m_{r,\theta}^{x,y}$. This term recapitulates experimentally observed phenomena in the normal cortex and in the LPZ, whereby the efficacy of the horizontal inputs is either modulated by coincident activation from within the RF (as in the normal cortex) (Hirsch and Gilbert, 1991) or else is facilitated to suprathreshold levels when the vertical input is silenced (as in the LPZ) (Darian-Smith and Gilbert, 1994; Calford et al., 2003). The term mediates a drastic and switch-like change in neuronal response properties after cortical reorganization. As the amount of retinal deterioration in a neuron's classical RF increases, the neuron rapidly becomes highly

responsive to lateral inputs from neighboring cortical cells, regardless of the visual input within its classical RF.

In addition to the excitation and inhibition mediated by the lateral connections, the simulated cortical activity is subject to a gain control mechanism mediated by surround inhibition. The firing of the model neurons is inhibited by the overall cortical activity as a function of the distance between the inhibited RFs and the RFs of the active cells in the surrounding cortex.

The total membrane depolarization, $S_{r,\theta}^{x,y,t}$, is therefore computed via the recurrence relation:

$$S_{r,\theta}^{x,y,t} = \frac{n_{r,\theta}^{x,y} + \gamma_1 \{ \xi G_{(1-p)}(1 - m_{r,\theta}^{x,y}) + [(f_{r,\theta}^{x,y})^2 + (h_{r,\theta}^{x,y})^2] G_p(m_{r,\theta}^{x,y}) \} \times \sum_{[(i,j) \in L_{x,y}]} \sum_{[\alpha \in \Phi]} g(S_{r,\alpha}^{i,j,t-1}) K_{i,j,\alpha}(x, y, \theta)}{1 + \gamma_2 \sum_{[(i,j) \in V_{x,y}]} \sum_{[\alpha \in \Phi]} g(S_{r,\alpha}^{i,j,t-1}) \exp(-d(x, y, i, j)^2 / b^2)} \quad (15)$$

The membrane depolarization $S_{r,\theta}^{x,y,0}$ at the initial time point (right after stimulus onset) is the classical response $n_{r,\theta}^{x,y}$, and the membrane potentials converge to a steady state in $t \approx 10$ iterations of network activity. The static

nonlinearity $g(\cdot)$ that converts the membrane depolarization to an instantaneous firing rate is given by ³

$$g(x) = \begin{cases} 0 & \text{for } x < 0 \\ x & \text{for } 0 \leq x \leq x_{thr} \\ x_{thr} & \text{for } x > x_{thr} \end{cases} \quad (16)$$

In Eq. 15, $f_{r,\theta}^{x,y}$ and $h_{r,\theta}^{x,y}$ are the convolution responses obtained from the Gabor filters of neuron $\tilde{n}_{r,\theta}^{x,y}$, except that each filter is normalized so that its L2-norm equals 1. $L_{x,y}$ is the set of all spatial positions from which neurons in the peri-LPZ laterally contact the cells with RF center (x,y) ; $V_{x,y}$ is the set of all points in the visible image within 2° of the point (x,y) ; Φ is the set of neuronal orientation preferences in the model, $\{0^\circ, 22.5^\circ, \dots, 157.5^\circ\}$; $d(x,y,i,j)$ is the spatial distance between the points (x,y) and (i,j) in degrees of arc subtended at the eye; $b = 1$ degree of arc subtended at the eye; and $G_x(m)$ is a logistic function that modulates the efficacy of horizontal connections as a function of the occluded fraction of the postsynaptic RFs

$$G_x(m) = \frac{1}{1 + (1/\rho_1 - 1)\exp(-\rho_2(1-m))} - \rho_1 \quad (17)$$

The function $G_x(m)$ is a one-dimensional logistic curve that decreases from 1 to 0 as the variable m increases from 0 to 1. Its parameters, ρ_1 and ρ_2 , are set to satisfy two constraints: 1.) that $G_x(m)$ maps its domain into the range $[0,1]$;

and 2.) that the midpoint of the function lies at the coordinate $m = x$ [i.e., so that $G_x(x) = 0.5$].

The parameters referenced above are: $x_{thr} = 20$; $\xi = 10$; $\gamma_1 = 0.001$; $\gamma_2 = 0.002$; and $p = 0.75$. For neurons in the peri-LPZ, $L_{x,y} = V_{x,y}$. For cells in the LPZ (i.e., all neurons in the set $\{\tilde{n}_{r,\theta}^{x,y} \mid (x,y) \in \bigcup_{i=1}^k S_i\}$), $L_{x,y}$ is the set of all periscotoma positions within 1° of the scotoma border, excluding all points (i,j) such that the cortical distance between the neurons representing the points (x,y) and (i,j) exceeds 3.5 mm, the radial extent of the lateral projections (Stettler et al., 2002).^{4,5} Note that the magnitude and spatial extent of the contextual interactions in the real cortex are strongly dependent upon the characteristics of the visual stimulus (Kapadia et al., 2000; Li et al., 2006). For instance, the modulatory effect of the lateral interactions in normal cortex is essentially nil for simple stimuli shown at high contrast. So whenever the model input $I(x,y)$ is specified as a black-and-white artificial image (as in Figs. 5 and 7), we use $\gamma_1 = \gamma_2 = 0$ for neurons in the peri-LPZ, and for LPZ neurons we reduce the spatial radius in $L_{x,y}$ from 1° to 0.5° .

In the numerator of *Eq. 15*, the term on the right represents the lateral inputs to neuron $\tilde{n}_{r,\theta}^{x,y}$, modulated postsynaptically by the term

$$\xi G_{(1-p)}(1 - m_{r,\theta}^{x,y}) + [(f_{r,\theta}^{x,y})^2 + (h_{r,\theta}^{x,y})^2]G_p(m_{r,\theta}^{x,y})$$

When the afferent input to neuron $\tilde{n}_{r,\theta}^{x,y}$ is severely ablated,

then $G_{(1-p)}(1 - m_{r,\theta}^{x,y}) \approx 1$ and $G_p(m_{r,\theta}^{x,y}) \approx 0$, so that the lateral interactions are scaled by the factor ξ . On the other hand, for a neuron with a healthy classical RF, $G_{(1-p)}(1 - m_{r,\theta}^{x,y}) \approx 0$ and $G_p(m_{r,\theta}^{x,y}) \approx 1$, so that the lateral interactions are modulated by the strength of the stimulus within the classical RF. The denominator of *Eq. 15* represents divisive surround inhibition (Carandini et al., 1997), the spatial characteristics of which were set to match measurements of its extent in V1 (Sceniak et al., 2001).

Note that the gating term

$$\xi G_{(1-p)}(1 - m_{r,\theta}^{x,y}) + [(f_{r,\theta}^{x,y})^2 + (h_{r,\theta}^{x,y})^2]G_p(m_{r,\theta}^{x,y})$$

which describes the efficacy of the lateral connections as a function of the retinal damage $m_{r,\theta}^{x,y}$, is the mathematical expression of cortical reorganization in the model. It is here that the preexisting scaffold of horizontal connections is converted from a modulatory force in the normal

cortex—whereby the connections modulate neural responses to stimuli in the RF center—to a driving force that engenders suprathreshold activity in the LPZ.

ILLUSTRATION OF PERCEPTUAL FILL-IN. For each point

$(x, y) \in \bigcup_{i=1}^k S_i$, the converged firing rates $g(S_{r,\theta}^{x,y,10})$ can be assembled into five covariance matrices [after (Guy and Medioni, 1996)], one for each RF size. Let \hat{v}_θ be the unit vector whose direction matches the orientation preference θ , and let $(\hat{v}_\theta)_x$ and $(\hat{v}_\theta)_y$ be the x - and y -components of the vector, respectively. Then the covariance matrix for RF size r at the scotoma point (x,y) is

$$\mathbf{C}_{xy}^r = \begin{bmatrix} \sum_{\theta \in \Phi} [(\hat{v}_\theta)_x g(S_{r,\theta}^{x,y,10})]^2 & \sum_{\theta \in \Phi} (\hat{v}_\theta)_x (\hat{v}_\theta)_y g(S_{r,\theta}^{x,y,10})^2 \\ \sum_{\theta \in \Phi} (\hat{v}_\theta)_x (\hat{v}_\theta)_y g(S_{r,\theta}^{x,y,10})^2 & \sum_{\theta \in \Phi} [(\hat{v}_\theta)_y g(S_{r,\theta}^{x,y,10})]^2 \end{bmatrix} \quad (18)$$

where r can be any of the five RF sizes simulated in the network (i.e.,

$r \in \{0.2^\circ, 0.4^\circ, 0.6^\circ, 0.8^\circ, 1.0^\circ\}$). The singular value decomposition (SVD)

of the covariance matrix for each RF size r yields the two eigenvectors

$(\vec{p}$ and $\vec{p}^\perp)$ of the matrix, together with their corresponding eigenvalues (λ_1 and λ_2):

$$\text{SVD}(\mathbf{C}_{\mathbf{xy}}^r) = \begin{bmatrix} \vec{p}^T \\ (\vec{p}^\perp)^T \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} \vec{p} & \vec{p}^\perp \end{bmatrix} \quad (19)$$

For each lesioned point (x,y) , we chose the covariance matrix (and its associated RF size r) that possessed the largest difference between its two eigenvalues. The predicted orientation at (x,y) was then \vec{p} , the eigenvector from this chosen matrix with the largest corresponding eigenvalue, and our measure of the predicted saliency around (x,y) was just the difference between the two eigenvalues of the matrix, $\lambda_1 - \lambda_2$ [after (Guy and Medioni, 1996)]. (Alternatively, if all the responses from neurons with different RF sizes were merged into one covariance matrix $\mathbf{C}_{\mathbf{xy}}$, and the orientation and saliency in the image were predicted from that matrix, the resultant predictions were quantitatively similar.) The products of the unit eigenvectors $\vec{p}_{x,y}$ and the saliencies $(\lambda_1 - \lambda_2)_{x,y}$, corresponding to each point (x,y) on the lesioned retinae, constitute a vector field of predicted orientations and saliencies everywhere in the occluded retinal image (Fig. 5F).

We computed our illustrations of perceptual fill-in via a function minimization that was inspired by, and closely related to, an anisotropic

diffusion process (Perona and Malik, 1990). Let the illustrated perception,

defined over a domain in $\mathbb{N} \times \mathbb{N}$, be $P(x,y) : [1,A] \times [1,B] \rightarrow \mathbb{R}$, with

$P(x,y) = I^{V1}(x,y)$ for all $(x,y) \notin \bigcup_{i=1}^k S_i$. Then we find the image P that

minimizes the function

$$F = \sum_{(x,y)} \frac{1}{T^{x,y}} \left\{ C_N^{x,y} \left([\Delta_N^{x,y}]^2 + [\Delta_S^{x,y}]^2 \right) + C_{NE}^{x,y} \left([\Delta_{NE}^{x,y}]^2 + [\Delta_{SW}^{x,y}]^2 \right) \right. \\ \left. + C_E^{x,y} \left([\Delta_E^{x,y}]^2 + [\Delta_W^{x,y}]^2 \right) + C_{NW}^{x,y} \left([\Delta_{NW}^{x,y}]^2 + [\Delta_{SE}^{x,y}]^2 \right) \right\} \quad (20)$$

where

$$C_N^{x,y} = \exp(q((\lambda_1 - \lambda_2)_{x,y} \vec{p}_{x,y} \cdot [0 \ 1]^T)^2) \quad (21)$$

$$C_{NE}^{x,y} = \exp(q((\lambda_1 - \lambda_2)_{x,y} \vec{p}_{x,y} \cdot [1/\sqrt{2} \ 1/\sqrt{2}]^T)^2) \quad (22)$$

$$C_E^{x,y} = \exp(q((\lambda_1 - \lambda_2)_{x,y} \vec{p}_{x,y} \cdot [1 \ 0]^T)^2) \quad (23)$$

$$C_{NW}^{x,y} = \exp(q((\lambda_1 - \lambda_2)_{x,y} \vec{p}_{x,y} \cdot [-1/\sqrt{2} \ 1/\sqrt{2}]^T)^2) \quad (24)$$

$$\Delta_N^{x,y} \equiv P(x, y-1) - P(x, y) \quad (25)$$

$$\Delta_{NE}^{x,y} \equiv P(x+1, y-1) - P(x, y) \quad (26)$$

$$\Delta_E^{x,y} \equiv P(x+1, y) - P(x, y) \quad (27)$$

$$\Delta_{SE}^{x,y} \equiv P(x+1, y+1) - P(x, y) \quad (28)$$

$$\Delta_S^{x,y} \equiv P(x, y+1) - P(x, y) \quad (29)$$

$$\Delta_{SW}^{x,y} \equiv P(x-1, y+1) - P(x, y) \quad (30)$$

$$\Delta_{\text{W}}^{x,y} \equiv P(x-1, y) - P(x, y) \quad (31)$$

$$\Delta_{\text{NW}}^{x,y} \equiv P(x-1, y-1) - P(x, y) \quad (32)$$

$$T^{x,y} = 2 \left(C_{\text{N}}^{x,y} + C_{\text{NE}}^{x,y} + C_{\text{E}}^{x,y} + C_{\text{NW}}^{x,y} \right) \quad (33)$$

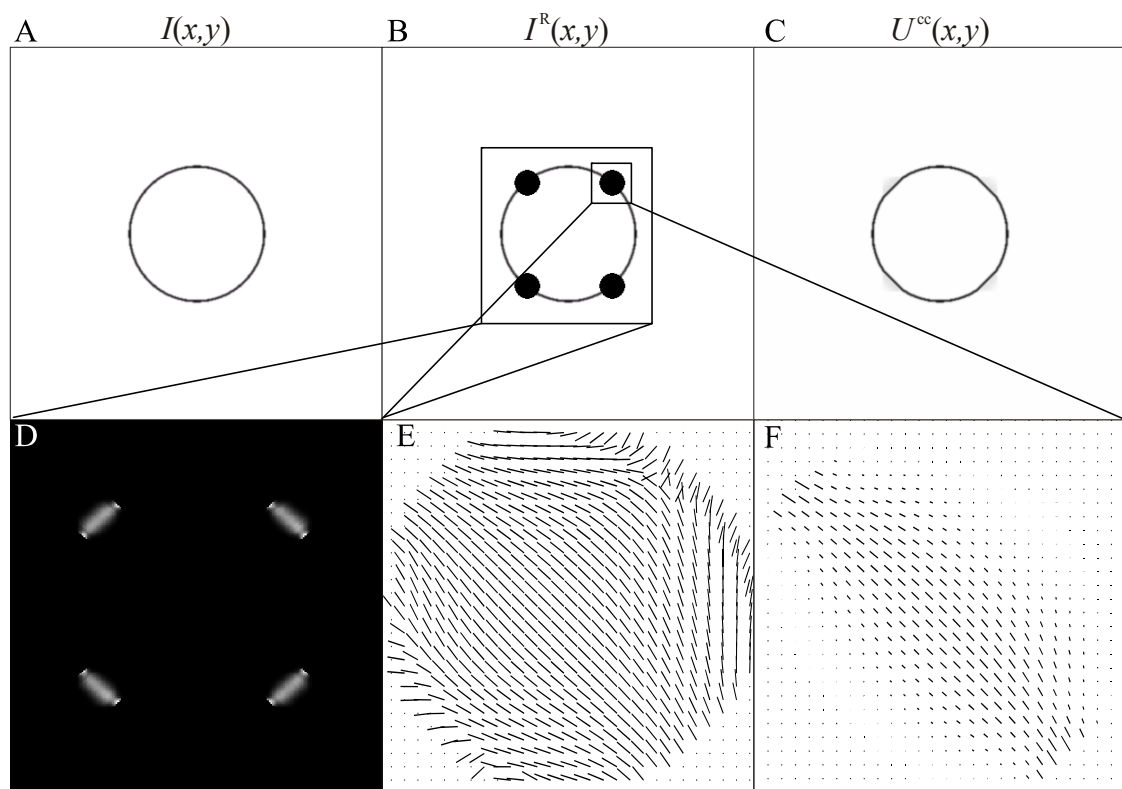
F is the average of the local brightness differences in the image $P(x,y)$ along four directions (north-south, northeast-southwest, east-west, northwest-southeast), weighted by a set of four coefficients (C_{N} , C_{NE} , C_{E} , C_{NW}) that measure how well each direction matches $\vec{p}_{x,y}$, the predicted orientation in the underlying image. The minimization of this function selectively reduces differences in the image brightness along the predicted contours in the scotomata, thereby filling-in the image along those contours. (Note that this minimization is not intended to explicitly simulate the physiological activity of any neural network. It is merely intended to illustrate the perceptual consequences of the cortical reorganization.)

The superscripts in the above equations represent pixel coordinates, while the subscripts represent image directions. For instance, NE denotes the “northeast” (top right) direction, so $\Delta_{\text{NE}}^{x,y}$ is the difference in the illustrated perception between the pixel (x,y) and the pixel immediately above it and to the right. For every scotoma point, the image is selectively filled-in along

the directions with the largest coefficients. The coefficient $C_X^{x,y}$, corresponding to direction X (e.g., $X = NE$) at scotoma point (x,y) , is determined by the product of three quantities: 1.) the predicted saliency at (x,y) ; 2.) the angular agreement between the predicted orientation at (x,y) and the direction X ; and 3.) the parameter q —which controls the direction specificity of the perceptual fill-in. Large values of the parameter q highlight differences in the coefficients between different directions. They therefore engender a sharp form of fill-in, whereby the image is filled-in along the directions that match the predicted orientations in the image. On the other hand, small values of q produce the opposite effect: they yield a diffuse, blurry perception, whereby the brightness differences in the image are minimized along all directions at once. We set q to be relatively large ($q \in \{10^4, 10^5\}$) in order to produce a highly directional mode of fill-in.

Figure 5. A demonstration of perceptual fill-in, using the circle. (A) a circle as the input image, $I(x,y)$, spanning the central 20° of the visual field. (B) the retinal image, $I^R(x,y)$, signaled by the damaged retina of the simulated subject. (C) $U^{cc}(x,y)$, the illustrated perception of the subject with the macular degeneration evident in (B) when viewing the image in (A), computed with the co-circular association field (AF). Here and elsewhere, we denote the illustrated perception arising from a co-circular AF by $U^{cc}(x,y)$. [The perceptions arising through a collinear field or through a non-directional fill-in process are denoted by $U^{cl}(x,y)$ and $U^{iso}(x,y)$, respectively.] The average angular discrepancy (i.e., image reconstruction error) in the perceptual completion of the circle is $D = 5.1^\circ$. See the *Results* for a description of this average angular difference, D . (In the figures that follow, the image reconstruction errors for each simulation are reported parenthetically as $D = x^\circ$). (D) the saliency map computed from the firing rates in the LPZ (bright white indicates high saliency; dark gray indicates low saliency). The lobes of high saliency indicate where contour facilitation within the LPZ signals the presence of a continuous contour passing through the lesions. (E) the (normalized) vector field of predicted orientations $\vec{p}_{x,y}$ at each position underneath the indicated scotoma. (F) the vector field of predicted orientations and saliencies at each position in the indicated scotoma. The magnitudes of the vectors signal the saliency $(\lambda_1 - \lambda_2)_{x,y}$ corresponding to the predicted orientation at each scotoma point (x,y) . The perceptual illustration in (C) was obtained by filling-in the lesions along the vector fields in each scotoma. See the text in the *Methods* section accompanying *Eqs. 18* and *19* for a description of how the predicted orientations and their corresponding saliencies are derived from the neural responses in the LPZ.

Figure 5



2.4 RESULTS

Our simulations show that cortical reorganization mediated by geometric horizontal connections can explain the perceptual fill-in phenomenon. Even in the face of severe retinal deterioration, the simulated activity in the LPZ can engender the robust perceptual completion of disrupted visual input. Figures 5-9 show several input images, their projections onto atrophied retinae, the corresponding visual signal transmitted from the LGN to V1, and the illustrated perceptual completion. Under the worst conditions of the modeled retinal damage, the perceptual fill-in image contains regions which are locally blurred or distorted (e.g., Fig. 6), but it reliably recreates the overall structure of the underlying image. Moreover, the modeled perceptual fill-in recapitulates the major observations from psychophysical studies on subjects suffering from retinal damage. Subjects report that their perception is characterized by some haziness and distortion, but that contours are typically completed and that their perception is sufficiently natural so that the presence and location of the lesions often goes unnoticed (Craik, 1966; Gerrits and Timmerman, 1969; Schuchard, 1993, 1995; Zur and Ullman, 2003). The results in Figs. 5-9 encompass all of these phenomena. Notice that the locations and shapes of the retinal lesions are generally not evident

Figure 6. A full set of simulations on an aerial view of Rockefeller University. (A) the input image, $I(x,y)$, spanning the central 15° of the visual field. (B) the retinal image, $I^R(x,y)$, signaled by the deteriorated photoreceptor lattice when the simulated subject fixates on the center of the input image. (C) the visual signal transmitted to V1, $I^{V1}(x,y)$, after the integration of the image within the RFs of retinal ganglion cells and LGN neurons. (D and E) the illustrated perception of a subject with the macular degeneration evident in (B), computed with either a co-circular [D: $U^{cc}(x,y)$] or a collinear [E: $U^{cl}(x,y)$] pattern of connectivity. In (D), the regions encircled in black are not reconstructed as well as they are in (E). They provide examples of how a collinear AF can engender better perceptual completion than a co-circular field where the visual scene is dominated by straight lines. The average angular discrepancies in the co-circular (D) and collinear (E) perceptual completion images are $D = 17.1^\circ$ and $D = 15.7^\circ$, respectively (see the *Results*). (F) $U^{iso}(x,y)$, an illustration of non-directional perceptual completion, computed by filling-in the image with isotropic diffusion (see the *Results*). Here, (C) shows the visual information that would be available to the observer if no cortical reorganization occurred. (D) and (E) show the observer's perception if the reorganization occurred in V1. (F) shows the observer's perception if the fill-in occurred without deference to the geometric relationships of the AF.

Figure 6





Figure 7. The Dayan caricature. (A) the input image, falling within the central 18.75° of the visual field. (B) the retinal image. (C) the visual signal sent to V1. (D) the illustrated perception, computed with the co-circular AF ($D = 11.1^\circ$). (E) the illustrated perception arising from a collinear AF ($D = 10.8^\circ$). The image regions circled in gray indicate where the collinear AF is subtly outmatched by the co-circular field. They point to the general phenomenon that the co-circular geometry can be better suited for the fill-in of contours that change sharply (relative to the size of the lesions under which they are hidden) and of junctions between intersecting curves. (The numerical gauge of perceptual fill-in reports a slightly smaller reconstruction error for the collinear field, but that difference is somewhat artifactual. It reflects the fact that the collinear field tends to be more accurate than the co-circular field along the sides of contours. That difference is not perceptually salient, however, because the perceptual fill-in is strongly directional only along the length of the contours themselves, not along the abutting image regions.)

Figure 7






Figure 8. The Einstein photograph. Clockwise, from *top* to *bottom*: the input image (16° in width), retinal image, illustrated perception, and visual signal transmitted to V1 ($D = 25.9^\circ$).

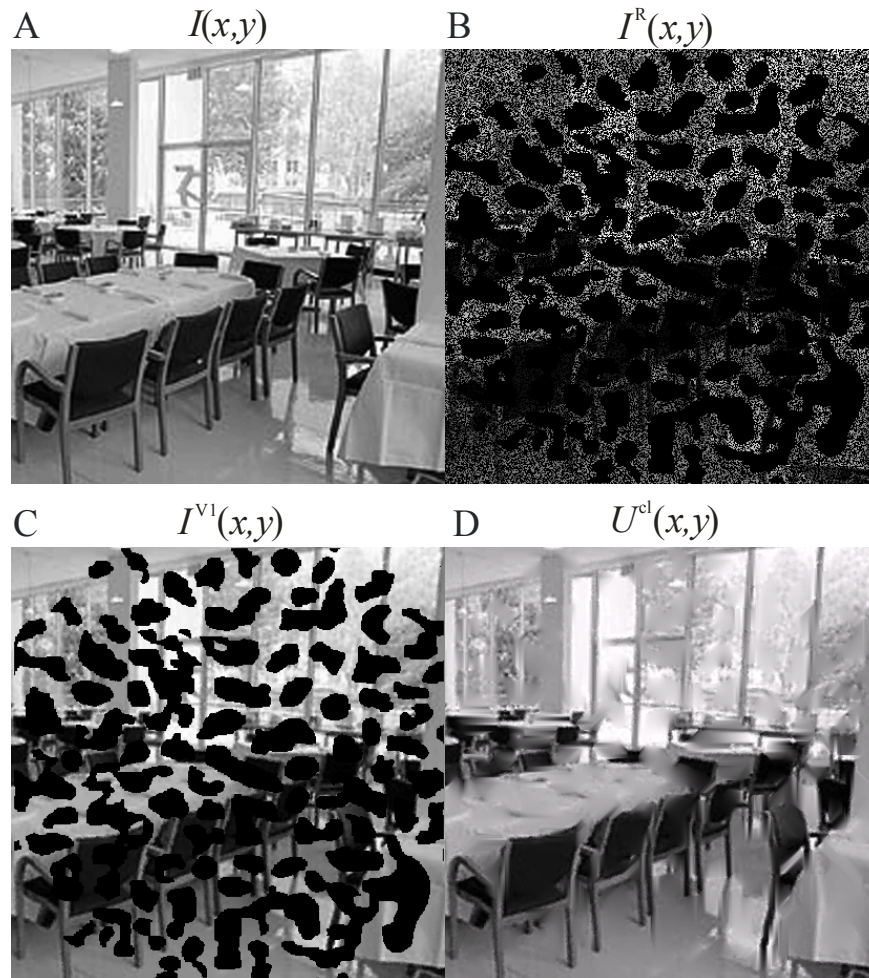


Figure 9. The dining room. Clockwise, from *top* to *bottom*: the input image (spanning 15°), retinal image, illustrated perception, and the visual signal sent to V1 ($D = 23.6^\circ$).

in the filled-in images, and that many of the globally salient features of the underlying visual scene are preserved. The model parameters used to generate these results were not optimized to produce the most accurate image reconstructions, but were constrained for consistency with available physiological data on factors like RF size and average orientation bandwidth, the spatial characteristics of contextual interactions, and the lateral extent of the horizontal connections.⁵ We also found that the quality of the perceptual fill-in is robust against perturbations of the chosen parameter values.

The precise shape of the AF can noticeably influence the perceptual completion of the visual scene. In Fig. 6, we show how the perceptual fill-in of a city landscape differs when the neural responses in the LPZ are driven through either the co-circular or collinear AF. Since the geometry of the collinear field corresponds more closely to the buildings in the image than does the co-circular field, it leads to a slightly more accurate completion of the skyscraper at the left of the image (Fig. 6). A co-circular field, on the other hand, can outperform the collinear geometry when curvature and junctions between intersecting lines are prominent features in the visual scene (Fig. 7). More generally, these results suggest how top-down modulation of the contextual interactions in early visual cortical areas may

be manifested in perceptual completion. Object expectation for an observer viewing a visual scene like the one depicted in Fig. 6 could conceivably lead to a dynamic, top-down reshaping of the AF to match the predominantly collinear geometry of the environment. Similarly, when looking at a face like Einstein's in Fig. 8, a higher-level face representation might reshape early contextual interactions to match expected facial characteristics. Although top-down influences on the AF are outside the scope of this model, there is some compelling psychophysical evidence demonstrating the importance of expectation in perceptual fill-in among patients with retinal damage (Schuchard, 1995).

Figure 10 offers a comparison between the perceptual fill-in obtained from our model and the image reconstructions generated by filling-in the scotomata with a non-directional process. While our model uses the simulated LPZ activity to fill-in the lesions along specific directions, the alternative method fills-in each scotoma pixel through isotropic diffusion. In isotropic diffusion, the image along the outside edges of the lesions “flows” non-directionally into the scotomata, so each pixel is filled-in with an average of the pixel values in the image surrounding the lesions. (The isotropic diffusion process is equivalent to performing the function

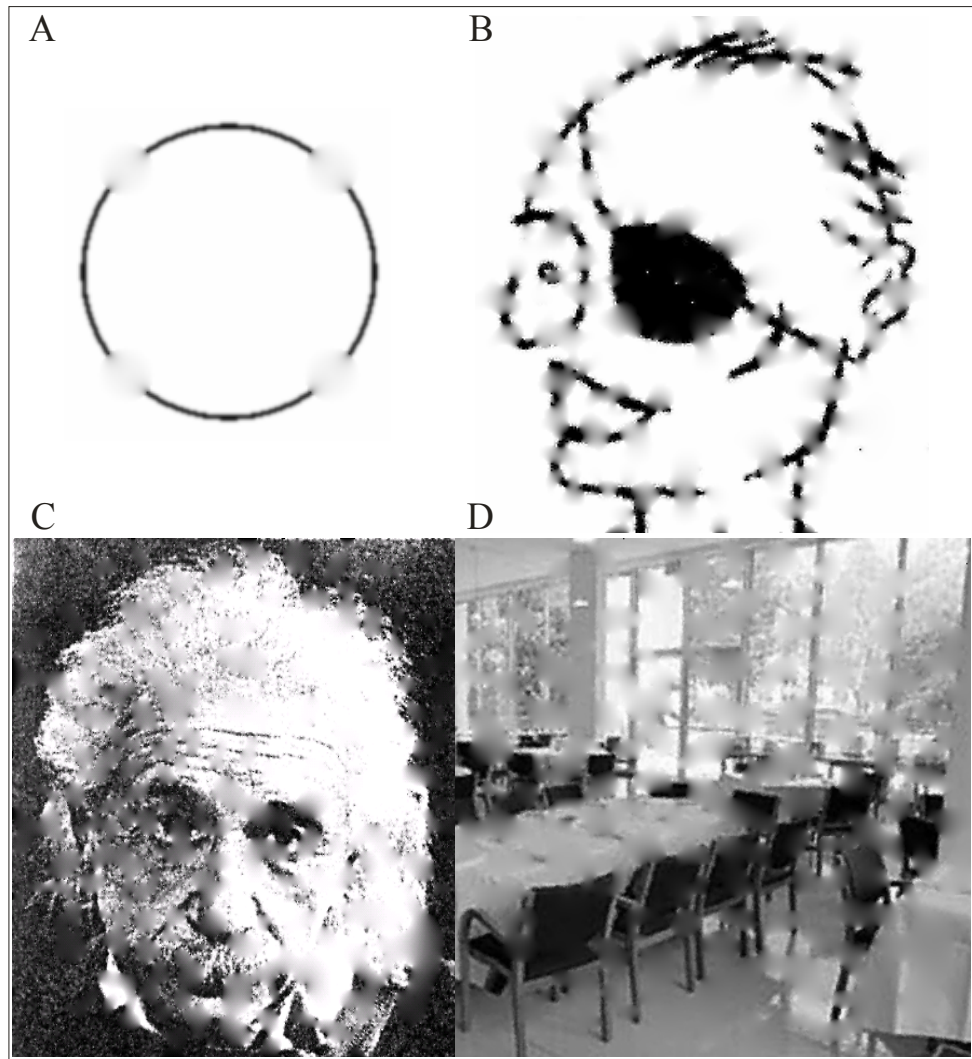


Figure 10. Non-directional image reconstructions. (A-D) results from applying an isotropic diffusion mechanism to the inputs shown in Figures 5, 7, 8, and 9.

minimization described in the *Methods*, but with the parameter q set to zero.) The comparisons in Fig. 10 clearly show the subjective disadvantage of a perceptual completion mechanism that does not employ prior knowledge of the statistical correlations in images (see the *Discussion*). The isotropic diffusion produces results that are characterized exclusively by a hazy, diffuse-looking form of fill-in. And the presence of the retinal lesions is much more evident in the image reconstructions generated by the non-directional method than in the perceptual fill-in from the model. The discrepancy is greatest for scenes that are characterized by a set of salient contours and that lack textured surfaces. While the non-directional filling-in process can produce reasonable reconstructions of certain images, it cannot capture many of the basic psychophysical observations reported by subjects with mild or intermediate retinal degeneration, particularly the completion of smooth contours and the inconspicuousness of the lesions.

We also analyzed the simulated ensemble activity in the LPZ quantitatively. We computed a measure of the angular difference between the image orientations underneath the retinal lesions and the orientations predicted by the model

$$D = \frac{\sum_{(x,y)} \Psi_{x,y} \min\left(\left|\theta_{x,y}^I - \theta_{x,y}^P\right|, \pi - \left|\theta_{x,y}^I - \theta_{x,y}^P\right|\right)}{\sum_{(x,y)} \Psi_{x,y}} \quad (34)$$

where $\theta_{x,y}^I$ is the actual orientation in the image around each “missing” image point (x,y) ; $\Psi_{x,y}$ is a measure of the contrast along that orientation in the underlying image; $\theta_{x,y}^P$ is the corresponding neuronal prediction for the orientation around (x,y) ; and $0 \leq \theta_{x,y}^I, \theta_{x,y}^P \leq \pi$. [$\Psi_{x,y}$ and $\theta_{x,y}^I$ are the “oriented energy” and the “dominant orientation” obtained from the method described in (Freeman and Adelson, 1991).] D , then, is the average of the angular difference between the real and predicted image orientations, weighted by the contrast of the actual orientation in the underlying image. This measure gauges how well the neural activity in the LPZ, despite the absence of visual input, signals the orientation of contours that are present in the underlying image. The average value of D , taken over all the simulation results shown here, is 15.6° , with a standard deviation of 7.4° and a range from 5.1° to 25.9° . That is significantly better than the expectation from random activity in the LPZ, which is 45° . The specific value of D for each simulation result is presented in the corresponding figure captions. Quantitative differences in the value of D resulting from whether the simulations were run with co-circular or collinear connections tended to be

small, although subtle differences in the illustrated perceptual fill-in can be observed (see Figs. 6 and 7). (The two connectivity patterns engender similar network behavior because collinearity is a special case of co-circularity, occurring when the radius of circular curvature is infinite, and because the strongest interactions in both fields lie along the collinear axis.) We also found that the angular discrepancy between the real and predicted image orientations tended to be quite small wherever the underlying image contained a salient contour, but larger elsewhere (data not shown). Not surprisingly, the image reconstructions were more accurate (the value of D was smaller) for the artificial images consisting of only salient contours (like the images in Figs. 5 and 7) than for the photographic images. The perceptual fill-in is mediated by a network of connections that normally underlies contour integration and contour saliency, so the fill-in mechanism specializes in the completion of lines and edges rather than surfaces and textures.

Both methods of evaluation—the subjective evaluation of the simulated perceptual fill-in and the quantitative analysis of LPZ activity—indicate that the neural activity in the LPZ can accurately describe a retinal image

composed of clearly delineated contours, in spite of severe damage encompassing 75% of the simulated retina.

2.5 DISCUSSION

We used AMD as a model to show how reorganization in V1 may engender perceptual fill-in, but the work described here pertains more to the general linkages between cortical reorganization and functional recovery than to AMD per se. Here, we have shown that cortical reorganization via the sprouting of V1 horizontal connections can engender perceptual fill-in by co-opting a preexisting set of connections ordinarily involved in the integration of local stimulus features into globally salient contours. The model makes the novel prediction that if the horizontal connections really underlie perceptual fill-in, then the reorganized RFs should shift along the directions of the geometric association field (either along collinear or circular directions). These results may represent a more general phenomenon, whereby the connectivity and functional architecture that subserve normal integrative mechanisms are advantageously recruited during cortical reorganization following CNS lesions.

Our model is based upon known patterns of connectivity in the visual cortex. The mathematical descriptions of the model neurons and their interconnections, presented in the *Methods*, are derived from a substantial body of experimental work in the normal and reorganized cortex, including anatomical, electrophysiological, optical, psychophysical, and statistical studies (Adelson and Bergen, 1985; Gilbert and Wiesel, 1989; Hirsch and Gilbert, 1991; Darian-Smith and Gilbert, 1994, 1995; Das and Gilbert, 1995a; Bosking et al., 1997; Kapadia et al., 2000; Geisler et al., 2001; Sceniak et al., 2001; Sigman et al., 2001; Stettler et al., 2002; Chisum et al., 2003; Zur and Ullman, 2003). The model is also inspired by a range of work in the theoretical literature on intermediate-level vision and image reconstruction (Sha'ashua and Ullman, 1988; Parent and Zucker, 1989; Perona and Malik, 1990; Guy and Medioni, 1996; Williams and Jacobs, 1997; Ballester et al., 2001).

With its strong experimental and theoretical foundations, our model accounts for many psychophysical observations reported by human subjects with retinal damage. The geometric relationships between the interconnected neurons in our model offer a parsimonious explanation for the robust perceptual completion of lines across retinal lesions. Furthermore, the

limited lateral extent of the long-range horizontal connections is consistent with the observation that lines are perceptually completed across scotomata spanning a few degrees or less, but lines interrupted by larger scotomata spanning 10° or more are often perceived with gaps (Craig, 1966; Gerrits and Timmerman, 1969; Schuchard, 1993, 1995; Zur and Ullman, 2003). Although our model does not address the perceptual fill-in of textures, it does reproduce the general perceptual features described by individuals with retinal damage. Human visual perception is likely determined by neuronal activity throughout the whole visual cortex, but our model simulates specific components of perceptual fill-in that may arise from a specific set of neurons in V1. The model predicts that perceptual fill-in of the lines and edges that form the structure of visual scenes is mediated by the horizontal connections in V1, and that the fill-in of other visual attributes (like texture) involves similar neuronal networks.

The fundamental basis of the model, and the value of cortical reorganization in V1, derives from the long-range statistical correlations that characterize natural images. In natural visual scenes, locally oriented image regions tend to have a co-circular geometric relationship: given an orientation α around point (x,y) , the most probable orientation θ_p around any point (i,j) is co-

circular to the orientation α at (x,y) (Fig. 4E) (Geisler et al., 2001; Sigman et al., 2001). The conditional probability of observing θ_p at (i,j) is strongest when θ_p is collinear with the orientation at the given point (x,y) , but the conditional probability decreases as the curvature of the tangent circle connecting the two points increases and as the distance between the two points increases. This character of natural scenes, combined with anatomical and psychophysical evidence (Field et al., 1993; Geisler et al., 2001; Sigman et al., 2001), provides the basis for the geometry of the association field used by the model. In a network exposed to natural images, neurons with co-circular RFs will tend to fire in coincidence more often than other neurons, in accordance with natural statistics. The principle of Hebbian plasticity, in turn, implies that neurons with co-circular RFs should have the strongest excitatory connection strengths, and that excitatory connection strengths between co-circular RFs should fall off according to distance and radius of curvature. The result is that the neurons in the LPZ are, in some sense, optimal predictors of the image orientations underlying the scotomata, since they are activated by connections which reflect the most likely stimulus configurations. But the usefulness of the cortical reorganization also depends on the perceptual interpretation of the recovered LPZ activity. The model cells represent “labeled lines” for specific regions of visual space, and while

their input structures (their RFs) change during cortical reorganization, the interpretation of their outputs by subsequent cortical stages does not. The activity of the LPZ neurons, which are really responding to stimuli falling outside the scotoma, is interpreted as a reflection of the visual scene falling within the original RF positions.

Still, the assertion that topographic reorganization occurs in V1 is not without controversy (Calford et al., 2005; Smirnakis et al., 2005). In opposition to a large body of electrophysiological and intrinsic imaging evidence for the reorganization in adult V1 (Gilbert et al., 1990; Kaas et al., 1990; Chino et al., 1992; Chino et al., 1995; Das and Gilbert, 1995a; Calford et al., 2000), Smirnakis et al. (2005) have put forward fMRI data suggesting that the LPZ border, as delineated by the BOLD signal, does not contract between the first 2-3 h following the retinal lesions and the ensuing 18-30 wk. An explanation for the apparent lack of reorganization in their study comes from the nature of the BOLD signal used in fMRI studies (for a detailed discussion, see the rebuttal in Calford et al., 2005). The BOLD signal reflects cortical inputs, including subthreshold activation, rather than outputs, as represented by spiking activity; and measuring cortical inputs can lead to a false determination of the LPZ boundary. Interestingly, the

electrode recordings reported by Smirnakis et al. from the cortex *outside* their BOLD-defined LPZ yielded a pattern of clustered, highly-overlapping RF positions known to be signature of cortical reorganization in the LPZ, potentially because the recordings were actually from reorganized cortex that the authors mistook for normal cortex. In any case, although we speculate that the network involved in cortical reorganization is implemented in V1, the model does not require that the remapping occurs in V1. The critical prediction of the model is only that the reorganization occurs in cortical areas where the association field is represented.

In general, any laterally connected network whose connections have been pruned by experience-dependent (Hebbian) learning to extract some stimulus correlation may be able to optimally estimate its characteristic operation when its input has been ablated—given that the network responds by a process of axonal sprouting and synaptogenesis that maintains the initial rules of connectivity (Das and Gilbert, 1995a). The model presented in this paper suggests a mode of reorganization that has an adaptive value in recovering visual perception, and it might represent a more general role of functional architecture—one that encodes information about the statistical structure of sensory input that is normally used to interpret the sensory

environment, but which becomes beneficial in restoring functionality to disrupted cortical regions.

Our model establishes a tentative link between major components of the perceptual fill-in phenomenon and a specific network of laterally connected neurons in V1. Correspondingly, it makes the experimentally verifiable prediction that RFs in the LPZ should shift according to a co-circular AF. If neurons in the LPZ are activated through horizontal connections from the peri-LPZ, then the post-lesion LPZ neurons will inherit the orientation preferences of the pre-synaptic cells that most effectively activate them. Moreover, if the horizontal connections tend to link co-circular or collinear neurons, then any LPZ neuron which maintains its original preferred orientation should be activated by collinear cells from the peri-LPZ. Any LPZ neuron that changes its preferred orientation after the lesion should be activated by co-circular cells from the peri-LPZ. Therefore, if a particular cell maintains the same preferred orientation before and after the retinal lesion, then the cell's RF should shift along the (linear) axis that runs through the original RF and that is parallel to the preferred orientation. And for any cell that changes its preferred orientation after the lesion, its new RF should shift along a circular arc to a position that brings into co-circular

alignment with the original RF position and preferred orientation. Testing these predictions promises to yield a better understanding not only of cortical reorganization, but of the precise functional relationships between laterally interconnected cortical neurons in general.

2.6 ACKNOWLEDGMENTS

We thank D. Zur for involvement early in this research. We also thank W. Li and V. Piëch for helpful comments and technical suggestions.

2.7 FOOTNOTES

1. Since the reader will tend to direct his or her gaze throughout the simulation images, we used the same spatial resolution (i.e., pixel size and image clarity) throughout each image, even though the density of the photoreceptor mosaic, and therefore the spatial resolution of our vision, decreases with retinal eccentricity. This simplifies the presentation and implementation of the model, allows the reader to make saccades when perceiving the results, and it partially mimics the real-life perception of a subject with retinal lesions, whereby spatial attention and saccadic eye

movements typically compensate for differences in resolution across the image.

2. Although the energy model does not capture certain biophysical properties of complex cells (e.g., it squares the outputs of simple cells rather than rectifying them), it does capture the required functional attributes of those cells. The model presented in this paper is not a detailed biophysical model, but a numerical neural network designed to capture the fundamental properties of V1 neurons (like RF size and orientation tuning) and to show how those properties, when embedded in a network of lateral connections, give rise to cortical reorganization.

3. Like many neural network models, we used a piecewise linear function to convert the membrane depolarizations into instantaneous firing rates. We chose the simplest and most straightforward model, in which the neuronal firing rate is a linear function of the membrane potential, the minimum firing rate is 0 Hz, and the maximum firing rate is a finite number. The size of our network precludes continuous-time simulations of spike trains. Therefore we took the common approach of simulating our network with instantaneous firing rates that were updated at discrete iterations of network activity.

4. The cortical distance constraint is applied only in the LPZ, where it is most relevant, because the computational time needed to implement it in the peri-LPZ is prohibitive. The cortical distance computation uses Schwartz's $\log(z + a)$ conformal mapping (Schwartz, 1980), but the size of our model network requires the Fourier convolution theorem to compute the recurrent activity in the peri-LPZ. Since the convolution theorem can only be applied with a spatially homogeneous convolution kernel, and since the cortical mapping procedure would introduce spatial inhomogeneities into the convolution kernel, the cortical distance constraint cannot be practically implemented in the peri-LPZ. For our model neurons in the LPZ, however, the computations are more tractable. Cells in the LPZ receive lateral connections but do not send them out, so the implementation of *Eq. 15* for these neurons is simpler and does not require the convolution theorem.

5. The model parameters were chosen to fit data from the macaque. In the case of the physiological response properties of V1 neurons (like orientation tuning width and RF size), we gave the human V1 cells in our model the same properties as macaque neurons (Dow et al., 1981; De Valois et al., 1982a; De Valois et al., 1982b). In the case of anatomical parameters, like

the cortical magnification factor in V1 and the lateral extent of the horizontal connections, we took the data from the macaque and scaled it to the human (Schwartz, 1980; Stettler et al., 2002). We assumed that the relationship between cortical size, cortical magnification factor, and the lateral extent of the horizontal projections (which determines the visuotopic spread of the lateral connections) was the same in the monkey and in the human.

3 SHAPE SELECTIVITY AND ITS TOP-DOWN MODULATION IN PRIMARY VISUAL CORTEX

Justin N. J. McManus, Wu Li, and Charles D. Gilbert

3.1 SUMMARY

In primary visual cortex (V1), lateral interactions mediated by horizontal connections are thought to follow a specific geometric pattern that underlies the integrative and dynamic properties of V1 neurons. Using an automated search algorithm, we directly mapped the geometry of these lateral interactions in behaving monkeys. We found that, over the population, the optimal interactions linked contour components with collinear and co-circular geometries, giving V1 neurons selectivity for contour shapes and an early role in the perception of object form. Moreover, individual neurons showed optimal responses to different shapes, and their shape preferences shifted when the animal was given different shape cues. Our results point to a general model of cortical function, whereby hardwired horizontal connections establish a broad domain of potential associations, and top-down influences dynamically gate these linkages to perform specific functions pertaining to the immediate task.

3.2 INTRODUCTION

The cortical processing of sensory information is guided by the spatial and temporal context surrounding the sensory stimulus, as well as by the cognitive state of the observer. In visual perception, the appearance of local image regions is determined by the organization of the surrounding visual scene (Albright and Stoner, 2002). These contextual influences shape our interpretation of the color, brightness, size, location, and form of both real and illusory stimuli. The neural correlates of these contextual effects on perception have been traced throughout visual cortex: matching patterns of neural activity have been found for the salience of objects that “pop-out” from the background; the judgment of relative depth; figure-ground segregation and border ownership; the perception of both brightness and color; perceptual grouping; illusory contours; and the perception of three-dimensional “shape from shading” (Zeki, 1983a, b; von der Heydt et al., 1984; Schein and Desimone, 1990; Knierim and Van Essen, 1992; Lamme, 1995; Rossi et al., 1996; Bradley and Andersen, 1998; Rossi and Paradiso, 1999; Sugita, 1999; Nothdurft et al., 2000; Zhou et al., 2000; Kinoshita and Komatsu, 2001; Super et al., 2001b; Lee et al., 2002; Wachtler et al., 2003; Li et al., 2006). In visual areas with small to intermediate receptive fields

(RFs), these contextual influences are manifested by the interaction between stimuli falling within the classical RF (cRF) and stimuli in the extra-classical RF surround. Many laboratories have investigated the basic attributes of these center-surround interactions, including their strength, sign (excitatory versus inhibitory), contrast dependence, selectivity, symmetry, and spatial properties (Maffei and Fiorentini, 1976; Fries et al., 1977; Nelson and Frost, 1978; Allman et al., 1985; Knierim and Van Essen, 1992; Li and Li, 1994; Kapadia et al., 1995; Sillito et al., 1995; Zipser et al., 1996; Levitt and Lund, 1997; Walker et al., 1999; Kapadia et al., 2000; Li et al., 2001). Some of these studies have examined how contextual interactions may give rise to higher-order vision, even in primary visual cortex (V1). For instance, Lamme and colleagues have addressed the role that V1 neurons may play in form perception, through medial axis coding, and object recognition, via the prerequisite process of figure-ground segregation (Lamme, 1995; Lee et al., 1998).

Beyond the global physical characteristics of a stimulus, the cognitive context in which the scene is viewed can shape the observer's perception. Spatial and feature-based attention improve an observer's ability to detect and discern the features of visual stimuli (Posner et al., 1977; Bashinski and

Bacharach, 1980; Rossi and Paradiso, 1995; Handy et al., 1996; Lee et al., 1999; Reynolds and Chelazzi, 2004; Maunsell and Treue, 2006; Wegener et al., 2008). Neurophysiology experiments have found signatures of space-, feature-, and object-oriented attention throughout the visual cortex, including V1 (Moran and Desimone, 1985; Mountcastle et al., 1987; Haenny et al., 1988; Spitzer et al., 1988; Maunsell et al., 1991; Motter, 1993, 1994a, b; Treue and Maunsell, 1996; Chelazzi et al., 1998; Roelfsema et al., 1998; Ito and Gilbert, 1999; McAdams and Maunsell, 1999b; Reynolds et al., 1999; Treue and Trujillo, 1999; Recanzone and Wurtz, 2000; Reynolds et al., 2000; Chelazzi et al., 2001; Roelfsema and Spekreijse, 2001; Reynolds and Desimone, 2003). The view that has emerged from these studies is that attention serves to increase the gain of neural responses to the attended attribute or stimulus, without changing the basic tuning properties of neurons (Reynolds and Chelazzi, 2004). More recent work has demonstrated that attention can modify a particular set of contextual interactions in V1 (Li et al., 2004), and that feature-based attention can alter the spectral tuning properties of V4 neurons (David et al., 2008).

The goal of the present study is to understand the higher-order, integrative properties of cortical processing in V1, in light of the contextual and

attentional influences that shape neural responses throughout the visual cortex. We focus on a set of orientation-specific contextual interactions between locally oriented image segments, interactions thought to be subserved by the long-range horizontal connections running parallel to the cortical surface. In V1, it has been suggested that the horizontal axon collaterals (mediating contextual effects), and their interplay with top-down connections (mediating attentional modulation), are responsible for a broad range of integrative and dynamic visual functions (Gilbert and Sigman, 2007), such as contour integration and contour saliency (Roelfsema et al., 1998; Li and Gilbert, 2002; Chisum et al., 2003; Hess et al., 2003; Li et al., 2006), perceptual fill-in (Darian-Smith and Gilbert, 1994; Das and Gilbert, 1995b; McManus et al., 2008), and perceptual learning (Crist et al., 2001; Li et al., 2004). Still, the geometry of these lateral contextual interactions, the very feature thought to endow them with broad functional significance, has never been fully characterized.

Gestalt psychologists, as early as the 1920s, formulated the law of “good continuation” to describe how the visual system binds local line segments into a perceptually unified contour (Wertheimer, 1923). Good continuation, and the other perceptual grouping phenomena described by the Gestalt

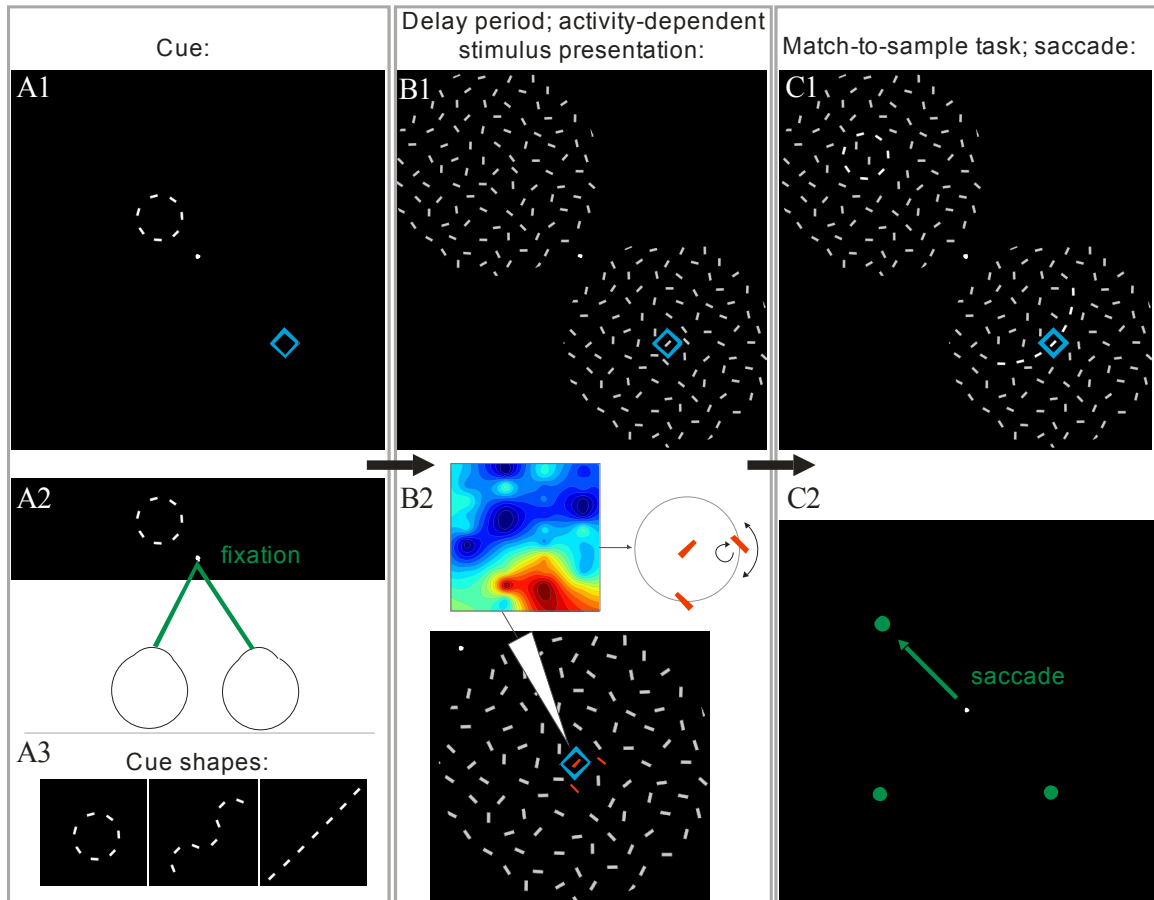
school, are some of the most fundamental contextual effects that shape our perception of object form. Recently, the law of good continuation has been mathematically recapitulated with the notion of an association field (Field et al., 1993; Hess et al., 2003) (AF). The AF is a vector field that specifies the position and orientation of local pairs of line segments that are most readily linked by the visual system (see Fig. 2). These perceptual linkages between line segments may be paralleled by lateral (contextual) interactions in V1, whereby neurons with corresponding receptive field (RF) properties facilitate each other's responses through horizontal connections running between them. Electrophysiological experiments showing contextual facilitation between collinear lines have revealed a portion of the AF in V1 (Kapadia et al., 1995; Kapadia et al., 2000; Li et al., 2006). Furthermore, anatomical (Rockland and Lund, 1982; Bosking et al., 1997; Stettler et al., 2002; Chisum et al., 2003), statistical (Geisler et al., 2001; Sigman et al., 2001), and modeling (Williams and Jacobs, 1997; Li, 1998; Ben-Shahar and Zucker, 2004; Ernst et al., 2004; McManus et al., 2008) studies suggest that the full AF is likely to be implemented in V1 by the long-range horizontal connections. Concomitant with the notion of a geometrically specific AF has been the implicit assumption that the field is static, representing a hard-wired set of connections. Previous results, however, have shown that the

contextual interactions in V1 are subject to dynamic top-down control (Ito and Gilbert, 1999; Crist et al., 2001; Li et al., 2004, 2006), which implies that the geometry of the AF may change dynamically according to perceptual task and expectation.

Here, we measure the geometry of the contextual interactions in V1, and we survey the functional consequences of its task-dependent modulation. We devised an optimization algorithm that searches for the geometric configuration of line elements that maximizes single unit responses. In order to simultaneously examine the topology of the AF and its modulation by task, we trained three monkeys (Monkeys A, B, and C) to detect the presence of a cued contour shape flashed within a patch of random bar elements (Fig. 11). By recording the activity of single neurons during the task, we could algorithmically construct contours that maximally activated the recorded neuron under different task conditions. We observed broad patterns of contextual facilitation that confer, even within V1, selectivity for complex shapes, and that change dynamically with perceptual task.

Figure 11. Trial and task design. Top row: the sequence of stimuli that comprised each trial; bottom row: schematic depictions of the trial events. The cyan rectangle represents the classical RF (minimum response field) of the recorded neuron. (A1 and A2) The monkey initiated each trial by fixating a spot in the center of the monitor (A2); a cue contour was then presented near the fixation point. The same cue was repeatedly presented to the monkey for ~1,000 consecutive trials. (B1 and B2) Following the cueing portion of the trial, two identical, circular fields of line segments were drawn on the computer screen. The fields contained a single geometric stimulus (red bars in B2) surrounded by randomly oriented and positioned line elements. The geometry of the embedded stimulus contour was determined by an optimization algorithm, which searched for the arrangement of contour bars that would maximize the recorded neuron's response. The neural activity recorded during this trial period, averaged over multiple trials for each tested stimulus, was used to measure the neuron's selectivity and to guide the generation of increasingly effective stimuli (B2; see the main text and Fig. 12 for details). (C1 and C2) After a random duration of geometric stimulus exposure, the bar elements in each field were abruptly rearranged to form a salient contour. The monkey performed a delayed match-to-sample task by making a saccade toward the direction of the field containing the cued contour, which could be in either field, or toward an alternative direction if neither field contained the cue. (A3) The three contour shapes that served as potential cues for our experiments. The cue contours and the geometric stimuli were always rotated so that the central bar, which fell in the classical RF, matched the recorded neuron's preferred orientation.

Figure 11



3.3 RESULTS

3.3.1 CONTOUR DETECTION TASK

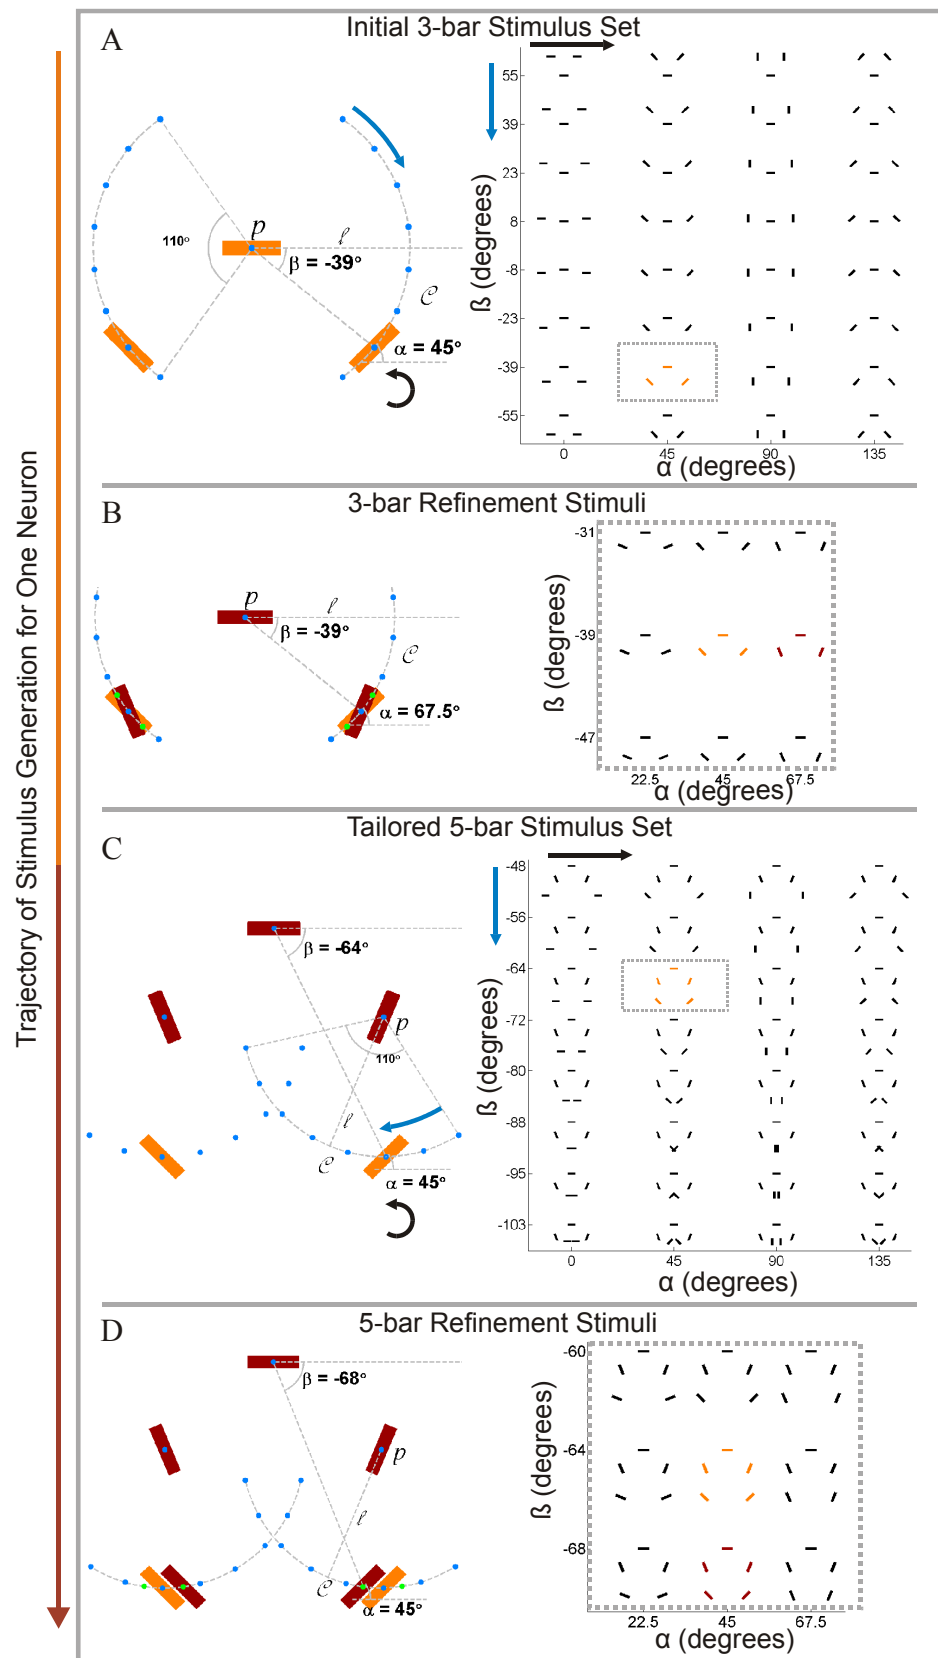
We recorded extracellularly from single units in the superficial layers of V1, from trained monkeys performing a contour detection task. Each experiment comprised hundreds of trials of the detection task, whereby a monkey determined the location of a cued contour in a delayed-match-to-sample paradigm (Fig. 11). During the delay period over the course of these trials, we recorded a neuron's responses to geometric stimuli, while the monkey was expecting the cued shape. Before every experiment, we selected a 7-bar contour (either a closed circle, a sinusoid, or a straight line; Fig. 11A3) to serve as the cue in all of the ensuing trials. The cue was presented next to the fixation point at the outset of each trial (Fig. 11A1), followed by two "contextual patches" of oriented line elements, displayed during the delay period (Fig. 11B1). One of the two contextual patches was centered over the RF of the recorded neuron; the position and orientation of the central bar in that patch were always set to match the neuron's RF location and preferred orientation. The second of the contextual patches was identical to the first, but was positioned in the opposite visual field quadrant. Most of the line segments in each patch were randomly positioned and oriented, but a few of

them were aligned in a particular geometric configuration with the bar at the patch center. We refer to these geometrically aligned elements as a contour or a geometric stimulus. The array of bars into which the geometric stimuli were embedded is called the “contextual background”. In each trial, during the delay period (Fig. 11B1), an embedded geometric stimulus was chosen from a set of contour shapes created by the contour optimization algorithm (see Fig. 12 and the following section for details). The neural responses to these embedded contours, each repeated over multiple trials, were used by the algorithm to determine the preferred contour shape, upon which the generation of new stimulus sets was based (Fig. 12). With each new stimulus set, the algorithm built up contours that elicited progressively stronger responses, leading to an approximation of the recorded neuron’s preferred stimulus shape. At the end of the trial, a 7-bar contour was flashed in each contextual patch, just before the stimuli were extinguished (Fig. 11C). In two-thirds of the trials, the cued contour was briefly embedded in either patch with equal probability, while a distracter contour (an open circular arc) was embedded in the other. The animal’s task was to make a saccade toward the location where the cued contour was flashed (Fig. 11C2). In the other third of the trials (catch trials), the distracter contour was flashed in both patches; the task in that case was to make a saccade toward a third,

alternative location. We interleaved the geometric stimuli between presentations of the cue, at the beginning of each trial, and the behaviorally relevant contours, at the end, in order to engage top-down mechanisms during the delay period when neural responses were measured.

Figure 12. Optimization algorithm for stimulus generation. (A-D) the geometric stimuli that were created by the optimization algorithm for an example neuron; the cell's responses to these stimuli are shown in Figure 3. The right side of each panel illustrates the stimuli that were generated during each stage of the optimization algorithm; the left side depicts how these stimuli were created. All stimuli are defined by two coordinates: α , the orientation of the bar on the right arm of the contour; and β , the polar angle (position) of the right contour bar with respect to the central bar in the classical RF (cRF). The stimuli always had bilateral symmetry, so the position and orientation of the bar on the right arm of the contour determine the geometry of the corresponding bar on the left arm. The letters p , ℓ , and \mathcal{C} refer to variables that are defined in the *Experimental Procedures*. (A) The initial stimulus set was identical for all neurons. It consisted of 32 distinct 3-bar stimuli. The central bar was centered in the cRF; the bars outside the minimum response field took one of 8 positions (left: blue dots), with 4 possible orientations at each position (0° , 45° , 90° , 135°), making for 32 different stimuli (right). The stimulus from this set that elicited the strongest response from the neuron is shown in orange ($\alpha = 45^\circ$, $\beta = -39^\circ$). (B) the stimulus set generated after the best contour was selected from among the original stimuli. This set of refinement stimuli consisted of 8 new stimuli, in addition to the best stimulus from the previous stage. The new stimuli took one of three positions, indicated by the green points and the preferred blue position from the preceding stage (left), and one of three bar orientations: $\{22.5^\circ, 45^\circ, 67.5^\circ\}$. Closely spaced in both position and orientation, the stimuli were chosen to refine the estimate of the cell's preferred 3-bar stimulus. The final estimate (red) had the same bar position as the orange contour, but a bar orientation of 67.5° rather than 45° . (C) the 5-bar stimuli generated during the third stage of the optimization algorithm, using the optimum 3-bar stimulus as a seed configuration to search for longer contours. The stimulus set was constructed as in (A), but by using the position and orientation of the bar on the right side of the optimum 3-bar contour, rather than the central bar in the RF, as the basis for the relative positioning and orientation of the 4th bar (and the symmetrically positioned 5th bar). Orange: the best configuration from this stimulus set. (D) the 5-bar refinement stimuli, created around the best stimulus from the previous stage; red: the improved estimate for this cell's 5-bar optimum.

Figure 12



3.3.2 STIMULUS GENERATION ALGORITHM

Figure 12 traces a typical sequence of geometric stimuli created by the optimization algorithm for a single neuron. The stimulus trajectory unfolded from the recorded neuron's activity while Monkey C was cued to detect the sinusoid target (the corresponding neural responses from this cell are shown in Fig. 13). The stimuli are depicted without the surrounding field of randomly oriented lines, and they are drawn in a horizontal orientation (the actual stimuli were rotated to the cell's preferred orientation). The optimization program was initialized with a set of 3-bar stimuli (Fig. 12A), consisting of the central bar in the cRF and pairs of symmetric line elements positioned just outside the minimum response field. The initial stimulus set, with 32 different contours, was intended to broadly but comprehensively sample the space of 3-bar stimuli. All of the neurons in our experiments, under all task conditions, were exposed to this same region of the stimulus space at the beginning of the optimization algorithm. The right side of Fig. 12A shows these 32 stimuli, arrayed in a 2-dimensional coordinate system that describes the geometry of each stimulus (see the legend). The stimuli were repeatedly shown to the neuron over a block of randomly interleaved trials, during the delay portion of each trial. After 20 repetitions of each stimulus, the orange contour in Fig. 12A was found to be the best (i.e., most

facilitatory) among this initial stimulus set. The optimization algorithm then generated a second, smaller set of stimuli to refine its estimate of the neuron's preferred contour within the surrounding stimulus space (gray rectangle; Fig. 12, A and B). Following another 20 repetitions of each of these stimuli, the algorithm settled on one (highlighted in red, Fig. 12B) as the best 3-bar configuration. The next stimuli generated by the algorithm, a set of 5-bar contours, were based on this optimum 3-bar seed (see the *Experimental Procedures*). Panel C shows these new stimuli, drawn in the coordinate system that defines their shape. These 5-bar stimuli all share the same three bars at their center, but represent different attempts by the algorithm to find the smooth, optimal extension of the 3-bar seed. The stimulus highlighted in orange (Fig. 12C, right) elicited the highest neural response among all the other initial 5-bar stimuli. The gray rectangle (Fig. 12, C and D) marks the region of the stimulus space, in the neighborhood of the orange contour, that was searched during the subsequent iteration of the algorithm. Panel D (right) provides a magnified view of the area, with the contours created during the final stage drawn in the stimulus space. As before, the algorithm's refinement step revealed a response peak for a stimulus (red) that differed slightly from the original estimate (orange). The best stimulus for this neuron consisted of pairs of line segments that were

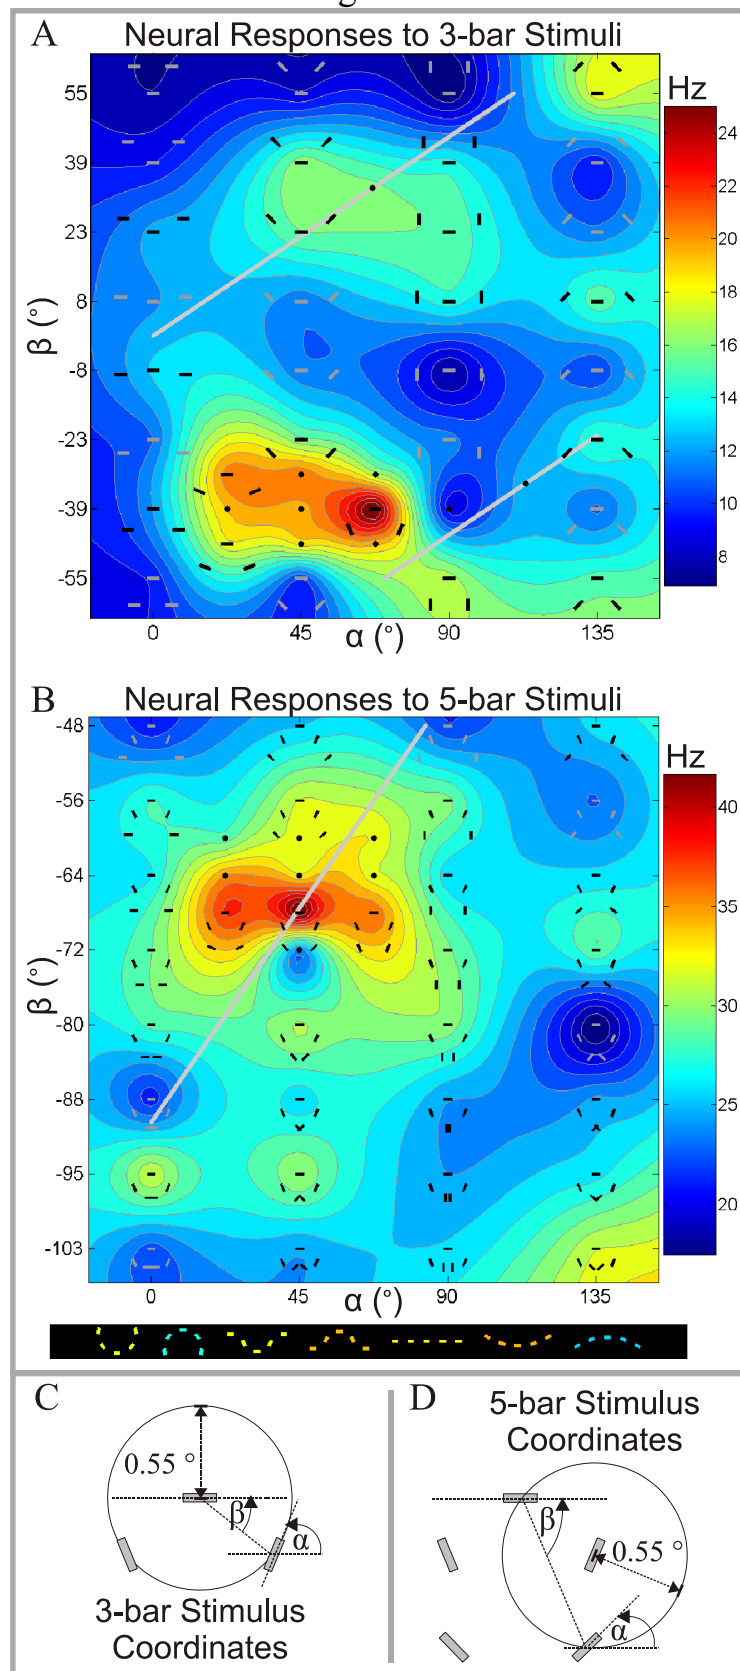
each approximately co-circular to the bar in the cRF (see Figs. 2 and 13), and which together formed an approximation to a closed circle.

3.3.3 GEOMETRIC TUNING SURFACES

To visualize neural responses to the sequential stimulus sets generated by the optimization routine, we constructed three-dimensional “tuning surfaces” (the higher-dimensional analog of tuning curves; see the *Experimental Procedures*). The surfaces are displayed as heat maps (e.g., Fig. 13): the two-dimensional coordinates in the plane define a range of stimulus geometries, as described in Fig. 12, and the surface height—indicating the neural response to each stimulus—is depicted with a color scale. In our experiments, as described above, separate sets of stimuli were successively generated during an automated search for each neuron’s optimum contour. These stimulus sets were composed of increasingly long contours as the experiment progressed. In the response plots for each neuron (e.g., Figs. 13-16), the separate surfaces (heat maps) show the responses to 3-, 5-, and, in some cases, 7-bar contours of various geometries. The peaks on each surface indicate the contour shapes that gave the highest firing rates, among other

Figure 13. Representation of neural shape selectivity. The responses of a single neuron (from Monkey C) to the dynamically generated stimuli shown in Fig. 12. (A) the neural response function, represented as a heat map, for the 3-bar stimuli. (B) the response map for the generated 5-bar stimuli. (A and B) The heat maps were created by plotting the neural response to each tested stimulus and interpolating the data over the stimulus plane with Hardy's multiquadratics. The magnitude of the neural response is given in Hertz. As in Fig. 12, the stimuli are defined by a pair of angular coordinates (α , β), and the neural responses are plotted as a function of these defining stimulus parameters. For ease of visualization, the tested geometric stimuli are drawn over the data; for each contour, the position of the central bar corresponds to the (α , β) coordinates of the stimulus. The contours from both the initial and refinement stimulus sets are drawn, but wherever the positions of stimuli in the α - β plane are too close to clearly show all the stimuli, the most effective stimuli are drawn explicitly while the others are indicated with points. (The intensity of the stimuli in the tuning surfaces was adjusted simply to improve their visibility.) The gray "co-circularity lines" are the loci of stimuli, with various radii of curvature, whose outermost contour bars are co-circular with the central bar in the cRF. This neuron exhibited clear peaks of facilitation near the co-circularity lines on the 3- and 5-bar heat maps; the optimum 5-bar contour for the cell resembled a closed circular contour. These data were recorded while the monkey performed the wave detection task. (B) The contours below the 5-bar surface are segments taken from the center of the targets and distracters. These segments were included in the 5-bar stimulus set (see the *Experimental Procedures*) along with the stimuli created by the optimization program; the neural response to each contour is indicated by the color of the stimulus bars. (C and D) schematics illustrating the meaning of the stimulus coordinates for 3- and 5-bar stimuli, using the optimum 3-bar (C) and 5-bar (D) contours as examples. Throughout this work, the geometric stimuli are drawn as though they were horizontally oriented, but the actual stimuli were rotated to match the preferred orientation of the recorded neuron.

Figure 13



contours of the same length but different geometries. In most of our experiments, we compared the 5-bar stimuli that arose during the optimization program with an additional set of 5-bar contours. These were segments taken from the middle of the cue shapes (the closed circle, wave, and line) and the distracter (open circle). Whenever they were included, the neural responses to these stimuli, as well as their mirror images, are indicated by the correspondingly colored shapes below the 5-bar tuning surface (e.g., Fig. 13B). The gray lines drawn over the surfaces are the stimuli whose outermost contour bars are co-circular with the central bar in the RF (for a description of co-circularity, see Fig. 2). In order to facilitate the interpretation of these figures, we have drawn the tested geometric stimuli over their corresponding coordinates in the plane. Wherever the stimulus positions in the plane are too tightly clustered to draw all the stimuli, the stimuli that elicited the highest firing rate are drawn; the positions of nearby stimuli that were also presented to the animal are indicated with points.

Figure 13 illustrates a single neuron's responses to the dynamically created stimuli shown in Fig. 12, when the animal was cued to look for the wave cue. In this case, the optimization routine was run until it converged on the cell's

preferred 5-bar contour, so the figure contains two surfaces, separately representing the cell's selectivity for 3- and 5-bar stimuli. At the bottom, Panels C and D schematize the (α , β) coordinates in which the surfaces are plotted. The stimuli used in our experiments were always mirror symmetric about the central bar in the RF, since the size of the stimulus space is too large to search comprehensively without this constraint. In order to completely describe the whole contour, it is therefore sufficient to specify the position and orientation of the bars on just one side of the stimulus. To enable comparison across neurons, the stimuli are always plotted in a coordinate system where the central bar is horizontally oriented, even though the actual stimuli were rotated to match the cell's preferred orientation. The β coordinate, plotted along the vertical axis, describes the position of the bar at the right edge of each contour. Since adjacent contour bars were always separated by a center-to-center distance of 0.55 degrees of arc, the position of the rightmost bar is defined by its polar angle (β) with respect to the central bar. The α coordinate, plotted along the horizontal axis, describes the angular orientation of the rightmost bar with respect to the horizontal.

The surface in Fig. 13A depicts the cell's responses to 3-bar stimuli. (Note the correspondence between the stimuli drawn over the neural response and

the contours drawn in Fig. 12, A and B.) An island of activation is prominent near the lower left corner of the heat map (for $(\alpha, \beta) \in [22.5^\circ, 67.5^\circ] \times [-31^\circ, -47^\circ]$), with a peak of near co-circular facilitation (red stimulus in Fig. 12B). This cell also expressed lower levels of facilitation for circular contours of high curvature [e.g., the stimuli at $(135^\circ, 55^\circ)$; $(90^\circ, -55^\circ)$] and for circular arcs of somewhat lower curvature (the local maximum, in green, along the co-circularity line). After discovering the global maximum on the 3-bar surface, the optimization algorithm constructed the 5-bar tuning function in Panel B. The neural selectivity for closed contours was recapitulated in the 5-bar surface, with local and global maxima occurring at stimulus coordinates $(0^\circ, -95^\circ)$, $(135^\circ, -103^\circ)$, and $(45^\circ, -68^\circ)$. Interestingly, the peak response occurred for a stimulus, at $(45^\circ, -68^\circ)$, whose outermost contour bars are precisely co-circular to the bar in the classical RF (they are tangent to opposite sides of the same circle). Below Panel B, the contours drawn over the black background are the 5-bar segments of the cue and distracter shapes (see the *Experimental Procedures*). The color of these shapes is directly proportional to the response they elicited, according to the same color scale used by the heat map. These stimuli (referred to here as “5-bar target contours”) were included as smooth *a priori* additions to the 5-bar stimulus set, but none elicited the strongest neural response. In fact, we may

have overestimated the neural response to the 5-bar target contours, compared to the stimuli on the 5-bar response surface. All the contours on the response surface, but not the predefined target contours, share the same three contour bars at their center. If neurons experience any adaptation to the repeated presentation of this configuration, then the optimized stimuli will experience a competitive disadvantage relative to the target contours. Therefore, the selectivity of V1 neurons for the optimized stimuli, compared to the target contours, may be even greater than reported here.

Figures 14-16 show additional examples of the response functions collected from different neurons under each of the three task conditions, when the animal was cued to detect the line (Fig. 14), circle (Fig. 15) and wave (Fig. 16) shapes. The tuning surfaces in these plots reveal the diverse patterns of geometric selectivity we encountered. Neurons were highly selective for stimulus components outside the classical RF, and for shapes beyond the regime of previously reported collinear interactions (Kapadia et al., 1995; Kapadia et al., 2000; Li et al., 2006). Across the neuronal populations from all three monkeys, individual cells showed a preference for a relatively narrow range of contour shapes, out of a broader landscape of surveyed geometries. Different cells expressed distinct modes of selectivity,

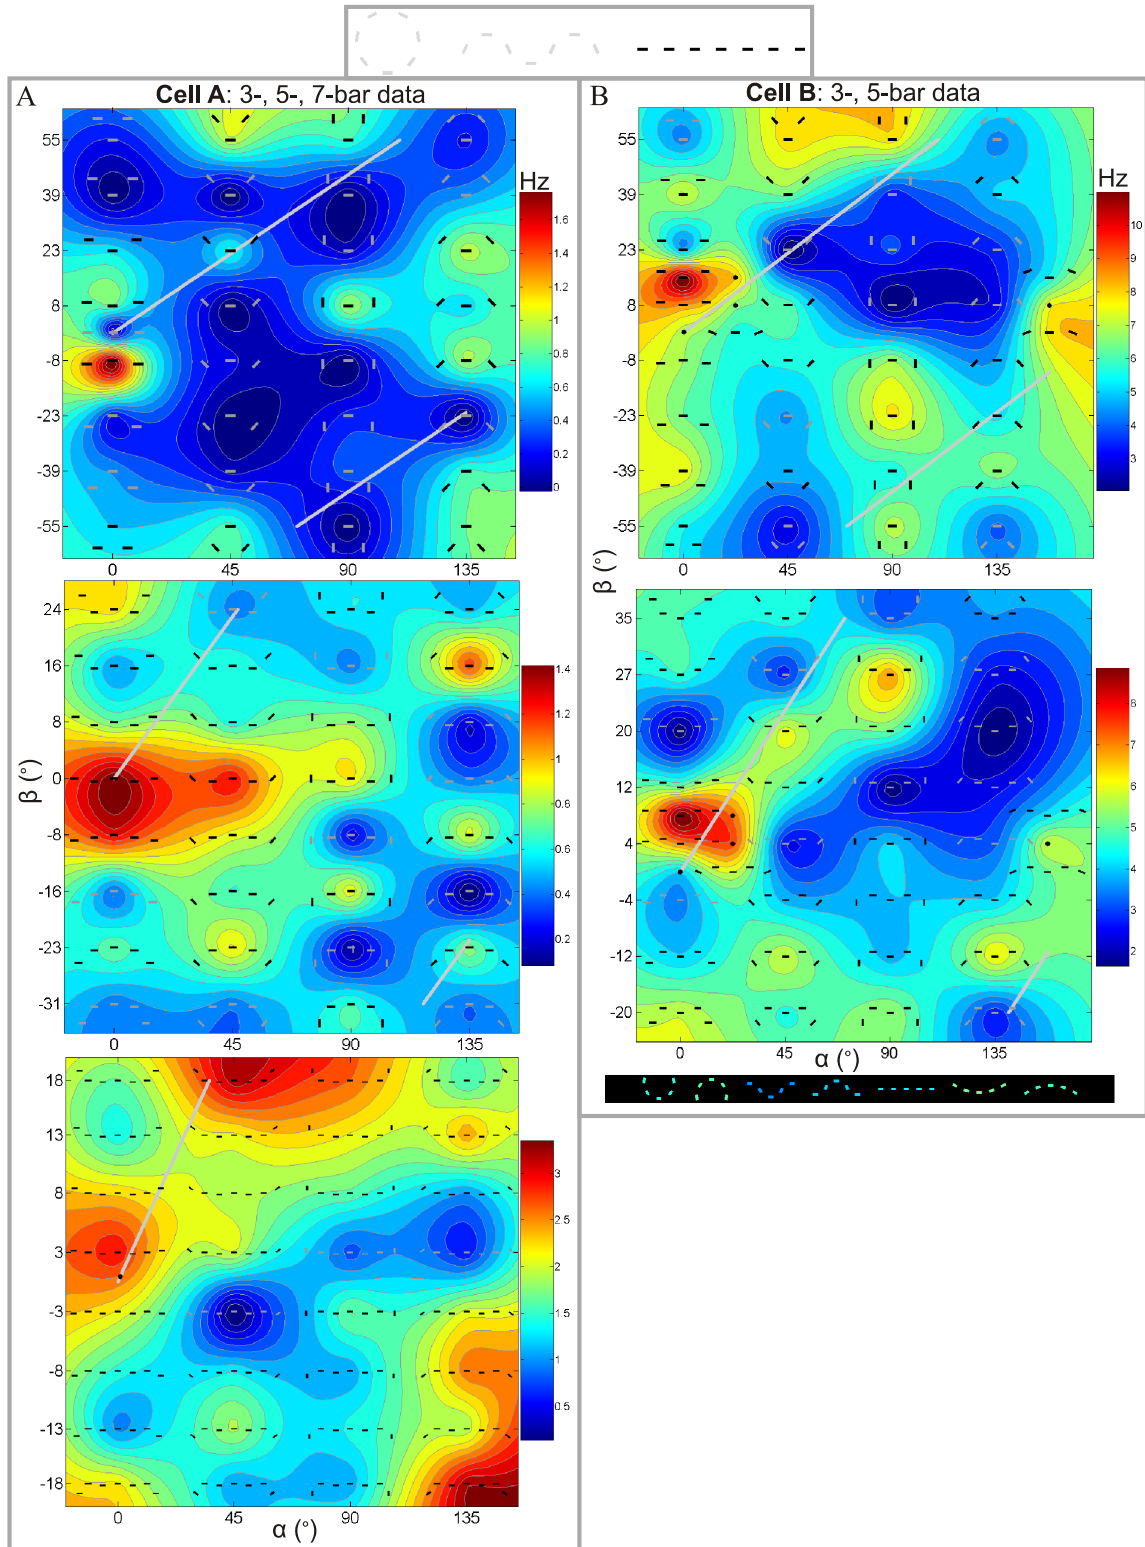
corresponding to optimum contours with various degrees of curvature and smoothness. This shape selectivity, via the expression of specific contextual interactions, was a ubiquitous property of the recorded V1 neurons. It is known from previous work that this geometric facilitation represents a release from the inhibition imposed by the contextual background (Li et al., 2006). We performed a series of eye tracking analyses (see the *Supplemental Experimental Procedures*) to demonstrate that the geometric selectivity was not an artifact of eye movements (Fig. 22, *Supplemental Figures*) and that eye movements did not cause systematic changes in neural firing rates (Fig. 23, *Supplemental Figures*).

Figure 14 recapitulates the well-known pattern of collinear facilitation in V1, but it also reveals the *selectivity* of the facilitation by comparing it to the responses obtained from other geometries. Panels A and B show the tuning surfaces obtained from two cells that were recorded during the line detection task (using the linear contour as the cue). Both neurons had peaks near the coordinate $(0^\circ, 0^\circ)$ on all of their tuning surfaces, corresponding to stimuli with co-aligned, iso-oriented bars. While many neurons were selective for perfectly collinear contours under this task condition, these cells preferred contours with small lateral offsets between adjacent contour bars. Fig. 14A

shows the tuning of one cell to the 3-, 5-, and 7-bar stimulus sets that were tailored to the neural responses. After the optimization routine settled on the most facilitatory 3-bar stimulus at coordinate (0° , -8°), two additional bars were appended to its ends in various configurations, creating a set of 5-bar stimuli used to search for the optimal extension of the contour. The peak at collinearity on the 5-bar surface illustrates the cell's response to the optimum configuration of bars at the ends of these contours, which extended well beyond the cRF. The local maximum at (0° , 3°) on the 7-bar surface demonstrates selectivity for the position and orientation of contour components several cRF diameters away from the RF center. For stimuli beyond 5 bars in length, however, neural responses to differences at the contour ends were often less selective, so we did not attempt to construct contours longer than 7 bars. The broader selectivity for longer contours was likely due, in part, to the increasing similarity between the competing contour shapes that were generated as the algorithm progressed. The cell in panel B shows a similar pattern of collinear facilitation, from an experiment in which the optimization algorithm was stopped after 5 bars. This cell was particularly sensitive to small changes in the lateral positioning of the contour bars: although the optimal contour for the cell was approximately collinear, the neuron responded very weakly to perceptually similar contours,

Figure 14. Shape selectivity for two neurons recorded during the line detection task. Left, from top to bottom: the responses of one cell (Cell A, Monkey A) to the 3-, 5-, and 7-bar stimuli that were generated by the optimization routine. Right: another example neuron (Cell B, Monkey C), for which the optimization routine was stopped after finding the optimum 5-bar contour. Both columns of data were recorded during the delay period after the monkey had been cued to detect the linear target; the neurons displayed a corresponding preference for linear geometries that was apparent for 3- and 5-bar stimuli. The labels “Cell A” and “Cell B” apply only to this figure; in other figures, the same labels refer to different neurons.

Figure 14

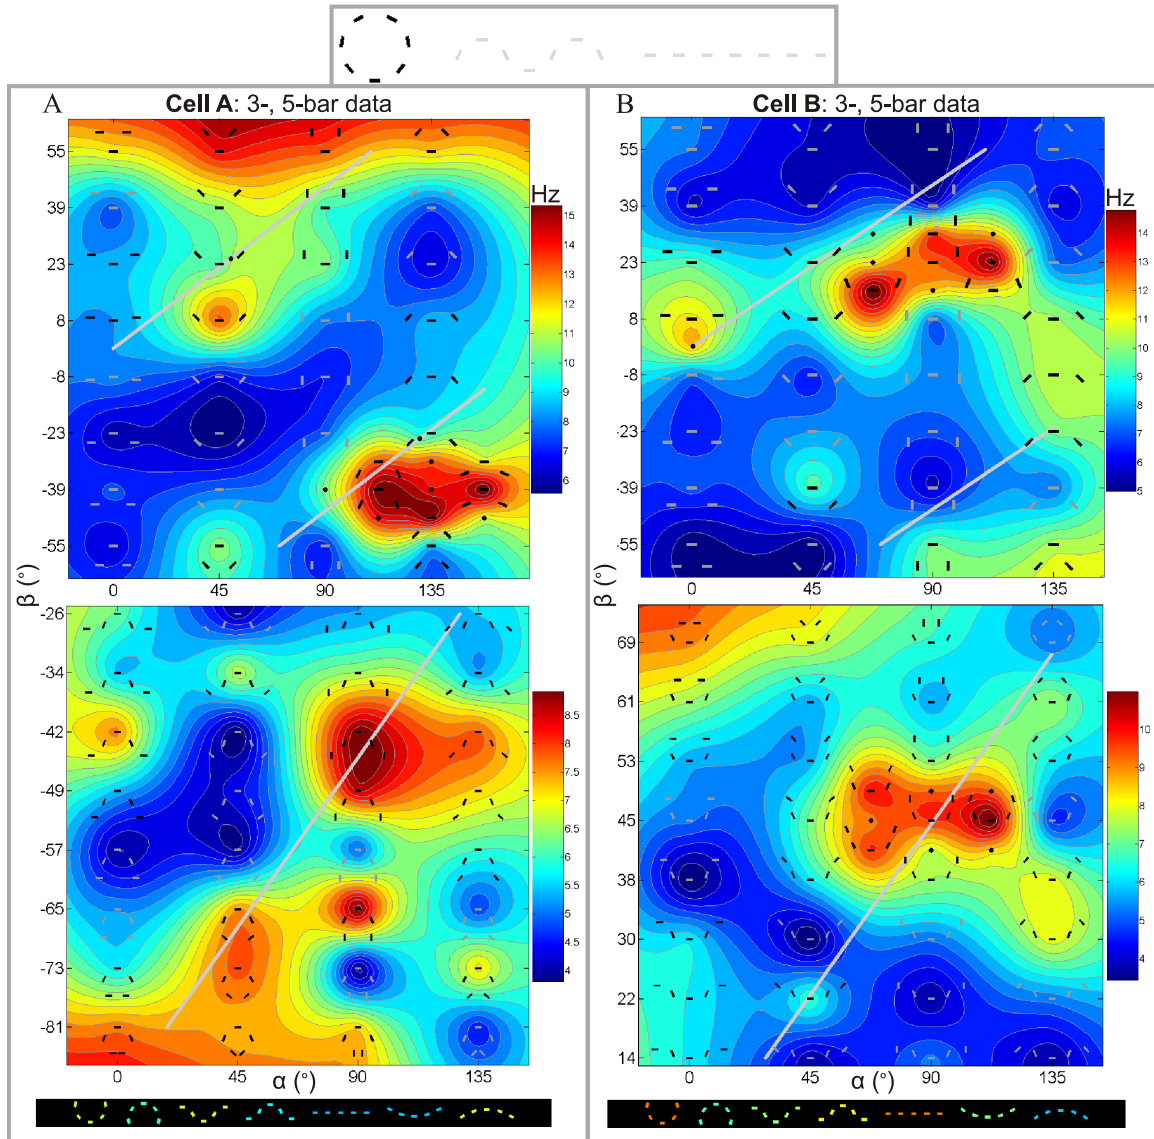


including the 3- and 5-bar segments from the strictly linear target to be detected.

Figures 15 and 16 return to a regime of facilitation characterized by curved geometries. Fig. 15 describes distinct patterns of circular facilitation from two cells recorded during the circle detection task. Both of the cells preferred curved, nearly co-circular stimuli, with each neuron favoring different radii of curvature. In its 3-bar tuning surface, Cell A expressed two broad regions of robust activation, each corresponding to a set of circular contours in roughly opposite orientations. This neuron also had two peaks along the co-circularity line on its 5-bar surface. Both peaks corresponded to stimuli whose outermost contour bars were co-circular to the bar in the RF. Close to one peak, all the contour bars in each stimulus [e.g.: $(45^\circ, -65^\circ)$; $(45^\circ, -73^\circ)$] rested along nearly the same circular contour; but close to the other peak, the outermost contour bars lie along a broader circular arc than the innermost three bars [e.g.: $(90^\circ, -42^\circ)$; $(90^\circ, -49^\circ)$]. Cell B, in its 3-bar tuning surface, has a prominent peak of facilitation along the co-circularity line, $(\alpha, \beta) \in [67.5^\circ, 112.5^\circ] \times [16^\circ, 31^\circ]$, from which lesser lobes of facilitation stretch outward toward other portions of the line. The 5-bar surface for this cell is punctuated by a central peak of co-circularity, with a

Figure 15. Circular optima under the circle detection task. The 3-bar (top) and 5-bar (bottom) tuning surfaces of two neurons recorded while Monkeys B (left) and C (right) were engaged in the circle detection task. Both cells expressed clear selectivity for circular geometries, but they were maximally activated by largely non-overlapping regions of stimulus space. In their 5-bar tuning functions, each cell expressed selectivity for both open and closed circular contours.

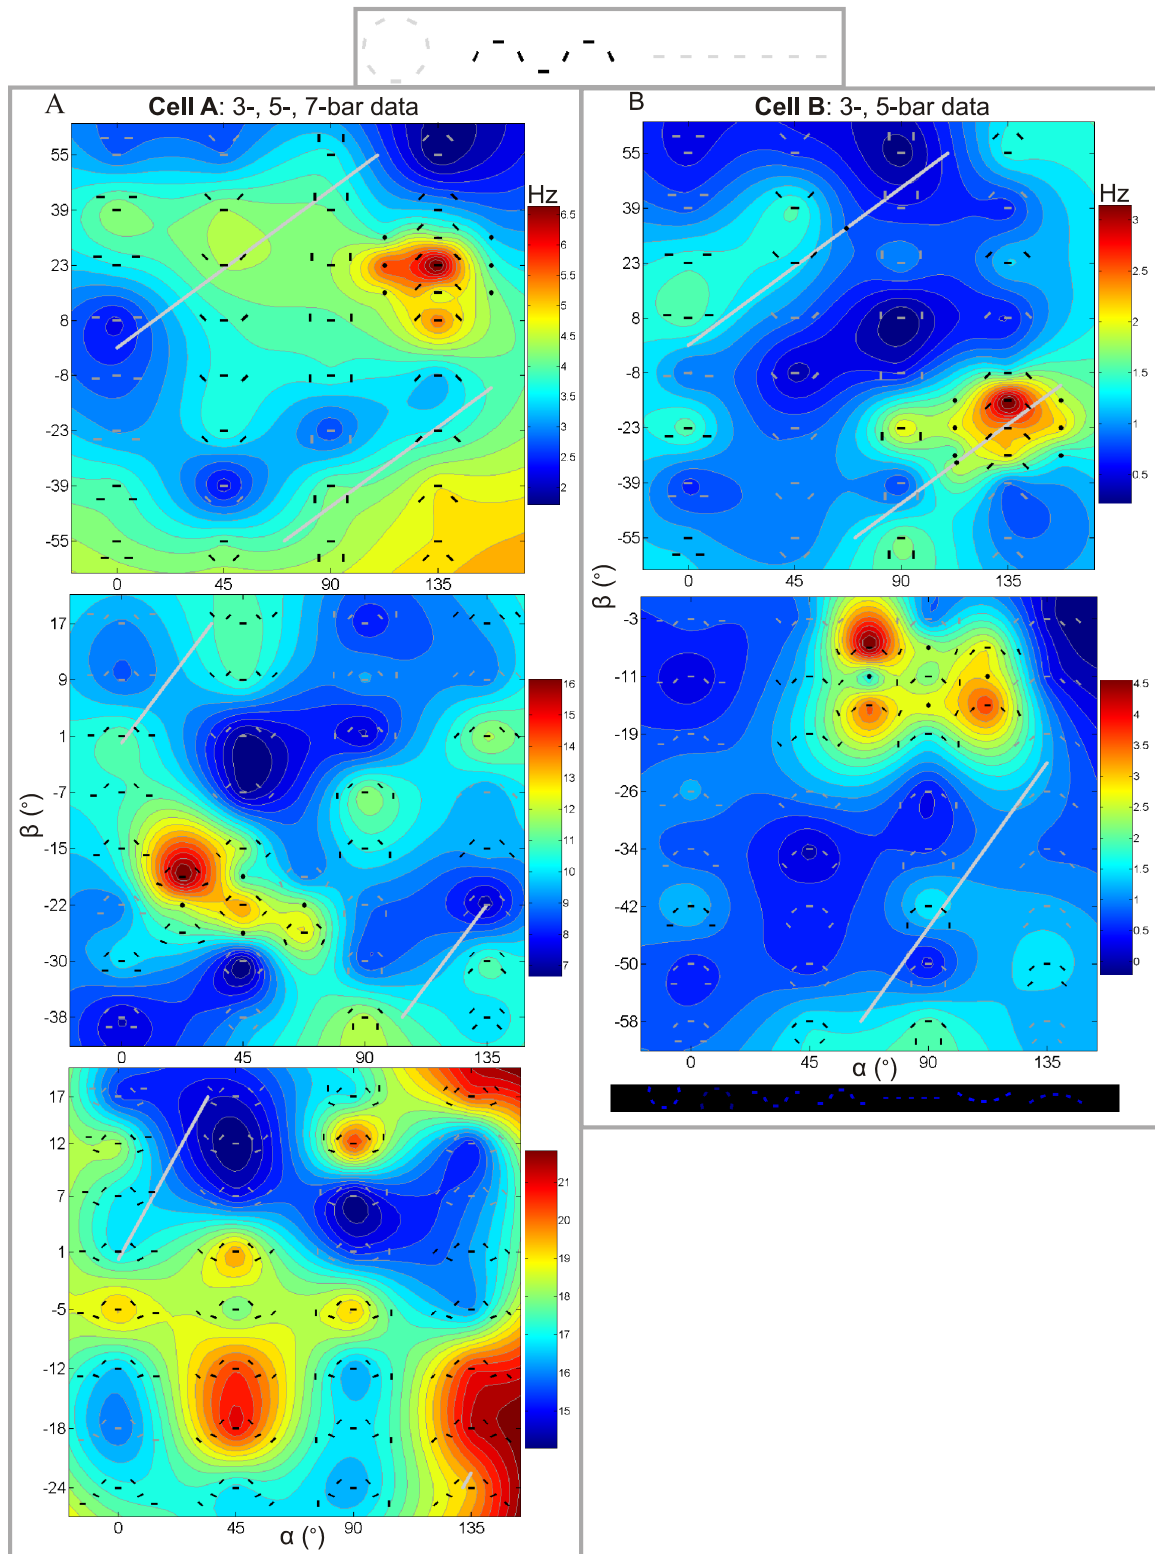
Figure 15



radius of curvature that matches the cued target. Neural responses increased at the upper left edge of the surface near the point (0° , 69°), showing selectivity for contour closure. The horizontal panel below the figure demonstrates that the cell responded strongly to both the circular and linear shapes. In particular, the response to the 5-bar line mirrors the local maximum seen for the same geometry on the 3-bar surface [at coordinate (0° , 0°)]. Although the optimal contour for this cell was circular, it was not exclusively selective for circular arcs. Figure 18E provides an example of a cell that preferred co-circular interactions with a smaller radius of curvature than the two cells in Fig. 15 (also see Fig. 21, *Supplemental Figures*). The examples described above demonstrate characteristic features of the data. First, neuronal response optima were seen as one or more islands of activation—regions of the stimulus space representing contours with a particular spatial scale, radius of curvature, and orientation. Moreover, neurons often contained secondary, suboptimal peaks of facilitation, which were geometrically distinct from the primary regions of activation. A final observation is that neurons contained rotational asymmetries in their responses: neurons were often selective for a contour of a particular orientation, and only weakly responsive to the same contour rotated by 180° .

Figure 16. Heat maps with wave-like optima under the sinusoid detection task. The stimulus evolution and neural responses for two cells recorded while Monkeys A (left) and B (right) performed the sinusoid detection task. The geometric optima expressed by these cells are complex configurations of line segments, whose positions and orientations suggest contours with reversing directions of curvature.

Figure 16



The final class of geometric selectivity we observed is depicted in Figure 16, which shows the responses of two cells recorded during the sinusoid detection task. Both cells were selective for perceptually “wave-like” shapes that exhibited changes in direction and curvature. The 3-bar surface for Cell A has a feature that distinguishes it from the heat maps in Figs. 13-15: the global maximum is far from either co-circularity lines. The same characteristic is also apparent in the 5- and 7-bar data. Despite the undulating shape of this cell’s optimum contours, the selectivity of the cell is as sharp as the line-selective neuron in Fig. 14A, recorded under the line detection task. These data highlight the more basic observation, seen throughout our data, that the smoothest contours do not always elicit the strongest responses. The 3-bar selectivity for Cell B is seen as a narrow peak for circular stimuli of intermediate curvature, but the 5-bar optima—which out-compete the stimuli in the bottom panel—are “rippled” or disjointed. Certain modes of circular facilitation also included disjointed (but not necessarily wave-like) stimuli, in which the line segments were individually co-circular with the bar in the RF, but did not collectively fall along the same contour (Fig. 15A; Fig. 21B). Another common feature of the data is illustrated by the 7-bar heat map in Fig. 16A (the bottom panel). Contours that fold back on themselves, like the shapes in the upper-right corner of this

surface, can elicit strong activity by mimicking the structure of their innermost contour bars with the bars at their outermost edges (also see Fig. 21B).

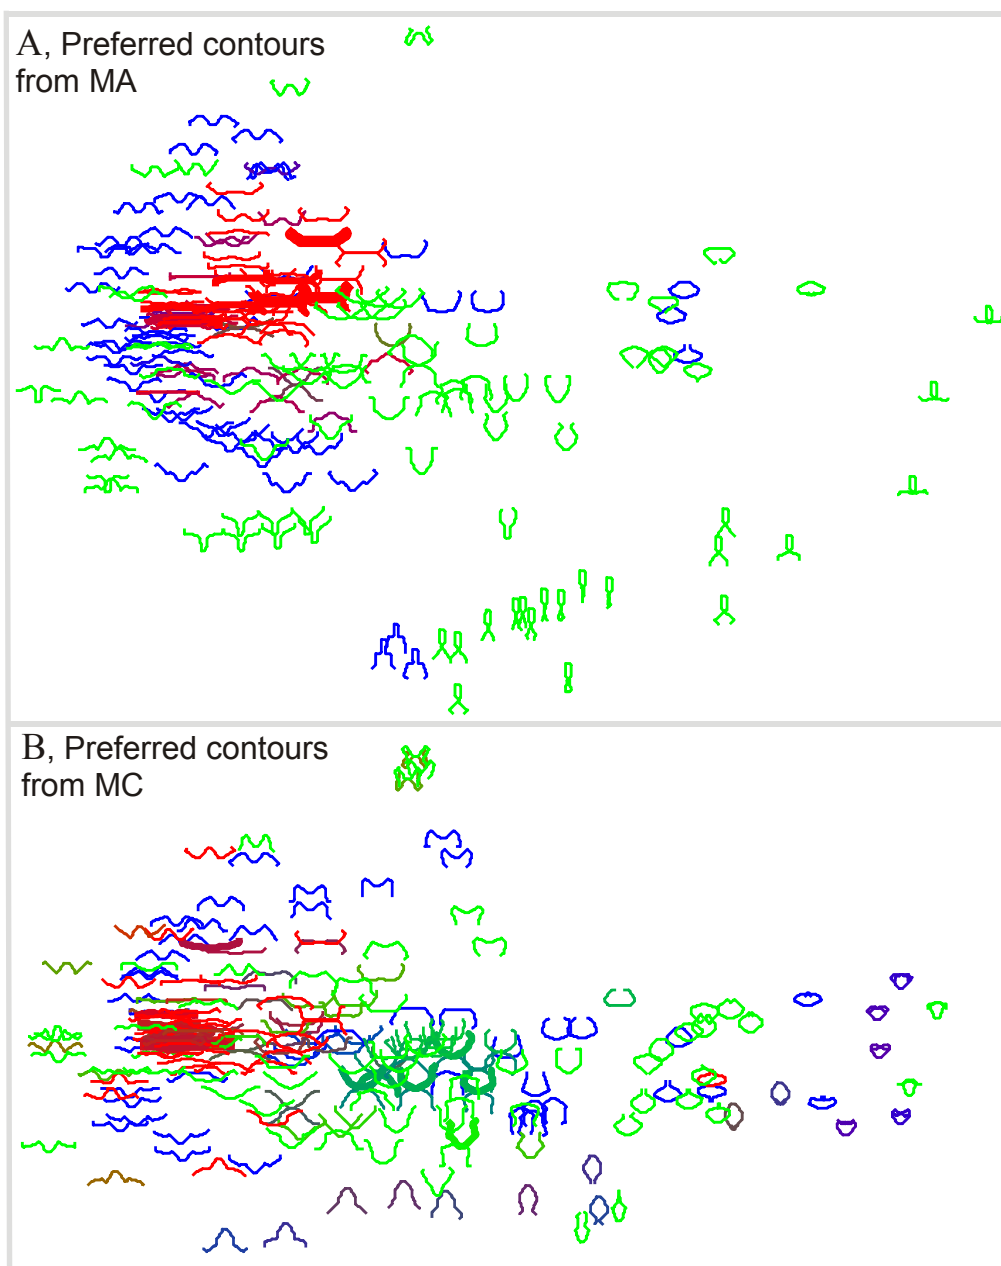
3.3.4 SHAPE SELECTIVITY AND TASK DEPENDENCE OVER THE POPULATION

The examples of collinear and curved facilitation shown in Figs. 13-16 were recorded from single cells while the monkey was cued to detect a particular target shape (either the line, circle, or wave) throughout the experiment. In Figure 17, we plot the most facilitatory 5-bar stimuli from all neurons in Monkeys A and C, to show how the preferred geometry varied as a function of the cued target shape. Here, the stimuli are represented as continuous curves by directly connecting the ends of adjacent contour bars in the original stimuli. We used a multidimensional scaling (MDS) approach known as Sammon's mapping (Kohonen, 2001) to project the 5-bar stimuli, which exist in a high dimensional space, onto the two-dimensional plane (see the *Experimental Procedures* for details). Similar stimuli (stimuli that are easily deformed into one another) are mapped to nearby points on the plane, while dissimilar stimuli are mapped to distant locations. The stimuli

are additionally color-coded to indicate the task condition under which they elicited strong neuronal responses: the red, green, and blue stimuli correspond to the line, circle, and wave detection tasks, respectively. The clustering of stimuli with the same color indicates that neurons tended to guide the optimization algorithm towards similar contours under the same task condition. For both monkeys shown here, neurons preferred near-collinear, elongated contours during the line detection task (shown in red). Conversely, neurons steered the optimization toward curved contours during the circle and wave detection tasks (shown in green and blue, respectively). The distribution of preferred stimulus shapes obtained from Monkey B (data not shown) closely resembled the distribution from Monkey C (Fig. 17B), except that the task dependence was weaker. That is, the plot generated from Monkey B contains similar contours, but the colors of the contours do not cluster. Neurons from Monkey B also underwent changes in their geometric tuning as a function of perceptual task (see below), but those changes were apparent over the averaged population response, rather than in the optimum shapes preferred by individual neurons.

Figure 17. Neurons preferred different 5-bar contours under different task conditions. (A) the preferred 5-bar stimuli generated by all the neurons recorded from Monkey A, pooled together. (B) the preferred stimuli generated by the population of cells from Monkey C. The stimuli are drawn as continuous curves by linking the ends of adjacent line segments with straight lines. To construct these plots, the firing rate of each cell was linearly mapped into the interval $[0,1]$ (0, weakest response; 1, response to optimum contour). All stimuli from each cell that elicited a response ≥ 0.7 are displayed; stimuli that were derived from more than one neuron are shown in bold. The position at which each stimulus is drawn was determined by Sammon's mapping, a multidimensional scaling (MDS) procedure that plots high-dimensional data into arbitrary two-dimensional coordinates. When applied to our data, Sammon's mapping projects structurally similar stimuli to nearby locations in a two-dimensional plane, while plotting dissimilar stimuli at distant coordinates (see the *Experimental Procedures* for details). The color of the stimulus contours indicates the task condition under which they were generated (red: line detection task; green: circle task; blue: sinusoid task). The tight clustering of red stimuli demonstrates that neurons preferred similar, elongated contours under the line detection task. Under the circle and sinusoid tasks, the optimal stimuli for different neurons were more diverse, consisting of groups of variously curved contours. The set of contours preferred under the line detection task is largely separate from the set of near-optimal contours under the two other task conditions.

Figure 17




To determine how perceptual task might reshape the whole tuning surfaces, rather than just their peaks, we analyzed the population responses to 3-bar stimuli. Since the optimization program was always seeded with the same set of 3-bar stimuli, the neural responses to this region of the stimulus space can be directly compared across cells and task conditions. We mapped the height of each 3-bar tuning surface into the interval $[0,1]$, in order to normalize for differences in neuronal firing rates. We then grouped the surfaces that were recorded while the monkeys performed each of the three detection tasks, and averaged the data over each group, to obtain three population tuning surfaces. We performed permutation tests between these mean tuning surfaces to determine whether any differences between them were statistically significant (see the *Experimental Procedures*). For all three monkeys, the differences between the average 3-bar tuning surfaces under the circle and wave tasks were not statistically significant ($p \geq 0.2$). We therefore pooled the data from those two task conditions, and compared the mean of the merged data with the mean surface under the line detection task. The difference between the mean tuning surfaces under the line and circle/wave tasks was statistically significant (Monkey A: total number of surfaces, $n = 53$; $p = 4 \times 10^{-5}$; Monkey B: $n = 63$; $p = 0.007$; Monkey C: $n = 62$; $p = 0.003$). The averaged tuning surfaces from Monkeys A and C under

the line and circle/wave detection tasks are presented in Figure 18A-D. In all three animals, the average 3-bar tuning surface contained a peak of collinear facilitation at $(0^\circ, 0^\circ)$ under the line detection task, which was absent under the other task conditions. Conversely, the circle and wave detection tasks elicited peaks of facilitation along the co-circularity line (but significantly away from the collinearity point at the origin), which were absent or weakly expressed under the line detection task.

These results demonstrate that, over the population, neurons dynamically changed their responses to the same geometric stimuli as a function of perceptual task. The traditional experimental design is to demonstrate this task dependency by recording from the same neurons under different task conditions, rather than averaging the data over separate groups of cells recorded under different task conditions. Since the construction of 5- or 7-bar optimum contours often required thousands of trials and hours of recording time, it was not possible to run the optimization algorithm to completion under different task conditions for the same cell in behaving monkeys. We did, however, generate 3-bar tuning surfaces for a subset of cells that were recorded under two different task conditions. Consistent with the results obtained from averaging the data over the population, we found

Figure 18. Neuronal heat maps are reshaped by task. (A and C) the mean 3-bar heat maps averaged over all the neurons recorded from Monkeys A and C during the circle and wave detection tasks. (B and D) the mean 3-bar heat maps averaged over all the neurons from Monkeys A and C during the line detection task. To normalize for differences in firing rate, the heat map of each cell was mapped into the range $[0,1]$ before being averaged with the rest of the data. Since each map is combined from many cells, the set of all stimuli generated by the whole group of cells is depicted over the data. The 32 stimuli from the initial stimulus set are drawn, as are the stimuli overlying the highest regions of the mean response. Weaker stimuli, or stimuli that are too closely clustered to draw, are indicated with gray points. Note the relative abundance of different refinement stimuli created by cells under the circle and wave tasks compared to sparsity of unique contours preferred under the line detection task. Over the population of cells from both monkeys, the mean neuronal response was highest for circular stimuli under the circle and wave tasks; under the line detection task, the network expressed a narrower selectivity, this time for the linear geometry. Arrows on the color scale in (C) and (D): response ($R = 0.40$) to a one-bar “contour” (one bar optimally positioned and oriented in the RF, and embedded in the complex background; see the *Experimental Procedures*). Sample sizes: (A), $n = 17$; (B), $n = 36$; (C), $n = 25$; (D), $n = 37$. (E and F) the 3-bar heat maps for a single neuron from Monkey A under two different task conditions. The neuron’s geometric tuning was first recorded under the circle task (E) and then under the line task (F). The cell underwent a dramatic change in its geometric tuning when the cue was switched from the circle to the line. In (E), the cell preferred small contours with a very high radius of curvature. In (F), the cell preferred linear contours and circular arcs with broad curvature. Note how the cell, under both task conditions, drove the optimization routine into a region of the stimulus space outside the boundaries of the initial stimulus set ($0^\circ \leq \alpha \leq 135^\circ$; $-60^\circ < \beta < 60^\circ$). Under the circle task, the cell guided the algorithm into the region where $\beta < -60^\circ$; under the line task, the search extended into the region where $\alpha > 135^\circ$.

Figure 18



that the same neurons were able to dynamically change their tuning surfaces according to the cued shape. We recorded from 14 neurons (8 from Monkey A, 6 from Monkey B) while the animal carried out two separate blocks of several hundred trials each. We switched the cue between the two blocks, and the neural responses under the two task conditions gave rise to two distinct tuning surfaces. Fig. 18E-F illustrates the results from one such experiment, in which Monkey A was cued to detect first the circle (E) and then the line (F) while we recorded from the same neuron. Coincident with the change in the cued target, the AF underwent a dramatic shift, from a co-circular to a collinear pattern of facilitation. For both monkeys, the shift in the 3-bar tuning surfaces for the same neuron paralleled the shift seen over the neural population, when different task conditions were used during recordings from different neurons. The analyses of population tuning surfaces described earlier include these 14 neurons.

3.3.5 TIME COURSE OF GEOMETRIC SELECTIVITY

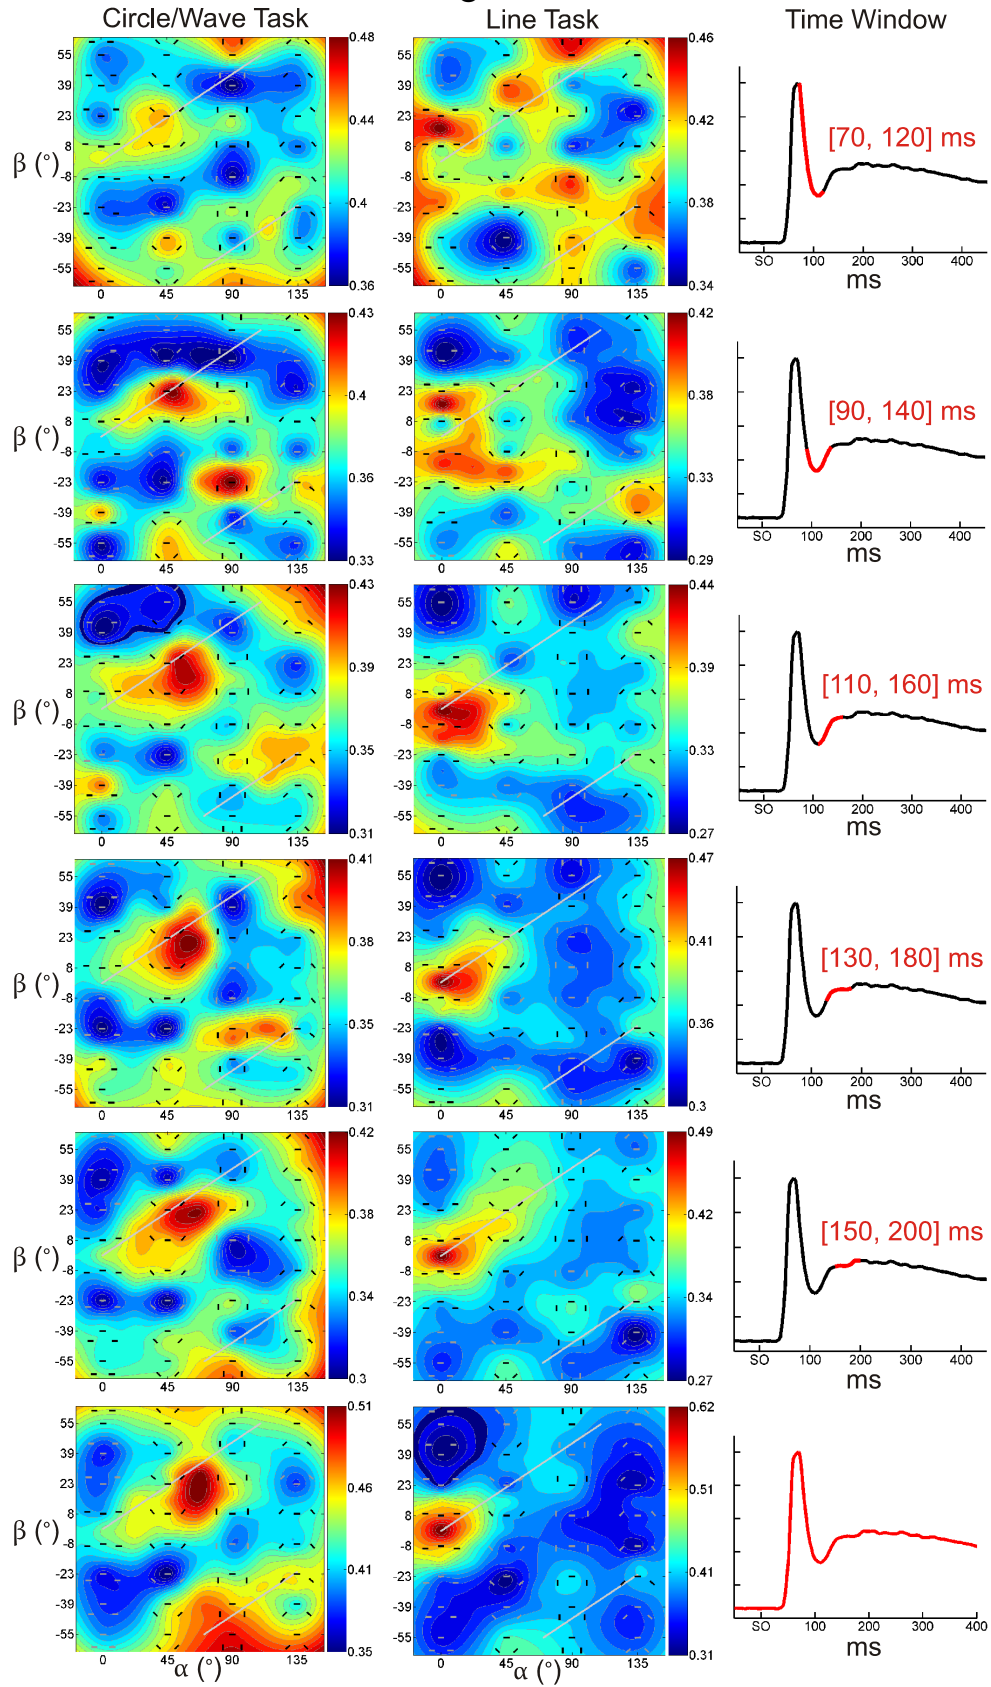
Lastly, we analyzed the temporal dynamics of shape selectivity in V1. In order to determine how quickly neural responses became geometrically selective after the stimulus onset, we performed a nonparametric two-way

ANOVA (Friedman's test) on a dataset that was amenable to statistical analysis. The data consisted of 22 neurons from Monkey C, for which each geometric stimulus was presented 40 times. For these cells, the stimulus set consisted of the 32 3-bar stimuli normally used to seed the optimization algorithm. We performed the statistical test on a sliding 50-ms window, which we swept across the time course of the neural response, to determine when the tuning surfaces took on a statistically non-planar shape. When the left edge of the spike-counting window was moved past 50 ms (corresponding to spikes collected 50-100 ms after stimulus onset), the p-value for the null hypothesis (that the tuning surfaces were flat) fell below 0.1. Past 61 ms, the p-value fell below 0.01; past 65 ms, it fell below 0.00001; past 72 ms, the p-value reached an asymptote of less than 10^{-10} . (When applied to the same data, the parametric two-way ANOVA yielded virtually identical results.) This analysis demonstrates that an initial mode of geometric selectivity was in place between 60 and 110 ms following stimulus onset. More precisely, if the statistical power of the test reaches saturation when the time window includes only selective portions of the response, then selectivity should first arise at about 72 ms. This time corresponds to the peak of the neuronal onset response (see Figs. 19 and 20 for the shape of the PSTH).

The statistical test above suggests that neurons respond differentially to at least some regions of the stimulus space within about 70 ms. However, the analysis does not describe how the pattern of geometric selectivity might evolve over the ensuing tens of milliseconds, nor does it give insight into the task-dependency of the selectivity. To investigate these features, we plotted the evolution of 3-bar shape selectivity over the population, as a function of task condition. We pooled the 3-bar tuning functions from all three monkeys and separated them according to the task condition under which they were recorded. Figure 19 shows the mean heat maps recorded during either the circle or wave task (left column) and during the line detection task (middle column), plotted using various time windows. When the entire duration of each trial is used to construct the mean heat maps (bottom row), the task-dependent differences between the surfaces are highly significant ($p < 4 \times 10^{-5}$; permutation test). To measure the temporal onset of these differences, we moved a sliding 50-ms window across the neural response and plotted the population activity within that window under each task condition. We found distinct patterns of selectivity between the task conditions that began to develop in the window between 70 and 120 ms following the stimulus onset. The shape of the tuning functions evolved over the next tens of

Figure 19. Temporal evolution of geometric selectivity. Left column: the 3-bar tuning surfaces, in different time windows, averaged from 107 different experiments from all three monkeys. The data in this column were collected during either the circle or the wave task. Middle column: the 3-bar heat maps, within different time windows, pooled over 72 line detection experiments from each monkey. Right column: the PSTH obtained from pooling the spikes elicited by all 3-bar stimuli from all recorded neurons in each monkey. The region of the PSTH used to construct each pair of response surfaces is highlighted in red. “SO” denotes stimulus onset. The p-values for the differences between the bottom four pairs of surfaces are, from bottom to top: $p < 4 \times 10^{-5}$; $p = 0.028$; $p = 0.053$; $p = 0.036$ (permutation test). The tuning surfaces were normalized and averaged as in Fig. 18A-D. For clarity, only the stimuli in the initial stimulus set are drawn over each surface (the refinement stimuli are not drawn).

Figure 19



milliseconds, reaching maturity in the time window between 110 and 160 ms. Within this window, peaks of collinear and co-circular facilitation were already apparent on the corresponding response surfaces ($p = 0.036$; permutation test). The results suggest that geometric selectivity starts to mature at 110 ms; this time corresponds precisely to a turning point in the peri-stimulus time histogram (PSTH), shown in the right column of Fig. 19. The PSTH is characterized by a sharp onset response, followed by an inhibitory dip induced by the suppressive contextual background (Li et al., 2006). Notably, the geometric response appears to mature during the excitatory rebound from the inhibitory dip.

We also analyzed the time course of 5-bar geometric selectivity. Since different neurons tended to steer the optimization program into different regions of the stimulus space for 5-bar stimuli, we could not directly average the 5-bar tuning surfaces across cells. However, we often included the 5-bar segments from the cue and distracter shapes as activity-independent additions to the stimulus sets. Since these same stimuli were used across a large subset of the data, we compared the mean neural responses to these stimuli, averaged over data from all three monkeys, as a function of task condition (Fig. 20). In particular, we plotted the time course of the neural

response to the 5-bar line in red, and we traced the mean response to the circular and wave shapes in blue. Figure 20A shows the neural responses recorded from 38 neurons in the line detection task; Fig. 20B plots the activity of 67 different neurons recorded under the circle or wave cues. The task-dependent geometric selectivity for the 5-bar line emerged rapidly during the rebound from inhibition, with the same time course as the 3-bar shape selectivity. The reciprocal selectivity for curved shapes under the circle/wave detection task is not apparent in Fig. 10B because the shapes preferred by neurons under those task conditions were usually very different than the cue shapes themselves. Moreover, the blue trace represents the average response over all the curved targets, which masks strong responses of individual neurons to one cue shape or another.




Figure 20. Time course of 5-bar shape selectivity. (A and B) peri-stimulus time histograms (PSTHs): the response to the 5-bar line is shown in red, and the mean response to the circle, wave, and distracter is shown in blue, under task conditions when the animal looked for the line (A), or for the circle or wave (B). SO marks stimulus onset (time zero); the spikes elicited by each stimulus, over a population of neurons, were pooled into 20 millisecond time bins and fit with a smooth interpolant. The data in Panel A are averaged over 38 cells that were recorded during the line detection task; panel B plots the pooled responses from 67 cells recorded during either the circle or wave detection tasks. Under the line detection task—but not under the circle or wave tasks—neurons were preferentially facilitated by the 5-bar line, compared to the segments from the other cues and the distracter. This task-dependent selectivity emerged rapidly following the rebound from contextual inhibition. The highlighted portions of the red and blue PSTHs in (A) are significantly different ($p = 0.0039$; one-sided Wilcoxon signed-rank test; time window: [100, 160] ms).

3.4 DISCUSSION

V1 has long been recognized as a geometric processor, responsible for parsing the visual scene into its component lines and edges (Hubel and Wiesel, 1962; Marr, 1982). More recent evidence suggests that V1 may also merge these components into perceptually unified wholes (Hess et al., 2003; Li et al., 2006) via the AF. In the standard model, the pattern of neuronal interconnections underlying the AF echoes the statistics of natural images (Geisler et al., 2001; Sigman et al., 2001; Hess et al., 2003; McManus et al., 2008), in which collinear and co-circular contours predominate. The strongest connections in the AF are thought to link neurons whose RFs fall along these linear and circular arcs. While the AF was originally coined as a psychophysical concept (Field et al., 1993), our study directly shows, for the first time, that a network with matching properties exists in V1. Over the population, the strongest contextual interactions we observed fell on, or very near, the co-circularity line (Figs. 18-19). Our data also parallel some psychophysical evidence that the AF may depend on higher-order cognitive influences (Schuchard, 1995). The network in V1 never expressed the complete range of lateral interactions from the full AF at once (which includes all collinear and co-circular interactions). Rather, the neural

population expressed subsets of these interactions—peaks of facilitation that shifted their position along the co-circularity line when the perceptual task was changed.

As an important nuance to the standard AF theory, the tuning surfaces expressed by individual neurons in V1—even under the same task condition—were not just random perturbations of the population mean, but represented fundamentally different modes of facilitation. In particular, different cells responded to a diverse range of curved optimal contours, including circles and sinusoids of various shapes, whereas previous results have focused on the responses to straight line geometries. The cortical strategy may be to ensure that the population activity follows a narrow pattern that is appropriate for encoding the saliency of smooth contours, while maintaining a richer repertoire of shape selectivity at the level of individual neurons. This “hidden” diversity of responses may be used to accommodate mechanisms of object recognition and scene segmentation (Ullman, 2007; Epshtein et al., 2008). The population activity seems well-suited for extracting smooth contours, since it contains optima for circular geometries, but individual cells often preferred shapes with considerably more curvature or with locally disjointed line segments. This class of

optimal contours may have arisen when the first-order contextual interactions (between the bar in the RF and the contour bars outside it) dominated the second-order interactions (between the contour segments outside the RF and each other). For instance, in Fig. 15A, the contextual contour bars in the 5-bar optimum are all co-circular with the bar in the classical RF, but they are not mutually co-circular with each other. In addition, neurons linked by horizontal connections may express different geometric preferences, which could also explain why some neurons prefer disjointed contours over smooth ones. If interconnected neurons have different patterns of geometric connections, then the interactions between them may cause kinks or discontinuities in the optimal contour for each cell.

It is important to introduce some caveats pertaining to the construction of “optimal” contours in our experiments. Throughout this work, we refer to the best stimuli created by the stimulus generation program as optimal contours, but they are not necessarily optimal over the entire “universe” of contour shapes. They were the most effective stimuli created under the constraints of the “optimization” program, but a rigorous search for a truly optimum stimulus would require an impossibly exhaustive search of all possible stimuli. The assumption of mirror symmetry, the stepwise

procedure for building successively longer contours, and the weak smoothness constraint that are implicit in the method could all prevent the detection of truly optimal contours for some neurons. However, it is not clear that neurons with complicated response properties, like those in V1, must have a single optimal stimulus. Indeed, many of the cells in our population have 3- and 5-bar tuning functions with more than one prominent peak of facilitation. Any plausible method of searching a sufficiently large stimulus space for an optimum will necessarily be incomplete, and the notion that neurons have a single optimal stimulus may be overly simplistic. The basic value of the optimization program is that it guides the stimulus generation into regions of the stimulus space where the responses of a given neuron are informative. It can reliably detect subsets of complex stimuli that strongly activate neurons, even if it cannot detect all such stimuli. The traditional approach of using static stimulus sets, particularly for complex stimuli, suffers from the drawback that the chosen stimuli may not yield any insight into the selectivity of an arbitrary neuron. Neurons were selective for shapes with a limited range of curvature and orientation, so using a predefined set of stimuli would sample most neurons' responses in a flat region of their tuning surface. As an extreme example, consider the cell in Fig. 16B. It has clear tuning for wave-like shapes in the 5-bar tuning surface,

which were discovered by the optimization algorithm, but it did not respond to any of the contours used as the static 5-bar stimulus set (bottom inset).

Our optimization program can explore a much broader range of stimuli than previous approaches applied to V1, and its generality reveals new sets of contextual interactions that could be missed by other approaches.

The interplay between perceptual task and geometric selectivity raises interesting questions. Previous experiments have shown that neuronal facilitation in V1 by collinear contours only emerges after animals have been trained on a contour detection task (Li et al., 2008). It remains to be seen what effect perceptual learning would have on neuronal selectivity for geometric shapes, in comparison with a pre-training default state. Another question pertains to the speed with which geometric preferences can be reset. If the AF dynamics have general implications for visual computations like object recognition and segmentation, then they must occur rapidly (on a trial-by-trial basis) (Ullman, 2007; Epshtein et al., 2008). Lastly, the nature of our behavioral paradigm has important implications. Our measurements of geometric selectivity were taken during the delay between the cue and the upcoming behavioral discrimination. This feature of the data puts them in the context of work on expectation, which is an understudied facet of visual

cognition (Summerfield and Egner, 2009), and short-term memory. Recent studies have shown that during such a delay period, V1 is involved in encoding short-term memory about the cued target (Super et al., 2001a; Vidyasagar and Pigarev, 2007; Harrison and Tong, 2009). In fact, neural activity throughout the visual cortex may be broadly related to memory and top-down control (Schlack and Albright, 2007; O'Herron and von der Heydt, 2009). The resemblance of the dynamically built contour shapes to the cues in our study could be an indication of either working memory or task-specific top-down control (Offen et al., 2009), including the specific influence of expectation for the cued shape. Finally, differences in cognitive engagement before the targets were presented, during the delay period, may account for the differences we observed in the extent of the top-down modulation between the monkeys.

We speculate that the shape selectivity in V1 arises from its geometrically tuned horizontal projections, and that the task dependency derives from a top-down gating of those connections (Gilbert and Sigman, 2007). While the prevailing view of top-down interactions emphasizes their role in gain control (Reynolds et al., 2000; Martinez-Trujillo and Treue, 2002; Reynolds and Heeger, 2009) or attentional competition (Reynolds et al., 1999), they

may also provide a fundamental gating, or input selection, mechanism that enables cortical areas to act as adaptive processors (Gilbert and Sigman, 2007). Anatomical and physiological evidence implicate the horizontal connections as the substrate for the geometric contextual interactions (Bosking et al., 1997; Schmidt et al., 1997; Sincich and Blasdel, 2001; Angelucci et al., 2002; Stettler et al., 2002; Chisum et al., 2003; Hess et al., 2003; Angelucci and Bressloff, 2006; Li et al., 2006; Gilbert and Sigman, 2007). The task dependency of those interactions, shown here, suggests that feedback projections may inhibit some sets of lateral interactions and/or activate others, thereby establishing different network states with different geometric optima. The time course of geometric facilitation is consistent with a mechanism whereby task-dependent shape selectivity is generated intrinsically in V1. In V4, for instance, attention does not enhance neural responses to cued shapes or spatial locations until about 160-170 ms following stimulus exposure (Chelazzi et al., 2001; Buffalo et al., 2010). This suggests that the task-dependent selectivity in V1 does not echo the responses in higher visual cortical areas, since the selectivity emerges first in V1. Instead, the early onset of the geometric selectivity is consistent with other contextual interactions in V1 (Lee et al., 1998). Furthermore, the gradual, task-specific evolution of the response surfaces is predicted by a

recent network model, in which the temporal dynamics derive from lateral interactions intrinsic to V1 (Piëch et al., 2010, in press).

One alternative explanation for our data is that spatial- or even feature-based attention could explain the results, by selectively boosting the neural responses to salient stimuli in the RF that resemble the cued target. Under this view, V1 neurons would respond strongly to specific contours only because the monkey notices them during a specific task condition, and then allocates attention either to the spatial location of the RF, to the stimulus shape as a whole, or to some feature of the stimulus shape. This interpretation fundamentally differs from our own, which holds that the shape selectivity is generated from intrinsic network interactions in V1 and is modulated by top-down influences.

There are several compelling reasons to favor our interpretation over the alternative. First, V1 neurons express task-dependent selectivity in the very first stage of the optimization program, for 3-bar stimuli. It is known from previous results that contours of this length, in the same range of retinal eccentricities, are not perceptually salient (Li et al., 2006). Even in a task that requires the monkey to directly detect the presence of a 3-bar linear

contour embedded in a contextual background, the detection performance is barely above chance level (Li et al., 2006). Second, even if the monkey notices particular contour shapes, and then allocates attention to these, different neurons should experience attentional enhancement to the *same* shapes—the shapes that the monkey notices. In our data, however, neurons that preferred co-circular contours under the circle detection task, for example, did not prefer the *same* circular contours. Instead, each neuron preferred different subsets of a larger family of circular contours. For instance, some neurons were strongly selective for small, closed circles while others preferred broad arcs. This was true of neurons recorded across days, but it was also true of simultaneously recorded neurons, which were picked up from the same electrode (data not shown). Moreover, the optimal contours for neurons tended to resemble the targets only in a general sense (curved versus linear) rather than in precise form or orientation. If the geometric selectivity arose from the salience of stimuli that resemble the cue, then the best shapes should have been the cues themselves. Taken together, we believe these arguments rule out the concern about conventional attentional effects.

The finding that task-dependent cognitive influences, including the expectation for particular shapes, can reshape neural responses in V1 is a significant departure from the current paradigm of attentional influences. Expectation, and the various attentional mechanisms that are often bound with it, may play a significant role in our experiments. We investigated the modulation of neural responses, not to the target shapes the monkeys were cued to detect, but rather to arbitrary contours that were presented over the RF while the monkey was primed to perform the detection. This experimental design gives our study relevance to theories of predictive coding and “countercurrent stream” models of cortical function (Bar et al., 2001; Bar et al., 2006; Ullman, 2007; Epshtein et al., 2008; Summerfield and Egner, 2009). Only a few studies (David et al., 2008; Ghose and Bearl, 2010) have reported changes in neural tuning properties under a similar task design, but those experiments were done in V4 and MT, where the influence of spatial- and feature-based attention is less controversial. Our experiments were not designed to isolate the cognitive influences of expectation, but our task design makes the shape of upcoming stimuli predictable.

Interestingly, the studies that have shown task-dependent sharpening of neural tuning have done so in the context of expectation (David et al., 2008;

Ghose and Bearl, 2010), while theories that emphasize gain control mechanisms are based on spatial- and feature-based attentional studies (Reynolds and Chelazzi, 2004; Maunsell and Treue, 2006; Reynolds and Heeger, 2009). These differences may result from fundamental differences in the functions that attention and expectation serve. While attention may be important for focusing the brain's limited computational resources on behaviorally relevant stimuli, expectation may be useful for resolving ambiguities arising from the projection of a four-dimensional world onto the two-dimensional retinal surface (Summerfield and Egner, 2009).

Correspondingly, attention studies tend to reveal gain control and stimulus selection mechanisms (Reynolds and Chelazzi, 2004; Reynolds and Heeger, 2009); expectation studies suggest modifications of neural tuning and, as shown in the current study, network behavior. Our results suggest that perceptual task shapes the lateral interactions mediating the association field, thereby promoting perceptual grouping and altering shape selectivity in a way that cannot be explained by a normalization of pre-existing tuning. A pair of related theories from Bar and Ullman (Bar et al., 2001; Bar et al., 2006; Ullman, 2007; Epshtein et al., 2008) suggest that top-down influences may bias the sensory processing in early cortical areas toward the detection of expected stimuli. In the context of objection recognition, this process may

be crucial for the detection of the detailed features that characterize particular objects (e.g., the eyes, nose, and mouth on a face) (Ullman, 2007; Epshtein et al., 2008). These ideas provide compelling hints about how the task-dependent shape selectivity in V1 participates in broader mechanisms of visual processing. Whatever purpose the network interactions in V1 may serve, our results demonstrate a new level of complexity in V1, and they highlight the role of V1 as an adaptive integrator.

3.5 EXPERIMENTAL PROCEDURES

3.5.1 ELECTROPHYSIOLOGY

We performed extracellular recordings from single units in V1 of three adult male monkeys (*Macaca mulatta*). We used glass-coated tungsten microelectrodes within an impedance range of 0.8 to 2.5 M Ω at 1KHz. All recordings were carried out in the superficial layers of the striate cortex, corresponding to eccentricities between 0.85° and 6.58°. Before beginning the optimization algorithm, we measured the RF position and orientation preference of the recorded neuron using grating and/or bar stimuli; these parameters were later used to position and orient the geometric stimuli. All procedures were in accordance with the National Institutes of Health Guide

for the Care and Use of Laboratory Animals and with the approval of the Institutional Animal Care and Use Committee at the Rockefeller University.

3.5.2 DATA ACQUISITION AND STIMULUS DESIGN

The stimuli were created by a visual stimulus generator (VSG2/5, Cambridge Research Systems) at a viewing distance of 138 cm. They were displayed on a NANA O monitor (FlexScan F2-21) at a resolution of 1024 x 769 pixels and with a frame refresh rate of 105 Hz. Neuronal activity was recorded with a spike sorting and acquisition system (Plexon, Inc.), and the stimulus-elicited spikes from each trial were simultaneously counted by a computer board (PCI-DAS1002, Measurement Computing Corp.) for online data analysis.

The stimuli in our experiments employed a previously reported paradigm (Kapadia et al., 1995; Li and Gilbert, 2002; Li et al., 2006), whereby a contour configuration is embedded in a contextual background of line elements. Each stimulus consisted of two circular arrays of 0.55° -by- 0.55° compartments, into which oriented line elements were positioned. The line elements in each patch were 0.2° -by- 0.05° in dimension and were displayed

at 50% Michelson contrast, with a luminance of 18 cd/m^2 . The arrays were 11 or 13 compartments in diameter, and the bar elements were partitioned into two sets: 1.) the contextual line elements, which were randomly positioned and oriented within their respective compartments; and 2.) the contour elements, whose position and orientation were brought into a particular pattern of geometric alignment with the bar in the central grid cell. The contours that were generated during the automated search for each neuron's optimum shape were composed of these contour elements.

3.5.3 TASK DESIGN

We trained three monkeys in a two alternative forced-choice (2AFC) task design with catch trials. The monkeys initiated each trial by gazing at a 0.08° fixation point, while their eye position was monitored at 30 Hz with an infrared tracking system. After a 500-600 ms pre-fixation epoch, the monkeys were required to maintain their gaze within a window (0.8° - 1.0° in diameter) around the fixation point. After the pre-fixation period, the monkeys were presented with one of three contours (Fig. 11A1) to serve as a cue for the forthcoming task: either a closed circle, a sinusoidal wave, or a straight line (Fig. 11A3). The experiments were conducted in large blocks of

many hundreds of trials (generally more than a thousand) in which the cued contour was always the same. (The presentation of the cue at every trial was therefore redundant, but was included to reinforce the monkey's familiarity with the cued shape and his attention to the task.) After cuing the monkey for 380 ms, and following an additional blank frame (9.5 ms), the contextual patches were drawn on the computer monitor (Fig. 11B1). Embedded in the center of each patch was a geometric stimulus, taken from a set of stimuli whose length and shape were adjusted at discrete time points during the recording session (for details see below). The neuronal responses recorded during this stimulus presentation epoch were used by the optimization algorithm to tailor the stimulus set to the neuronal activity. After a variable period of stimulus exposure for each trial (ranging from 191 to 1429 ms), the embedded stimuli were abruptly replaced by two 7-bar "target" contours (Fig. 11C1), which were shown for 48 to 191 ms. (The range of stimulus exposure times, and the duration of target presentation, were parameters that we varied between experiments. We anticipated that making the time of target presentation both short and unpredictable might increase the monkeys' attentiveness during the stimulus presentation, before the behavioral discrimination. Changing those parameters, however, had no apparent effect on neuronal responses, so we used a constant period of stimulus exposure,

475 ms, for Monkey C.) To form the target contour in each patch, seven line elements from the patch, including lines from both the geometric stimulus and the contextual background, were transiently rearranged. One of these flashed contours was the cue presented at the beginning of the trial, while the other was a distracter contour. (For most of the experiments, the same distracter was used for all trials and under all task conditions. In other experiments, we either used random distracter shapes or distracters that were randomly “jumbled” or distorted versions of the cue, but we never observed any dependence of the animals’ behavior or the neural responses on the nature of the distracter.) At the end of each trial, the two stimulus patches were replaced by two dots (saccade targets); the monkey’s task was to make a saccade toward the dot location where the cued contour was flashed. In most of the experiments, one-third of the trials were catch trials, in which the same distracter contour was flashed in both patches. These trials were included to ensure that the monkey attended to the spatial location of both contextual patches; the correct decision during a catch trial was to look toward a third saccade target, which was present in all trials together with the other two saccade targets (Fig. 11C2).

In the majority of our experiments, we allowed the monkey to perform the task with the same cue for the entire recording session, thereby enabling the progressive construction of long contours. In a smaller set of experiments, we probed a neuron's geometric selectivity while the monkey was cued on a particular target shape, and then repeated the optimization algorithm after cueing the monkey on a different contour.

3.5.4 AUTOMATED STIMULUS GENERATION

The optimization algorithm we devised is a stepwise procedure for constructing the symmetric configuration of line elements that maximizes the response of the recorded neuron. At the beginning of each experiment, all cells were shown the same set of stimuli, which was identical from cell to cell, except for a rotation about the center of the RF (the stimuli for each cell were rotated to orient them with respect to each cell's preferred orientation). The shapes of the stimuli shown to different cells would only diverge later in the experiment, when different cells would drive the stimulus generation in separate directions.

At the outset of the recording session, the optimal configuration contained only one contour element: the central bar in the cRF, surrounded by contextual lines. (The position and orientation of the central bar were always matched to the recorded neuron's RF properties.) The algorithm then attempted to find the optimal positions and orientations of two adjacent contextual bars, each spaced a center-to-center distance of 0.55° from the central bar. The contextual bars were always placed symmetrically on either side of the bar in the RF; once the position and orientation of one bar are specified, the geometry of its mirror-symmetric counterpart is also determined. After settling on the optimal arrangement of the first two adjacent bars, the algorithm made their configuration, together with the central bar, permanent throughout the remainder of the experiment. Every succeeding stimulus shown to the monkey then contained the optimal 3-bar configuration at its core. The next iteration of the algorithm searched for the best arrangement of the next two bar elements, each of which was spaced 0.55° from either end of the already established contour (Fig. 12). Proceeding in this way, the algorithm could execute an arbitrary number of iterations, each time extending the number of bars in the constructed contour. But time constraints and diminishing neuronal selectivity for distal contour

elements obliged the termination of the program after it settled on the optimal 5- or 7-bar contour.

At each iteration, the algorithm generated an initial set of stimuli that was intended to comprehensively sample the stimulus space for the next two line segments to add to the contour. The routine repeatedly tested each of these stimuli. After 5 or 7 randomly interleaved repetitions of each stimulus, the algorithm removed half of the stimuli from the test set, eliminating those stimuli that elicited the lowest firing rate from the recorded neuron. The remaining stimuli were tested for 15 to 21 additional trials, and the stimulus that elicited the highest firing rate was selected. The algorithm then carried out a “refinement” step, constricting the region of the stimulus space it would search around the selected stimulus. The refinement step tested a new, smaller set of stimuli which were closer to the selected stimulus than the stimuli in the previous set. After testing each of these new stimuli 20 or 28 times, together with the best stimulus from the previous set, the program selected the most effective stimulus and could execute another refinement step by further constricting the stimulus space around the current optimum. At each of these stimulus space constrictions, the differences in the position and orientation of the contour elements in the test stimuli were progressively

decreased until a good approximation to the optimal configuration was found. In our experiments, we configured the optimization routine to use between 0 to 2 refinement steps per iteration (usually 1) before fixing the best configuration and moving on to the next set of contour bars.

The initial stimulus set at each iteration contained 32 automatically generated stimuli, with 4 different bar orientations at each of 8 spatial locations on each side of the contour (Fig. 12). Let ℓ be the collinear axis that extends outward from the bar at the right edge of the existing contour; let the point p be the center of that bar. Now consider the circular arc \mathcal{C} , such that: the distance between p and \mathcal{C} is 0.55° ; \mathcal{C} is bisected by ℓ ; and the endpoints of \mathcal{C} form a central angle of 110° with p . Then the 8 spatial locations on the right side of the contour were evenly spaced between the endpoints of \mathcal{C} , covering a broad range of positions; and the 4 orientations at each position were 0° , 45° , 90° , and 135° , relative to the orientation of the bar in the RF. The bars on the left end of the contour were mirror-symmetric to the bars on the right edge.

Let q and θ be the optimum position and orientation of the outermost contour bar from a previous stimulus set, S . Further, let d be the difference in

polar angle between adjacent bar positions, and let Δ denote the smallest difference between different bar orientations, for the stimuli in S . If Ω represents the set of all previously tested stimuli, then the new stimulus set generated by a refinement step would be:

$$R = \{(\theta + \Delta, q + d), (\theta, q + d), (\theta - \Delta, q + d), (\theta + \Delta, q), (\theta - \Delta, q), (\theta + \Delta, q - d), (\theta, q - d), (\theta - \Delta, q - d)\} / \Omega$$

If R was the empty set, then the algorithm would replace it with:

$$R = \{(\theta + \Delta/2, q + d/2), (\theta, q + d/2), (\theta - \Delta/2, q + d/2), (\theta + \Delta/2, q), (\theta - \Delta/2, q), (\theta + \Delta/2, q - d/2), (\theta, q - d/2), (\theta - \Delta/2, q - d/2)\}$$

3.5.5 *A PRIORI* STIMULUS ADDITIONS

In addition to the stimuli that were generated by the dynamic optimization routine, we also included several geometrically pertinent shapes in our stimulus sets. When generating the first set of 3-bar stimuli, we often included the innermost three bars of the cued target shape (and the cue rotated by 180°). For Monkey C (but not A or B), we also included in the initial stimulus set a 1-bar stimulus, consisting only of the bar in the cRF, surrounded by the contextual background. This stimulus was identical to the

3-bar stimuli, except that none of the random bars from the contextual background were brought into a particular geometric alignment with the bar in the RF. Lastly, when generating the first set of 5-bar stimuli, we included the innermost five bars (middle segments) of all three cues, the false target distracter, and their 180° rotations.

3.5.6 THREE-DIMENSIONAL TUNING SURFACES

We rendered the tuning surfaces (heat maps) by first plotting the neural response to each stimulus as a set of three-dimensional points, and then interpolating between those points using Hardy's multiquadratics (Amidror, 2002). The stimuli were mapped onto the plane by specifying each stimulus as a function of its two outermost line segments. Since the stimuli were all mirror-symmetric, only one of the two outermost contour bars was needed to specify each stimulus. In order to compare the tuning surfaces between neurons, we transferred the stimuli into a Cartesian coordinate system, in which the central bar in the RF was collinear with the horizontal axis. The outermost bar on the right arm of any stimulus in this coordinate system was then chosen to specify the stimulus geometry. Each stimulus was specified by two angular coordinates: the orientation of this rightmost contour bar

with respect to the horizontal (α , Figs. 12 and 13); and the polar angle of the bar's position relative to the center of the RF (β , Figs. 12 and 13). The gray lines drawn over the data are the loci of all stimuli along the surface whose outermost contour bars are co-circular to the bar in the RF.

For experiments during which only 3-bar stimuli were presented to the monkey, the neuronal response consists of a single three-dimensional surface (shown as a heat map). For experiments that included 5- or 7-bar stimuli, the neural response includes additional surfaces which separately correspond to the 5- and 7-bar stimuli. Each surface is generated by mapping only the outermost contour bars onto the plane.

Finally, the 5-bar heat maps also include, in a panel beneath the surface data, the neural responses to the 5-bar segments of the cued target and distracter shapes (and their mirror images, obtained by rotating these stimuli 180°). These are plotted below the surface because they generally fall outside the region of the stimulus space searched by the optimization program. The response to each 5-bar target segment is indicated by its color, according to the same color scale used for the response surface.

3.5.7 5-BAR POPULATION ANALYSIS WITH SAMMON’S MAPPING

Sammon’s mapping projects high-dimensional objects, like our 5-bar stimuli, into two dimensions by minimizing the following cost function:

$$E_S = \frac{1}{\sum_i \sum_{j>i} d_{ij}} \sum_i \sum_{j>i} \frac{(d_{ij} - d'_{ij})^2}{d_{ij}}$$

Here, d_{ij} is our measure of distance between stimulus i and stimulus j ; and d'_{ij} is the Euclidean distance between the two-dimensional points that represent stimuli i and j .

We define the distance between any two stimuli as the amount of “energy” required to deform one stimulus configuration into the other. Variants of this “elastic deformation energy” are used extensively in the image processing literature, particularly in shape recognition and classification. We represent the stimuli as one-dimensional, piecewise linear functions of arc length, t .

Let $\vec{S}_i(t) : [0, 1] \mapsto \mathbb{R}^2$ be the function representing stimulus i , which consists of the contour bars that make up the original stimulus, interconnected by straight line segments that adjoin adjacent bars. The center of each stimulus

is always taken to reside at the origin, and each stimulus is rotated to the same orientation. Let the difference between stimuli i and j be represented by $\vec{D}_{ij}(t) = \vec{S}_i(t) - \vec{S}_j(t)$, and let $(\vec{D}_{ij}(t))_x$ and $(\vec{D}_{ij}(t))_y$ be the x - and y -components of the difference function $\vec{D}_{ij}(t)$. Then our measure of the elastic deformation distance between stimuli i and j is

$$d_{ij} = \int_0^1 (\vec{D}_{ij}(t))_x^2 + (\vec{D}_{ij}(t))_y^2 dt + \int_0^1 (d(\vec{D}_{ij}(t))_x / dt)^2 + (d(\vec{D}_{ij}(t))_y / dt)^2 dt$$

3.5.8 PERMUTATION STATISTICS

Let $\{\vec{F}_1, \vec{F}_2, \dots, \vec{F}_N\}$ and $\{\vec{G}_1, \vec{G}_2, \dots, \vec{G}_M\}$ be two sets of tuning surfaces that were averaged together to form the mean surfaces \vec{T}_1 and \vec{T}_2 . Here, each surface is represented as a vector that samples the surface height at the stimulus locations that were tested, over all the experiments, from the corresponding monkey(s). The permutation test between \vec{T}_1 and \vec{T}_2 was performed by pooling all the individual surfaces into the set

$\{\vec{F}_1, \vec{F}_2, \dots, \vec{F}_N, \vec{G}_1, \vec{G}_2, \dots, \vec{G}_M\}$, and then randomly reassigning each surface into one of two subsets, of length N and M . The surfaces in these randomly permuted subsets were averaged together, and a test statistic computed on

these mean surfaces was compared to the statistic computed on \vec{T}_1 and \vec{T}_2 .

The random permutation and averaging procedure was performed 25,000 times; the p-value was reported as the probability that the test statistic from the permuted data was as small as that from the original data sets.

Let $\vec{T}_{1,i}$ and $\vec{T}_{2,i}$ be mean neural response to stimulus i . Additionally, let D_i be the distance between stimulus i and the nearest co-circular stimulus. If D_i is computed as the elastic deformation energy between stimulus i and its nearest co-circular neighbor, then our test statistic is:

$$TS = \sqrt{\sum_{\{i \mid D_i \leq \delta\}} (\vec{T}_{1,i} - \vec{T}_{2,i})^2}$$

This test statistic measures the Euclidean distance between the mean tuning surfaces along the co-circularity lines. We used $\delta = 0.2$ (arbitrary units) to focus the statistical comparisons close to the co-circularity lines, but the results from the permutation tests were not sensitive to this parameter. The components of the neural response vectors near the co-circularity line were linearly mapped into the range $[0,1]$, so that

$$\max_i \vec{T}_{1,i} = \max_i \vec{T}_{2,i} = 1; \min_i \vec{T}_{1,i} = \min_i \vec{T}_{2,i} = 0$$

for all i such that $D_i \leq \delta$.

3.5.9 TWO-WAY ANOVA

We applied Friedman's test (Matlab Statistics Toolbox, The MathWorks, Inc.) to examine the time course of geometric selectivity. Friedman's test is a nonparametric alternative to the standard two-factor ANOVA; it avoids the normality assumption, but it still assumes that all the data derive from distributions with the same variance. Since the variance in neuronal spiking is proportional to the mean response, we applied Friedman's test to the square root of the neuronal spike counts, rather than to the spike counts themselves. When applied to data that are distributed according to a Poisson or a Gamma distribution, the square root operation is a variance stabilizing transformation: it adjusts the data so that samples with different means will have approximately the same variance. The transformed spike counts in a given time window were then grouped according to two "factors": the stimulus that elicited the spikes (factor A) and the neuron from which the spikes were recorded (factor B). In our data set, we had 32 "levels" of factor A (i.e., 32 different stimuli), 22 levels of factor B, and 40 observations per treatment (i.e., 40 trials per stimulus, for each neuron). Friedman's test corrects for the extraneous variability introduced by the different levels of

factor B, and tests for the presence of a “main effect” due to factor A, over all the supplied observations. In the context of our data, this means that Friedman’s test corrects for the variability of neural responses from cell to cell, and tests the null hypothesis that all 32 stimuli elicit the same neural response (within a given time window).

3.6 SUPPLEMENTAL FIGURES

Figure 21. Facilitation by co-circular stimuli with a small radius of curvature. (A and B) additional examples of circular facilitation, from two separate cells recorded during the circle detection task. Cell A was largely silent to the geometric stimuli embedded in the contextual background, but it expressed sharp selectivity for co-circular 3-bar stimuli, with a local optimum that suggested a preference for stimulus closure. In an extension of this 3-bar optimum, the optimization routine converged on a closed contour from only five contour bars. Cell B had an even more extreme preference for stimulus closure. Having converged on a closed circular contour from only three bars, the optimization algorithm generated a set of 5-bar configurations with stimulus bars positioned just outside the 3-bar circle. The neuron's preferred contextual geometries clustered around the 5-bar co-circularity line: the cell was selectively facilitated by contour bars that were approximately co-circular to the central bar in the RF, even though they did not form a smooth extension of the 3-bar optimum. While these preferred 5-bar configurations did not form smooth contours, they nevertheless elicited more facilitation than the 5-bar segments from the cues and the distracter (although the line also elicited strong facilitation).

Figure 21

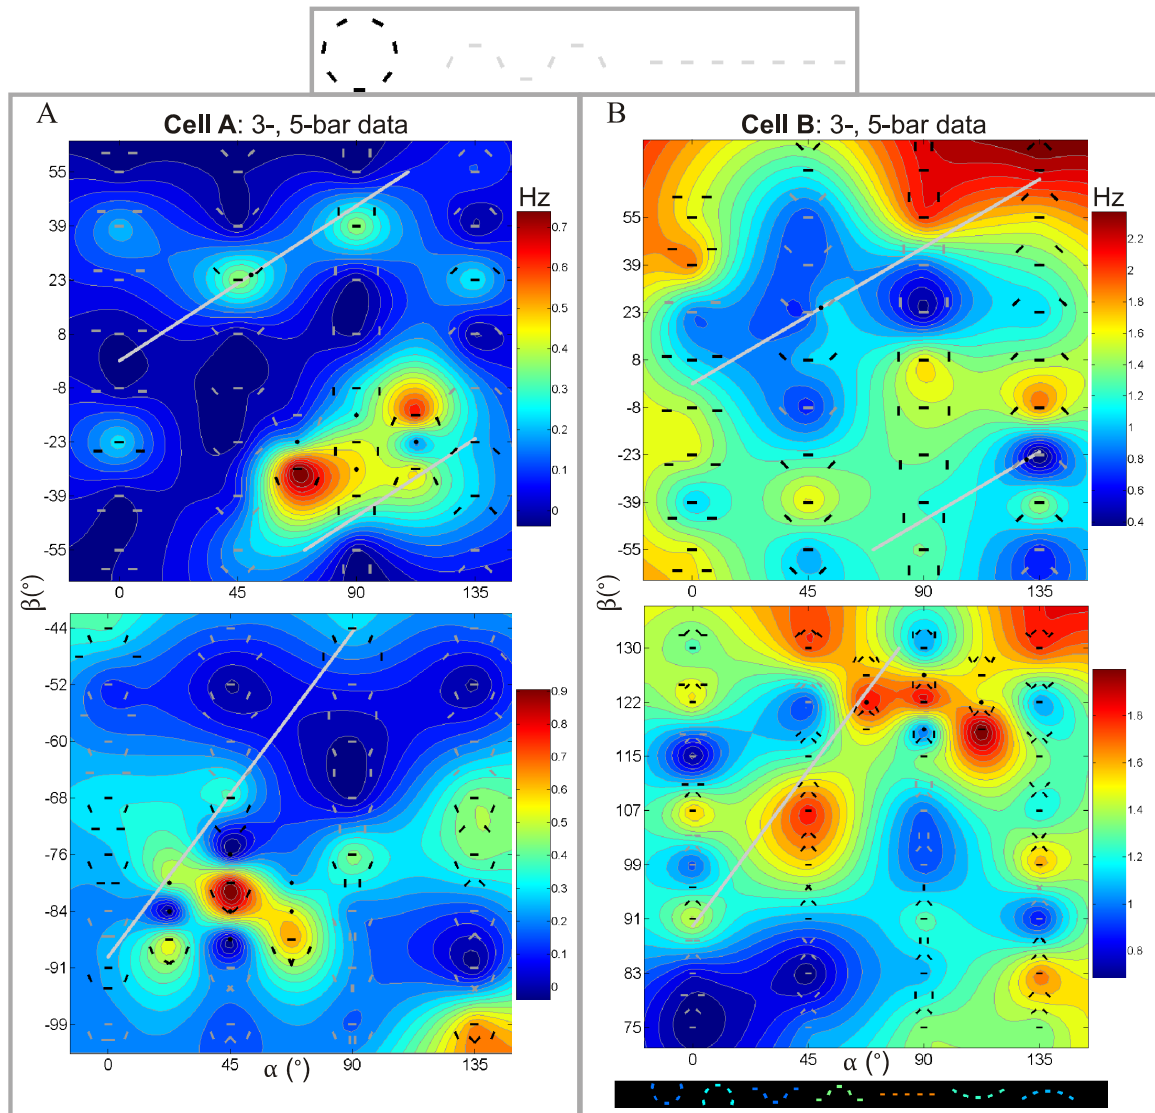


Figure 22. The eye position distributions during the presentation of the strongest and weakest 5-bar stimuli, for all three monkeys. Left column: each distribution was formed by pooling the eye traces recorded during the display of the 5 most effective stimuli for each cell; right column: eye traces pooled from a corresponding number of the weakest stimuli. White dot: fixation point (FP); outermost white circle: fixation window, radius = 0.5° ; inner circle: drawn for reference, radius = 0.2° . Mean eye movement speed (eye trace length/time, averaged over trials) in $^\circ/\text{sec}$: A1, 1.43; A2, 1.41; B1, 1.30; B2, 1.35; C1, 1.22; C2, 1.24. The monkeys tended to focus their fixation within $\sim 0.2^\circ$ of the FP; their behavior was the same for both strong and weak stimuli.

Figure 22


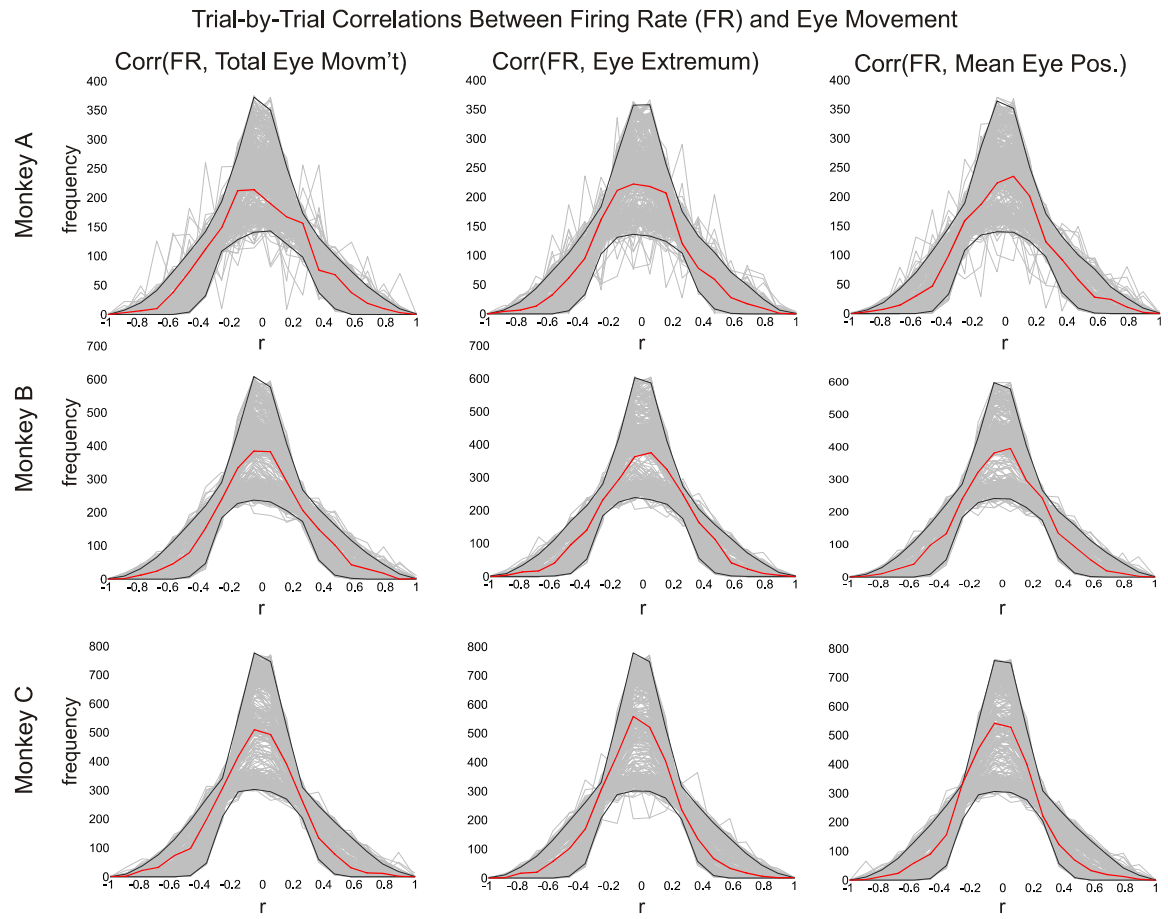


Figure 23. Pearson correlation coefficients between neuronal responses to geometric stimuli and three features of the monkeys' eye traces. Rows: data for each monkey; columns: correlations between neural activity and different attributes of the eye movements. Red line: distribution of observed correlations between sets of stimulus responses and the corresponding eye traces (see below). Gray lines: chance distributions obtained by randomly reshuffling the order of the eye traces with respect to the neural responses, then re-calculating the correlation coefficients. Each set of data was randomly reshuffled 500 times. The distributions computed from these reshuffled data were ranked according to their peak height (a simple measure of their overall shape); the random distributions falling between the 2.5th and 97.5th percentile are displayed. Black lines: 95% confidence bounds for chance correlations, computed for each point along the distributions. Left column: correlations between neural responses and the total length of the eye position trajectory in a trial, normalized by the recording time (i.e., eye movement speed). Middle column: activity correlations with the maximum displacement of the eye position away from the fixation point within a trial. Right column: firing rate correlations with the distance of the mean eye position from the fixation point. The correlations were all computed in the following way. Let $\mathbf{n}_i = \{n_{i,1}, n_{i,2}, \dots, n_{i,S}\}$ be the set of neural responses to the i -th presentation of each contour in a given stimulus set ($n_{i,j}$ is the response to the j -th stimulus during trial i). \mathbf{n}_i is therefore a single-trial slice through the measured neural response function, consisting of a series of S trials recorded in close temporal succession (while the firing rate was stationary). The distributions show the Pearson correlation coefficients between each of these single-trial series and the simultaneously recorded eye traces, from all experiments in which the optimization algorithm constructed 5-bar contours. Here, the Pearson coefficient measures the relationship between eye movements and the neural selectivity for one stimulus over another. We found no systematic correlations between any feature of the monkeys' eye movements and their neural responses, beyond the chance correlations we would expect if the relationship between response rate and eye movement were completely random.

Figure 23



3.7 SUPPLEMENTAL EXPERIMENTAL PROCEDURES

3.7.1 EYE TRACKING ANALYSES

We performed two analyses to confirm that the neural shape selectivity was not an artifact of systematic biases in the monkeys' eye movements during the stimulus presentation, within the delay period of a trial (e.g., Fig. 11B). First, we plotted the eye position distributions for each monkey, using separate sets of trials corresponding to the strongest and weakest 5-bar stimuli (Fig. 22). The distributions are two-dimensional spatial plots, showing the probability that the monkeys' gaze fell within each region of the fixation window. In our experiments, the fixation windows ranged from 0.4° to 0.5° in radius (trials with eye movements outside this window were aborted), but the majority of the gaze directions fell within 0.2° of the fixation point. The shape and spatial extent of the eye position distributions did not differ between presentations of optimal and suboptimal contours, so the differences in neural responses to those stimuli cannot be attributed to systematic deviations in eye position. Moreover, the total amount of eye movement in each set of trials was virtually identical, indicating that monkeys did not make more microsaccades or re-fixations during one set of stimulus conditions than in the other. The precision of our eye tracking

system is sufficient to detect saccades as small as 0.05° in either the vertical or horizontal dimensions (Li et al., 2006).

The concern about stimulus-related differences in fixation is based on the assumption that those differences engender systematic differences in the neural response. In our second analysis, we therefore computed the trial-by-trial correlations between the neural response and several distinct features of eye movements for each monkey (Fig. 23). We found no correlations beyond chance level between neuronal firing rates and any feature of the monkeys' eye movements, including the total length of the eye position trajectory, the maximum displacement of the gaze direction from the fixation point, and the mean displacement of the eye position from the fixation point. The results demonstrate that, even if there were small differences in eye position or movement between stimuli, they would not have biased our measurements of neural selectivity.

3.8 ACKNOWLEDGMENTS

We thank Valentin Piëch and Shimon Ullman for helpful discussions.

4 FUTURE DIRECTIONS

4.1 INTRODUCTION

The theoretical and experimental work described in this thesis provide fertile ground for further investigation, not just to extend the conclusions presented here, but to answer exciting questions about the circuit diagram underlying cortical processing and visual perception. Extracellular recording techniques reveal the astounding computations the cortex performs, but *how* it carries out those calculations is a baffling question. Here, I outline numerous experiments, drawing on a broad range of established and emerging techniques, to address this challenging question.

4.2 RELATING AMD PSYCHOPHYSICS TO COMPUTATIONAL SIMULATIONS

Toward the goal of understanding the cortical algorithms run by the brain, one of the fundamental steps is to unravel the tangle of cortical connections that give rise to neural responses. In this vein, extending the computational model of cortical reorganization and testing its predictions promises to yield new insights. Progress in this direction should start by sharpening our

understanding of visual perception in AMD and its correspondence to the computational model. The deduction from the model, that the horizontal connections dictate rules of sensory integration and perceptual grouping, rests on the similarity between the simulated fill-in and the perception described by AMD subjects. But this correspondence is vague, since the psychophysical literature on AMD perception is sparse and incomplete. There are only a handful of studies that report the perception of human subjects with retinal damage, and those studies focus on the perceptual fill-in of artificially simple stimuli like lines, gratings, and uniform arrays of dots (Craik, 1966; Gerrits and Timmerman, 1969; Schuchard, 1993, 1995; Zur and Ullman, 2003). Even the perceptual fill-in phenomenon itself is intrinsically culpable for the dearth of psychophysical data, since retinal lesions that are small enough to be patched by cortical reorganization often go unnoticed. By the time subjects are diagnosed with AMD, their retinal lesions have become so large they are no longer effectively filled-in.

Developing a focused picture of perceptual fill-in will require a random sample of elderly subjects (≥ 65 years of age) for a longitudinal human psychophysics study. The advances contributed by the study would come from finding subjects at an early to intermediate stage of the disease, when

retinal lesions are effectively filled-in, and from comparing their perception of natural images to the simulations produced by a computational model. The model would simulate each subject's perception, using the pattern of binocular deterioration measured from his or her retinae. An iterative process—of presenting the simulation results to the patients, asking them to report features that coincide or disagree with their own perception, and adjusting the model parameters to match simulation with perception—would be used to refine the model. Furthermore, manipulations of the input images outside the retinal lesions could be introduced, and the predicted perceptual changes could be compared to the reported visual consequences. Since expectation is known to profoundly influence perceptual fill-in, the experimental manipulations of the subjects' perceptions should include not only physical changes to the images but also cognitive changes to the observers' expectations. Still, the computational model is only intended to reproduce the perceptual features generated by the reorganization in V1. The contributions of the rest of the visual system, which are not captured by the model, will undoubtedly exert important influences on perception. The interpretation of the results from these experiments should therefore be tempered by acknowledging the limited scope of the model. On the other hand, repetitive transcranial magnetic stimulation (rTMS) (Fitzgerald et al.,

2006) could potentially be used to determine the cortical origins of specific perceptual features. For instance, if the reported perception of an AMD subject contains features that cannot be explained by the V1 model, the cause could either be the contribution of extrastriate cortical areas or a missing component of the model in V1. To disambiguate between these two alternatives, one could observe whether the discordant feature disappears when specific extrastriate areas are transiently silenced by rTMS.

4.3 IMPROVING THE COMPUTATIONAL MODEL

Alongside efforts to establish a more rigorous link between the observed and simulated perceptual fill-in, the theoretical model itself could be improved by incorporating insights gleaned from the electrophysiological recordings in V1. The existing models of lateral connectivity and network interactions in V1, including the model of cortical reorganization, assume that all neurons express the same pattern of connections, based on the shape of the AF. The connectivity kernels in these models are also typically static—they represent hardwired sets of connections. The experimental recordings in V1, however, revealed several layers of added complexity that are not captured by these simple models. Importantly, different neurons express different

patterns of connectivity. And at any moment in time, even over the population of all recorded neurons, the network in V1 expresses only a subset of the interactions described by the AF. Moreover, these subsets shift dynamically as a function of visual experience and cognitive state. While the population activity is consistent with the Gestalt principles of continuity and smoothness, many individual neurons, under specific cognitive states, prefer stimulus shapes with surprising levels of curvature and discontinuity. Improved network models of V1, whether they seek to explain cortical reorganization or normal patterns of neuronal activity, must include these features. In particular, a rigorous attempt should be made to model how different behavioral states might dynamically reshape the effective lateral connections. During our simulations of perceptual fill-in, we observed patterns of simulated neural activity and perceptual fill-in that would be clearly improved by introducing some of these features.

4.4 TESTING THE THEORETICAL MODEL WITH MULTIELECTRODE ARRAYS

The most fundamental assumption of the cortical reorganization model is that neurons in the LPZ are activated by neurons in the peri-LPZ, via the

same horizontal connections that mediate normal sensory integration. While it is known that the RFs of LPZ neurons shift from inside to outside retinal lesions during cortical reorganization, it is not known whether those RF shifts follow any particular geometric rules. If the central idea of the computational model is correct, then the RFs of LPZ neurons should move toward subsets of their AFs. Moreover, both the theoretical and experimental work suggest that the sets of horizontal connections responsible for activating LPZ neurons (i.e., the active subsets of their AFs) should shift according to behavioral state. Accordingly, the post-lesion RF positions and orientation preferences for LPZ neurons should change dynamically as a function of perceptual task and expectation.

The key to testing the core tenet of the model is to record from the same cells, or at least from the same cortical volume, before and after making retinal lesions. The recent emergence of technology for chronic multielectrode recordings makes the methodology for this experiment straightforward. Multielectrode arrays can be chronically implanted in macaque V1 and used to characterize the response properties of neurons near each electrode (Nordhausen et al., 1996; Suner et al., 2005; Kelly et al., 2007). Because successful array implantations allow the experimenter to

record from neighboring cells for months to years, the arrays could track the same group of cells both before and after retinal lesions ablate their bottom-up input. The stimuli that drive cortical responses in the LPZ could then be directly compared with the contextual stimuli that modulated their responses before the retinal lesions. The same contours that maximally facilitate neuronal responses before the lesions should preferentially drive neural activity afterward. Merging this approach with appropriate behavioral tasks could provide a powerful tool for directly relating LPZ activity to perceptual fill-in. For instance, a monkey with a retinal lesion could be trained to match a perceptually filled-in stimulus with a similar image from a set of simulation results. Recording from LPZ neurons while the monkey performs the task would enable comparisons between the animal's perception and the underlying activity in the LPZ. Even further, direct electrical (Salzman et al., 1990; Moore and Fallah, 2001) or optical manipulations (Zhang et al., 2007; Zhang et al., 2010) of neural responses in the LPZ could be correlated with the changes they induce in the monkey's perceptual judgments. These ideas are analogous to the human psychophysics suggested earlier, except that: 1.) a behavioral task (rather than verbal communication) would be used to query the observer's perception; and 2.) the method allows direct access to local populations of V1 neurons in the LPZ.

4.5 A DIRECT TEST OF THE ADAPTIVE PROCESSOR THEORY

The enhanced efficacy of the horizontal connections after retinal lesions can also be leveraged to test general theories of cortical function. The work in this thesis reinforces the burgeoning theory that top-down connections can rapidly rewire the lateral connectivity in any brain region, adaptively switching the computations the area performs as a function of behavioral demands (Gilbert and Sigman, 2007). This “adaptive processor” theory represents a dramatic departure from the current understanding of how top-down interactions modulate sensory processing, but it has never been directly tested. While the standing paradigm holds that top-down influences only change the gain on neural responsiveness, by uniformly amplifying a neuron’s responses to all its inputs (Reynolds and Chelazzi, 2004; Reynolds and Heeger, 2009), this new theory posits that top-down interactions can selectively amplify/suppress neural responses to specific subsets of inputs. This added control would overwhelmingly bolster the computational power of feedback connections, but the extra layer of complexity it entails makes its implementation difficult to envision. The whole theory rests on the idea that top-down interactions can rapidly activate or subdue existing

connections between neurons, an assumption that can be examined in the context of cortical reorganization. If feedback connections can dynamically alter the effective lateral connectivity in the LPZ, then the top-down modulation of particular lateral inputs should induce clear shifts in the basic RF attributes of LPZ neurons, like position and orientation preference. (Recall that, in the LPZ, these basic RF properties likely depend on inputs from the horizontal connections.) Moreover, the top-down modulation of neuronal connections could be assessed with well-established physiological methods for measuring the connection strength between neurons (Luo et al., 2008). Consider two distinct subsets of cells in the peri-LPZ, populations A and B, that send horizontal projections to a postsynaptic neuron in the LPZ. Imagine further that the horizontal connections onto the postsynaptic cell are gated, via top-down connections, by the cognitive and/or visual context of a stimulus, and that two distinct contexts exclusively activate either one set of connections or the other. Specifically, assume that “context A” activates the horizontal connections from cell population A onto the postsynaptic cell, and “context B” activates the connections from subpopulation B. Then stimulating neural population A (either through direct or stimulus-induced methods) should enhance the response of the postsynaptic neuron under “context A” but not “context B”, and vice versa. Inhibiting the neurons in

either peri-LPZ subpopulation should have a corresponding effect, and both manipulations should have consequences on the perceptual fill-in of the stimulus overlying the lesion. The idea here is that an expectation, set up either by the global attributes of the stimulus or by the monkey's behavioral task, would cause top-down influences to selectively activate/deactivate specific sets of horizontal connections, and that the set of currently active connections onto any cell could be measured by observing which presynaptic neurons can elicit activity in the postsynaptic cell. These experiments would provide a powerful test of the theory that top-down interactions gate the lateral inputs onto cells in V1. Their essential components involve experimental control over expectation and neural responses, but not necessarily cortical reorganization. Analogous experiments could also be performed in normal monkeys with intact retinæ.

4.6 LOCALIZING THE CORTICAL ORIGIN OF THE TASK-DEPENDENT AF

Perhaps the most likely alternative to the interpretation of V1 as an adaptive processor is the possibility that sensory integration, shape selectivity, and the cognitive modulation of these are achieved not in V1, but in extrastriate

visual areas. Whenever an interesting neural response property is observed, particularly when it exceeds the computational tasks for which the cortical area is thought to be responsible, the question arises as to whether the property emerges through local, intrinsic activity or whether it merely echoes computations performed elsewhere. To date, the methods used to determine the cortical origin of specific computations have been largely restricted to timing analyses and procedures to temporarily or permanently silence cortical regions that may play a role in the computation (Lamme et al., 1998). Timing analyses based on data recorded from a single cortical area are often problematic because the relative latencies involved in intrinsic network activity compared to feedback interactions are controversial (Lamme et al., 1998; Girard et al., 2001; Hupe et al., 2001; Smith et al., 2006; Gilbert and Sigman, 2007). The cortical silencing approach, whether it employs cooling, pharmacological methods, or permanent lesions, is hampered by logical interpretive difficulties. Even if silencing a particular cortical area abolishes a response property recorded from another area, the result only indicates that cortical activity in the silenced area is requisite, not sufficient, for the recorded response. The problem is particularly severe when one considers the distributed nature of cortical processing, whereby most computations likely involve complex interactions between multiple

areas. Silencing any of the involved regions would degrade or abolish the dependent patterns of activity in the other areas, but that result should not be interpreted to mean that the silenced cortical compartment is the seat of the computation.

More promising methods of localizing the origins of cortical algorithms are becoming prominent. One novel idea is to use simultaneous recordings of local field potentials (LFPs) and spiking activity, obtained from the same electrode (Nielsen et al., 2006; Monosov et al., 2008). Since LFPs are thought to derive from aggregate input and synaptic activity in a local cortical area (Gustafsson, 1984; Mitzdorf, 1985, 1987; Buzsaki and Kandel, 1998; Kamondi et al., 1998; Juergens et al., 1999; Logothetis et al., 2001; Cruikshank et al., 2002; Kaur et al., 2004; Logothetis and Wandell, 2004; Kreiman et al., 2006; Chen et al., 2007; Rasch et al., 2009), and spiking activity is the cortical output, simultaneous measurements of these provide a picture of how local circuits transform their inputs into outputs. Applying this idea to extracellular recordings from V1, like those described in Chapter 3, would provide a strong test of the idea that sensory integration and shape selectivity are emergent properties of recurrent network interactions in V1. If this idea is correct, then shape selectivity in V1 should first emerge in the

form of spiking activity and only later, if at all, in the LFP signal. Another approach to determine the cortical origin of shape selectivity is to simultaneously record from multiple areas across the cortical hierarchy in the visual system. If our model of sensory integration holds true, then the earliest appearance of task-dependent contour selectivity should be in V1, rather than the extrastriate areas of the temporal stream (V2, V4, and IT) traditionally associated with higher form perception. An ambitious experimental design would also include paired recordings from subcortical areas. The pulvinar nuclei of the thalamus, for instance, are thought to play an important role in attention (Robinson and Petersen, 1992; Grieve et al., 2000; Shipp, 2003) and are known to have distributed connections with the visual cortex (Shipp, 2003; Kaas and Lyon, 2007). A rigorous demonstration that shape selectivity emerges first in V1 should also rule out an earlier effect in the pulvinar thalamus.

4.7 RESOLVING THE SPEED OF AF PLASTICITY AND ITS ROLE IN OBJECT PERCEPTION

Beyond the anatomical origins of V1 response properties, interesting questions remain about the properties themselves and about their role in

visual perception. We have proposed that the shape selectivity in V1, because of its diversity and plasticity, could form the neural substrate for computational procedures that are crucial for object recognition. Theoretical work suggests that a complete recognition of objects requires top-down modulation of neural responses, according to the expectation that the stimulus features encoded by the neurons are present in the image (Ullman, 2007; Epshtein et al., 2008). Chapter 3 demonstrates that V1 cells possess two of the necessary attributes to participate in this process: they exhibit tuning for “intermediate-level” visual features that are important in form processing, and their responses to these features can be profoundly reshaped by top-down processes. These expectation-induced shifts in neural selectivity must also be very rapid (on timescales of ~ 100 ms) if V1 neurons are truly engaged in this component of object recognition, but the experimental design employed in Chapter 3 did not allow us to resolve the speed of the selectivity shifts. The stimulus selectivity of the recorded neurons was assessed during blocks of many hundreds of consecutive trials obtained under the same task condition, over the span of tens of minutes to more than one hour. Deriving a more precise bound on the speed of selectivity shifts will require a behavioral task that can rigidly and dynamically control a monkey’s expectations over time periods of a second

or less. This is a fundamentally challenging proposition. It means designing a task that allows the experimenter to cleanly switch the monkey's expectations from one trial to the next, without any cognitive contamination from preceding trials. We attempted to detect rapid changes in neural selectivity by randomly switching the cue shape in successive trials, but we found evidence that this task design set up a simultaneous expectation for all the cue shapes rather than a dynamically shifting expectation from trial to trial (data not shown). The monkey's expectations seemed to bleed through from one trial to the next. The requirement that the monkey's cognitive state during any trial be distinguishable from preceding trials may preclude any design in which the same cues are used intermittently during the course of an experiment. A compelling alternative is an experimental regime in which neural responses are obtained from single trials of a particular behavioral condition, rather than averaged over hundreds of repetitions. It would emulate natural viewing conditions, and it could be achieved by averaging over a population of simultaneously recorded neurons, rather than averaging the responses from the same neuron over many trials. Methods for simultaneously monitoring large populations of neurons are now coming into prominence, including multielectrode arrays and 2-photon functional imaging (Hires et al., 2008; Luo et al., 2008; Mank and Griesbeck, 2008;

Dreosti et al., 2009; Grewe and Helmchen, 2009; Tian et al., 2009). (Other methods also interrogate large populations of neurons, like fMRI and optical imaging, but the signal-to-noise ratio of these techniques is too low to avoid averaging over trials and their relationship to underlying neural activity is poorly understood.)

Our original experiments explored the impact of cognitive states that were established *before* the stimuli were presented. But expectations about a visual scene can also be generated by the scene itself, after an initial stage of cursory sensory processing. It will also be important to investigate this mode of stimulus-driven expectation and its influence on the network activity in V1. In fact, these stimulus-driven effects are the focus of experimental and theoretical investigations into the role of expectation in object recognition (Bar et al., 2001; Bar et al., 2006; Fenske et al., 2006; Ullman, 2007; Epshtein et al., 2008). Moreover, studying expectations that are generated by the global characteristics of an image, rather than by the nuances of a behavioral task, may facilitate sharper temporal transitions between the expectations experienced by an observer from moment to moment. Using stimulus-driven changes in expectation is therefore another approach to resolve the speed of the plastic changes in V1 shape selectivity. (This idea

leads immediately to the use of stimuli derived from natural environments. Natural scenes are rich in structure, replete with recurring stimulus correlations that can engender strong expectations after years of visual experience.)

Even a detailed account of neural response properties and their anatomical substrates would be incomplete without an understanding of how those properties relate to overall cortical function. It has been proposed that the cortex uses expectations derived from learned stimulus correlations to resolve ambiguities in the sensory input (Summerfield and Egner, 2009). We have further proposed that the task-dependency of shape selectivity in V1 may contribute to this process. A major question, then, is whether the dynamic shape selectivity expressed by V1 neurons actually augments an observer's ability to detect or recognize objects in ambiguous or nebulous visual scenes. Information theoretic analyses or other computational approaches, based on the data in Chapter 3, might determine whether V1 responses could be harnessed for improved discrimination. An ideal experimental test would involve precise functional manipulations of the inputs to V1 neurons, to transiently ablate their ability to undergo expectation-dependent shifts in selectivity. Experiments of this kind may

soon be within reach, through the development of genetic and optical tools to functionally manipulate neural circuits *in vivo* (see below).

4.8 STUDYING THE MECHANISMS OF TOP-DOWN CONTROL

The impressive top-down control over sensory processing is now clearer than ever, but the mechanisms by which cognitive influences exert themselves remain unclear. Recent experiments and modeling work, however, are producing tantalizing hints. Biochemically, Herrero et al. have shown that the neurotransmitter acetylcholine, operating through postsynaptic muscarinic receptors, mediates much of the response enhancement in V1 under spatial attention (Herrero et al., 2008). A major question is whether acetylcholine, or another neurotransmitter(s), underlies the more complicated cognitive modulation of higher-order selectivity in V1. That question can be readily addressed by repeating the experiments in Chapter 3 under various pharmacological manipulations. Other investigators have approached the mechanism of top-down modulation from the perspective of dendritic integration (De Meyer and Spratling, 2009). De Meyer and Spratling have proposed a computational model in which the top-down gating of horizontal connections occurs through a nonlinear

integration of synaptic inputs within the apical dendrites of pyramidal neurons. In this model, contextual and cognitive influences enhance pyramidal cell responses only when there is a coincidence of feedback and lateral inputs onto the apical dendrite.

However, neither the description of the acetylcholine contribution to attentional modulation, nor the theoretical model of dendritic integration, explain how specific subsets of horizontal connections may be gated. In fact, the Spratling model assumes that top-down influences uniformly gate the full spectrum of horizontal inputs onto any neuron. Presumably, it could be extended to include the specific modulation of subsets of horizontal inputs, via the integration of distinct horizontal and top-down inputs on separate branches of the apical dendrites. In this model, the top-down-induced plasticity of a cell's geometric tuning should be related to the structural complexity and compartmentalization of its apical dendrites—a theory that may eventually be testable with two-photon functional imaging [e.g., see (Jia et al., 2010)]. Another possibility, but not necessarily a mutually exclusive one, is that different sets of horizontal connections may be gated by different sets of neurotransmitters and postsynaptic receptors along the dendritic tree. In any case, the notion of specific top-down gating of the horizontal

connections requires either a physical and/or neurochemical overlap of functionally interacting top-down and lateral connections on the postsynaptic dendrites.

An alternative theory, which also explains the top-down modulation of higher-order V1 responses, could be implemented with the same machinery that underlies the well-established rule of attentional gain control. It involves modulating the gain on specific sets of presynaptic cells that send projections to a postsynaptic neuron, rather than modulating the connections themselves at the postsynaptic dendrite. The gain changes, which would work just as they do in space-, object-, and feature-based attention, could be mediated by fluctuating UP and DOWN states (Gilbert and Sigman, 2007) in the membrane potential, or by other mechanisms. This model is fundamentally different from the one advocated in this thesis, whereby top-down connections directly gate the horizontal inputs, and there is evidence against it (Motter, 1993; Ito and Gilbert, 1999; McAdams and Maunsell, 1999a). Nevertheless, its relative simplicity and its mechanistic relationship to already established rules of attentional control make it an important competitor that should be ruled out with careful experiments (e.g., see the section, *A Direct Test of the Adaptive Processor Theory*.)

Emerging technologies provide unprecedented potential for studying mechanisms of top-down control, by enabling direct measurements of the interacting circuit components in V1. The ultimate experimental goal is to monitor the output of a single V1 neuron while simultaneously measuring the full repertoire of its inputs. Combining single cell electroporation with two-photon functional imaging and viral transgenesis (Callaway, 2008) may make this possible. Rabies virus is known to spread exclusively in the retrograde direction from infected neurons to the presynaptic cells that contact them (Callaway, 2008; Luo et al., 2008). Moreover, the virus can be engineered so that the spread of the infection traverses only a single synaptic step (Wickersham et al., 2007). Using the method of single cell electroporation (Kitamura et al., 2008) or AAV infection (Stettler et al., 2002) to inject DNA encoding a cell-surface receptor (EnvA receptor TVA) and the rabies glycoprotein (RG) into a single neuron, and then infecting the cortex with an EnvA pseudotyped rabies virus that has its RG gene replaced with eGFP, all of the neuron's presynaptic partners could be selectively labeled. Moreover, if SyGCaMP (Dreosti et al., 2009) is used in place of eGFP, one could also monitor the activity of all the presynaptic neurons, by imaging the dendritic tree of the postsynaptic cell. And by using two-photon

guided electrode recordings (or another optical probe), one could simultaneously measure the cell's spiking output, together with all of its inputs. Labeling V1 neurons with eGFP and extrastriate feedback neurons with eRFP (using AAV-mediated transfection, for instance) would further allow the experimenter to distinguish between synaptic inputs arising from intrinsic connections versus corticocortical feedback. Finally, the RF properties and even higher-order tuning characteristics of the presynaptic cells could be measured and related to their functional influence on the postsynaptic cell.

This approach involves a combination of techniques that are difficult even when applied individually, but the potential power of the method is undeniable. For the first time, it would enable an investigator to study the function of the whole cortex within the microcosm of a single neuron's dendrites. The relative contributions of lateral and feedback projections in a cortical area, and the carefully choreographed cacophony of inputs onto single neurons, could be directly observed, potentially in awake and behaving primates. Moreover, the method could provide the first full understanding of the function of any single neuron, by describing its input-output transformations with extraordinary clarity. Still, major technical

hurdles remain. Perhaps the tallest of these is a limitation of two-photon microscopy. In its fullest realization, the method entails the simultaneous imaging of the entire span of a neuron's dendritic tree, but dendrites are distributed over a cortical depth of more than 50 microns. Conventional two-photon microscopy, however, can only image within a single depth plane at once. Fortunately, a recently developed technology, the acousto-optic deflector (AOD), enables very rapid imaging of specific points distributed at arbitrary coordinates in the cortical volume (Reddy et al., 2008; Grewe et al., 2010). This technology allows the experimenter to approach the idealized goal of recording all of a neuron's inputs and outputs simultaneously. Its use in the awake, behaving monkey, however, will require parallel advances in methods to eliminate or compensate for cortical movements.

4.9 CONCLUSION

The ideas described throughout this chapter can be combined and expanded in effective ways. For instance, the rabies virus approach outlined above (or similar variants of the idea) could be used to express opsins like channelrhodopsin-2 or halorhodopsin (Zhang et al., 2007; Zhang et al., 2010) in the presynaptic cells that contact a specific postsynaptic neuron. This

optogenetic approach would enable the investigator to either activate or inhibit specific subsets of lateral or feedback neurons and to observe how those manipulations influence the postsynaptic cell under various cognitive states. Various pharmacological manipulations could be applied in concert, to determine which neurotransmitters are responsible for various pathways of neuronal communication. The use of natural scene stimuli, or at least stimuli with naturalistic levels of complexity, is another recurring idea that can be integrated with the other methods in this chapter. Of course, many details about the proposed experimental designs have been omitted; the purpose of this chapter is not to provide a comprehensive catalogue of all the interesting or important work that might follow this thesis. Rather, it is meant to identify the most important questions that remain about the data in Chapters 2 and 3, and to point the reader in interesting intellectual and methodological directions. Many of the experiments will require multiple steps of ingenious innovation, but the technologies they rely upon are under rapid development. If the experiments can be carried out, they promise profound discoveries about the nature of distributed cortical processing.

5 REFERENCES

Adelson, E.H., and Bergen, J.R. (1985). Spatiotemporal Energy Models for the Perception of Motion. *Journal of the Optical Society of America a-Optics Image Science and Vision* 2, 284-299.

Ahmed, B., Anderson, J.C., Douglas, R.J., Martin, K.A.C., and Nelson, J.C. (1994). Polyneuronal Innervation of Spiny Stellate Neurons in Cat Visual-Cortex. *Journal of Comparative Neurology* 341, 39-49.

Albrecht, D.G., and Hamilton, D.B. (1982). Striate Cortex of Monkey and Cat - Contrast Response Function. *Journal of Neurophysiology* 48, 217-237.

Albright, T.D., Desimone, R., and Gross, C.G. (1984). Columnar Organization of Directionally Selective Cells in Visual Area Mt of the Macaque. *Journal of Neurophysiology* 51, 16-31.

Albright, T.D., and Stoner, G.R. (2002). Contextual influences on visual processing. *Annual Review of Neuroscience* 25, 339-379.

Allman, J., Miezin, F., and 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.

Amidror, I. (2002). Scattered data interpolation methods for electronic imaging systems: a survey. *Journal of Electronic Imaging* 11, 157-176.

Angelucci, A., and 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. *Visual Perception, Part 1, Fundamentals of Vision: Low and Mid-Level Processes in Perception* 154, 93-120.

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

Assad, J.A., and Maunsell, J.H.R. (1995). Neuronal Correlates of Inferred Motion in Primate Posterior Parietal Cortex. *Nature* 373, 518-521.

Baker, C.I., Peli, E., Knouf, N., and Kanwisher, N.G. (2005). Reorganization of visual processing in macular degeneration. *Journal of Neuroscience* 25, 614-618.

Balan, P.F., and Gottlieb, J. (2006). Integration of exogenous input into a dynamic salience map revealed by perturbing attention. *Journal of Neuroscience* 26, 9239-9249.

Ballester, C., Bertalmio, M., Caselles, V., Sapiro, G., and Verdera, J. (2001). Filling-in by joint interpolation of vector fields and gray levels. *Ieee Transactions on Image Processing* 10, 1200-1211.

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience* 5, 617-629.

Bar, M., Kassam, K.S., Ghuman, A.S., Boshyan, J., Schmidt, A.M., Dale, A.M., Hamalainen, M.S., Marinkovic, K., Schacter, D.L., Rosen, B.R., and Halgren, E. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences of the United States of America* 103, 449-454.

Bar, M., Tootell, R.B.H., Schacter, D.L., Greve, D.N., Fischl, B., Mendola, J.D., Rosen, B.R., and Dale, A.M. (2001). Cortical mechanisms specific to explicit visual object recognition. *Neuron* 29, 529-535.

Bashinski, H.S., and Bacharach, V.R. (1980). Enhancement of Perceptual Sensitivity as the Result of Selectively Attending to Spatial Locations. *Perception & Psychophysics* 28, 241-248.

Ben-Shahar, O., and Zucker, S. (2004). Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Computation* 16, 445-476.

Bichot, N.P., Schall, J.D., and Thompson, K.G. (1996). Visual feature selectivity In frontal eye fields induced by experience in mature macaques. *Nature* 381, 697-699.

Bichot, N.P., Thompson, K.G., Rao, S.C., and Schall, J.D. (2001a). Continuous processing in macaque frontal cortex during visual search. *Neuropsychologia* 39, 972-982.

Bichot, N.P., Thompson, K.G., Rao, S.C., and Schall, J.D. (2001b). Reliability of macaque frontal eye field neurons signaling saccade targets during visual search. *Journal of Neuroscience* 21, 713-725.

Bisley, J.W., and Goldberg, M.E. (2003). Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299, 81-86.

Bisley, J.W., and Goldberg, M.E. (2010). Attention, Intention, and Priority in the Parietal Lobe. *Annual Review of Neuroscience* 33, 1-21.

Blaser, E., Pylyshyn, Z.W., and Holcombe, A.O. (2000). Tracking an object through feature space. *Nature* 408, 196-199.

Blatt, G.J., Andersen, R.A., and Stoner, G.R. (1990). Visual Receptive-Field Organization and Cortico-Cortical Connections of the Lateral Intraparietal Area (Area Lip) in the Macaque. *Journal of Comparative Neurology* 299, 421-445.

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

Bradley, D.C., and Andersen, R.A. (1998). Center-surround antagonism based on disparity in primate area MT. *Journal of Neuroscience* 18, 7552-7565.

Buffalo, E.A., Fries, P., Landman, R., Liang, H.L., and Desimone, R. (2010). A backward progression of attentional effects in the ventral stream. *Proceedings of the National Academy of Sciences of the United States of America* 107, 361-365.

Burke, W. (1999). Psychophysical observations concerned with a foveal lesion (macular hole). *Vision Research* 39, 2421-2427.

Buschman, T.J., and Miller, E.K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* 315, 1860-1862.

Buzsaki, G., and Kandel, A. (1998). Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat. *Journal of Neurophysiology* 79, 1587-1591.

Calford, M.B., Chino, Y.M., Das, A., Eysel, U.T., Gilbert, C.D., Heinen, S.J., Kaas, J.H., and Ullman, S. (2005). Neuroscience - Rewiring the adult brain. *Nature* 438, E3-E3.

Calford, M.B., and Tweedale, R. (1988). Immediate and Chronic Changes in Responses of Somatosensory Cortex in Adult Flying-Fox after Digit Amputation. *Nature* 332, 446-448.

Calford, M.B., and Tweedale, R. (1991). Acute Changes in Cutaneous Receptive-Fields in Primary Somatosensory Cortex after Digit Denervation in Adult Flying Fox. *Journal of Neurophysiology* 65, 178-187.

Calford, M.B., Wang, C., Taglianetti, V., Waleszczyk, W.J., Burke, W., and Dreher, B. (2000). Plasticity in adult cat visual cortex (area 17) following circumscribed monocular lesions of all retinal layers. *Journal of Physiology-London* 524, 587-602.

Calford, M.B., Wright, L.L., Metha, A.B., and Taglianetti, V. (2003). Topographic plasticity in primary visual cortex is mediated by local corticocortical connections. *Journal of Neuroscience* 23, 6434-6442.

Callaway, E.M. (2008). Transneuronal circuit tracing with neurotropic viruses. *Current Opinion in Neurobiology* 18, 617-623.

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

Cavada, C., Company, T., Tejedor, J., Cruz-Rizzolo, R.J., and Reinoso-Suarez, F. (2000). The anatomical connections of the macaque monkey orbitofrontal cortex. A review. *Cerebral Cortex* 10, 220-242.

Chelazzi, L., Duncan, J., Miller, E.K., and Desimone, R. (1998). Responses of neurons in inferior temporal cortex during memory-guided visual search. *Journal of Neurophysiology* 80, 2918-2940.

Chelazzi, L., Miller, E.K., Duncan, J., and Desimone, R. (1993). A Neural Basis for Visual-Search in Inferior Temporal Cortex. *Nature* 363, 345-347.

Chelazzi, L., Miller, E.K., Duncan, J., and Desimone, R. (2001). Responses of neurons in macaque area V4 during memory-guided visual search. *Cerebral Cortex* 11, 761-772.

Chen, C.M., Lakatos, P., Shah, A.S., Mehta, A.D., Givre, S.J., Javitt, D.C., and Schroeder, C.E. (2007). Functional anatomy and interaction of fast and slow visual pathways in macaque monkeys. *Cerebral Cortex* 17, 1561-1569.

Chino, Y.M., Kaas, J.H., Smith, E.L., Langston, A.L., and Cheng, H. (1992). Rapid Reorganization of Cortical Maps in Adult Cats Following Restricted Deafferentation in Retina. *Vision Research* 32, 789-796.

Chino, Y.M., Smith, E.L., Kaas, J.H., Sasaki, Y., and Cheng, H. (1995). Receptive-Field Properties of Deafferentated Visual Cortical-Neurons after Topographic Map Reorganization in Adult Cats. *Journal of Neuroscience* 15, 2417-2433.

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

Chun, M.M., and Marois, R. (2002). The dark side of visual attention. *Current Opinion in Neurobiology* 12, 184-189.

Conway, B.R., Hubel, D.H., and Livingstone, M.S. (2002). Color contrast in macaque V1. *Cerebral Cortex* 12, 915-925.

Corbetta, M., and Shulman, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience* 3, 201-215.

Craik, K.J.W. (1966). On the Effects of Looking at the Sun. In *The Nature of Psychology*, S.L. Sherwood, ed. (Cambridge: Cambridge University Press), pp. 98-101.

Crist, R.E., Li, W., and Gilbert, C.D. (2001). Learning to see: experience and attention in primary visual cortex. *Nature Neuroscience* 4, 519-525.

Cruikshank, S.J., Rose, H.J., and Metherate, R. (2002). Auditory thalamocortical synaptic transmission in vitro. *Journal of Neurophysiology* 87, 361-384.

Damasio, A.R., Damasio, H., and Chui, H.C. (1980). Neglect Following Damage to Frontal-Lobe or Basal Ganglia. *Neuropsychologia* 18, 123-132.

Darian-Smith, C., and Gilbert, C.D. (1994). Axonal Sprouting Accompanies Functional Reorganization in Adult Cat Striate Cortex. *Nature* 368, 737-740.

Darian-Smith, C., and Gilbert, C.D. (1995). Topographic Reorganization in the Striate Cortex of the Adult Cat and Monkey Is Cortically Mediated. *Journal of Neuroscience* 15, 1631-1647.

Das, A. (1996). Orientation in visual cortex: A simple mechanism emerges. *Neuron* 16, 477-480.

Das, A., and Gilbert, C.D. (1995a). Long-Range Horizontal Connections and Their Role in Cortical Reorganization Revealed by Optical-Recording of Cat Primary Visual-Cortex. *Nature* 375, 780-784.

Das, A., and Gilbert, C.D. (1995b). Receptive-Field Expansion in Adult Visual-Cortex Is Linked to Dynamic Changes in Strength of Cortical Connections. *Journal of Neurophysiology* 74, 779-792.

David, S.V., Hayden, B.Y., Mazer, J.A., and Gallant, J.L. (2008). Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron* 59, 509-521.

Dayan, P., and Abbott, L.F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*, Vol 1 (Cambridge, Massachusetts: MIT Press).

De Meyer, K., and Spratling, M.W. (2009). A model of non-linear interactions between cortical top-down and horizontal connections explains the attentional gating of collinear facilitation. *Vision Research* 49, 553-568.

De Valois, R.L., Albrecht, D.G., and Thorell, L.G. (1982a). Spatial-Frequency Selectivity of Cells in Macaque Visual-Cortex. *Vision Research* 22, 545-559.

De Valois, R.L., Yund, E.W., and Hepler, N. (1982b). The Orientation and Direction Selectivity of Cells in Macaque Visual-Cortex. *Vision Research* 22, 531-544.

Dorris, M.C., and Glimcher, P.W. (2004). Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 44, 365-378.

Dow, B.M., Snyder, A.Z., Vautin, R.G., and Bauer, R. (1981). Magnification Factor and Receptive-Field Size in Foveal Striate Cortex of the Monkey. *Experimental Brain Research* 44, 213-228.

Dreosti, E., Odermatt, B., Dorostkar, M.M., and Lagnado, L. (2009). A genetically encoded reporter of synaptic activity in vivo. *Nature Methods* 6, 883-U122.

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General* 113, 501-517.

Eagleman, D.M. (2001). Visual illusions and neurobiology. *Nature Reviews Neuroscience* 2, 920-926.

Egly, R., Driver, J., and Rafal, R.D. (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General* 123, 161-177.

Epshtein, B., Lifshitz, I., and Ullman, S. (2008). Image interpretation by a single bottom-up top-down cycle. *Proceedings of the National Academy of Sciences of the United States of America* 105, 14298-14303.

Ernst, U.A., Mandon, S., Pawelzik, K.R., and Kreiter, A.K. (2004). How ideal do macaque monkeys integrate contours? *Neurocomputing* 58-60, 971-977.

Eysel, U.T. (1982). Functional Reconnections without New Axonal Growth in a Partially Denervated Visual Relay Nucleus. *Nature* 299, 442-444.

Eysel, U.T., Gonzalezaguilar, F., and Mayer, U. (1981). Late Spreading of Excitation in the Lateral Geniculate-Nucleus Following Visual

Deafferentation Is Independent of the Size of Retinal Lesions. *Brain Research* 204, 189-193.

Feldmeyer, D., Lubke, J., Silver, R.A., and Sakmann, B. (2002). Synaptic connections between layer 4 spiny neurone-layer 2/3 pyramidal cell pairs in juvenile rat barrel cortex: physiology and anatomy of interlaminar signalling within a cortical column. *Journal of Physiology-London* 538, 803-822.

Felleman, D.J., and Van Essen, D.C. (1991). Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cerebral Cortex* 1, 1-47.

Fenske, M.J., Aminoff, E., Gronau, N., and Bar, M. (2006). Top-down facilitation of visual object recognition: object-based and context-based contributions. *Visual Perception, Pt 2: Fundamentals of Awareness: Multi-Sensory Integration and High-Order Perception* 155, 3-21.

Field, D.J., Hayes, A., and Hess, R.F. (1993). Contour Integration by the Human Visual-System - Evidence for a Local Association Field. *Vision Research* 33, 173-193.

Fischer, B., and Kruger, J. (1979). Disparity Tuning and Binocularity of Single Neurons in Cat Visual-Cortex. *Experimental Brain Research* 35, 1-8.

Fitzgerald, P.B., Fountain, S., and Daskalakis, Z.J. (2006). A comprehensive review of the effects of rTMS on motor cortical excitability and inhibition. *Clinical Neurophysiology* 117, 2584-2596.

Foster, K.H., Gaska, J.P., Nagler, M., and Pollen, D.A. (1985). Spatial and Temporal Frequency-Selectivity of Neurons in Visual Cortical Areas V1 and V2 of the Macaque Monkey. *Journal of Physiology-London* 365, 331-363.

Freedman, D.J., and Assad, J.A. (2009). Distinct Encoding of Spatial and Nonspatial Visual Information in Parietal Cortex. *Journal of Neuroscience* 29, 5671-5680.

Freeman, W.T., and Adelson, E.H. (1991). The Design and Use of Steerable Filters. *Ieee Transactions on Pattern Analysis and Machine Intelligence* 13, 891-906.

Fries, W., Albus, K., and Creutzfeldt, O.D. (1977). Effects of Interacting Visual-Patterns on Single Cell Responses in Cats Striate Cortex. *Vision Research* 17, 1001-1008.

Fritz, J.B., Elhilali, M., David, S.V., and Shamma, S.A. (2007). Auditory attention - focusing the searchlight on sound. *Current Opinion in Neurobiology* 17, 437-455.

Fuster, J.M., Bauer, R.H., and Jervey, J.P. (1985). Functional interactions between inferotemporal and prefrontal cortex in a cognitive task. *Brain Research* 330, 299-307.

Gaffan, D., and Hornak, J. (1997). Visual neglect in the monkey - Representation and disconnection. *Brain* 120, 1647-1657.

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

Gerrits, H.J., and Timmerman, G.J. (1969). Filling-in Process in Patients with Retinal Scotomata. *Vision Research* 9, 439-442.

Gerrits, H.J.M., Dehaan, B., and Vendrik, A.J.H. (1966). Experiments with Retinal Stabilized Images . Relations between Observations and Neural Data. *Vision Research* 6, 427-&.

Ghose, G.M., and Bearl, D.W. (2010). Attention directed by expectations enhances receptive fields in cortical area MT. *Vision Research* 50, 441-451.

Giannikopoulos, D.V., and Eysel, U.T. (2006). Dynamics and specificity of cortical map reorganization after retinal lesions. *Proceedings of the National Academy of Sciences of the United States of America* 103, 10805-10810.

Gilbert, C.D., Hirsch, J.A., and Wiesel, T.N. (1990). Lateral Interactions in Visual-Cortex. *Cold Spring Harbor Symposia on Quantitative Biology* 55, 663-677.

Gilbert, C.D., and Sigman, M. (2007). Brain states: Top-down influences in sensory processing. *Neuron* 54, 677-696.

Gilbert, C.D., and Wiesel, T.N. (1979). Morphology and Intra-Cortical Projections of Functionally Characterized Neurons in the Cat Visual-Cortex. *Nature* 280, 120-125.

Gilbert, C.D., and Wiesel, T.N. (1983). Clustered Intrinsic Connections in Cat Visual-Cortex. *Journal of Neuroscience* 3, 1116-1133.

Gilbert, C.D., and Wiesel, T.N. (1989). Columnar Specificity of Intrinsic Horizontal and Corticocortical Connections in Cat Visual-Cortex. *Journal of Neuroscience* 9, 2432-2442.

Gilbert, C.D., and Wiesel, T.N. (1992). Receptive-Field Dynamics in Adult Primary Visual-Cortex. *Nature* 356, 150-152.

Girard, P., Hupe, J.M., and Bullier, J. (2001). Feedforward and feedback connections between areas V1 and V2 of the monkey have similar rapid conduction velocities. *Journal of Neurophysiology* 85, 1328-1331.

Gottlieb, J., and Goldberg, M.E. (1999). Activity of neurons in the lateral intraparietal area of the monkey during an antisaccade task. *Nature Neuroscience* 2, 906-912.

Gottlieb, J.P., Kusunoki, M., and Goldberg, M.E. (1998). The representation of visual salience in monkey parietal cortex. *Nature* 391, 481-484.

- Grewe, B.F., and Helmchen, F. (2009). Optical probing of neuronal ensemble activity. *Current Opinion in Neurobiology* 19, 520-529.
- Grewe, B.F., Langer, D., Kasper, H., Kampa, B.M., and Helmchen, F. (2010). High-speed in vivo calcium imaging reveals neuronal network activity with near-millisecond precision. *Nature Methods* 7, 399-U391.
- Grieve, K.L., Acuna, C., and Cudeiro, J. (2000). The primate pulvinar nuclei: vision and action. *Trends in Neurosciences* 23, 35-39.
- Grinvald, A., Lieke, E., Frostig, R.D., Gilbert, C.D., and Wiesel, T.N. (1986). Functional Architecture of Cortex Revealed by Optical Imaging of Intrinsic Signals. *Nature* 324, 361-364.
- Grosov, D.H., Shapley, R.M., and Hawken, M.J. (1993). Macaque-V1 Neurons Can Signal Illusory Contours. *Nature* 365, 550-552.
- Grossberg, S., and Mingolla, E. (1985). Neural Dynamics of Perceptual Grouping - Textures, Boundaries, and Emergent Segmentations. *Perception & Psychophysics* 38, 141-171.
- Gustafsson, B. (1984). Afterpotentials and Transduction Properties in Different Types of Central Neurons. *Archives Italiennes De Biologie* 122, 17-30.
- Guy, G., and Medioni, G. (1996). Inferring global perceptual contours from local features. *International Journal of Computer Vision* 20, 113-133.
- Haenny, P.E., Maunsell, J.H.R., and Schiller, P.H. (1988). State Dependent Activity in Monkey Visual-Cortex .2. Retinal and Extraretinal Factors in V4. *Experimental Brain Research* 69, 245-259.
- Haenny, P.E., and Schiller, P.H. (1988). State Dependent Activity in Monkey Visual-Cortex .1. Single Cell-Activity in V1 and V4 on Visual Tasks. *Experimental Brain Research* 69, 225-244.

Handy, T.C., Kingstone, A., and Mangun, G.R. (1996). Spatial distribution of visual attention: Perceptual sensitivity and response latency. *Perception & Psychophysics* 58, 613-627.

Harrison, S.A., and Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458, 632-635.

Hawken, M.J., and Parker, A.J. (1987). Spatial Properties of Neurons in the Monkey Striate Cortex. *Proceedings of the Royal Society of London Series B-Biological Sciences* 231, 251-288.

Heilman, K.M., and Valenstein, E. (1972). Frontal Lobe Neglect in Man. *Neurology* 22, 660-&.

Heinen, S.J., and Skavenski, A.A. (1991). Recovery of Visual Responses in Foveal V1 Neurons Following Bilateral Foveal Lesions in Adult Monkey. *Experimental Brain Research* 83, 670-674.

Herrero, J.L., Roberts, M.J., Delicato, L.S., Gieselmann, M.A., Dayan, P., and Thiele, A. (2008). Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* 454, 1110-1114.

Hess, R.F., Hayes, A., and Field, D.J. (2003). Contour integration and cortical processing. *Journal of Physiology-Paris* 97, 105-119.

Hires, S.A., Tian, L., and Looger, L.L. (2008). Reporting neural activity with genetically encoded calcium indicators. *Brain Cell Biology* 36, 69-86.

Hirsch, J.A., and Gilbert, C.D. (1991). Synaptic Physiology of Horizontal Connections in the Cats Visual-Cortex. *Journal of Neuroscience* 11, 1800-1809.

Hubel, D.H., and Livingstone, M.S. (1987). Segregation of Form, Color, and Stereopsis in Primate Area-18. *Journal of Neuroscience* 7, 3378-3415.

Hubel, D.H., and Wiesel, T.N. (1959). Receptive Fields of Single Neurones in the Cats Striate Cortex. *Journal of Physiology-London* 148, 574-591.

Hubel, D.H., and Wiesel, T.N. (1962). Receptive Fields, Binocular Interaction and Functional Architecture in Cats Visual Cortex. *Journal of Physiology-London* 160, 106-154.

Hubel, D.H., and Wiesel, T.N. (1968). Receptive Fields and Functional Architecture of Monkey Striate Cortex. *Journal of Physiology-London* 195, 215-243.

Hubel, D.H., and Wiesel, T.N. (1974). Uniformity of Monkey Striate Cortex - Parallel Relationship between Field Size, Scatter, and Magnification Factor. *Journal of Comparative Neurology* 158, 295-306.

Hubener, M., Shoham, D., Grinvald, A., and Bonhoeffer, T. (1997). Spatial relationships among three columnar systems in cat area 17. *Journal of Neuroscience* 17, 9270-9284.

Hupe, J.M., James, A.C., Girard, P., Lomber, S.G., Payne, B.R., and Bullier, J. (2001). Feedback connections act on the early part of the responses in monkey visual cortex. *Journal of Neurophysiology* 85, 134-145.

Imig, T.J., and Adrian, H.O. (1977). Binaural Columns in Primary Field (A1) of Cat Auditory-Cortex. *Brain Research* 138, 241-257.

Ipata, A.E., Gee, A.L., Bisley, J.W., and Goldberg, M.E. (2009). Neurons in the lateral intraparietal area create a priority map by the combination of disparate signals. *Experimental Brain Research* 192, 479-488.

Ipata, A.E., Gee, A.L., Goldberg, M.E., and Bisley, J.W. (2006a). Activity in the lateral intraparietal area predicts the goal and latency of saccades in a free-viewing visual search task. *Journal of Neuroscience* 26, 3656-3661.

Ipata, A.E., Gee, A.L., Gottlieb, J., Bisley, J.W., and Goldberg, M.E. (2006b). LIP responses to a popout stimulus are reduced if it is overtly ignored. *Nature Neuroscience* 9, 1071-1076.

Ito, M., and Gilbert, C.D. (1999). Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron* 22, 593-604.

Jagadeesh, B., Wheat, H.S., and Ferster, D. (1993). Linearity of Summation of Synaptic Potentials Underlying Direction Selectivity in Simple Cells of the Cat Visual-Cortex. *Science* 262, 1901-1904.

Jia, H.B., Rochefort, N.L., Chen, X.W., and Konnerth, A. (2010). Dendritic organization of sensory input to cortical neurons in vivo. *Nature* 464, 1307-1312.

Johnson, E.N., Hawken, M.J., and Shapley, R. (2001). The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience* 4, 409-416.

Juergens, E., Guettler, A., and Eckhorn, R. (1999). Visual stimulation elicits locked and induced gamma oscillations in monkey intracortical- and EEG-potentials, but not in human EEG. *Experimental Brain Research* 129, 247-259.

Julesz, B. (1981). Textons, the Elements of Texture-Perception, and Their Interactions. *Nature* 290, 91-97.

Kaas, J.H., Krubitzer, L.A., Chino, Y.M., Langston, A.L., Polley, E.H., and Blair, N. (1990). Reorganization of Retinotopic Cortical Maps in Adult Mammals after Lesions of the Retina. *Science* 248, 229-231.

Kaas, J.H., and Lyon, D.C. (2007). Pulvinar contributions to the dorsal and ventral streams of visual processing in primates. *Brain Research Reviews* 55, 285-296.

Kamondi, A., Acsady, L., Wang, X.J., and Buzsaki, G. (1998). Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: Activity-dependent phase-precession of action potentials. *Hippocampus* 8, 244-261.

Kapadia, M.K., Gilbert, C.D., and Westheimer, G. (1994). A Quantitative Measure for Short-Term Cortical Plasticity in Human Vision. *Journal of Neuroscience* 14, 451-457.

Kapadia, M.K., Ito, M., Gilbert, C.D., and 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.

Kapadia, M.K., Westheimer, G., and Gilbert, C.D. (2000). Spatial distribution of contextual interactions in primary visual cortex and in visual perception. *Journal of Neurophysiology* 84, 2048-2062.

Kastner, S., and Ungerleider, L.G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience* 23, 315-341.

Kaur, S., Lazar, R., and Metherate, R. (2004). Intracortical pathways determine breadth of subthreshold frequency receptive fields in primary auditory cortex. *Journal of Neurophysiology* 91, 2551-2567.

Keck, T., Mrsic-Flogel, T.D., Afonso, M.V., Eysel, U.T., Bonhoeffer, T., and Hubener, M. (2008). Massive restructuring of neuronal circuits during functional reorganization of adult visual cortex. *Nature Neuroscience* 11, 1162-1167.

Kelly, R.C., Smith, M.A., Samonds, J.M., Kohn, A., Bonds, A.B., Movshon, J.A., and Lee, T.S. (2007). Comparison of recordings from microelectrode arrays and single electrodes in the visual cortex. *Journal of Neuroscience* 27, 261-264.

- Kinoshita, M., and Komatsu, H. (2001). Neural representation of the luminance and brightness of a uniform surface in the macaque primary visual cortex. *Journal of Neurophysiology* 86, 2559-2570.
- Kiorpes, L., and Bassin, S.A. (2003). Development of contour integration in macaque monkeys. *Visual Neuroscience* 20, 567-575.
- Kitamura, K., Judkewitz, B., Kano, M., Denk, W., and Hausser, M. (2008). Targeted patch-clamp recordings and single-cell electroporation of unlabeled neurons in vivo. *Nature Methods* 5, 61-67.
- Knierim, J.J., and 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.
- Kohonen, T. (2001). *Self-Organizing Maps*, Third Edition edn (Berlin: Springer-Verlag).
- Kovacs, I., Kozma, P., Feher, A., and Benedek, G. (1999). Late maturation of visual spatial integration in humans. *Proceedings of the National Academy of Sciences of the United States of America* 96, 12204-12209.
- Kreiman, G., Hung, C.P., Kraskov, A., Quiroga, R.Q., Poggio, T., and DiCarlo, J.J. (2006). Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. *Neuron* 49, 433-445.
- Kusunoki, M., Gottlieb, J., and Goldberg, M.E. (2000). The lateral intraparietal area as a salience map: the representation of abrupt onset, stimulus motion, and task relevance. *Vision Research* 40, 1459-1468.
- Lamme, V.A.F. (1995). The Neurophysiology of Figure Ground Segregation in Primary Visual-Cortex. *Journal of Neuroscience* 15, 1605-1615.

Lamme, V.A.F., Super, H., and Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology* 8, 529-535.

Lampl, I., Anderson, J.S., Gillespie, D.C., and Ferster, D. (2001). Prediction of orientation selectivity from receptive field architecture in simple cells of cat visual cortex. *Neuron* 30, 263-274.

Landisman, C.E., and Ts'o, D.Y. (2002). Color processing in macaque striate cortex: Electrophysiological properties. *Journal of Neurophysiology* 87, 3138-3151.

Latto, R., and Cowey, A. (1971). Visual Field Defects after Frontal Eye-Field Lesions in Monkeys. *Brain Research* 30, 1-&.

Lee, D.K., Itti, L., Koch, C., and Braun, J. (1999). Attention activates winner-take-all competition among visual filters. *Nature Neuroscience* 2, 375-381.

Lee, T.S., Mumford, D., Romero, R., and Lamme, V.A.F. (1998). The role of the primary visual cortex in higher level vision. *Vision Research* 38, 2429-2454.

Lee, T.S., and Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* 98, 1907-1911.

Lee, T.S., Yang, C.F., Romero, R.D., and Mumford, D. (2002). Neural activity in early visual cortex reflects behavioral experience and higher-order perceptual saliency. *Nature Neuroscience* 5, 589-597.

Lennie, P., Krauskopf, J., and Sclar, G. (1990). Chromatic Mechanisms in Striate Cortex of Macaque. *Journal of Neuroscience* 10, 649-669.

Levay, S., and Gilbert, C.D. (1976). Laminar Patterns of Geniculocortical Projection in Cat. *Brain Research* 113, 1-19.

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

Li, C.Y., and Li, W. (1994). Extensive Integration Field Beyond the Classical Receptive-Field of Cats Striate Cortical-Neurons - Classification and Tuning Properties. *Vision Research* 34, 2337-2355.

Li, W., and Gilbert, C.D. (2002). Global contour saliency and local colinear interactions. *Journal of Neurophysiology* 88, 2846-2856.

Li, W., Piech, V., and Gilbert, C.D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience* 7, 651-657.

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

Li, W., Piech, V., and Gilbert, C.D. (2008). Learning to link visual contours. *Neuron* 57, 442-451.

Li, W., Thier, P., and Wehrhahn, C. (2000). Contextual influence on orientation discrimination of humans and responses of neurons in V1 of alert monkeys. *Journal of Neurophysiology* 83, 941-954.

Li, W., Thier, P., and Wehrhahn, C. (2001). Neuronal responses from beyond the classic receptive field in VI of alert monkeys. *Experimental Brain Research* 139, 359-371.

Li, Z.P. (1998). A neural model of contour integration in the primary visual cortex. *Neural Computation* 10, 903-940.

Livingstone, M.S., and Hubel, D.H. (1984). Anatomy and Physiology of a Color System in the Primate Visual-Cortex. *Journal of Neuroscience* 4, 309-356.

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412, 150-157.

Logothetis, N.K., and Wandell, B.A. (2004). Interpreting the BOLD signal. *Annual Review of Physiology* 66, 735-769.

Luo, L., Callaway, E.M., and Svoboda, K. (2008). Genetic dissection of neural circuits. *Neuron* 57, 634-660.

Lynch, J.C., and McLaren, J.W. (1989). Deficits of Visual-Attention and Saccadic Eye-Movements after Lesions of Parietooccipital Cortex in Monkeys. *Journal of Neurophysiology* 61, 74-90.

MacEvoy, S.P., and Paradiso, M.A. (2001). Lightness constancy in primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* 98, 8827-8831.

Maffei, L., and Fiorentini, A. (1976). Unresponsive Regions of Visual Cortical Receptive-Fields. *Vision Research* 16, 1131-1139.

Malach, R., Amir, Y., Harel, M., and Grinvald, A. (1993). Relationship between Intrinsic Connections and Functional Architecture Revealed by Optical Imaging and in-Vivo Targeted Biocytin Injections in Primate Striate Cortex. *Proceedings of the National Academy of Sciences of the United States of America* 90, 10469-10473.

Maljkovic, V., and Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory and Cognition* 22, 657-672.

- Mank, M., and Griesbeck, O. (2008). Genetically encoded calcium indicators. *Chemical Reviews* 108, 1550-1564.
- Marr, D. (1982). *Vision : a computational investigation into the human representation and processing of visual information* (San Francisco: Freeman).
- Martinez-Trujillo, J.C., and Treue, S. (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron* 35, 365-370.
- Martinez, L.M., and Alonso, J.M. (2001). Construction of complex receptive fields in cat primary visual cortex. *Neuron* 32, 515-525.
- Martinez, L.M., and Alonso, J.M. (2003). Complex receptive fields in primary visual cortex. *Neuroscientist* 9, 317-331.
- Maunsell, J.H.R., Sclar, G., Nealey, T.A., and Depriest, D.D. (1991). Extraretinal Representations in Area-V4 in the Macaque Monkey. *Visual Neuroscience* 7, 561-573.
- Maunsell, J.H.R., and Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences* 29, 317-322.
- McAdams, C.J., and Maunsell, J.H.R. (1999a). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience* 19, 431-441.
- McAdams, C.J., and Maunsell, J.H.R. (1999b). Effects of attention on the reliability of individual neurons in monkey visual cortex. *Neuron* 23, 765-773.
- Mcguire, B.A., Gilbert, C.D., Rivlin, P.K., and Wiesel, T.N. (1991). Targets of Horizontal Connections in Macaque Primary Visual-Cortex. *Journal of Comparative Neurology* 305, 370-392.

McManus, J.N.J., Ullman, S., and Gilbert, C.D. (2008). A computational model of perceptual fill-in following retinal degeneration. *Journal of Neurophysiology* 99, 2086-2100.

Merzenich, M.M., Kaas, J.H., Wall, J., Nelson, R.J., Sur, M., and Felleman, D. (1983a). Topographic Reorganization of Somatosensory Cortical Areas 3b and 1 in Adult Monkeys Following Restricted Deafferentation. *Neuroscience* 8, 33-55.

Merzenich, M.M., Kaas, J.H., Wall, J.T., Sur, M., Nelson, R.J., and Felleman, D.J. (1983b). Progression of Change Following Median Nerve-Section in the Cortical Representation of the Hand in Areas-3b and Area-1 in Adult Owl and Squirrel-Monkeys. *Neuroscience* 10, 639-665.

Merzenich, M.M., Nelson, R.J., Stryker, M.P., Cynader, M.S., Schoppmann, A., and Zook, J.M. (1984). Somatosensory Cortical Map Changes Following Digit Amputation in Adult Monkeys. *Journal of Comparative Neurology* 224, 591-605.

Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience* 24, 167-202.

Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience* 16, 5154-5167.

Mitzdorf, U. (1985). Current Source-Density Method and Application in Cat Cerebral-Cortex - Investigation of Evoked-Potentials and Eeg Phenomena. *Physiological Reviews* 65, 37-100.

Mitzdorf, U. (1987). Properties of the Evoked-Potential Generators - Current Source-Density Analysis of Visually Evoked-Potentials in the Cat Cortex. *International Journal of Neuroscience* 33, 33-59.

Monosov, I.E., Trageser, J.C., and Thompson, K.G. (2008). Measurements of simultaneously recorded spiking activity and local field potentials suggest that spatial selection emerges in the frontal eye field. *Neuron* 57, 614-625.

Moore, T., and Armstrong, K.M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421, 370-373.

Moore, T., and Fallah, M. (2001). Control of eye movements and spatial attention. *Proceedings of the National Academy of Sciences of the United States of America* 98, 1273-1276.

Moore, T., and Fallah, M. (2004). Microstimulation of the frontal eye field and its effects on covert spatial attention. *Journal of Neurophysiology* 91, 152-162.

Mooser, F., Bosking, W.H., and Fitzpatrick, D. (2004). A morphological basis for orientation tuning in primary visual cortex. *Nature Neuroscience* 7, 872-879.

Moran, J., and Desimone, R. (1985). Selective Attention Gates Visual Processing in the Extrastriate Cortex. *Science* 229, 782-784.

Motter, B.C. (1993). Focal Attention Produces Spatially Selective Processing in Visual Cortical Areas V1, V2, and V4 in the Presence of Competing Stimuli. *Journal of Neurophysiology* 70, 909-919.

Motter, B.C. (1994a). Neural Correlates of Attentive Selection for Color or Luminance in Extrastriate Area V4. *Journal of Neuroscience* 14, 2178-2189.

Motter, B.C. (1994b). Neural Correlates of Feature Selective Memory and Pop-out in Extrastriate Area V4. *Journal of Neuroscience* 14, 2190-2199.

Mountcastle, V.B. (1957). Modality and Topographic Properties of Single Neurons of Cats Somatic Sensory Cortex. *Journal of Neurophysiology* 20, 408-434.

Mountcastle, V.B., Motter, B.C., Steinmetz, M.A., and Sestokas, A.K. (1987). Common and Differential-Effects of Attentive Fixation on the Excitability of Parietal and Prestriate (V4) Cortical Visual Neurons in the Macaque Monkey. *Journal of Neuroscience* 7, 2239-2255.

Murthy, A., Thompson, K.G., and Schall, J.D. (2001). Dynamic dissociation of visual selection from saccade programming in frontal eye field. *Journal of Neurophysiology* 86, 2634-2637.

Nelson, J.I., and Frost, B.J. (1978). Orientation-Selective Inhibition from Beyond Classic Visual Receptive-Field. *Brain Research* 139, 359-365.

Nelson, R.J., Sur, M., Felleman, D.J., and Kaas, J.H. (1980). Representations of the Body-Surface in Postcentral Parietal Cortex of Macaca-Fascicularis. *Journal of Comparative Neurology* 192, 611-643.

Nielsen, K.J., Logothetis, N.K., and Rainer, G. (2006). Dissociation between local field potentials and spiking activity in macaque inferior temporal cortex reveals diagnosticity-based encoding of complex objects. *Journal of Neuroscience* 26, 9639-9645.

Nordhausen, C.T., Maynard, E.M., and Normann, R.A. (1996). Single unit recording capabilities of a 100 microelectrode array. *Brain Research* 726, 129-140.

Nothdurft, H.C., Gallant, J.L., and Van Essen, D.C. (2000). Response profiles to texture border patterns in area V1. *Visual Neuroscience* 17, 421-436.

O'Craven, K.M., Downing, P.E., and Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature* 401, 584-587.

O'Herron, P., and von der Heydt, R. (2009). Short-Term Memory for Figure-Ground Organization in the Visual Cortex. *Neuron* 61, 801-809.

Offen, S., Schuppeck, D., and Heeger, D.J. (2009). The role of early visual cortex in visual short-term memory and visual attention. *Vision Research* 49, 1352-1362.

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

Olshausen, B.A., and Field, D.J. (2006). What is the other 85% of V1 doing? In *23 Problems in Systems Neuroscience*, J.L. van Hemmen, and T.J. Sejnowski, eds. (New York, NY: Oxford University Press), pp. 182-214.

Parent, P., and Zucker, S.W. (1989). Trace Inference, Curvature Consistency, and Curve Detection. *Ieee Transactions on Pattern Analysis and Machine Intelligence* 11, 823-839.

Perona, P., and Malik, J. (1990). Scale-Space and Edge-Detection Using Anisotropic Diffusion. *Ieee Transactions on Pattern Analysis and Machine Intelligence* 12, 629-639.

Peters, A., and Payne, B.R. (1993). Numerical Relationships between Geniculocortical Afferents and Pyramidal Cell Modules in Cat Primary Visual-Cortex. *Cerebral Cortex* 3, 69-78.

Pettet, M.W., and Gilbert, C.D. (1992). Dynamic Changes in Receptive-Field Size in Cat Primary Visual-Cortex. *Proceedings of the National Academy of Sciences of the United States of America* 89, 8366-8370.

Poggio, G.F., and Fischer, B. (1977). Binocular Interaction and Depth Sensitivity in Striate and Prestriate Cortex of Behaving Rhesus-Monkey. *Journal of Neurophysiology* 40, 1392-1405.

Poggio, G.F., Motter, B.C., Squatrito, S., and Trotter, Y. (1985). Responses of Neurons in Visual-Cortex (V1 and V2) of the Alert Macaque to Dynamic Random-Dot Stereograms. *Vision Research* 25, 397-406.

Polat, U., and Sagi, D. (1994). Spatial Interactions in Human Vision - from near to Far Via Experience-Dependent Cascades of Connections. *Proceedings of the National Academy of Sciences of the United States of America* *91*, 1206-1209.

Posner, M.I., Snyder, C.R.R., and Davidson, B.J. (1977). Attention and Detection of Visual Signals. *Bulletin of the Psychonomic Society* *10*, 251-251.

Prenger, R., Wu, M.C.K., David, S.V., and Gallant, J.L. (2004). Nonlinear V1 responses to natural scenes revealed by neural network analysis. *Neural Networks* *17*, 663-679.

Priebe, N.J., and Ferster, D. (2005). Direction selectivity of excitation and inhibition in simple cells of the cat primary visual cortex. *Neuron* *45*, 133-145.

Raizada, R.D.S., and Grossberg, S. (2003). Towards a theory of the laminar architecture of cerebral cortex: Computational clues from the visual system. *Cerebral Cortex* *13*, 100-113.

Rajan, R., Irvine, D.R.F., Wise, L.Z., and Heil, P. (1993). Effect of Unilateral Partial Cochlear Lesions in Adult Cats on the Representation of Lesioned and Unlesioned Cochleas in Primary Auditory-Cortex. *Journal of Comparative Neurology* *338*, 17-49.

Ramachandran, V.S., and Gregory, R.L. (1991). Perceptual Filling in of Artificially Induced Scotomas in Human Vision. *Nature* *350*, 699-702.

Rasch, M., Logthetis, N.K., and Kreiman, G. (2009). From Neurons to Circuits: Linear Estimation of Local Field Potentials. *Journal of Neuroscience* *29*, 13785-13796.

Rasmusson, D.D. (1982). Reorganization of Raccoon Somatosensory Cortex Following Removal of the 5th Digit. *Journal of Comparative Neurology* *205*, 313-326.

Rasmusson, D.D. (1988). Projections of Digit Afferents to the Cuneate Nucleus in the Raccoon before and after Partial Deafferentation. *Journal of Comparative Neurology* 277, 549-556.

Reale, R.A., and Imig, T.J. (1980). Tonotopic Organization in Auditory-Cortex of the Cat. *Journal of Comparative Neurology* 192, 265-291.

Recanzone, G.H., and Wurtz, R.H. (2000). Effects of attention on MT and MST neuronal activity during pursuit initiation. *Journal of Neurophysiology* 83, 777-790.

Reddy, G.D., Kelleher, K., Fink, R., and Saggau, P. (2008). Three-dimensional random access multiphoton microscopy for functional imaging of neuronal activity. *Nature Neuroscience* 11, 713-720.

Reid, R.C., and Alonso, J.M. (1995). Specificity of Monosynaptic Connections from Thalamus to Visual-Cortex. *Nature* 378, 281-284.

Reid, R.C., Soodak, R.E., and Shapley, R.M. (1991). Directional Selectivity and Spatiotemporal Structure of Receptive-Fields of Simple Cells in Cat Striate Cortex. *Journal of Neurophysiology* 66, 505-529.

Reynolds, J.H., and Chelazzi, L. (2004). Attentional modulation of visual processing. *Annual Review of Neuroscience* 27, 611-647.

Reynolds, J.H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience* 19, 1736-1753.

Reynolds, J.H., and Desimone, R. (2003). Interacting roles of attention and visual salience in V4. *Neuron* 37, 853-863.

Reynolds, J.H., and Heeger, D.J. (2009). The Normalization Model of Attention. *Neuron* 61, 168-185.

Reynolds, J.H., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron* 26, 703-714.

Robertson, D., and Irvine, D.R.F. (1989). Plasticity of Frequency Organization in Auditory-Cortex of Guinea-Pigs with Partial Unilateral Deafness. *Journal of Comparative Neurology* 282, 456-471.

Robinson, D.L., and Petersen, S.E. (1992). The Pulvinar and Visual Salience. *Trends in Neurosciences* 15, 127-132.

Rockland, K.S., and Lund, J.S. (1982). Widespread Periodic Intrinsic Connections in the Tree Shrew Visual-Cortex. *Science* 215, 1532-1534.

Roelfsema, P.R., Lamme, V.A.F., and Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395, 376-381.

Roelfsema, P.R., and Spekreijse, H. (2001). The representation of erroneously perceived stimuli in the primary visual cortex. *Neuron* 31, 853-863.

Rossi, A.F., and Paradiso, M.A. (1995). Feature-Specific Effects of Selective Visual-Attention. *Vision Research* 35, 621-634.

Rossi, A.F., and Paradiso, M.A. (1999). Neural correlates of perceived brightness in the retina, lateral geniculate nucleus, and striate cortex. *Journal of Neuroscience* 19, 6145-6156.

Rossi, A.F., Rittenhouse, C.D., and Paradiso, M.A. (1996). The representation of brightness in primary visual cortex. *Science* 273, 1104-1107.

Rust, N.C., Schwartz, O., Movshon, J.A., and Simoncelli, E.P. (2005). Spatiotemporal elements of macaque V1 receptive fields. *Neuron* 46, 945-956.

Salzman, C.D., Britten, K.H., and Newsome, W.T. (1990). Cortical Microstimulation Influences Perceptual Judgments of Motion Direction. *Nature* 346, 174-177.

Sanes, J.N., Suner, S., and Donoghue, J.P. (1990). Dynamic Organization of Primary Motor Cortex Output to Target Muscles in Adult-Rats .1. Long-Term Patterns of Reorganization Following Motor or Mixed Peripheral-Nerve Lesions. *Experimental Brain Research* 79, 479-491.

Sanes, J.N., Suner, S., Lando, J.F., and Donoghue, J.P. (1988). Rapid Reorganization of Adult-Rat Motor Cortex Somatic Representation Patterns after Motor-Nerve Injury. *Proceedings of the National Academy of Sciences of the United States of America* 85, 2003-2007.

Sceniak, M.P., Hawken, M.J., and Shapley, R. (2001). Visual spatial characterization of macaque V1 neurons. *Journal of Neurophysiology* 85, 1873-1887.

Schall, J.D., Hanes, D.P., Thompson, K.G., and King, D.J. (1995). Saccade Target Selection in Frontal Eye Field of Macaque .1. Visual and Premovement Activation. *Journal of Neuroscience* 15, 6905-6918.

Schein, S.J., and Desimone, R. (1990). Spectral Properties of V4 Neurons in the Macaque. *Journal of Neuroscience* 10, 3369-3389.

Schiller, P.H., and Tehovnik, E.J. (2005). Neural mechanisms underlying target selection with saccadic eye movements. *Progress in Brain Research* 149, 157-171.

Schlack, A., and Albright, T.D. (2007). Remembering visual motion: Neural correlates of associative plasticity and motion recall in cortical area MT. *Neuron* 53, 881-890.

Schmid, L.M., Rosa, M.G.P., and Calford, M.B. (1995). Retinal-Detachment Induces Massive Immediate Reorganization in Visual-Cortex. *Neuroreport* 6, 1349-1353.

Schmidt, K.E., Goebel, R., Lowel, S., and Singer, W. (1997). The perceptual grouping criterion of colinearity is reflected by anisotropies of connections in the primary visual cortex. *European Journal of Neuroscience* 9, 1083-1089.

Schuchard, R.A. (1993). Validity and interpretation of Amsler grid reports. *Archives of Ophthalmology* 111, 776-780.

Schuchard, R.A. (1995). Adaptation to macular scotomas in persons with low vision. *American Journal of Occupational Therapy* 49, 870-876.

Schwartz, E.L. (1980). Computational Anatomy and Functional Architecture of Striate Cortex - a Spatial-Mapping Approach to Perceptual Coding. *Vision Research* 20, 645-669.

Sclar, G., and Freeman, R.D. (1982). Orientation Selectivity in the Cats Striate Cortex Is Invariant with Stimulus Contrast. *Experimental Brain Research* 46, 457-461.

Sha'ashua, A., and Ullman, S. (1988). Structural Saliency: The Detection of Globally Salient Structures Using a Locally Connected Network. In *Second International Conference on Computer Vision* (Tarpon Springs, Florida, The IEEE Computer Society Press), pp. 321-327.

Shipp, S. (2003). The functional logic of cortico-pulvinar connections. *Philosophical Transactions of the Royal Society B-Biological Sciences* 358, 1605-1624.

Shmuel, A., Korman, M., Sterkin, A., Harel, M., Ullman, S., Malach, R., and Grinvald, A. (2005). Retinotopic axis specificity and selective clustering of feedback projections from V2 to V1 in the owl monkey. *Journal of Neuroscience* 25, 2117-2131.

Shulman, G.L., Ollinger, J.M., Akbudak, E., Conturo, T.E., Snyder, A.Z., Petersen, S.E., and Corbetta, M. (1999). Areas involved in encoding and

applying directional expectations to moving objects. *Journal of Neuroscience* 19, 9480-9496.

Sigman, M., Cecchi, G.A., Gilbert, C.D., and Magnasco, M.O. (2001). On a common circle: Natural scenes and Gestalt rules. *Proceedings of the National Academy of Sciences of the United States of America* 98, 1935-1940.

Sillito, A.M., Grieve, K.L., Jones, H.E., Cudeiro, J., and Davis, J. (1995). Visual Cortical Mechanisms Detecting Focal Orientation Discontinuities. *Nature* 378, 492-496.

Simoncelli, E.P., Pillow, J.W., Paninski, L., and Schwartz, O. (2004). Characterization of neural response with stochastic stimuli. In *The Cognitive Neurosciences*, M. Gazzaniga, ed. (Cambridge, MA: MIT Press), pp. 327-338.

Sincich, L.C., and Blasdel, G.G. (2001). Oriented axon projections in primary visual cortex of the monkey. *Journal of Neuroscience* 21, 4416-4426.

Smirnakis, S.M., Brewer, A.A., Schmid, M.C., Tolias, A.S., Schuz, A., Augath, M., Inhoffen, W., Wandell, B.A., and Logothetis, N.K. (2005). Lack of long-term cortical reorganization after macaque retinal lesions. *Nature* 435, 300-307.

Smith, M.A., Bair, W., and Movshon, J.A. (2006). Dynamics of suppression in macaque primary visual cortex. *Journal of Neuroscience* 26, 4826-4834.

Spitzer, H., Desimone, R., and Moran, J. (1988). Increased Attention Enhances Both Behavioral and Neuronal Performance. *Science* 240, 338-340.

Stepanyants, A., Martinez, L.M., Ferecsko, A.S., and Kisvarday, Z.F. (2009). The fractions of short- and long-range connections in the visual cortex.

Proceedings of the National Academy of Sciences of the United States of America *106*, 3555-3560.

Stettler, D.D., Das, A., Bennett, J., and Gilbert, C.D. (2002). Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron* *36*, 739-750.

Sugita, Y. (1999). Grouping of image fragments in primary visual cortex. *Nature* *401*, 269-272.

Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* *304*, 1782-1787.

Summerfield, C., and Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences* *13*, 403-409.

Suner, S., Fellows, M.R., Vargas-Irwin, C., Nakata, G.K., and Donoghue, J.P. (2005). Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *Ieee Transactions on Neural Systems and Rehabilitation Engineering* *13*, 524-541.

Super, H., Spekreijse, H., and Lamme, V.A.F. (2001a). A neural correlate of working memory in the monkey primary visual cortex. *Science* *293*, 120-124.

Super, H., Spekreijse, H., and Lamme, V.A.F. (2001b). Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nature Neuroscience* *4*, 304-310.

T'so, D.Y., Frostig, R.D., Lieke, E.E., and Grinvald, A. (1990). Functional-Organization of Primate Visual-Cortex Revealed by High-Resolution Optical Imaging. *Science* *249*, 417-420.

T'so, D.Y., and Gilbert, C.D. (1988). The Organization of Chromatic and Spatial Interactions in the Primate Striate Cortex. *Journal of Neuroscience* 8, 1712-1727.

Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience* 19, 109-139.

Thompson, K.G., and Bichot, N.P. (2005). A visual salience map in the primate frontal eye field. *Progress in Brain Research* 147, 251-262.

Thompson, K.G., Bichot, N.P., and Sato, T.R. (2005a). Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience. *Journal of Neurophysiology* 93, 337-351.

Thompson, K.G., Bichot, N.P., and Schall, J.D. (1997). Dissociation of visual discrimination from saccade programming in macaque frontal eye field. *Journal of Neurophysiology* 77, 1046-1050.

Thompson, K.G., Biscoe, K.L., and Sato, T.R. (2005b). Neuronal basis of covert spatial attention in the frontal eye field. *Journal of Neuroscience* 25, 9479-9487.

Thompson, K.G., Hanes, D.P., Bichot, N.P., and Schall, J.D. (1996). Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *Journal of Neurophysiology* 76, 4040-4055.

Thorell, L.G., Devalois, R.L., and Albrecht, D.G. (1984). Spatial-Mapping of Monkey V1-Cells with Pure Color and Luminance Stimuli. *Vision Research* 24, 751-769.

Tian, L., Hires, S.A., Mao, T., Huber, D., Chiappe, M.E., Chalasani, S.H., Petreanu, L., Akerboom, J., McKinney, S.A., Schreier, E.R., *et al.* (2009). Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nature Methods* 6, 875-U113.

Tomita, H., Ohbayashi, M., Nakahara, K., Hasegawa, I., and Miyashita, Y. (1999). Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature* 401, 699-703.

Toth, L.J., and Assad, J.A. (2002). Dynamic coding of behaviourally relevant stimuli in parietal cortex. *Nature* 415, 165-168.

Touryan, J., Felsen, G., and Dan, Y. (2005). Spatial structure of complex cell receptive fields measured with natural images. *Neuron* 45, 781-791.

Treisman, A., and Gormican, S. (1988). Feature Analysis in Early Vision - Evidence from Search Asymmetries. *Psychological Review* 95, 15-48.

Treisman, A.M., and Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology* 12, 97-136.

Treue, S., and Maunsell, J.H.R. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382, 539-541.

Treue, S., and Trujillo, J.C.M. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399, 575-579.

Tsunoda, K., Yamane, Y., Nishizaki, M., and Tanifuji, M. (2001). Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. *Nature Neuroscience* 4, 832-838.

Ullman, S. (2007). Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Sciences* 11, 58-64.

Vallar, G., and Perani, D. (1986). The Anatomy of Unilateral Neglect after Right-Hemisphere Stroke Lesions - a Clinical Ct-Scan Correlation Study in Man. *Neuropsychologia* 24, 609-622.

Vidyasagar, T.R., and Pigarev, I.N. (2007). Modulation of neuronal responses in macaque primary visual cortex in a memory task. *European Journal of Neuroscience* 25, 2547-2557.

Volchan, E., and Gilbert, C.D. (1994). Interocular Transfer of Receptive-Field Expansion in Cat Visual-Cortex. *Vision Research* 35, 1-6.

von der Heydt, R., Peterhans, E., and Baumgartner, G. (1984). Illusory Contours and Cortical Neuron Responses. *Science* 224, 1260-1262.

Wachtler, T., Sejnowski, T.J., and Albright, T.D. (2003). Representation of color stimuli in awake macaque primary visual cortex. *Neuron* 37, 681-691.

Walker, G.A., Ohzawa, I., and Freeman, R.D. (1999). Asymmetric suppression outside the classical receptive field of the visual cortex. *Journal of Neuroscience* 19, 10536-10553.

Wang, G., Tanaka, K., and Tanifuji, M. (1996). Optical imaging of functional organization in the monkey inferotemporal cortex. *Science* 272, 1665-1668.

Webster, M.A., and De Valois, R.L. (1985). Relationship between Spatial-Frequency and Orientation Tuning of Striate-Cortex Cells. *Journal of the Optical Society of America a-Optics Image Science and Vision* 2, 1124-1132.

Webster, M.J., Bachevalier, J., and Ungerleider, L.G. (1994). Connections of Inferior Temporal Areas Teo and Te with Parietal and Frontal-Cortex in Macaque Monkeys. *Cerebral Cortex* 4, 470-483.

Wegener, D., Ehn, F., Aurich, M.K., Galashan, F.O., and Kreiter, A.K. (2008). Feature-based attention and the suppression of non-relevant object features. *Vision Research* 48, 2696-2707.

Welch, K., and Stuteville, P. (1958). Experimental Production of Unilateral Neglect in Monkeys. *Brain* 81, 341-347.

Weliky, M., Bosking, W.H., and Fitzpatrick, D. (1996). A systematic map of direction preference in primary visual cortex. *Nature* 379, 725-728.

Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt. *Psychologische Forschung* 4, 301-350.

White, L.E., Bosking, W.H., and Fitzpatrick, D. (2001). Consistent mapping of orientation preference across irregular functional domains in ferret visual cortex. *Visual Neuroscience* 18, 65-76.

Wickersham, I.R., Lyon, D.C., Barnard, R.J.O., Mori, T., Finke, S., Conzelmann, K.K., Young, J.A.T., and Callaway, E.M. (2007). Monosynaptic restriction of transsynaptic tracing from single, genetically targeted neurons. *Neuron* 53, 639-647.

Williams, L.R., and Jacobs, D.W. (1997). Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation* 9, 837-858.

Xiao, D.K., Marcar, V.L., Raiguel, S.E., and Orban, G.A. (1997). Selectivity of macaque MT/V5 neurons for surface orientation in depth specified by motion. *European Journal of Neuroscience* 9, 956-964.

Yamahachi, H., Marik, S.A., McManus, J.N.J., Denk, W., and Gilbert, C.D. (2009). Rapid Axonal Sprouting and Pruning Accompany Functional Reorganization in Primary Visual Cortex. *Neuron* 64, 719-729.

Yang, Z.Y., Heeger, D.J., and Seidemann, E. (2007). Rapid and precise retinotopic mapping of the visual cortex obtained by voltage-sensitive dye imaging in the behaving monkey. *Journal of Neurophysiology* 98, 1002-1014.

Yen, S.C., Baker, J., and Gray, C.M. (2007). Heterogeneity in the responses of adjacent neurons to natural stimuli in cat striate cortex. *Journal of Neurophysiology* 97, 1326-1341.

Yoshimura, Y., Sato, H., Imamura, K., and Watanabe, Y. (2000). Properties of horizontal and vertical inputs to pyramidal cells in the superficial layers of the cat visual cortex. *Journal of Neuroscience* 20, 1931-1940.

Yoshioka, T., Blasdel, G.G., Levitt, J.B., and Lund, J.S. (1996). Relation between patterns of intrinsic lateral connectivity, ocular dominance, and cytochrome oxidase-reactive regions in macaque monkey striate cortex. *Cerebral Cortex* 6, 297-310.

Zeki, S. (1983a). Color Coding in the Cerebral-Cortex - the Reaction of Cells in Monkey Visual-Cortex to Wavelengths and Colors. *Neuroscience* 9, 741-765.

Zeki, S. (1983b). Color Coding in the Cerebral-Cortex - the Responses of Wavelength-Selective and Color-Coded Cells in Monkey Visual-Cortex to Changes in Wavelength Composition. *Neuroscience* 9, 767-&.

Zhang, F., Aravanis, A.M., Adamantidis, A., de Lecea, L., and Deisseroth, K. (2007). Circuit-breakers: optical technologies for probing neural signals and systems. *Nature Reviews Neuroscience* 8, 577-581.

Zhang, F., Gradinaru, V., Adamantidis, A.R., Durand, R., Airan, R.D., de Lecea, L., and Deisseroth, K. (2010). Optogenetic interrogation of neural circuits: technology for probing mammalian brain structures. *Nature Protocols* 5, 439-456.

Zhou, H., Friedman, H.S., and von der Heydt, R. (2000). Coding of border ownership in monkey visual cortex. *Journal of Neuroscience* 20, 6594-6611.

Zipser, K., Lamme, V.A.F., and Schiller, P.H. (1996). Contextual modulation in primary visual cortex. *Journal of Neuroscience* 16, 7376-7389.

Zur, D., and Ullman, S. (2003). Filling-in of retinal scotomas. *Vision Research* 43, 971-982.