

UCSF

UC San Francisco Electronic Theses and Dissertations

Title

Approaches to the function of object recognition areas of the visual cortex

Permalink

<https://escholarship.org/uc/item/5wv361bv>

Author

Caywood, Matthew Shields

Publication Date

2009

Peer reviewed|Thesis/dissertation

Approaches to the function of object recognition areas of the visual cortex

by

Matthew Shields Caywood

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Neuroscience

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

Dedication and Acknowledgements

I would like to thank my research advisor Michael Stryker, who made time to participate in many of these experiments, and whose input during the writing process was invaluable.

Michael taught me a great deal about doing neurophysiology, as did David Blake. I also learned from Hiroki Sugihara, Jianhua “J.C.” Cang, and Laila Dadvand.

Many of these experiments would have been impossible without the able surgical and veterinary assistance of Karen MacLeod. I received essential experimental support from Sheri Harris, and computer support from Sergei “C.P.” Rebrik.

Several people read and commented on drafts, including Melinda Owens, Cris Niell, and Marnie Riddle; thanks to them this dissertation is much better, and is less comma-laden.

Most of all I would like to thank my parents, Clarke and Mary Caywood, for all the opportunities they gave me.

Chapter 1 of this thesis is a reprint of the material as it appears in the Journal of Neurophysiology. This study was done during my studies at UCSF, in collaboration with Ben Willmore and David Tolhurst, colleagues from my former laboratory at the University of Cambridge, England, and was published in the Journal of Neurophysiology in 2004. Ben generated the model receptive fields from color photographic data using physiologically realistic preprocessing and independent component analysis, and made a number of significant intellectual contributions in project discussions. David supervised Ben's research and helped to revise and improve the paper.

Chapter 2 was submitted to a refereed conference, but has not been published. The computational topology methods used in this study were introduced by my collaborator Andrew Blumberg, a mathematician then at Stanford University, who wrote the section describing them and oversaw their execution and the statistical analysis of the results.

Approaches to the function of object recognition areas of the visual cortex

Matthew S. Caywood

Abstract

This dissertation describes recent theoretical and experimental efforts to understand the areas of the visual cortex and the neural coding underlying our ability to see and recognize objects. The first chapter is theoretical, describing a computer model of color processing in macaque primary visual cortex (V1). This previously published study shows that the apparent complexity of V1 neurons' responses to color grating stimuli may arise from the statistics of the natural visual world in combination with a learning rule under which neurons reduce the redundancy of their responses to natural stimuli. The second chapter is a hybrid of theoretical work and preliminary experimental data. We ask how one could begin to study the function of higher visual areas for which simple stimuli are insufficient. We suggest a "bottom-up" approach that starts with the visual area below the area of interest, and uses knowledge about the lower area to define a basis set of stimuli for studying the higher area. Since neurons in higher visual areas respond invariantly to stimuli, we introduce a computational topology method that allows us to rigorously study known invariances and identify unanticipated invariances. The third chapter is a significant new experimental contribution, a method for defining and describing the functional maps in individual visual areas of the cortex. We adapt the continuous intrinsic signal imaging technique of Kalatsky and Stryker to characterize multiple functional maps, including ocular dominance, direction, spatial frequency, and temporal frequency, in primary visual areas 17 and 18 and extrastriate visual area 21a of the cat. We show that the functional architecture of area 21a is distinct from earlier visual areas, and that its function is closely linked with other ventral stream visual areas involved in object recognition. Finally, we provide a systematic classification, inspired by machine learning theory, of attempts to study the selectivity of neurons in object recognition areas of the visual brain, and discuss what the functional architecture of a visual area implies about its neural responses.

Table of Contents

Introduction.....	1
Chapter 1. Independent components of color natural scenes resemble V1 neurons in their spatial and color tuning	6
Chapter 2. Detecting topological invariances of visual neuron receptive fields with persistent homology	22
Chapter 3. Functional maps of cat visual area 21a reveal a role in processing spatial form	44
Conclusion: Experimental approaches to the visual ventral stream	98

Introduction

The cortex is a several-millimeter-thick sheet of cells on the convoluted outer surface of the brain, responsible for complex human visual behaviors such as reading, recognizing objects and faces, and higher cognitive functions such as language use, planning and the processing of abstract information. Much of what is happening in your mind as you read this thesis takes place somewhere in your cortex. During the evolution of modern humans, the expansion of the cortex is believed to be responsible for increased intelligence and complex behavior (Marino, 2002; Sherwood et al., 2008).

In the quest to understand intelligence and the human brain, neuroscientists have concentrated their research efforts on specific areas of the cortex. Many of these areas represent input from the senses, since it is easy to change sensory input experimentally and study the correlates in the brain. For example, an experimenter can present simple video images to an animal while simultaneously monitoring the visual cortex. Another advantage of studying the visual cortex is its relatively simple arrangement into streams of connected visual areas – visual information from the eye goes first to a single area, the primary visual cortex (known as V1), then sweeps out along multiple streams (Ungerleider & Pasternak, 2004). This relatively simple hierarchy, which may pass information unidirectionally in some visual recognition tasks (Riesenhuber & Poggio, 2004) is likely to be easier to decipher than other brain circuits connected in recurrent loops.

The function of V1 is relatively well understood. It has been known for nearly 50 years that its neurons fire tiny electrical spikes tens of times a second in response to an oriented edge in a particular location in visual space – for example, when you view a pencil held vertically at arm’s length, hundreds or thousands of cells in V1 will fire (Hubel & Wiesel, 2004).

These cells are not distributed randomly throughout V1, but are organized into clusters called columns. Columns extend through the full thickness of the cortex, and are about 0.5 mm wide; cells within a column share similar response properties and stimulus preferences. Columns define the functional map of the area – for any visual stimulus, this map will predict where to find cells that will respond to it. Since mammals’ brains do not develop identically, no two animals will have identical maps, but certain organizing principles always hold (Hübener & Bonhoeffer, 2002). These principles include retinotopic organization (considering again the example of a pencil, its location in the two-dimensional image seen by the eyes directly corresponds to a location in the two-dimensional sheet of V1), orientation (pencil vertical or horizontal), spatial frequency (width of the pencil), direction (which way it is moving).

Although V1 has been thoroughly studied, vast areas of the visual cortex remain poorly understood since elementary stimuli such as the pencil evoke a much weaker response than they do in V1. There is not even agreement on the number of visual areas in the cortex; humans have 25-40; monkeys have fewer, and cats fewer still (Van Essen, 2004; Scannell et al., 1999). Thus, one open question is how to demarcate areas, by

identifying distinguishing functional or structural properties. Another is identifying areas' specific function: what unique role does a given area play? A final question is functional architecture: how are the neurons within an area organized, and what does this organization tell us about the function of the cortex?

This dissertation describes recent theoretical and experimental efforts to understand the neural coding underlying our ability to see and recognize objects. It is divided into three stand-alone papers and a concluding discussion of previous studies of object recognition with an exposition of this work's contribution to the field.

The first paper is theoretical, describing a computer model of the processing of color in V1 in the primate (macaque monkey). In order to understand neurons' complex and diverse visual responses, and the complex functional architecture of the cortex, it is important to identify whether simple developmental rules and statistical regularities in the natural world can interact to produce complexity. This study shows that the apparent complexity of V1 neurons' responses to color stimuli may arise from fairly simple causes: post-natal exposure to the natural visual world in combination with a rule under which neurons learn to respond differentially to natural stimuli.

The second paper is a hybrid of theoretical work and some preliminary experimental data. It asks how we could begin to study the function of higher visual areas for which simple stimuli are insufficient. We suggest a "bottom-up" approach that starts with the visual area below the area of interest, and uses knowledge about the lower area to define stimuli for studying the higher area. Since neurons in higher visual areas often

respond invariantly to stimuli (i.e. they are not sensitive to exact stimulus positioning), we propose a method of rigorously studying these invariances and identifying new ones. Although the experimental portion of this study is preliminary, the theoretical work points toward a possible future of neuroscience experiments that can be performed semi-automatically by computers, freeing experimenters to theorize more abstractly.

The final paper is a significant new experimental contribution, explaining a method for defining and describing the maps in individual visual areas of the cortex. I use the technique of intrinsic signal imaging to observe, through a high-speed camera, the metabolic response that correlates with local neuronal activity in the cortex. I characterize multiple visual areas of the cat simultaneously, specifically demarcating and characterizing area 21a, an intermediate visual area with previously uncertain boundaries; we show that its properties are distinct from earlier visual areas, and that its function is closely linked with other visual areas that recognize objects. We find that cortical maps in individual cats can be rapidly characterized, demonstrating that even in the anesthetized cat, optical imaging techniques can be used to study intermediate visual areas.

References

- Hubel, D. H., & Wiesel, T. N. (2004). *Brain and visual perception: The story of a 25-year collaboration*. New York: Oxford University Press.
- Hübener, M. and Bonhoeffer, T. (2002). Optical imaging of functional architecture in cat primary visual cortex. In B. R. Payne, & A. Peters (Eds.), *The cat primary visual cortex* (pp. 131-165). San Diego: Academic Press.
- Marino, L. (2002). Convergence of complex cognitive abilities in cetaceans and primates. *Brain, Behavior and Evolution*, 59(1-2), 21-32.
- Riesenhuber, M., & Poggio, T. (2004). How the visual cortex recognizes objects: The tale of the standard model. In L. M. Chalupa, & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1640-1653). Cambridge, MA: MIT Press.
- Scannell, J. W., Burns, G. A., Hilgetag, C. C., O'Neil, M. A., & Young, M. P. (1999). The connectional organization of the cortico-thalamic system of the cat. *Cerebral Cortex*, 9(3), 277-299.
- Sherwood, C. C., Subiaul, F., & Zawidzki, T. W. (2008). A natural history of the human mind: Tracing evolutionary changes in brain and cognition. *Journal of Anatomy*, 212(4), 426-454.
- Ungerleider, L. G., & Pasternak, T. (2004). Ventral and dorsal cortical processing circuits. In L. M. Chalupa, & J. S. Werner (Eds.), *The visual neurosciences* (pp. 541-562). Cambridge, MA: MIT Press.
- Van Essen, D. C. (2004). Organization of visual areas in macaque and human cerebral cortex. In L. M. Chalupa, & J. S. Werner (Eds.), *The visual neurosciences* (pp. 507-521). Cambridge, MA: MIT Press.

Independent Components of Color Natural Scenes Resemble V1 Neurons in Their Spatial and Color Tuning

Matthew S. Caywood,¹ Benjamin Willmore,² and David J. Tolhurst³

¹Program in Neuroscience, University of California, San Francisco, California 94143–0444; ²Department of Psychology, University of California, Berkeley, California 94720–1650; and ³Department of Physiology, University of Cambridge, Cambridge CB2 3EG, United Kingdom

Submitted 8 August 2003; accepted in final form 21 January 2004

Caywood, Matthew S., Benjamin Willmore, and David J. Tolhurst. Independent components of color natural scenes resemble V1 neurons in their spatial and color tuning. *J Neurophysiol* 91: 2859–2873, 2004. First published January 28, 2004; 10.1152/jn.00775.2003. It has been hypothesized that mammalian sensory systems are efficient because they reduce the redundancy of natural sensory input. If correct, this theory could unify our understanding of sensory coding; here, we test its predictions for color coding in the primate primary visual cortex (V1). We apply independent component analysis (ICA) to simulated cone responses to natural scenes, obtaining a set of colored independent component (IC) filters that form a redundancy-reducing visual code. We compare IC filters with physiologically measured V1 neurons, and find great spatial similarity between IC filters and V1 simple cells. On cursory inspection, there is little chromatic similarity; however, we find that many apparent differences result from biases in the physiological measurements and ICA analysis. After correcting these biases, we find that the chromatic tuning of IC filters does indeed resemble the population of V1 neurons, supporting the redundancy-reduction hypothesis.

INTRODUCTION

The mammalian visual system is believed to efficiently encode natural visual information. One way in which it might do this is by reducing the redundancy of the representation at successive stages of processing (Attneave 1954; Barlow 1959). Natural visual input contains features such as edges and homogenous color patches, which make the patterns of light falling on the retina highly redundant and which give rise to statistical dependencies between neighboring regions of the visual image (Field 1987). To reduce redundancy, the visual system might use these features as a basis for representing visual input (Barlow 1989).

Independent component analysis (ICA) (Comon 1994) is a widely used method for finding a redundancy-reducing encoding of data (such as natural visual scenes). Although the resulting independent components (ICs) are often only approximately independent, when ICA is applied to achromatic natural images, it produces ICs that are strikingly similar to the achromatic spatial receptive fields (RFs) of simple cells in primary visual cortex (V1) (Bell and Sejnowski 1997). Sparseness maximization (Olshausen and Field 1996) is conceptually similar to ICA and produces similar results; Ringach (2002) has argued that it better models the diversity of spatial RFs in V1. However, ICA

has been validated by many rigorous comparisons with V1: van Hateren and van der Schaaf (1998) quantified the comparison between achromatic ICs and macaque V1 simple cells and found that the distributions of all RF measurements except optimal spatial frequency (SF) were well matched. When a temporal dimension is added and simple cells are compared with the IC filters of achromatic natural movies, the SF similarity improves (van Hateren and Ruderman 1998).

Because the ICA model successfully explains the spatial tuning of V1 simple cells to achromatic stimuli, the model may also explain other response properties such as color tuning. To test this hypothesis, ICA has been applied to colored natural scenes to produce spatiochromatic ICs (Doi et al. 2003; Hoyer and Hyvärinen 2000; Taylor et al. 2000; Wachtler et al. 2001). However, these studies used varying methodologies and found conflicting sets of ICs. Additionally, none of the studies has quantitatively compared ICs to a standard set of physiological measurements of V1. Finally, a richer understanding of V1 color coding has recently been developing: the view that color sensitivity is infrequent in V1 and restricted to weakly orientation-tuned neurons in cytochrome oxidase blobs (Livingstone and Hubel 1984; Lennie et al. 1990) has given way to a more diverse picture, in which oriented cells also have a rich variety of color sensitivities (Conway 2001; Johnson et al. 2001).

In this study, our goal was to rigorously test the hypothesis that ICA can account for the chromatic tuning of neurons in V1. We have controlled for the methodological variations between previous color ICA studies and have chosen the most biologically realistic set of ICs for analysis. Treating the ICs as though they were real neuronal RFs, we use the cone-opponent grating stimuli of Derrington et al. (1984) to determine their color tuning. We then directly, quantitatively compare the color tuning with the Lennie et al. (1990) classic V1 data. To further test the hypothesis, we compare the structure of ICs' model cone inputs with recent V1 data from Conway (2001) and Johnson et al. (2001). At first sight, the color tuning of ICs of natural scenes does not look very similar to V1 neurons, but we will show that, in fact, there is considerable similarity, once appropriate ICA methodology is employed, and the limitations of real experimental protocols are considered.

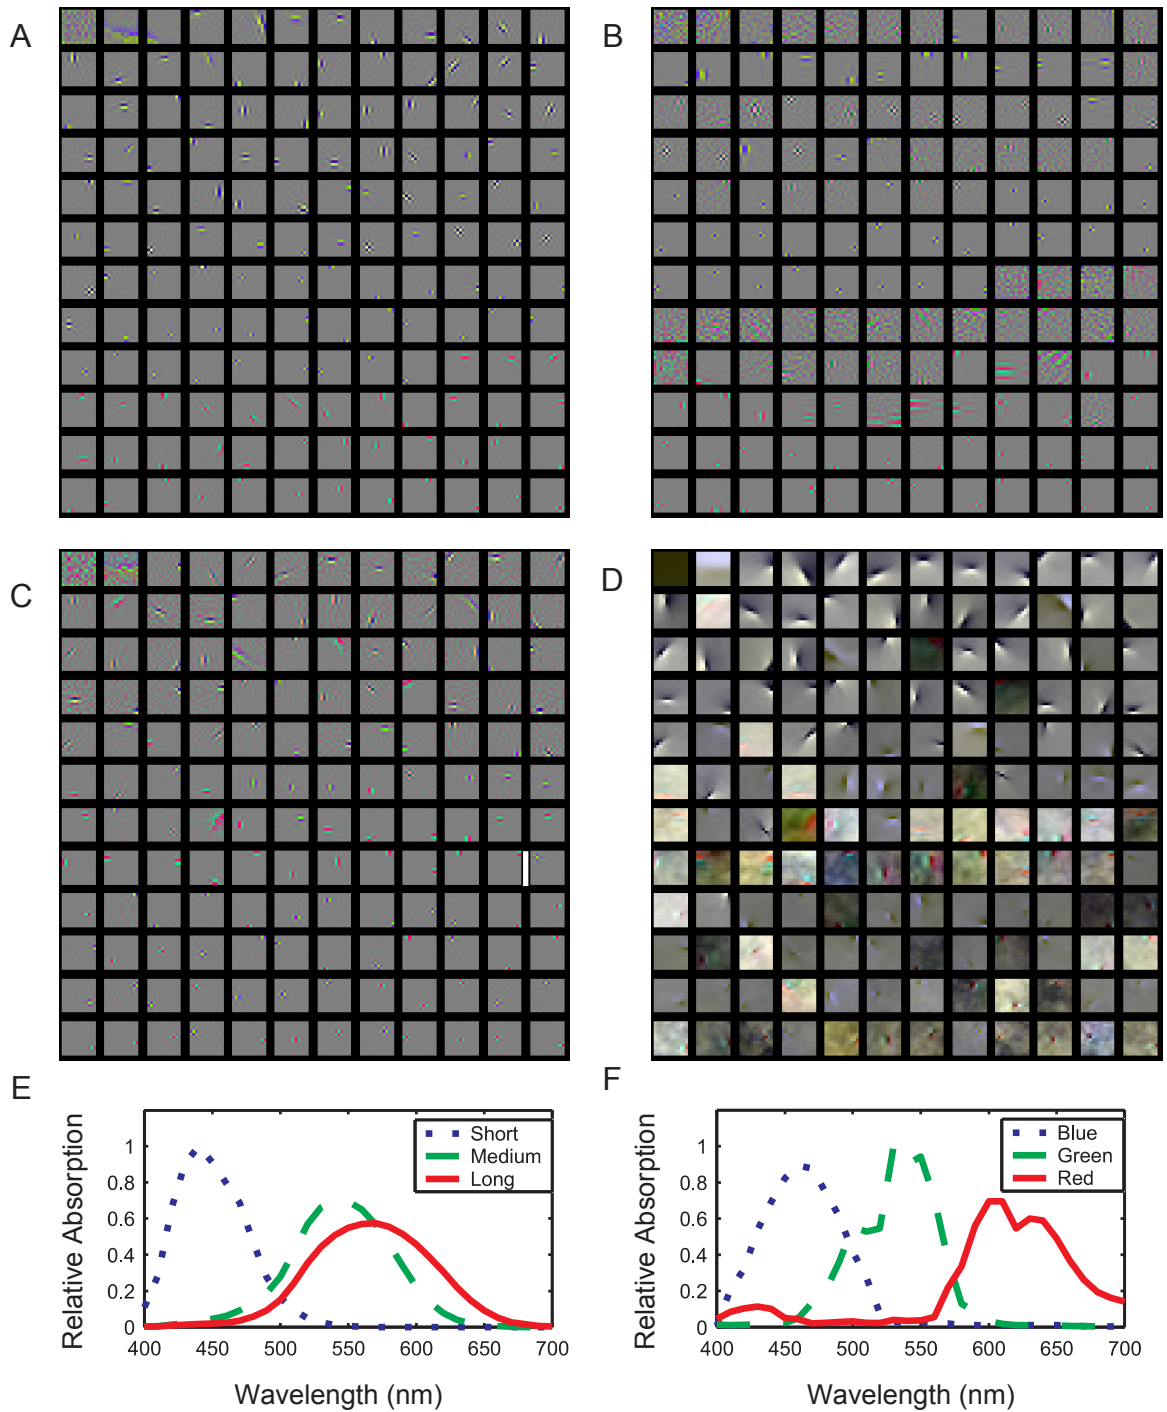


FIG. 1. (See text for full legend.) A and B: IC filters generated from RGB input (A) and from JPEG-encoded RGB input (B). C: IC filters generated from images encoded using human LMS cone sensitivities, shown in pseudocolor so that L is red, M green, and S blue. D: LMS IC basis functions, matched to the filters in C. E: absorption spectra of human L, M and S cones. F: R, G and B absorption spectra of a typical digital camera.

METHODS

Colored natural scenes

We analyzed a “hyperspectral” set of 25 distinct colored natural scenes (Párraga et al. 1998). Each 256×256 -pixel scene was photographed through 31 filters covering the human visible spectrum (centered at wavelengths 400, 410 . . . 700 nm, bandwidth: ~ 10 nm); thus it was not limited by the spectral sensitivity of a standard RGB (Red-Green-Blue) camera. Each plane was digitized to 12 bits of intensity. We aligned the color planes of each scene with subpixel accuracy by maximizing cross-correlation of adjacent planes; then, to ensure proper alignment, we averaged pixels together in 2×2 blocks.

We encoded the scenes in three ways that encompass previous work. Wachtler et al. (2001) transformed the Párraga et al. (1998) hyperspectral images into the human visual colorspace defined by the absorption spectra of L, M, and S cones. Hoyer and Hyvärinen (2000) used uncompressed images encoded in a digital camera’s RGB colorspace, while Tailor et al. (2000) used JPEG-compressed RGB images taken from the Internet.

In the LMS condition, we used the Smith-Pokorny cone-sensitivity curves (Fig. 1E) (Smith and Pokorny 1975) to transform the 31-plane hyperspectral scenes into 3-plane LMS space. These human psycho-

physical curves are consistent with, but more precise than, the physiological absorption spectra of macaque photoreceptors (Baylor et al. 1987). In the RGB condition, we converted the 31-plane hyperspectral scenes into 3-plane RGB images, using the sensitivity curves (Fig. 1F) of the red, green, and blue detectors of a typical digital camera (Nikon Coolpix 950; C.A. Párraga, unpublished observations). In the JPEG condition, we applied JPEG compression to these RGB-encoded images to mimic the Tailor et al. (2000) dataset. JPEG images were encoded at quality setting 90 [using Matlab’s (the MathWorks) IMWRITE function]. This is conservative compared with the quality setting of between 50 and 75, which we estimate (from the encoded file sizes) was used in the Tailor et al. (2000) image set. After performing ICA, we linearly transformed the resulting RGB ICs into the LMS colorspace for comparison with neurophysiological data.

We pseudorandomly extracted a set of 12×12 -pixel fragments from the scenes in each image set, excluding the calibration reference (a white card) from each scene. The number of fragments was 200 times the fragment dimensionality (i.e., 86,400 for 12×12 -pixel, 3-plane images).

In the LMS and RGB conditions, each image fragment was log-transformed. The log transform is commonly applied as ICA preprocessing (e.g., van Hateren and van der Schaaf 1998) because it

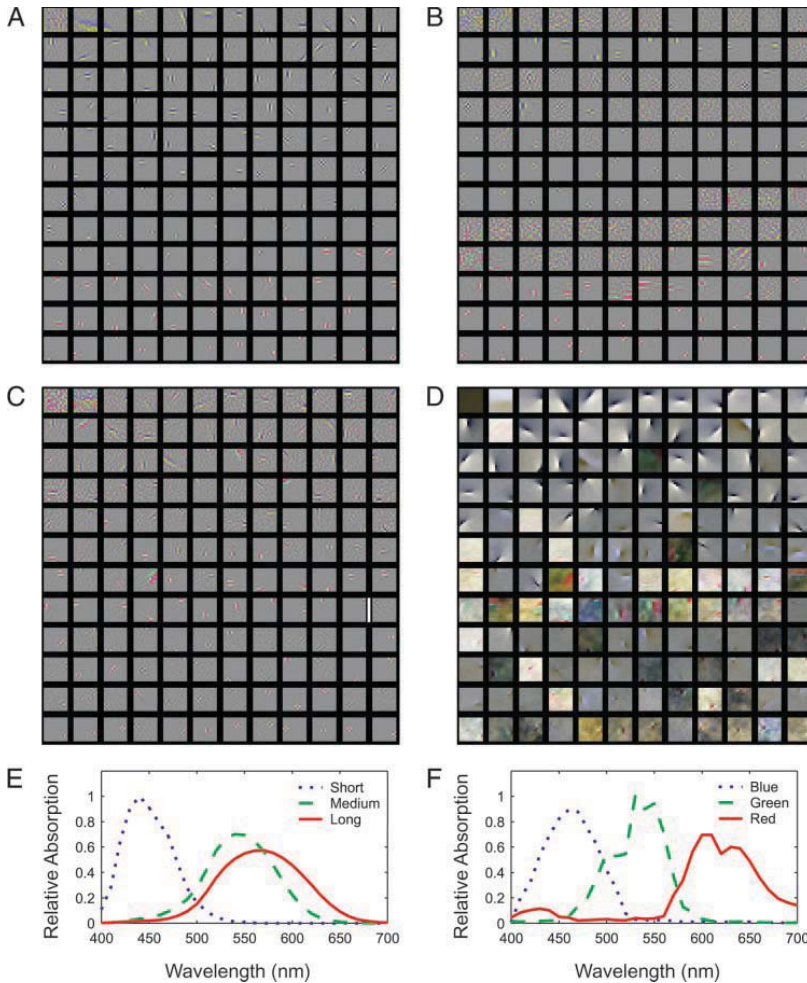


FIG. 1. Sample of independent components (IC) filters and IC basis functions generated by running IC analysis (ICA) on differently encoded input, and RGB (Red-Green-Blue) and LMS (Long-Medium-Short) absorption spectra used in encoding. ICs, which contain positive- and negative-valued pixels with arbitrary range, were normalized so that a value of 0 is displayed as medium intensity (0.5) and all pixels of an IC are between 0 and 1. *A* and *B*: IC filters generated from RGB input (*A*) and from JPEG-encoded RGB input (*B*). Filters were manually separated into approximate color opponency type (blue-yellow versus red-green). Within color groups, they are subdivided by spatial structure (full-field, Gabor-like, checkerboard (present in *B* only), and small noisy patches. Finally, within these groups, they are arranged in descending order by variance. *C*: IC filters generated from images encoded using human LMS cone sensitivities, shown in pseudocolor so that L is red, M green, and S blue. LMS filters are divided into 2 groups (short vertical white line): significant and artifactual, using the criteria of van Hateren and van der Schaaf (1998). Within groups, they are shown in decreasing order by variance. *D*: LMS IC basis functions, matched to the filters in *C*. *E*: absorption spectra of human L, M and S cones. *F*: R, G and B absorption spectra of a typical digital camera.

improves the convergence of ICA learning algorithms (see Willmore et al. 2000 for analysis of its effects). The raw luminance variation within our set of natural scenes, which covers three orders of magnitude, badly skews the data distribution and breaches the linear superposition model assumed by ICA (see next section). There are other remedies for this problem: while raw LMS images yield poorly converged ICs, standardizing the mean and variance of those images produces ICs similar to those of log-transformed images (data not shown). We chose to use the log transform, however, because it makes our results comparable with other studies.

From a functional point of view, the log transform has a clear biological correlate in retinal luminance adaptation (Field 1994; van Hateren and van der Schaaf 1998), a process that also compensates for the luminance changes (of many orders of magnitude) occurring in the natural environment. Our investigations of color tuning (in the following text) are conducted in the equivalent of laboratory conditions and always modulate cone contrast (at fixed mean luminance) rather than raw luminance. Therefore we did not apply log transformation to our “grating stimuli.”

In the JPEG condition, we did not apply the log transform, because our aim was to replicate the results of Taylor et al. (2000), where the transform was not used.

ICA: algorithm and preprocessing

Each image fragment r_j was converted into a one-dimensional vector, and these comprised the columns of the $432 \times 86,400$ data matrix R . The mean of each vector was set to zero. As a computational convenience for the ICA step, we whitened the data vectors; this transformation was inverted before the ICs were analyzed, so it has no effect on the final results. Our whitening method was based on principal component analysis, and produced a whitened data matrix

$$X = D^{-1/2}E^T R$$

where the columns of E are the eigenvectors of the covariance matrix RR^T (i.e., its principal components) and D is a diagonal matrix containing the corresponding eigenvalues.

The assumption behind the ICA algorithm is that each whitened image fragment x_j is composed of a weighted sum of a fixed set of underlying source vectors a_i , the activities of which are given by scalars s_i

$$x_j = \sum_{i=1}^n a_i s_i$$

This can be rewritten as a matrix equation, $x = As$, where A is the matrix of sources and s is the vector of activities.

These sources, also called basis functions, are identified as independent sources of variance of the image data, which produce image fragments by mixing together linearly. Along with the basis functions, one can also derive a set of IC filters that reverse the process: they “unmix” the image fragment to yield activity values s_i , which correspond to individual sources and which should also be maximally independent. The set of filters produced by ICA is directly analogous to a population of V1 neurons, each analyzing part of the incoming visual stimulus.

There are a number of ICA algorithms, which make slightly different assumptions about the distribution of the underlying sources, but for ICA of colored natural scenes most algorithms yield similar results (Wachtler et al. 2001). We used an implementation of the information-maximization ICA algorithm with natural gradient feature (Bell and Sejnowski 1997; Delorme and Makeig 2003). This algorithm uses gradient ascent to find vectors in the distribution that give us the IC filter matrix F , which is the inverse of the sources: $F = A^{-1}$.

IC filters were similar over multiple runs of the algorithm with

different random seeds and different sets of image fragments. As a control for artifacts due to misalignment of the spectral planes, we performed extra analyses where we subsampled the LMS images in blocks of 2×2 , 3×3 , and 4×4 before extracting fragments; at each block size, the character of filters was preserved, which shows that the color planes were sufficiently aligned to prevent artifacts.

Spatial tuning

To find the best spatial stimulus for each IC, we exploited the equivalence between the Fourier transform of an RF and its responses to grating stimuli. We took the two-dimensional discrete Fourier transform of each IC, identified the highest amplitude Fourier component, and computed its orientation and spatial frequency.

These measurements were limited in precision by the 12×12 patch size because only discrete values corresponding to integer harmonics are possible. This patch size enabled us to measure orientation with precision of $\sim 10^\circ$, and to measure SF with precision of 1 octave at low SFs (near 0.1 cycles/pixel) and < 0.3 octaves at high SFs (near 0.5 cycles/pixel).

We also assessed other spatial properties of each IC. Using the best achromatic stimulus, we computed orientation bandwidth, SF bandwidth, length, and width by measuring full width at half-maximum (FWHM) for band-pass functions or twice the half-width at half-maximum for low- and high-pass functions. SF bandwidth was computed as the log of the SF ratio between the half-heights in the amplitude spectrum along the radius corresponding to peak orientation following van Hateren and van der Schaaf (1998). Orientation tuning bandwidth was measured as the difference between half-heights along the circle corresponding to peak SF. For a small number of very high-SF oblique filters, some measurements fell outside the boundary of the Fourier transformed patch; in these cases, we took the measurement at the boundary, which caused a slight narrowing of their bandwidth. Aspect ratio was computed as the ratio of the length and width, where the length is the RF envelope FWHM in the low-pass direction and width is the envelope FWHM in the band-pass direction (Field and Tolhurst 1986).

Color tuning

Color tuning was quantified using the Lennie et al. (1990) cone-opponent stimuli in the “DKL” colorspace developed by Derrington et al. (1984) from the work of Macleod and Boynton (1979). In this three-dimensional colorspace, the xy plane is equiluminant, similar to the CIE colorspace, and so any changes along the x and y axes alter stimulus color without affecting luminance. Increasing x increases L-cone input while proportionally decreasing M-cone input (i.e., a stimulus becomes more “red” and less “green”), leaving S cones unaffected. Increasing y increases both L- and M-cone inputs while reducing S-cone input (resulting in more “yellow” and less “blue”). Increasing z increases input to all three cones, i.e., increases overall luminance. Lennie et al.’s white point was set to cone input values (0.311, 0.336, 0.353). Modulation from this point along the x axis to (1, 0, 0) corresponds to LMS cone input changes of (0.074, -0.14 , 0), and modulation along the y axis to (0, 1, 0) corresponds to changes of (0, 0, -0.84). Modulation along the z axis causes a change in luminance without a change in color so that modulation to (0, 0, 1) doubles all three cones’ inputs. Because of the overlap of L and M cone spectra (Fig. 1E), the effect of X axis modulation on L and M cones is small (compared with the effect of y axis modulation on S cones) so we followed Lennie et al. in scaling the x coordinate by 3.125.

In the DKL colorspace, as defined in Lennie et al. (1990), a linearly summing neuron’s color tuning can be described using only two parameters. To do this, we treat each set of coordinates (x, y, z) as a vector and describe it in spherical coordinates (θ, ϕ, ρ) . The azimuth θ is the preferred color within the equiluminant xy plane: A neuron which is purely sensitive to R-G variation lies along the X axis and has

an azimuth of 0° , whereas a neuron purely sensitive to B-Y variation lies along the y axis and has an azimuth of 90° . The elevation θ is the angle the vector makes with the xy plane, and it describes the neuron's response to luminance change relative to chromatic change. A neuron with elevation 0° responds best to chromatic stimuli, whereas a neuron with elevation 90° responds best to luminance stimuli. The stimuli are cone-opponent drifting gratings that are spatially modulated in color and luminance along the line defined by (ϕ, θ) and crossing the sphere defined by fixed ρ . Thus a cell's color tuning can be fully described using only azimuth and elevations in the upper hemisphere ($\theta > 0^\circ$).

We computed the color tuning of each IC by finding the elevation and azimuth values that produce the maximal response, at a resolution of 1° . We converted each (ϕ, θ) pair into (x, y, z) coordinates and then converted these coordinates into cone-specific luminance values (L, M, S). To measure an IC's response to a color grating, we multiplied each cone luminance by the IC's Fourier component for that cone, summed the components over cones, and then took the amplitude.

Complementary to this cone-opponency analysis, we measured the cone-isolating grating responses of single cone inputs. In each color plane, we found the highest amplitude Fourier component and measured its phase, orientation, and SF.

Comparing color tuning distributions

The color tuning of each IC was thus represented by a single point in the DKL color space, $q(\phi, \theta)$. The distribution of all 432 ICs, $Q(\phi, \theta)$, was then compared with the distribution $P(\phi, \theta)$ of the 96 neurons described by Lennie et al. (1990). To compare two DKL color tuning distributions $P(\phi, \theta)$ and $Q(\phi, \theta)$, we needed a metric valid for arbitrary two-dimensional distributions. We used the Kullback-Leibler (K-L) distance from P to Q , which is also known as the relative entropy of P with respect to Q

$$D(P\|Q) = \sum_{\phi, \theta} P(\phi, \theta) \log \frac{P(\phi, \theta)}{Q(\phi, \theta)}$$

K-L distance is not symmetric, that is, in general $D(P\|Q) \neq D(Q\|P)$. However, when comparing physiological distributions $P(\phi, \theta)$ to IC distributions $Q(\phi, \theta)$, we avoid this issue by consistently using it in one direction: $D(P\|Q)$.

We binned the color tuning distribution into $30 \times 30^\circ$ bins, centered at azimuths $-30, 0, \dots, 150^\circ$ and elevations $15, 45, \text{ and } 75^\circ$. Some of the distributions had a small number n_0 of zero bins (oriented cells had 7, IC filters had between 1 and 4 depending on condition); to avoid singularity, these bins must be assigned a count. For each distribution, we wanted to calculate a zero-bin count, which was as conservative (high) as possible. Therefore we computed the maximum probability of data in each zero bin, p_0 , such that the probability of no data falling in all n_0 observed zero bins was < 0.5 (because, in our single trial, we observed n_0 0 bins). The relative order of K-L distances was robust, and our results did not change substantially for other bin sizes, center locations, and zero bin counts between 0.01 and 1.

Measurements of K-L distance, like other information-theoretic measurements, cannot be assigned statistical confidence without extensive prior information. Therefore these distances must be referenced to another distance value. We chose our reference to be the distance between Lennie et al.'s nonoriented and oriented cell distributions (shown in Fig. 2, *C* and *D*), which represents a substantial change between distributions, and is generally believed to represent a difference between cell classes. Because of the asymmetry of K-L distance, $D(\text{oriented}\|\text{nonoriented}) = 0.97$ while $D(\text{nonoriented}\|\text{oriented}) = 1.69$, so we took as our reference the average, 1.33. However, using either individual distance value does not affect our results.

Artificiality in ICs

In our spatial and color tuning analyses, we considered the influence of artifactual ICs on the distribution of color tuning. Because ICA is a noisy optimization process with a finite data set, it will inevitably produce some artifactual ICs. However, using low variance as an indicator of artifactuality (as is commonly done) would introduce a bias against color because color is a smaller source of variance in natural images than luminance (Ruderman et al. 1998). Therefore we used the van Hateren and van der Schaaf (1998) criteria: artifactual ICs extend over only a few pixels and have nearly equal power in all four corners of the power spectrum. These criteria are entirely spatial and do not, a priori, introduce a color bias.

We also considered whether the limited size of image fragments could cause artifacts in the color-tuning data. If spatial or spatial frequency edge effects were a problem, we would expect to see more artifactual bias at the smaller size. Notably, we obtained identical results using two patch sizes: 8×8 and 12×12 (data not shown). This suggests that edge effects do not significantly alter color tuning.

Spatial-chromatic separability assumption

We replicated the Lennie et al. (1990) study of V1 neurons in two ways: under ideal conditions and under conditions that imitate their exact methods, including experimental limitations. In the ideal case, we simultaneously measured spatial tuning (optimal orientation and SF) and color tuning (azimuth and elevation). In the imitation case, we followed the Lennie et al. assumption that color and spatial tuning are separable (i.e., regardless of stimulus color, a cell will have the same optimal grating stimulus). Therefore in this case we first measured spatial tuning using achromatic gratings and then used the orientation and SF of the best achromatic grating to make chromatic gratings for color tuning measurement. Lennie et al. sometimes re-optimized spatial tuning with a colored grating and used colored gratings on occasions when achromatic ones did not work well. However, we excluded these heuristics, because they were not formalizable.

Effect of noise

A major difference between much modeling and real experiments is the presence of noise in the experiments. Therefore when imitating the Lennie et al. (1990) experiments, we tested the effect of noise on color tuning. We assumed that each IC's maximum amplitude of response modulation equaled 75 spikes/s, a value typical of Lennie et al.'s neuronal responses. For response modulations of that size, the amplitude SD is ~ 5.25 spikes/s (Levine 1995). We simulated 25 different experiments in which each grating stimulus was presented for 20 trials. To each trial's "grating response" (Fourier amplitude), we added amplitude noise with a SD of 5.25 and a mean of zero. In each experiment, we averaged responses over all 20 trials and took the best stimulus as the color tuning. Then to assess whether noise caused an overall bias, we computed the spherical mean of these color tuning measurements across 25 experiments (Fisher 1987).

RESULTS

Starting with a set of hyperspectral images (Párraga et al. 1998), we constructed three different image sets with three color planes each (LMS-encoded, RGB-encoded, and JPEG-compressed RGB-encoded). We took each data set, performed ICA, and obtained 432 IC basis functions and IC filters. Figure 1, *A-C*, shows equivalent samples of ICA filters for the RGB, JPEG, and LMS image sets, whereas Fig. 1*D* shows a corresponding sample of LMS basis functions (see figure legend). Many of the filters look similar to the elongated Gabor function RFs of V1 simple cells (e.g., Ringach 2002), and many have

apparent red-green (L-M) or blue-yellow [S-(L+M)] opponency. However, there are significant differences between the three sets of filters in spatial structure and color tuning.

JPEG encoding creates artifacts in ICs

Both Tailor et al. (2000) and Hoyer and Hyvärinen (2000) have published ICA results from RGB images. Because the Tailor et al. images were encoded using JPEG compression, which is lossy (i.e., does not preserve all information in the image), we investigated whether this had affected the ICs by directly comparing the effect of the different encodings on IC filters derived from our data set.

In the RGB condition, our IC filters (Fig. 1A) corresponded fairly well to the basis functions of Hoyer and Hyvärinen, whereas in the JPEG condition, our filters (Fig. 1B) were similar to those found by Tailor et al. We found that the spatial structure of the ICs found by Tailor et al. was indeed affected by JPEG compression of the source images: JPEG filters contained a large number of blue-yellow and red-green checkerboard patterns (e.g., rows 3–4, and 9–10 in Fig. 1B) that do not correspond to any RGB filters. To ensure that these artifacts were the results of JPEG compression, and not the absence of the log transform, we also ran ICA on non-log-transformed RGB images. The results (data not shown) are spatially similar to the log-transformed RGB ICs, not the JPEG ICs.

These artifacts result from JPEG’s compression algorithm, which divides images into 8×8 blocks and then discards information within blocks. Because the image fragments used for ICA are not aligned with block boundaries, the boundaries become artifacts. JPEG encoding also produces chromatic artifacts because it imitates an idealized human visual system by encoding images into luminance, red-green and blue-yellow color planes. Thus ICA of JPEG images reveals mechanisms of the JPEG compression algorithm rather than features of the natural visual world.

RGB encoding also biases ICs

We next investigated the effect of using uncompressed RGB-encoded images, compared with the more biologically realistic LMS encoding. Hoyer and Hyvärinen (2000) justify their use of RGB images from a digital camera by arguing that the colorspace choice will make no difference. Because RGB and LMS colorspace are related by a linear transform, they argue, the correlations in the data remain unchanged, thus ICA should produce the same ICs for both image sets. However, there are several nonlinearities in the digital camera’s imaging process, such as color balancing and gamma correction, and preprocessing also includes another nonlinear step, the log transform. Because all of these nonlinearities may cause discrepancies between the RGB and LMS conditions, we decided to measure their influence on spatial and chromatic tuning.

The spatial structures of IC filters in the RGB condition (Fig. 1A) and the LMS condition (Fig. 1C) were very similar, insofar as both sets contain mostly elongated Gabor-like features. There was noise in the LMS filters that was not present in the RGB filters, which was also true of the basis functions (data not shown). This may have resulted from the substantially greater overlap of L and M cone spectra (Fig. 1E) compared with R and G detectors (Fig. 1F): a given color image will have

very similar L and M activations, and thus image noise within the nonoverlapping spectral regions will have an exaggerated effect on ICs in which L and M are opponent.

To compare color tuning in the RGB condition with the LMS condition, we linearly transformed RGB filters into the LMS colorspace, using the matrix of dot products between the RGB and LMS spectra in Fig. 1, *F* and *E*. Color tuning was summarized as the elevation (luminance sensitivity) and azimuth (chromatic preference) of the drifting sinusoidal grating that evoked the greatest response (see METHODS). The color tuning of RGB filters (Fig. 2A) differed substantially from LMS filters’ color tuning (Fig. 2B): RGB filters fell into a small number of line-like, continuous clusters corresponding to yellow (30° azimuth), magenta (135°), and blue (90°). LMS filters clustered at red (0°) and blue-magenta (120°) and were much more diffusely clustered. The shift in the centers of azimuth clusters suggests that the difference between RGB and LMS is not merely due to noise and that, contrary to the claim of Hoyer and Hyvärinen (2000), RGB-encoded ICs are not appropriate after all for comparison with the visual system. Therefore an ICA model of the visual cortex should use the visual system’s LMS encoding of the chromatic information in natural scenes.

Insofar as RGB filters are comparable to the visual cortex, RGB encoding has one provisional advantage over LMS encoding: it reduces noise in the ICs, as explained above. Consequently, although RGB encoding produces biased IC filters, any similarity to V1 cells bolsters our claim (in the following text) that IC filters resemble V1 cells.

LMS IC filters have similar spatial tuning to simple cells and achromatic ICs

To examine whether color information affects the spatial structure of ICs, we compared the spatial properties of LMS IC filters with the spatial properties of achromatic IC filters (van Hateren and van der Schaaf 1998) and macaque simple cells (De Valois et al. 1982; Parker and Hawken 1988). We obtained the distributions of best orientation, orientation bandwidth, spatial frequency, spatial frequency bandwidth, and aspect ratio (envelope length divided by envelope width), using the best achromatic stimulus for each IC.

We identified artifactual ICs following the van Hateren and van der Schaaf (1998) criteria; they constituted 29% (124/432) of our ICs. However, some of these artifactual ICs had high variances (i.e., they were robust sources for the image data), so we also show, but do not discuss, spatial analyses for the full set of ICs.

The orientation distribution of nonartifactual filters (Fig. 3A, ■) was similar to the distribution of achromatic filters, with peaks at 0° (vertical), 45° , 90° (horizontal), and 135° . The orientation bandwidth peak (Fig. 3B, ■) was near 60° in our colored ICs, and the entire distribution seemed to be shifted toward higher values compared with achromatic ICs and older studies of macaque V1 (De Valois et al. 1982; Parker and Hawken 1988). Interestingly, more recent physiological studies have reported an orientation bandwidth distribution much like ours (Ringach et al. 2002); in fact, Ringach (2002) comments on the discrepancy between their data and achromatic ICA data, which color ICA apparently resolves. The distribution of color filter spatial frequencies (Fig. 3C) was qualitatively similar to achromatic filters in that it increased exponentially with

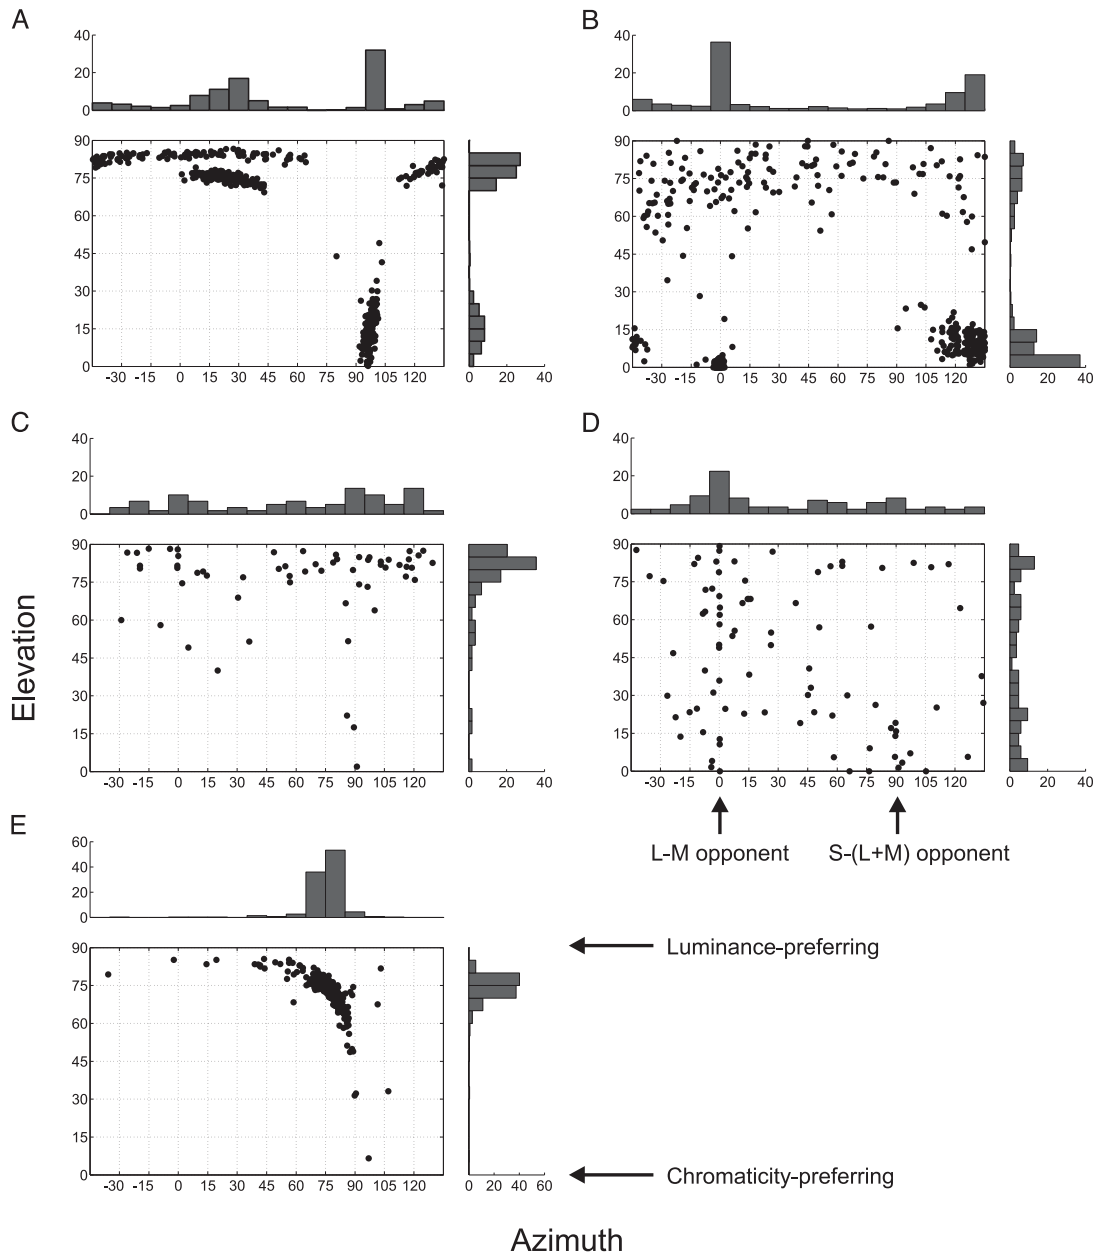
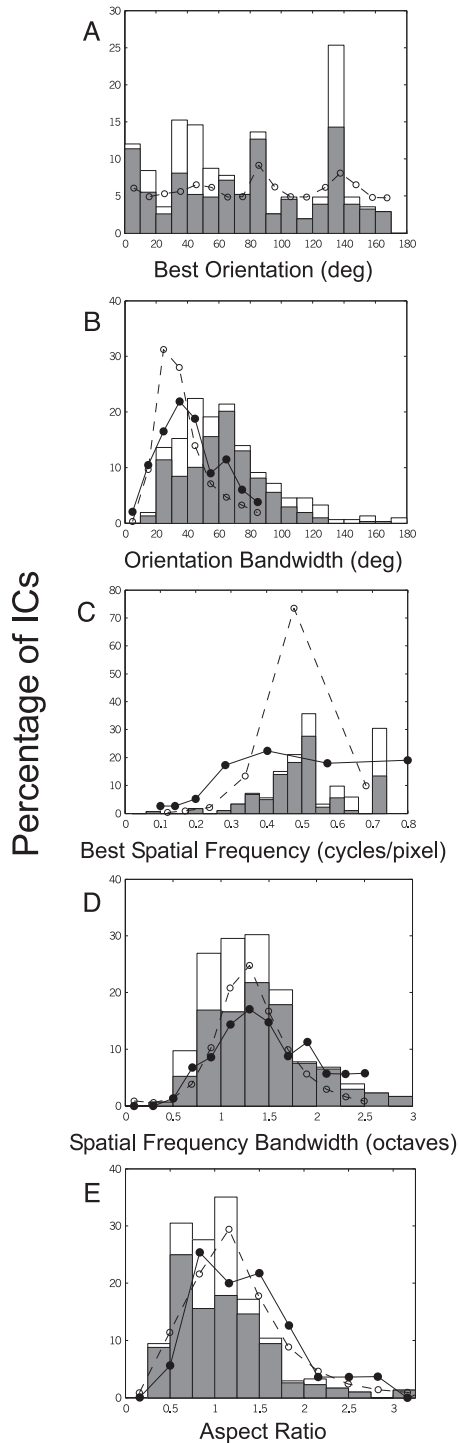


FIG. 2. The color tuning of IC filters and primary visual cortex (V1) cells is similar. The color tuning of each IC (*A*, *B*, and *E*) was measured in the DKL colorspace following the methods of Lennie et al. (1990) except that we optimized the grating stimulus over all spatial and chromatic gratings (the “ideal case”). The color tuning of V1 cells (*C* and *D*) is replotted from Lennie et al. and assumes that spatial and chromatic tuning are separable, with separate stages for finding the optimal achromatic and chromatic stimuli. Each plot’s marginal distributions indicate density, as a percentage of the total distribution. Following Lennie et al., azimuth is plotted between -45 and 135° to better depict clustering near 0° and is taken modulo 180° to emphasize color opponency rather than exact color preference. *A*: IC filters derived from RGB input and transformed into the LMS cone space. *B*: LMS filters. *C*: oriented cells. *D*: nonoriented cells. *E*: IC basis functions from LMS filters.



SF and peaked at the Nyquist frequencies for both horizontal/vertical and diagonal filters. However, color filters were more diverse in preferred SF than achromatic, and better resemble the physiological data. The spatial frequency bandwidth distribution of color filters (Fig. 3D) was similar to both achromatic filters (van Hateren and van der Schaaf 1998) and simple cells (De Valois et al. 1982). The aspect ratio distribution of color filters (Fig. 3E) was centered at 1 (envelope length equal to envelope width) as with achromatic filters and simple cells (van Hateren and van der Schaaf 1998). Overall, with artifactual ICs removed, the distributions of the spatial parameters of colored IC filters were similar to the distributions reported for achromatic ICs and fit the simple cell distributions about as well. The fit between the spatial properties of colored IC filters and simple cells was somewhat better for SF, marginally worse for aspect ratio; and while colored IC filters fit the most recent orientation tuning bandwidth data, achromatic IC filters fit the older data.

LMS IC filters, but not basis functions, have similar color tuning to oriented V1 cells

Because LMS IC filters are spatially similar to simple cells, we compared the filters' color tuning with V1. We used the data of Lennie et al. (1990), who divided V1 cells into three classes: oriented, nonoriented and complex. We first considered oriented cells, which roughly correspond to simple cells, but because many ICs are chromaticity sensitive, we also considered nonoriented cells, which Lennie et al. found to be most color sensitive.

Comparing the color tuning of LMS filters (Fig. 2B) to oriented V1 cells (Fig. 2C), we found similar classes of color tuning in both. The azimuth distributions were bimodal, with a red-green (L-M) opponent cluster near 0° azimuth and a slightly larger blue-yellow [S-(L+M)] opponent cluster near 90° . However, the azimuth distribution of chromaticity-tuned, S-cone-sensitive filters was actually centered near 120° ; thus filters are closer to (S+L)-M cone opponency (135°). In both distributions, we found chromaticity-tuned and luminance-tuned RFs. Chromaticity-tuned RFs (at lower elevations) were bimodal with very few intermediates. However, filters seemed substantially more chromaticity tuned than oriented V1 cells; a large number of L-M opponent filters with azimuths near 0° were purely chromatic, with elevations near 0° .

According to the Lennie et al. (1990) strict criteria, which require a nonoriented cell to respond better to full-field modulation than to any grating, only one of our IC filters (a full-field noise patch) was nonoriented. Nevertheless, a moderately large population of filters had a center-surround-like organization reminiscent of blob cells (Livingstone and Hubel 1984), so we also compared the filters to Lennie et al.'s

FIG. 3. Spatial tuning parameters of LMS IC filters resemble achromatic ICs and simple cells. The spatial properties of each IC were measured using each IC's best achromatic stimulus. \square , distributions for all ICs; \blacksquare , the distribution of ICs that were nonartifactual by the van Hateren and van der Schaaf (1998) criteria. \bullet and --- , the reported V1 distributions (De Valois et al. 1982; Parker and Hawken 1988); \circ and - - - , the achromatic IC distributions (van Hateren and van der Schaaf 1998). *A*: best orientation; 0° corresponds to vertical. *B*: bandwidth (full width at half-maximum) of orientation tuning. *C*: best spatial frequency. *D*: bandwidth of spatial frequency tuning. *E*: aspect ratio of the receptive field (RF) envelope.

nonoriented (presumed blob) cells (Fig. 2D). Like filters, non-oriented cells were basically bimodal in azimuth with a red-green (L-M) cluster near 0° and a blue-yellow [S-(L+M)] cluster near 90°; however, the nonoriented cells' L-M mode was dominant, and azimuths between 30 and 60° were more common. The elevation distributions of nonoriented cells and filters were similar: both were bimodal, although nonoriented cells were less sharply so.

Previously, LMS IC basis functions (rather than LMS filters) have been compared with V1 (Hoyer and Hyvärinen 2000; Lee et al. 2002). Although there are a priori reasons why basis functions should not be analyzed this way (see DISCUSSION), we have also observed that the color tuning of basis functions (Fig. 2E) was markedly different from both V1 cells and filters: the distribution contained only a single cluster near azimuth 75° and elevation 75°. This is somewhat surprising because basis functions are apparently diverse in color; however, their cone inputs were typically poorly balanced. Full-field modulation, which is a poor stimulus for V1 cells, was the best stimulus for 73% of basis functions and <1% of filters; this empirically confirms that IC filters are the correct comparison for V1 cells.

Quantitative comparison between IC filters and V1 shows significant similarity once experimental biases are removed

The color tuning of the LMS filters showed some qualitative similarities with the Lennie et al. (1990) V1 cells; still, the raw filters contained two notably chromatic groups, near 0 and 120° azimuth, which were not strongly evident in the raw V1 data. However, comparing raw color tuning is misleading because it fails to take account of three discrepancies between the V1 data and the filters. First, the raw ICA data included a subpopulation of artifactual low-variance ICs. Second, the V1 data assumed that spatial tuning can be determined achromatically before determining color tuning; that is, it assumes the best chromatic stimulus is spatially identical to the best luminance stimulus. This spatial-chromatic separability assumption might have biased measurements of color tuning. Third, the firing rates of V1 cells are subject to noise; this could also have biased estimates of cells' color tuning. Because we could not test the effects of noise and separability on actual neurons, we measured their effects on the color tuning of the model IC filters.

Using the van Hateren and van der Schaaf (1998) criteria, we found that 29% of color ICs are artifactual compared to the 25% of achromatic ICs they report. In terms of color tuning, these artifactual ICs (Fig. 4A) tended to have very low elevations. Eliminating them made the color tuning of the remaining filters more similar to V1 in addition to improving spatial tuning.

To test what difference the spatial-chromatic separability assumption makes for our dataset, we compared chromatic tuning when spatial tuning was measured achromatically with chromatic tuning when spatial and chromatic tuning were simultaneously determined. Measuring spatial tuning achromatically did bias color tuning. It caused elevations to increase and azimuths to decrease; that is, filters apparently became more luminance sensitive and red-green opponency became more prominent (Fig. 4B). Thus the physiological data's higher elevations, as well as its azimuth cluster at 90° rather than 120° (Fig. 2, C and D), may have been partly caused by this separability assumption.

Noise in neuronal responses is another possible source of bias. In Lennie et al.'s spherical coordinate system, there is a nonlinear relationship between elevation and stimulus contrast. In essence, because z-axis modulation strength dominates the x and y axes, cone contrast changes little with changes at low elevations and much more with changes at medium elevations. For most cells, contrast determines the firing rate (and therefore noise level). Because noise increases with elevation, stimuli with higher elevations are more likely to be experimentally misidentified as the optimal stimulus, which could cause a systematic upward elevation bias.

We measured the distance between the color tuning of noiseless IC filters and their mean color tuning with noise added and found that noise biased the color tuning toward higher elevations (Fig. 4C), causing a mean elevation overestimate of 7.6° and a number of severe overestimates. Thus some of the higher elevations seen in the physiological data (Fig. 2, C and D) compared with the noiseless LMS IC filters (Fig. 2B) may have arisen from noise in the physiological responses.

Thus the assumption of spatial-chromatic separability, and the presence of noise in real V1 neuronal responses, both result in experimental biases that make cells appear dissimilar to IC filters. Furthermore, the inclusion of low-variance artifactual ICs in the comparison with real neurons also makes the IC filters look dissimilar to the real cells. Figure 4D shows how the color tuning of our IC filters (from Fig. 2B) changes once we account for these three biases; it shows more similarity to the color tuning distributions of real V1 neurons (Fig. 2, C and D).

To make the comparison between IC filters and V1 cells more rigorous, we measured the Kullback-Leibler (K-L) distance between the color tuning distribution of filters and the distributions of oriented cells (Fig. 2C), nonoriented cells (Fig. 2D), and both cell types pooled (i.e., all noncomplex V1 cells). K-L distances were normalized relative to a reference distance, which we took to be the distance between the distributions of oriented and nonoriented cells (see METHODS).

We measured these normalized K-L distances for raw IC filters, filters corrected for each individual bias, filters corrected for noise and separability, and filters corrected for all three biases. Although raw IC filters best resembled the pooled V1 cells, they were not especially similar to this group or to the oriented or nonoriented groups since all K-L distances were near 1 (Table 1, column 1). However, every single bias correction improved the fit between filters and all types of cells (Table 1, columns 2–4); correcting the noise bias had the largest effect. Noise and separability corrections appeared to be most important; subsequent to these, the effect of removing artifactual IC filters was only moderate (Table 1, column 5). When all three biases were taken into account (Table 1, column 6), the filters revealed a strong similarity to the pooled cells and substantial similarity to oriented and nonoriented cell groups.

Even after all biases were corrected, some residual differences between filters and V1 cells remained, contributing to the K-L distance. The low-elevation S-cone cluster fell at higher azimuths in filters than cells (120 rather than 90°) even after correction of the spatial-chromatic separability assumption. Also, there were nonoriented cells, but not filters, at intermediate azimuths (30–75°).

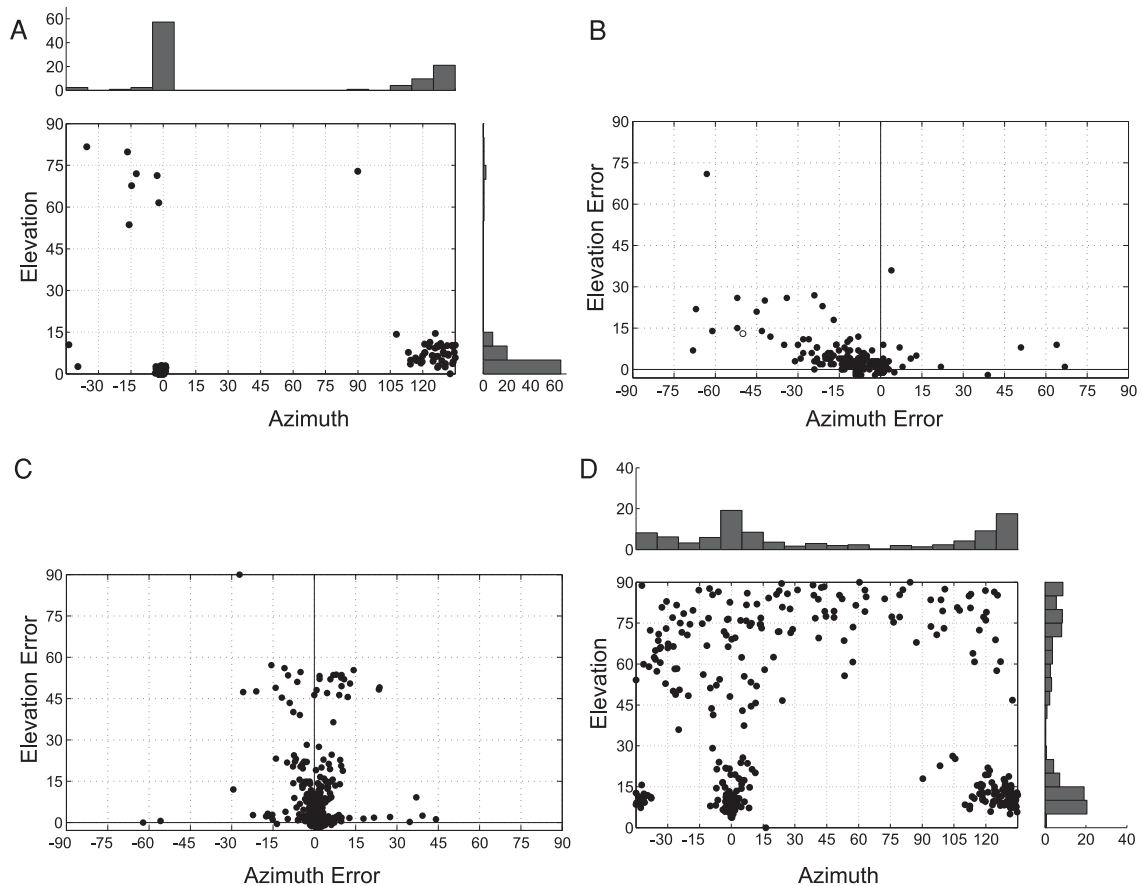


FIG. 4. Experimental biases strongly affect the comparison between ICs and physiological data, but correction greatly improves the correspondence with V1. Three sources of bias were considered: low variance artifactual ICs, the assumption of spatial-chromatic separability, and physiological noise. *A*: color tuning (plotted as in Fig. 2) of the LMS filters that were artifactual according to the van Hateren and van der Schaaf (1998) criteria. Marginal distributions indicate density as a percentage of the total distribution. *B*: inferred bias in physiological color tuning data due to assuming separability, inferred from comparing measurements of filter color tuning made under ideal and experimental conditions. The open circle represents the example in Fig. 7. *C*: inferred bias in physiological data due to noise. The displacement shown is between the mean of 25 color tuning estimates and the ideal tuning. *D*: distribution of LMS IC filters after all biases—artifactual ICs, separability, and noise—are taken into account. Marginal distributions are as in *A*.

Ultimately, in all analyses of our filters, we found that the closest match (smallest K-L distance) was with the pooled V1 cells (Table 1) rather than with either subpopulation, such as the more chromaticity-preferring nonoriented cells or the more Gabor-like oriented cells. It seems that IC filters are a heterogeneous group most similar in color tuning to the pool of oriented and nonoriented V1 cells.

Dependence of chromaticity on orientation selectivity is consistent with experimental observations

Because it has been widely believed that responsiveness to purely chromatic stimuli is concentrated among nonoriented cells (Livingstone and Hubel 1984; Lennie et al. 1990), we examined the relation between orientation tuning bandwidth

and chromaticity sensitivity (elevation in DKL space). Figure 5 shows that, indeed, filters with broad orientation tuning tend to be chromaticity sensitive and have low elevations. Filters with narrow orientation tuning are more diverse but have a tendency toward higher elevations. Nonartifactual filters are primarily responsible for this observation since almost all artifactual filters have elevations near zero. Because most filters have relatively sharp orientation tuning, we can thus account for the physiological observation that nonoriented cells tend to be more chromatic without inferring that oriented cells are necessarily achromatic. In fact, our results are also consistent with the Johnson et al. (2001) recent report that 79% of color-responsive neurons, stimulated with cone-isolating gratings (see following text), had oriented RFs. Because most filters are well tuned for orientation, it emerges that most chromaticity-preferring filters do have fairly narrow orientation tuning.

TABLE 1. Normalized Kullback-Leibler distances between color tuning distributions of V1 cells and IC filters

V1 Cell Class	IC Group					All corrections
	Raw	Non-artifactual	Separability corrected	Noise corrected	Separability + noise	
Oriented	1.02	0.84	0.88	0.89	0.72	0.58
Nonoriented	0.88	0.79	0.84	0.64	0.50	0.48
Pooled	0.79	0.66	0.71	0.59	0.44	0.37

All distances are normalized relative to the Kullback-Leibler distance between oriented and nonoriented cell classes (see METHODS). Pooled class contains both oriented and nonoriented cells. All corrections indicates that high variance independent components (ICs) have been corrected for both separability and noise. V1, primary visual cortex; IC, independent component.

Filters fall into two strongly double-opponent clusters

To compare ICs to more recent physiological data (Conway 2001; Johnson et al. 2001), we simulated the LMS IC filters' responses to cone-isolating stimuli. These stimuli avoid some of the confounds of cone-opponent stimuli by modulating only one cone type at a time. For instance, in an L-cone isolating grating, the light stripe increases L-cone excitation relative to background, whereas the dark stripe decreases L-cone excitation. Although physiological cone-isolating measurements are often noisy and difficult to calibrate because of the overlap between cone absorption spectra (Conway 2001), in our model, this measurement is trivial: we simply look at the filter's three color planes, which represent the individual LMS inputs.

For each IC filter, we measured its response amplitude to the optimal L-, M-, and S-isolating gratings, which indicates the relative strength of different cone inputs. We also measured its response phase to each grating because this determines the spatial relationship between the cone inputs. For example, a cell could have L and M inputs with identical best orientation and SF, and if those inputs are in phase with each other (0° apart) the cell will be (L+M) or yellow sensitive. However, if they are 180° apart in phase, it will be L-M or red-green double-opponent. Figure 7A shows a schematic RF containing a double-opponent component (Livingstone and Hubel 1984; Michael 1978), which responds fairly specifically to a border

between two colors (yellow and blue) because its inputs from the corresponding cones are opposite in sign.

The response amplitude of all our IC filters to L-, M-, and S-isolating gratings of optimal orientation and SF is shown in Fig. 6A (○). Filters fell into three clusters: one elongated cluster with high S response, some M response, and little L response; another elongated cluster with L and M response but little S response; and a diffuse cluster near zero. These roughly correspond to the modes of the cone-opponency data in Fig. 2B.

To quantify the double-opponency of the different clusters, we compared the response phases of cone pairs. Among strongly L- and M- responsive filters (L and M amplitudes >100), L and M cones were almost always 180° apart in phase (94%, Fig. 6B). Moreover, even though the optimal cone-isolating stimuli were not constrained by orientation and SF, 94% of L and M cone-isolating stimuli did have identical orientations and 95% did have identical SFs. Thus for these double-opponent ICs, the spatiochromatic tuning to an L-cone-isolating stimulus predicts the tuning to the M-cone-isolating stimulus, which is consistent with Johnson et al. (2001) physiological SF measurements.

Conway (2001) reported that, in a population of weakly oriented L-M cells studied with cone-isolating pixel stimuli, most S responses were in phase with M responses. These cells could therefore be called red-cyan [i.e., L-(M+S)]. Among our L-M responsive filters, the amplitude of S response indeed correlates with M response (lower cluster in Fig. 6A), but the M and S inputs were phase opponent (Fig. 6C) when measured at the M-isolating grating's best orientation and SF. Thus our filters tend to show magenta-green [(L+S)-M] opponency, which was atypical but occasionally seen in Conway's data. This tendency of our filters may arise from the spectral overlap of M and S cones, which ICA will attempt to decorrelate even within L-M opponent IC filters.

Although there is little physiological data on S-cone double-opponency in V1, we measured its prevalence in our filters. Among S-responsive filters (S amplitude >20 units), we found double-opponency between M and S when measured at the S-isolating grating's best orientation and SF (Fig. 6D). We might have expected that these S-cone driven filters would be opponent to yellow (L+M). However, in the majority of filters, L and S actually have similar phase (Fig. 6E), whereas the remainder are phase opponent. This is probably also due to the spectral overlap described in the preceding paragraph. Thus our strongly S-responsive filters are also mostly magenta-green opponent, although a few are red-cyan. This prediction of our model has not yet been explicitly tested, although it does appear consistent with Lennie et al.'s data.

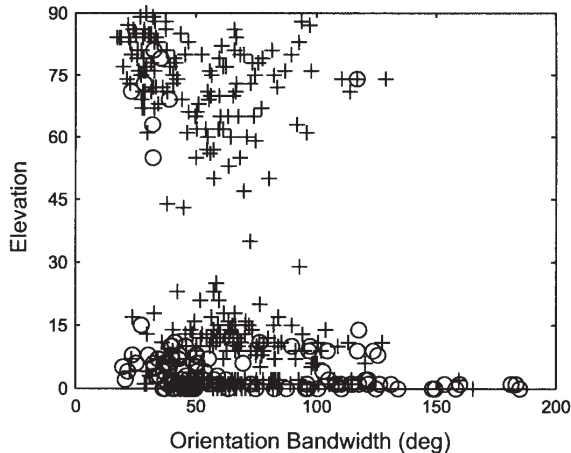


FIG. 5. Width of orientation tuning correlates with preference for chromaticity. Color tuning elevation, measured with optimal chromatic gratings, is shown as a function of orientation tuning bandwidth. ●, nonartifactual LMS IC filters; ○, artifactual filters.

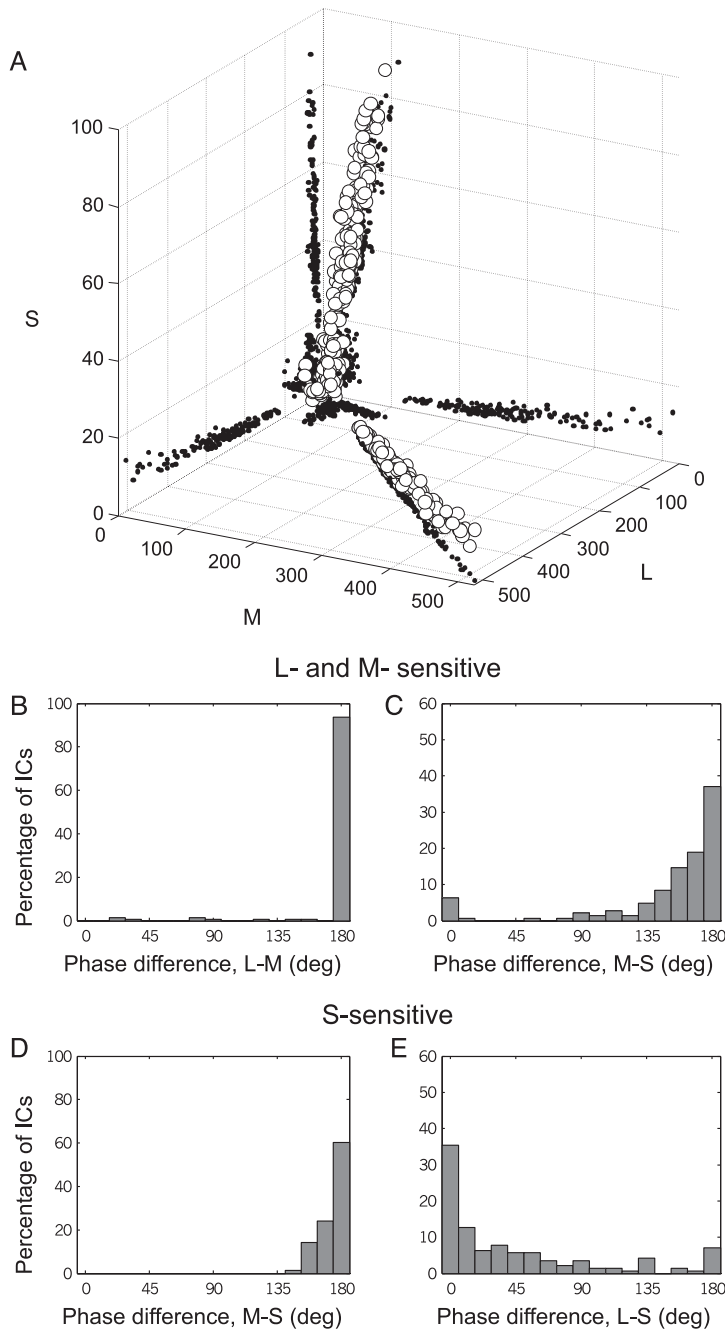


FIG. 6. LMS IC filter responses to cone-isolating gratings fall into distinct double-opponent groups. LMS IC filter responses were measured in each color plane, using cone-isolating gratings, and the spatial relationship between cone inputs was determined. *A*: response amplitudes of each filter to its optimal L, M and S-isolating gratings (\circ), \bullet , the projections of the \circ onto the cardinal ($L = 0, M = 0, S = 0$) planes. Axes are in arbitrary units; the L and M axes are much greater in magnitude than the S axis because of the overlapping absorption spectra of L and M cones. To make those 2 dimensions independent, the filter must be stronger along those axes. *B-E*: phase difference between the optimal cone-isolating gratings for pairs of cones. A peak at 180° indicates phase opponency, whereas a peak at 0° indicates phase coherence. Among the strongly L- and M-sensitive filters, *B* shows the phase difference between M and L and *C* the difference between M and S. Among the strongly S-sensitive filters, *D* shows the phase difference between S and M and *E* the difference between S and L.

DISCUSSION

The colored IC filters of LMS-encoded natural scenes are similar to primate V1 cells in their color tuning and are spatially similar to both simple-cell RFs and achromatic ICs. This is a surprising result, given that V1 cells have often been reported to be mostly achromatic, especially those with ori-

ented RFs (Hubel and Wiesel 1968; Lennie et al. 1990; Livingstone and Hubel 1984), whereas ICs have been reported to be unbiologically chromatic and double-opponent (Tailor et al. 2000). However, it emerges naturally from a biologically realistic set of LMS IC filters when care is taken with the ICA methodology to make proper comparisons between V1 cells




FIG. 7. Assuming separability may cause severe biases when measuring double-opponent cells. The problem with assuming separability of spatial and chromatic RFs is illustrated, schematically and with an example IC filter. For clarity, we describe stimuli using RGB terms, even though the filter is LMS-encoded. *A*: schematic double-opponent cell with 4 subfields. The *leftmost subfield* is inhibited by red, and the *rightmost* is excited by green. The *middle 2 subfields* are yellow-blue double-opponent; the *middle left subfield* is excited by red and green (R+ and G+) and inhibited by blue (B-), whereas the *middle right* is inhibited by red and green (R- and G-) and excited by blue (B+). *B*: a filter from our data set (open circle in Fig. 4B) that is modeled by the schematic cell. It is shown in full color (*left*), and then separated into color planes for clarity (*right*). In all figures, medium gray indicates 0 input. *C*: the true best cone-opponent grating. Note that it is spatially aligned with the filter and shows the influence of its L and M cone inputs. *D*: false “best” cone-opponent grating, assuming spatial-chromatic separability. Note that the grating is completely misaligned and shows little evidence of the filter’s strong L input.

and ICA output and when the reality of physiological experiments (such as response noise) is taken into account.

It is methodologically essential to use LMS encoding of natural scenes (as was also done by Doi et al. 2003; Lee et al. 2002; Wachtler et al. 2001); this accurately represents the chromatic information available to the cortical, experience-dependent stages of the visual system. Colored images encoded in the usual ways for transmission over, say, the internet (RGB or even JPEG encoding) produce filters that are not appropriate for comparison; moreover, they are unlike real cells. Additionally, it is important to compare IC filters, not basis functions, with the physiological data. Although some previous studies (Hoyer and Hyvärinen 2000; Lee et al. 2002) have focused on basis functions, there is no theoretical justification for this (see following text); additionally, we find that basis functions are dominated by

properties (such as their strong response to full-field illumination) that have no counterpart in visual neurons.

The other major factor in this comparison is the correction of biases both in the IC analysis and the physiological data. A simplistic, uncorrected comparison between the IC filters and V1 cells (Lennie et al. 1990) suggests only a weak similarity (Fig. 2, *B* and *C*). The ICs appear to be more chromatically tuned than the real neurons. However, we find that this comparison is skewed by several biases. Some of the IC filters are highly chromatic noisy structures; when these are eliminated, the remaining filters better resemble simple cells, spatially and chromatically. Similarly, we estimate that the Lennie et al. V1 data are biased by their assumption of spatial-chromatic separability and by the effect of cortical noise.

Cortical noise is a particularly important source of bias in the Lennie et al. physiological data. We find that, in a simulated neurophysiological experiment, luminance (high-elevation) stimuli, which modulate all cones in concert, tend to produce noisier responses than chromatic (low-elevation) stimuli, which modulate contrast between two cone types. This skews measurements of color tuning, so that even cells’ averaged responses make them appear much less chromaticity sensitive than they actually are. Any method of finding the best color direction will be swayed by this noise, especially if it interpolates between stimuli (as do many experiments) or picks the best stimulus (as does our simulation). Johnson et al. (2001) speculate that one reason for the lack of chromaticity-tuned cells in Lennie et al.’s study is that their stimuli only weakly modulate chromatic neurons. For modulation strength to directly influence tuning in the DKL colorspace, however, a nonlinearity is necessary; we suggest that noise is a very significant source of nonlinearity.

Analysis of these biases suggests that simple cells, which are generally not considered the substrate for color coding, are in fact more chromatically sensitive than has often been depicted. This result is independent of the correctness of our model and aids in reconciling the Lennie et al. study with a growing body of recent neurophysiology (Conway 2001; Johnson et al. 2001).

The similarity between IC filters and V1 cells suggests that chromatic and spatial information are distributed across cells in V1 in a way that is compatible with theoretical principles of independence and sparse representation. In other words, the spatial and chromatic response properties of neurons in V1 can be accounted for by the redundancy-reduction hypothesis (Atneave 1954; Barlow 1959).

Some systematic differences between IC filters and V1 cells remain. Even the corrected filters shown in Fig. 4D are somewhat more chromaticity sensitive than the pool of oriented and non-oriented cells in V1. Filters are also probably more double-opponent than simple cells. Some limitations of our study, which may account for these differences, are discussed in the following text.

Limitations of the IC analysis

ICA is not a guaranteed method for finding an independent code for a data set. It is designed to discover independent sources of variability (basis functions), assuming these are linearly superposed to form the observed data. Natural images break this assumption, because they are not formed by a linear

superposition of sources (occlusion would be a better model than superposition). Under these circumstances, ICA can only discover sources that appear (by their non-Gaussianity) to be independent; these may be unrelated to the underlying structure of the data set.

Another limitation of our study is that the number of IC filters is capped by the dimensionality of the input data, which can artificially limit our ability to capture the structure of the input distribution. This limitation may become more critical as dimensionality increases as when going from achromatic to chromatic images. Generating an overcomplete basis set (e.g., Lewicki and Sejnowski 2000; Olshausen and Field 1996) from colored natural scenes would likely change the distribution of spatial tuning and may also change the distribution of color tuning.

Although we consider the effects of noise on neurophysiological experiments, we do not consider its effects on the redundancy-reduction hypothesis. Atick and Redlich (1990) have shown that, under noisy conditions, the RFs need to be more correlated (less independent) to preserve information about the image. We do not yet have an ICA model that allows us to take this into account.

Finally, ICA produces a strictly linear encoding; real neurons exhibit output nonlinearities and contrast normalization (e.g., Heeger 1992a,b), which are not part of our ICA model. There is evidence that nonlinearities play a role in V1 color processing (De Valois et al. 2000; Hanazawa et al. 2000). However, these nonlinearities primarily seem to sharpen a cell's tuning across color stimuli (De Valois et al. 2000). We surmise that this kind of sharpening nonlinearity might enhance the effect of noise and separability bias in studies of V1.

Filters versus basis functions

To compare the ICs of natural visual input with the RFs of real neurons, it is necessary to decide whether the IC filters or the IC basis functions are comparable to real RFs. For achromatic images, the basis functions and filters are very similar to one another, and this decision is not very significant (van Hateren and van der Schaaf 1998). For color data, however, there are substantial differences between the spatial and chromatic structure of the basis functions and filters (compare the color tuning of filters in Fig. 2B to the equivalent basis functions in 2E), which makes the choice critical. There has been some disagreement about this question in the color ICA literature: Hoyer and Hyvärinen (2000) and Lee et al. (2002) focus on basis functions, whereas Tailor et al. (2000) and Doi et al. (2003) focus on filters.

The mathematical formulation of ICA is unambiguous that IC filters are the correct choice for comparison with neuronal RFs. Basis functions represent sources of variability in the data set; they are notional image patches that can be superposed to generate the data set. IC filters, on the other hand, are simply linear filters designed to extract information from the data set; in this way, they are completely compatible with the popular linear-filtering model of simple cells (e.g., Field 1987).

Retinal cone mosaic

There is a discrepancy between our results and those of Doi et al. (2003), who also analyzed LMS IC filters but used a

model featuring a fixed preprocessing stage mimicking the distribution of cones in the retinal mosaic. Their results, however, compare rather poorly to their choice of V1 data (Hanazawa et al. 2000), in both the prevalence of color selectivity (54% in V1 vs. 5% in their filters) and color tuning distribution (broad in V1 vs. extremely peaked in their filters). In addition, their model's retinogeniculate stage developed many fewer S-preferring cells (3%) than koniocellular/parvocellular LGN (12%). The likely explanation for these differences is that Doi et al. only modeled the fovea (actually foveola), so that just 3% of the cone inputs to their model cortex were S-type. Most primate studies of color, including the data they use for comparison, have been done in the periphery, where the proportion of S cones is much higher: Lennie et al. (1990) studied cells within the central 3° of the visual field, and in fact Hanazawa et al.'s cells were between 0.8 and 2.4°. Within the central 2° (excluding the foveola), S cones represent $\geq 10\%$ of the population (de Monasterio et al. 1985). Therefore the Doi et al. results may be most valid for minute foveolar RFs that are acutely sensitive to the structure of the cone mosaic.

Our present model, which assumes that S-cone density does not constrain RFs' spatial structure, compares better with existing physiological data. However, if the model were to incorporate a retinal mosaic stage with a normal perifoveal concentration of S cones, we would expect an increase in the proportion of luminance-preferring filters (given the Doi et al. results), and possibly improvement in the fit between the model and V1.

Double-opponency and separability

An interesting feature of our IC filters is that the vast majority is double-opponent. This double-opponency has two components: first, each cone type's excitatory and inhibitory inputs balance across the RF; second, the inputs from opponent cones are exactly opposite in phase. The balance of excitatory and inhibitory inputs is a natural outcome for ICA primarily because unbalanced filters would be correlated. The phase opponency, however, arises from a fixed property of the early visual system: the overlap in the absorption spectra of L and M cones, as well as S and M cones, forces the filters to decorrelate those input pairs as much as possible. The color statistics of natural scenes may also contribute significantly to opponency (Lee et al. 2002).

Double-opponency seemed rare in Lennie et al. (1990)'s study. This is likely to be an underestimate because they used achromatic gratings to measure the optimal SF and orientation of the cells in their sample. This reflects an assumption of spatial-chromatic separability and generates especially misleading results for double-opponent cells. Figure 7A shows a schematic double-opponent cell the preferred stimulus of which contains a yellow-blue edge. When an achromatic grating (of optimal orientation and SF) is presented at any position, its light stripe will drive the yellow (R+G) and blue (B) inputs in one subfield equally, causing them to cancel out. Its dark stripe will do the same, and thus the response will not modulate. As a result, the estimates of preferred SF and orientation will be inaccurate; achromatic gratings cannot be used to estimate the tuning of double-opponent cells.

This effect can be demonstrated in our ICA data. An IC filter

that responds poorly to achromatic gratings, much like the schematic cell in Fig. 7A, is shown in Fig. 7B. Its true best color grating stimulus, which has the filter's dominant orientation and shows the influence of its L and M inputs, is shown in Fig. 7C. However, if one assumes separability and stimulates this filter with an achromatic grating, the orientation and SF measurement is confounded by the double-opponency, and the resulting "best" grating (Fig. 7D) bears no relation to the RF; it seems primarily influenced by spatial noise, and gives rise to arbitrary color tuning (with error shown by the open circle in Fig. 4B).

More recent studies, using reverse correlation (Conway 2001) and cone-isolating stimuli (Johnson et al. 2001), have found a population of double-opponent cells in V1. However, it is likely that the nearly obligatory double-opponency of our filters exceeds that of cortical cells. One possible reason is that the cortex, being subject to noise, does not decorrelate its RFs to the same degree as our noise-free IC filters. Another possible reason is that our image preprocessing excludes some known properties of the precortical visual system, such as chromatic aberration, the lower spatial resolution of the S-cone system, and the scattered color tuning of LGN cells. It may be that such considerations also help explain why our filters show magenta-green opponency rather than the red-cyan reported by Conway (2001).

Nonoriented cells

It has often been supposed that nonoriented cells in V1 are the substrate for color tuning rather than oriented simple and complex cells (Livingstone and Hubel 1984). However, ICA generates few nonoriented, center-surround filters because oriented features are more significant sources of variance in natural scenes than center-surround features. This would seem to make ICA a poor model for the center-surround color-selective cells in cytochrome oxidase blobs. However, we did find many small, chromatically double-opponent filters, especially among the lower-variance ICs. These resemble the side-by-side double-opponent subunits found by Johnson et al. (2001) and the wide-band orientation tuned cells of Conway (2001), both of which probably occur in blobs. Also, consistent with the traditional understanding of blobs, we found that preference for chromaticity correlates with weakness of orientation tuning. ICA might therefore be able to partially model color selectivity within blobs, however a complete model for V1 will require additional organizing principles such as topography (Barrow et al. 1996; Hyvärinen et al. 2001).

Conclusions

The similarity between the color-independent components of natural scenes and RFs in V1 suggests that redundancy reduction (by ICA in particular) provides a plausible account of spatiochromatic RF structure in V1. This similarity is not apparent from inspection largely because some filters that appear colored are often very sensitive to luminance variations. It is likely that a similar effect is present in V1, where simple cells are conventionally thought to be luminance-dominated, but we find that they are likely to have substantial color tuning that would be evident if their inputs were visualized.

We predict that, as our understanding of V1 color coding

improves, the color sensitivity of oriented cells will assume a more important role. In place of a clear division between luminance and color coding, our results suggest that simple cells may multiplex spatial and color information. Similarly, our results highlight the importance of double-opponency, which is advantageous for coding of colored borders and, as we have shown, for redundancy reduction. Finally, the success of ICA in accounting for spatiochromatic RF structure in V1 suggests that redundancy reduction will also prove to be a fruitful hypothesis in other sensory systems.

ACKNOWLEDGMENTS

The authors are grateful to C. Alejandro Párraga, Gavin Brelstaff, T. Troscianko, and I. Moorhead for making "hyperspectral" natural scenes publicly available and to C. Alejandro Párraga for the Nikon camera sensitivity curves. We thank M. Kvale and C. Onufryk for helpful commentary.

GRANTS

This research was begun while M. S. Caywood was supported by a Frank Knox Fellowship to the University of Cambridge and was continued with support from an National Defense Science and Engineering Graduate Fellowship. B. Willmore was supported by a Biotechnology and Biological Sciences Research Council studentship.

REFERENCES

- Atick JJ and Redlich AN. Towards a theory of early visual processing. *Neural Comput* 2: 308–320, 1990.
- Attneave F. Some informational aspects of visual perception. *Psych Rev* 61: 183–193, 1954.
- Barlow HB. Sensory mechanisms, the reduction of redundancy, and intelligence. In: *The Mechanisation of Thought Processes*. London: HM Stationery Office, 1959, pp. 535–539.
- Barlow HB. Unsupervised learning. *Neural Comput* 1: 295–311, 1989.
- Barrow HG, Bray AJ, and Budd JM. A self-organizing model of "color blob" formation. *Neural Comput* 8: 1427–1448, 1996.
- Baylor DA, Nunn BJ, and Schnapf JL. Spectral sensitivity of cones of the monkey *Macaca fascicularis*. *J Physiol* 390: 145–160, 1987.
- Bell AJ and Sejnowski TJ. The "independent components" of natural scenes are edge filters. *Vis Res* 37: 3327–3338, 1997.
- Comon P. Independent component analysis—a new concept? *Signal Processing* 36: 287–314, 1994.
- Conway BR. Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1). *J Neurosci* 21: 2768–2783, 2001.
- Delorme A and Makeig S. EEGLAB: Matlab Toolbox for electrophysiological research. Swartz Center for Computational Neuroscience, Institute for Neural Computation, UCSD. www.sccn.ucsd.edu/eeelab, 1997–2003.
- Derrington AM, Krauskopf J, and Lennie P. Chromatic mechanisms in lateral geniculate nucleus of macaque. *J Physiol* 357: 241–265, 1984.
- de Monasterio FM, McCrane EP, Newlander JK, and Schein SJ. Density profile of blue-sensitive cones along the horizontal meridian of macaque retina. *Invest Ophthalmol Vis Sci* 26: 289–302, 1985.
- De Valois RL, Albrecht DG, and Thorell LG. Spatial frequency selectivity of cells in macaque visual cortex. *Vis Res* 22: 545–559, 1982.
- De Valois RL, Cottaris NP, Elfar SD, Mahon LE, and Wilson JA. Some transformations of color information from lateral geniculate nucleus to striate cortex. *Proc Natl Acad Sci USA* 97: 4997–5002, 2000.
- Doi E, Inui T, Lee TW, Wachtler T, and Sejnowski TJ. Spatiochromatic receptive field properties derived from information-theoretic analyses of cone mosaic responses to natural scenes. *Neural Comput* 15: 397–417, 2003.
- Field DJ and Tolhurst DJ. The structure and symmetry of simple-cell receptive-field profiles in the cat's visual cortex. *Proc R Soc London B Biol Sci* 228: 379–400, 1986.
- Field DJ. Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A* 4: 2379–2394, 1987.
- Field DJ. What is the goal of sensory coding? *Neural Comput* 6: 559–601, 1994.
- Fisher NI, Lewis T, and Embleton BJJ. *Statistical Analysis of Spherical Data*. Cambridge, UK: Cambridge Univ. Press, 1987.

- Hanazawa A, Komatsu H, and Murakami I.** Neural selectivity for hue and saturation of color in the primary visual cortex of the monkey. *Eur J Neurosci* 12: 1753–1763, 2000.
- Heeger DJ.** Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9: 181–97, 1992a.
- Heeger DJ.** Half-squaring in responses of cat striate cells. *Vis Neurosci* 9: 427–43, 1992b.
- Hoyer PO and Hyvärinen A.** Independent component analysis applied to feature extraction from colour and stereo images. *Network* 11: 191–210, 2000.
- Hubel DH and Wiesel TN.** Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195: 215–243, 1968.
- Hyvärinen A, Hoyer PO, and Inki M.** Topographic independent component analysis. *Neural Comput* 13: 1527–1558, 2001.
- Johnson EN, Hawken MJ, and Shapley R.** The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nat Neurosci* 4: 409–416, 2001.
- Lee TW, Wachtler T, and Sejnowski TJ.** Color opponency is an efficient representation of spectral properties in natural scenes. *Vis Res* 42: 2095–2103, 2002.
- Lennie P, Krauskopf J, and Sclar G.** Chromatic mechanisms in striate cortex of macaque. *J Neurosci* 10: 649–669, 1990.
- Levine MW.** Variability of responses to sinusoidal modulation. *Vis Neurosci* 11: 155–163, 1994.
- Lewicki MS and Sejnowski TJ.** Learning overcomplete representations. *Neural Comput* 12: 337–365, 2000.
- Livingstone MS and Hubel DH.** Anatomy and physiology of a color system in the primate visual cortex. *J Neurosci* 4: 309–356, 1984.
- MacLeod DI and Boynton RM.** Chromaticity diagram showing cone excitation by stimuli of equal luminance. *J Opt Soc Am* 69: 1183–1186, 1979.
- Michael CR.** Color vision mechanisms in monkey striate cortex: simple cells with dual opponent-color receptive fields. *J Neurophys* 41: 1233–1249, 1978.
- Olshausen BA and Field DJ.** Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381: 607–609, 1996.
- Parker AJ and Hawken MJ.** Two-dimensional spatial structure of receptive fields in monkey striate cortex. *J Opt Soc Am A* 5: 598–605, 1988.
- Párraga CA, Brelstaff G, Troscianko T, and Moorhead IR.** Color and luminance information in natural scenes. *J Opt Soc Am A* 15: 563–569, 1998.
- Ringach DL.** Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *J Neurophysiol* 88: 455–463, 2002.
- Ringach DL, Shapley RM, and Hawken MJ.** Orientation selectivity in macaque V1: diversity and laminar dependence. *J Neurosci* 22: 5639–5651, 2002.
- Ruderman DL, Cronin TW, and Chiao C.** Statistics of cone responses to natural images: implications for visual coding. *J Opt Soc Am A* 15: 2036–2045, 1998.
- Smith VC and Pokorny J.** Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm. *Vis Res* 15: 161–171, 1975.
- Tailor DR, Finkel LH, and Buchsbaum G.** Color-opponent receptive fields derived from independent component analysis of natural images. *Vis Res* 40: 2671–2676, 2000.
- van Hateren JH and Ruderman DL.** Independent component analysis of natural image sequences yields spatiotemporal filters similar to simple cells in primary visual cortex. *Proc R Soc Lond B Biol Sci* 265: 2315–2320, 1998.
- van Hateren JH and van der Schaaf A.** Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc R Soc Lond B Biol Sci* 265: 359–366, 1998.
- Wachtler T, Lee TW, and Sejnowski TJ.** Chromatic structure of natural scenes. *J Opt Soc Am A* 18: 65–77, 2001.
- Willmore B, Watters PA, and Tolhurst DJ.** A comparison of natural-image-based models of simple-cell coding. *Perception* 29: 1017–1040, 2000.

Detecting topological invariances of visual neuron receptive fields with persistent homology

Abstract

We introduce a new methodology for detecting and identifying arbitrarily complicated invariances in the receptive fields (RFs) of cortical neurons throughout the visual pathway, utilizing techniques of computational algebraic topology to obtain qualitative descriptions of these invariances. The portion of the parameterized stimuli space which causes a neuron to respond forms a manifold. By computing the persistent homology of this manifold we can identify invariances in the RF. To validate this methodology, we present a series of experiments with simulated cells and with real neurophysiological data recorded from the primary visual cortex. The method successfully detects various invariances in the data (corresponding to physically plausible invariances in the early stages of the visual pathway) and is able to handle noisy and sparse data.

Introduction

Visual processing in humans, primates, and cats involves two divergent networks: a dorsal stream that represents the location and movement of objects, and a ventral stream that recognizes the form and identity of objects. Knowledge about the ventral stream is strikingly incomplete. On the one hand, we know a great deal about the responses of the lowest stage (primary visual cortex or V1) to lines and edges, and a considerable amount about representation of faces in the highest stages (inferotemporal areas). On the other hand, relatively little is known about intermediate

visual areas such as V4 and TEO, which are presumed to respond to stimuli of intermediate complexity.

Two properties that are believed to progressively develop over the ventral stream are invariance and feature selectivity: neurons simultaneously increase in their preference for complex stimuli over simple stimuli, and in their tolerance for variation in the positioning of their preferred stimuli. Within V1, a minimal kind of position invariance develops — the initial layer of simple cells prefer an oriented bar at a specific position, while the subsequent layers of complex cells respond to oriented bars in several positions. Much richer invariances are observed to occur in the higher visual cortex, including translation and three-dimensional rotation invariance of entire objects (Wallis & Rolls, 1997). Again, little is known about the development of invariance and feature selectivity through the intermediate stages; for instance, fundamental uncertainties remain about the transition between two and three-dimensional object representations. The purpose of this paper is to introduce a new methodology for studying invariances present in intermediate levels of the visual pathway.

The invariances of the response surfaces of neurons in the ventral stream are reflected in the topological signatures of their response manifolds. We study these by parameterizing the stimulus space and then computing the persistent homology groups (Zomorodian & Carlsson, 2005) associated to the portion of the stimulus space to which a neuron responds. The parameterization of the stimulus space we employ follows a “bottom-up” approach, motivated by biological constraints and our experimental focus

on lower and intermediate visual areas. It is inspired by studies in V1, and in primate V4 (Pasupathy, 2006); however, our approach and algorithms generalize to arbitrary choices of parameterization and are suitable for use at various stages of the visual system.

We will illustrate our method with a simple example. Suppose that we are recording from a vertically oriented complex cell, and present it with bar stimuli of fixed length, width and y position, but varying orientation and x position. The total stimulus space can be parameterized as (x, θ) . The neuron is selective for orientation, but has a translation-invariant response along part of the x axis. Thus, the neuron will respond above threshold only to stimuli of the form (α, θ) for some range of $\alpha_1 \dots \alpha_2$ and $\theta_1 \dots \theta_2$.

We apply a window to the stimulus space, so that only $\alpha \in \alpha_1 \dots \alpha_2$ exist, and then “glue” the edges of the window so that $\alpha_1 = \alpha_2$. This makes the stimulus space cylindrical, and now the subspace of the stimuli space to which the neuron responds is a circle around the cylinder. The presence of this circle precisely reflects the translation-invariance of the RF. The techniques of computational algebraic topology allow us to reconstruct this circle even from noisy and incomplete data. More importantly, by computing the topological signatures of the positive response surface we can identify invariances which are considerably more complicated; for instance, invariances which involve multiple axes in the parameter space simultaneously.

To validate our algorithm, we simulated a series of neurons with specified invariances in their RFs that reflect known invariances in primary visual cortex and higher visual areas,

specifically, translation, rotation, and scale invariances. Our algorithm successfully detected all of these invariances. It was also successful in detecting the translation invariance of real V1 complex cells, and in distinguishing them from simple cells. The experiments with real neurophysiological data indicate that the algorithms are sufficiently robust to cope with the noisy and sparse samples produced by typical experimental regimes. However, the true benefits of a topological approach will likely be manifested when studying the very complicated invariances that are likely to turn up in the later stages of the visual pathway. The parameter spaces that stimulate such neurons are high-dimensional (consisting of multiple independent bars, for example) and as a consequence extremely difficult to visualize. To this end, we studied simulated neurons with complicated RFs and invariances. In this case also, our algorithm was able to detect the invariances.

2 Methods

2.1 Parametric representation of visual receptive fields (RFs)

It is known that neurons in intermediate visual areas integrate numerous V1 inputs. By analogy with V1, which builds RFs similar to Gabor functions by conjoining circularly symmetric RFs, we would expect higher areas such as V4 to conjoin its Gabor-like inputs from V1. Thus, the characteristic parameter space over which its invariant response manifolds are defined is comprised of wavelet-like stimuli, as opposed to the pixelated image space of reverse correlation. We identify the response surface of a neuron within a wavelet-like basis space to explore the topology of these neurons' RFs.

The stimulus parameter space is parameterized by B^n , where B is a single wavelet subunit and n is the number of (unordered) subunits in the stimulus. In bar stimuli, each subunit has 5 parameters: x position, y position, orientation, length, and width. When the parameterized stimuli are rendered as images, bar contrast is set to +1. A Euclidean distance metric is used within the parameter space, in which the dimensions are weighted inversely based on the range; also, the orientation is regarded as an intrinsically circular parameter. In physiological experiments, the spatial extent of a cell's RF is not completely known a priori; this is especially true for neurons with spatial invariance. Therefore, the spatial extent of stimuli was set to roughly 3x the size of the RF, in simulations and experiments.

2.2 Neural responses

In physiological experiments, we recorded from neurons in V1 (area 17) of the cat visual cortex. The neuron's dominant eye was determined using square grating stimuli.

Subsequently, subunit stimuli from B^1 were presented to the dominant eye, with flash duration 200 ms and inter-stimulus interval 100 ms. Spikes were isolated using an analog window discriminator, and recorded from 35 ms after stimulus onset to 35 ms after offset, accounting for visual transmission delays. Between batches of approximately 100-200 stimuli, square gratings were presented, to monitor changes in the neuron's ongoing responsiveness; however no changes were observed.

Our simulations were designed to be roughly equivalent to experimental conditions, except that simulated neuron responses were continuous values rather than discrete

spike counts. Simulated neurons either responded linearly based on the product of its RF with a stimulus, or took the maximum of several linear input neurons. Noise was always added, with variance = mean.

In order to apply the topological algorithms, we binarized the response data, that is, transformed the data from a response surface to a distribution of high-response points. To do this, we thresholded the responses as a fraction of the maximum response, beginning at 0.15 and increasing it until a persistent homology feature was seen or a maximum of 0.5 is reached. For experimental data, the threshold was set to 0.25 of maximum; this corresponded to all responses with 2+ spikes.

2.3 Topology and persistent homology

The techniques of algebraic topology assign algebraic quantities (e.g. numbers or groups) to geometric objects, and study the properties of the objects via the more tractable algebraic data. These associated algebraic structures are insensitive to continuous deformations of the geometric object. This latter property makes this approach seem felicitous for analyzing the qualitative properties of the neuron's receptive fields.

A standard topological invariant of a space X is the collection of homology groups. Given a choice of coefficients R (in our case, the field $Z/2$, i.e. binary integers), for each natural number there is a homology group $H_n(X;R)$. Loosely speaking, the n th homology group measures the number of holes of dimension n in the space X . When R is a field, e.g. rational numbers (Q) or integers modulo p (Z/p), the homology groups are in fact vector

spaces and it is sensible to speak of their rank. The rank of the n th homology group is called the n th Betti number. Generators of the homology groups (called cycles) have a geometric interpretation. For instance, for a torus T the rank of $H_1(T)$, the first homology group of T , is 2. One cycle corresponds to a circle wrapped through the hole of the torus, and the other to a circle that goes around the torus. More generally, in our situation the cycles will correspond to invariant submanifolds of the positive response surface.

Traditionally, homology groups of a space X are computed based on some analytic presentation of X , such as a set of equations. In recent years, there has been interest in the problem of computing homology groups from a set of points which have been sampled from some underlying space. Typically, one hopes to recover the homological information of the underlying space from the computations done using the sampled points.

There are now a variety of algorithms and techniques for computing topological features from such a point cloud. The basic idea is to build a simplicial complex from the points and estimate the topological features of the underlying space from the topological features of the simplicial complex. Recall that a simplicial complex is a set of simplices (a 0-simplex or vertex is a point, a 1-simplex is a line segment, a 2-simplex is a triangle, and so forth) along with instructions for gluing these simplices along edges. To build a simplicial complex from a set of sampled points, roughly one proceeds by choosing a fixed radius ϵ , and assigning points which are “close together” to the same

simplex. It is possible to prove that in the limit of large numbers of points the simplicial complex built from the points faithfully captures the topological information of the underlying space.

In practice, the correct choice of ϵ for a given underlying space and sampling regime is in principle unknowable (and extremely sensitive to noise). To deal with this problem, much effort now focuses on a related invariant called “persistent homology” (Edelsbrunner et al., 2002; Zomorodian & Carlsson, 2005). Here, the idea is to look at how the homology of the simplicial complexes associated to a set of points changes as ϵ is varied. Features which appear and remain in the persistent homology over long intervals of variation of ϵ reflect topological features in the underlying space.

It is important to note that homological data will not completely identify the underlying manifold; in general, no such classification of manifolds exists in dimensions 3 and higher. Instead, persistent homology tells us about submanifolds embedded in the response surface which encode the invariances of the receptive field. In fact, this “limitation” is actually an advantage, as the problem of manifold learning is extremely difficult whereas computing persistent homology is quite tractable.

2.4 Windowing

Neurons are invariant over only a certain range; for instance, a translation-invariant neuron will respond invariantly over only a certain spatial extent. Given our choice of parameterization of the stimulus space, certain invariances naturally give rise to topological equivalence classes; specifically, a fully orientation-invariant neuron will give

rise to a manifold which is circular in the orientation parameter. But most neurons are only invariant over a restricted range of parameter space, such as the translation-invariant cell in the example above. Thus, in order for different invariances to give rise to different topological equivalence classes, we must apply a window in the parameter space, and connect the boundaries of the window (by adjusting the metric on the space). We employed two broad classes of windows: “cylindrical/toroidal” windows, where we identified (“glued”) edges in the parameter space, and “spherical” windows, where we collapsed circles in parameter space to a point. Windows were determined by choosing a subspace of the parameter space in which the window would be applied, and progressively shrinking the windows until features appeared in the persistent homology.

2.5 Technical details of our experiments

We employed “weak witness” complex techniques (Carlsson & de Silva, 2004; Zomorodian & Carlsson, 2005) for the computation of persistent homology of point cloud data, and in particular the Plex 2.0 implementation (Perry et al., 2005). This approach takes as input a set of points in some ambient space M and a distance metric. The distance metric reflected the any windowing that was applied. We first apply a density filter, removing outlying points with k th nearest-neighbor distances in the lowest 10-25% ($k = 25$). A greedy optimization algorithm selects a set of 20-50 “landmark points” as the basis for the complexes to be built. Both of these tunable parameters are initially set to a minimum value, and then increased until persistent homology features are found.

Based on the landmark points, the weak witness complex is built and the persistent homology is computed. As one would expect, the results are sensitive to the choice of the number of landmark points, and so we used rough heuristics based on the number of points and then automatically varied the number of landmark points to find sets which led to significant topological features (as discussed below).

2.6 Invariance detection

In order to interpret the results of the persistent homology computations, we performed control experiments when possible, but we also employed an absolute measure of significance. Associated to a particular choice of landmark points is a “mesh size” m , which reflects the coverage of the space of points by the landmarks (like a measure of sampling density). On theoretical grounds, features in persistent homology which are present at values of ϵ roughly proportional to m are likely to correspond to real underlying topological features. Pragmatically, experience with the Plex suite by a variety of research groups suggests that features which appear at values of $\epsilon > 0.1m$ are likely to be significant.

Next, there is a statistical question of interpreting the significance of the results.

Essentially, there is a problem in nonparametric statistics — to determine to what degree the persistent homology computations based on a particular choice of landmark points reflect the persistent homology of the entire set. Given the nature of the setup, the appropriate procedure appears to be to compute the mean over the “topological significance” scores discussed above over repeated trials (choices of landmark points).

Invariance type was detected based on which windows produced cycles (in the new windowed space), relative to the (unwindowed) controls. It would be extremely useful to be able to extract representatives of these detected cycles for visualization purposes. Although this is currently not possible with the Plex toolkit, recent work (Zomorodian & Carlsson, 2007) is likely to lead to this capability in the near future.

3 Invariance in simulated simple and complex neurons

Initially, we tested whether we could identify known invariances in simulated cells modeling the cells which arise in the first stages of the visual pathway. We compared the topology of a (noninvariant) simple cell and a modestly translation-invariant complex cell (modeled by taking the maximum response over four linear simple cells, arranged in quadrature), using a cylindrical spatial window (Figure 1). Each data set contained 10,000 stimulus-response pairs, with stimuli chosen from a uniform distribution in all parameters. This was a fairly sensitive test, as the width over which the complex cell was invariant was only 3-4x the width of the simple cell's ON subfield.

Denote the length of the longest bar (corresponding to a cycle in the first homology group) by b_1 . In the majority of the complex cell runs, $b_1 > 0.1m$. Similarly, in almost all simple cell runs, $b_1 < 0.1m$. The distributions of b_1 (over 100 runs with randomly chosen landmarks) are easily distinguishable, indicating that the complex cell contains invariance not present in the simple cell. When the spatial window is removed in the control case, the invariance disappears; this confirms that the invariance is translational. In addition, although we cannot currently visualize the actual homology cycles, we can

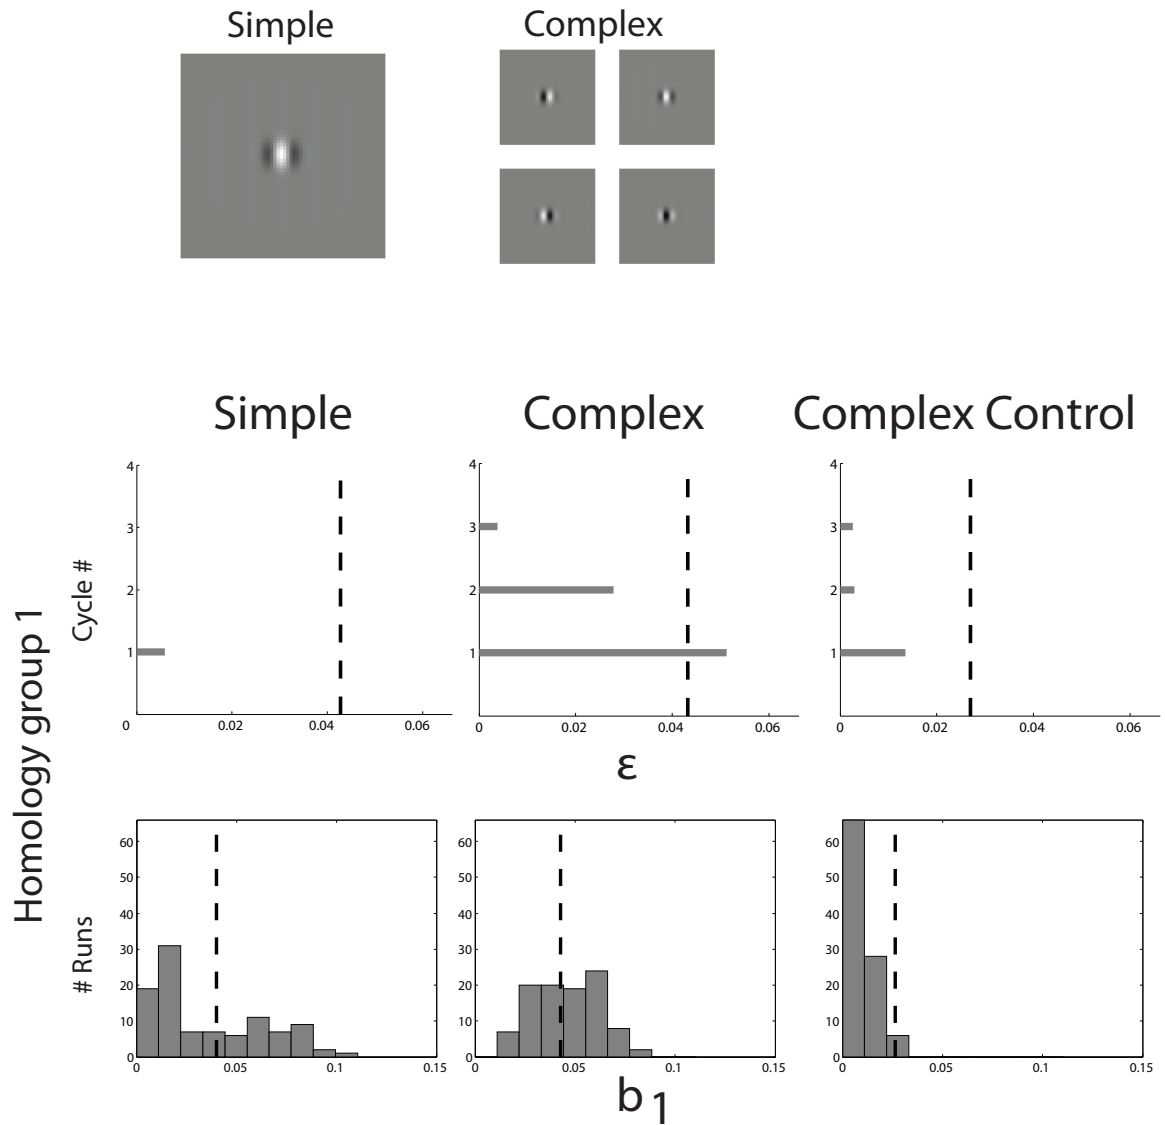


Figure 1: Receptive field components of a simple cell and a complex cell, representative bar code plots (top) and longest-bar (b_1) histograms (bottom). Simulated simple cell with windowing (left), complex cell with windowing (middle) and without (right). Dashed line in all figures indicates $0.1m$, in single run (top) or averaged over all runs (bottom).

confirm translation invariance by inspecting the distribution of landmark points (not shown).

Occasionally, simple cells yielded large b_1 values. This is a problem with data density, as the simple cells are more feature selective than the complex cells and produce sparser data for the same number of uniformly distributed stimuli. In sparse data, occasional spurious circles can arise when building complexes. Note that obtaining the distinction between the cells required averaging over many runs of the persistent homology computation (which takes on the order of a second per iteration). Also note that in windowed runs, the mesh size parameters m and ϵ were larger than in control runs. Windowing tends to eliminate sparse, non-responsive regions of space from consideration, and decreases the distance between many pairs of points near the newly wrapped or collapsed boundary.

4 Invariances in real V1 simple and complex neurons

We tested whether the algorithm was robust enough to work on real data recorded from cat V1, in the presence of physiological noise and with the number of stimuli (1200-1500) limited by recording time and experimental conditions (Figure 2). Stimuli were chosen by an adaptive sampling process; they were initially randomly chosen from a uniform distribution, but as information about the neuron's response accumulated, they were chosen from a distribution biased towards responsive stimuli. Since complex cells are translation invariant over a limited spatial range, we compared the translation invariance of a simple cell (with a classical odd-symmetric linear RF as determined by

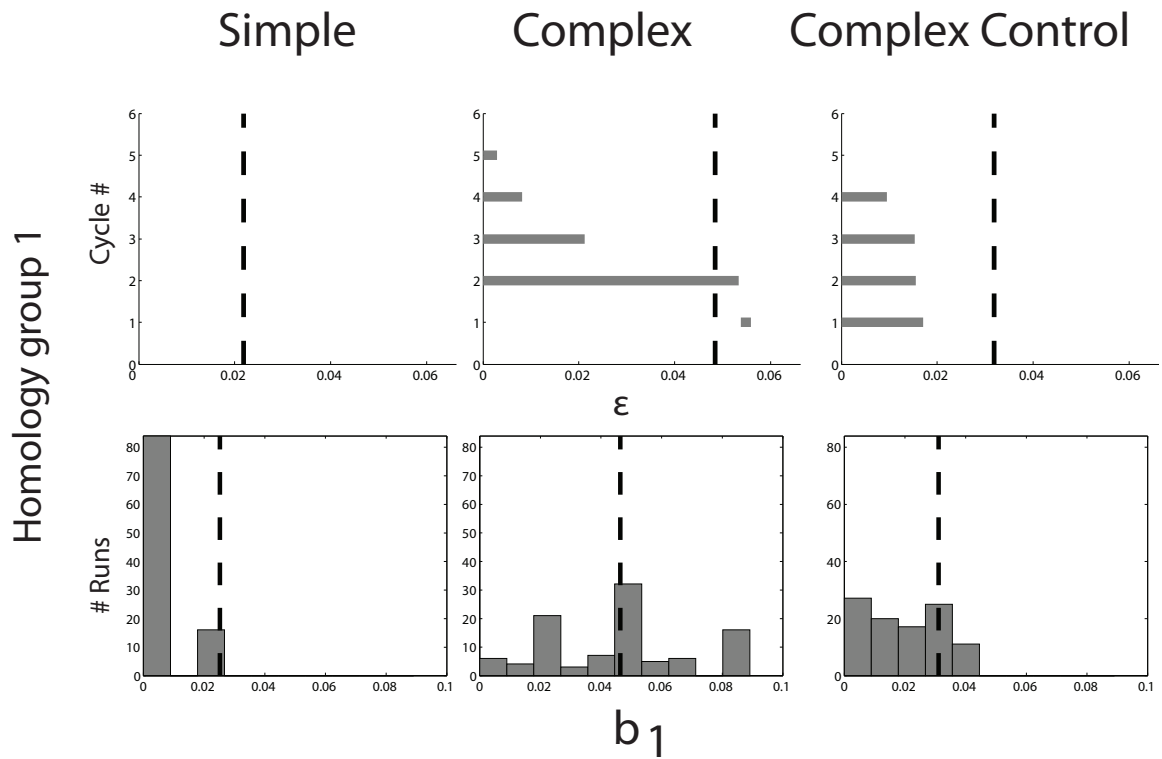


Figure 2: Bar codes (top) and b_1 histograms (bottom) for a real simple cell (left), complex cell with windowing (middle), and complex cell without windowing (right).

noise mapping, and F1/DC modulation ratio = 0.8) and a complex cell (without a linear RF, F1/DC = 0.09).

The complex cell showed a cycle in the first homology group, when windowed along its orthogonal axis. In the control non-windowed case, no such cycles were found above the noise level, indicating that the cycle derives from translation invariance. The simple cell typically showed no cycles for any value of ϵ , indicative of point topology. Overall, this data set improved on the synthetic data set above, mainly because of the efficacy of adaptive stimulus presentation. This is encouraging, since in later stages of the visual pathway adaptive stimulus presentation techniques will be necessary due to the high dimension of the stimulus space.

5 Higher-order invariances

We also tested our algorithm on simulated rotation- and scale-invariant neurons which occur in higher areas of the visual pathway, although they are not present in V1. We modeled these simulated neurons similarly to the complex cell, by taking the maximum response over four (orientation) and twelve (scale) linear simple cells.

Again, we were able to directly detect the invariances in the “bar code” plots (Figure 3). The fully (180°) orientation invariant RF has a natural cycle in the first homology group in the orientation dimension and does not require windowing. The partially (45°) orientation invariant RF has a cycle when windowed, but not in the control case.

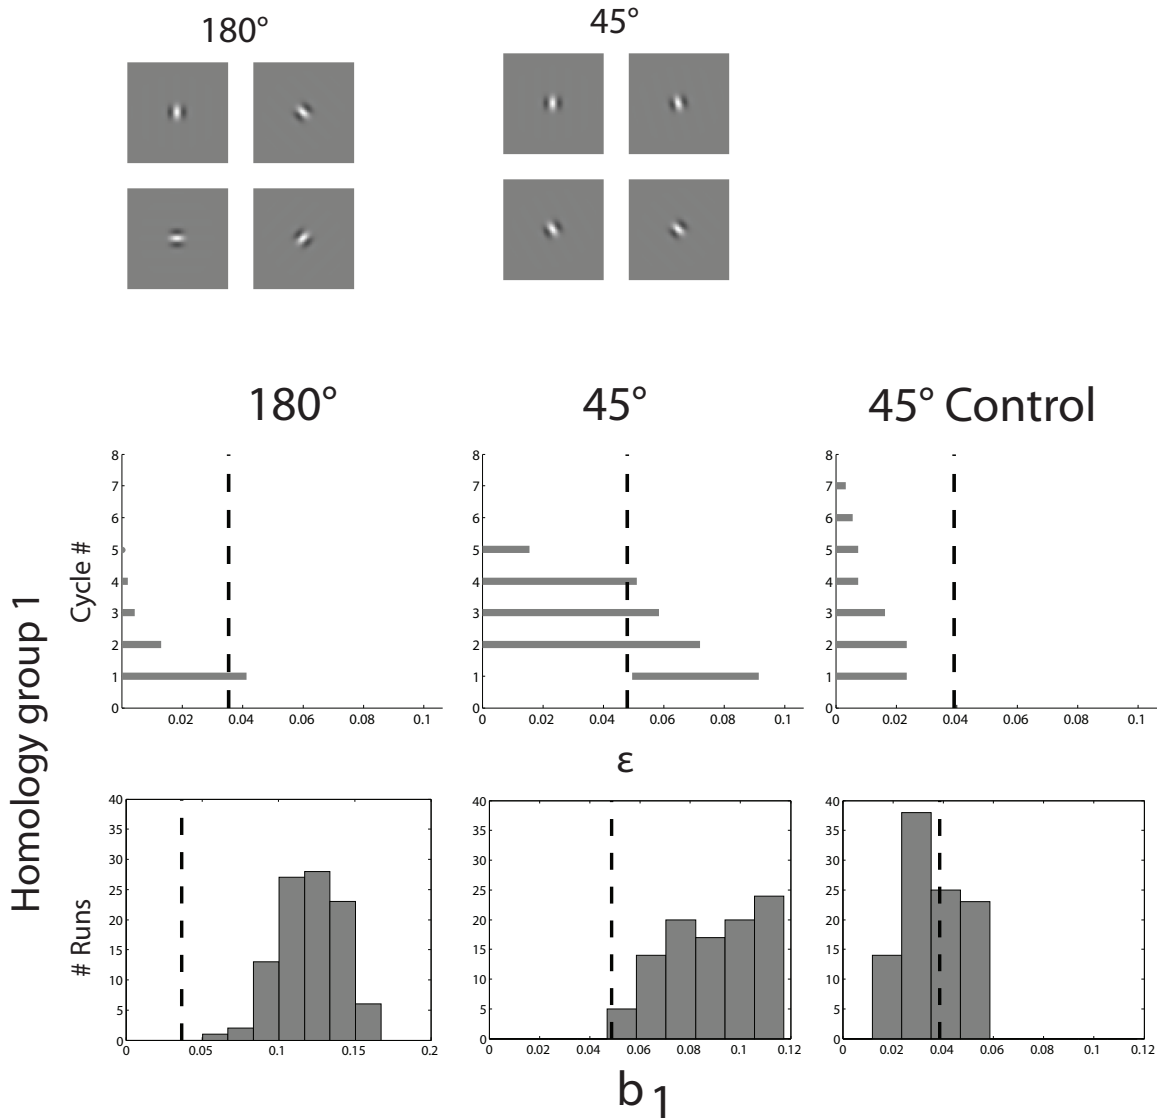


Figure 3: RF components of fully (180°) rotation invariant neuron and partially (45°) orientation invariant neuron. Bar code plots (top) and b_1 histogram (bottom) of fully rotation invariant neuron (left), partially rotation invariant neuron without windowing (middle) and with (right).

The scale invariant RF presents an interesting case (Figure 4). In the first homology group, it has two cycles (corresponding to length and width invariance). However, based on that information alone, these two cycles could be disjoint, or they could indicate the presence of a torus. We can look at the second homology group, where (after projecting onto the “length-width” subspace) we find a long bar at ϵ values overlapping those in the first group. This confirms that the homology is generated by this torus in response space. Thus, we can use windowing and analysis of higher homology groups to identify invariances known to occur in later stages of the visual cortex.

6 Highly feature selective RFs (with simple stimuli)

It is possible to use topological information to help identify the feature selective properties of neuron’s RF. In particular, topological properties can distinguish between neurons which are selective to different vertex and edge geometries, since vertices of different degree will give rise to different numbers of “holes” between edge segments when a spatial boundary around the response region is collapsed to a point. Indeed, we could discriminate between vertices of degree 2 and degree 3 based on the presence of cycles after collapsing a spherical boundary (Figure 5).

This information can be obtained even when more components are present in the RF than the stimulus, provided the nonlinearity of feature selectivity is not too extreme. This can be advantageous because representations with fewer components are lower dimensional, and can thus be explored in a reasonable amount of time, using optimization methods. We may also be able to exploit the fact that as we increase the

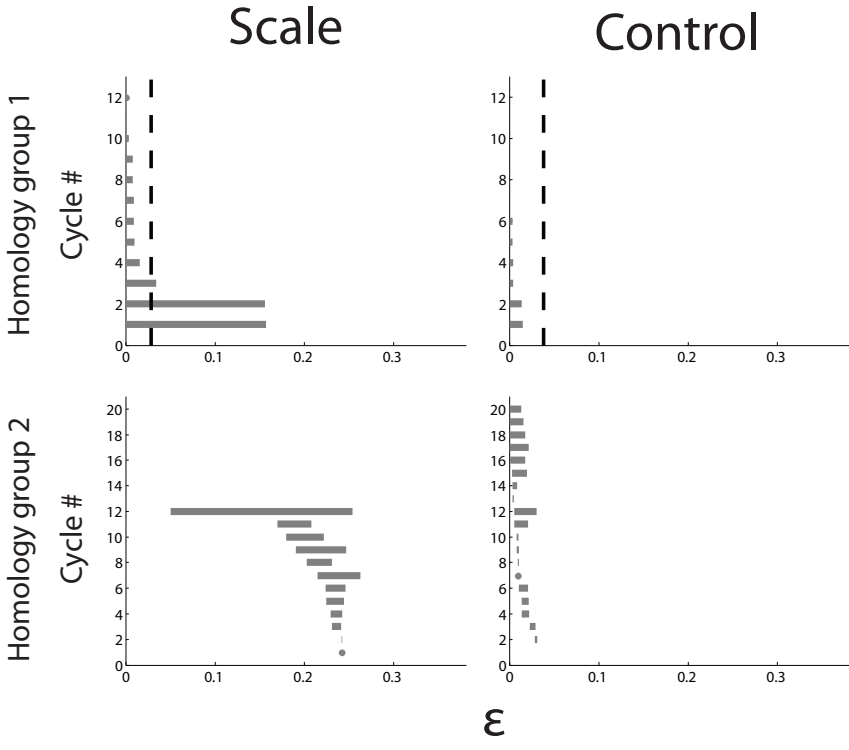



Figure 4: Scale invariant RF components, its bar code plot and b1 histogram (left), and controls (right). Top histogram indicates first Betti number, bottom histogram second.

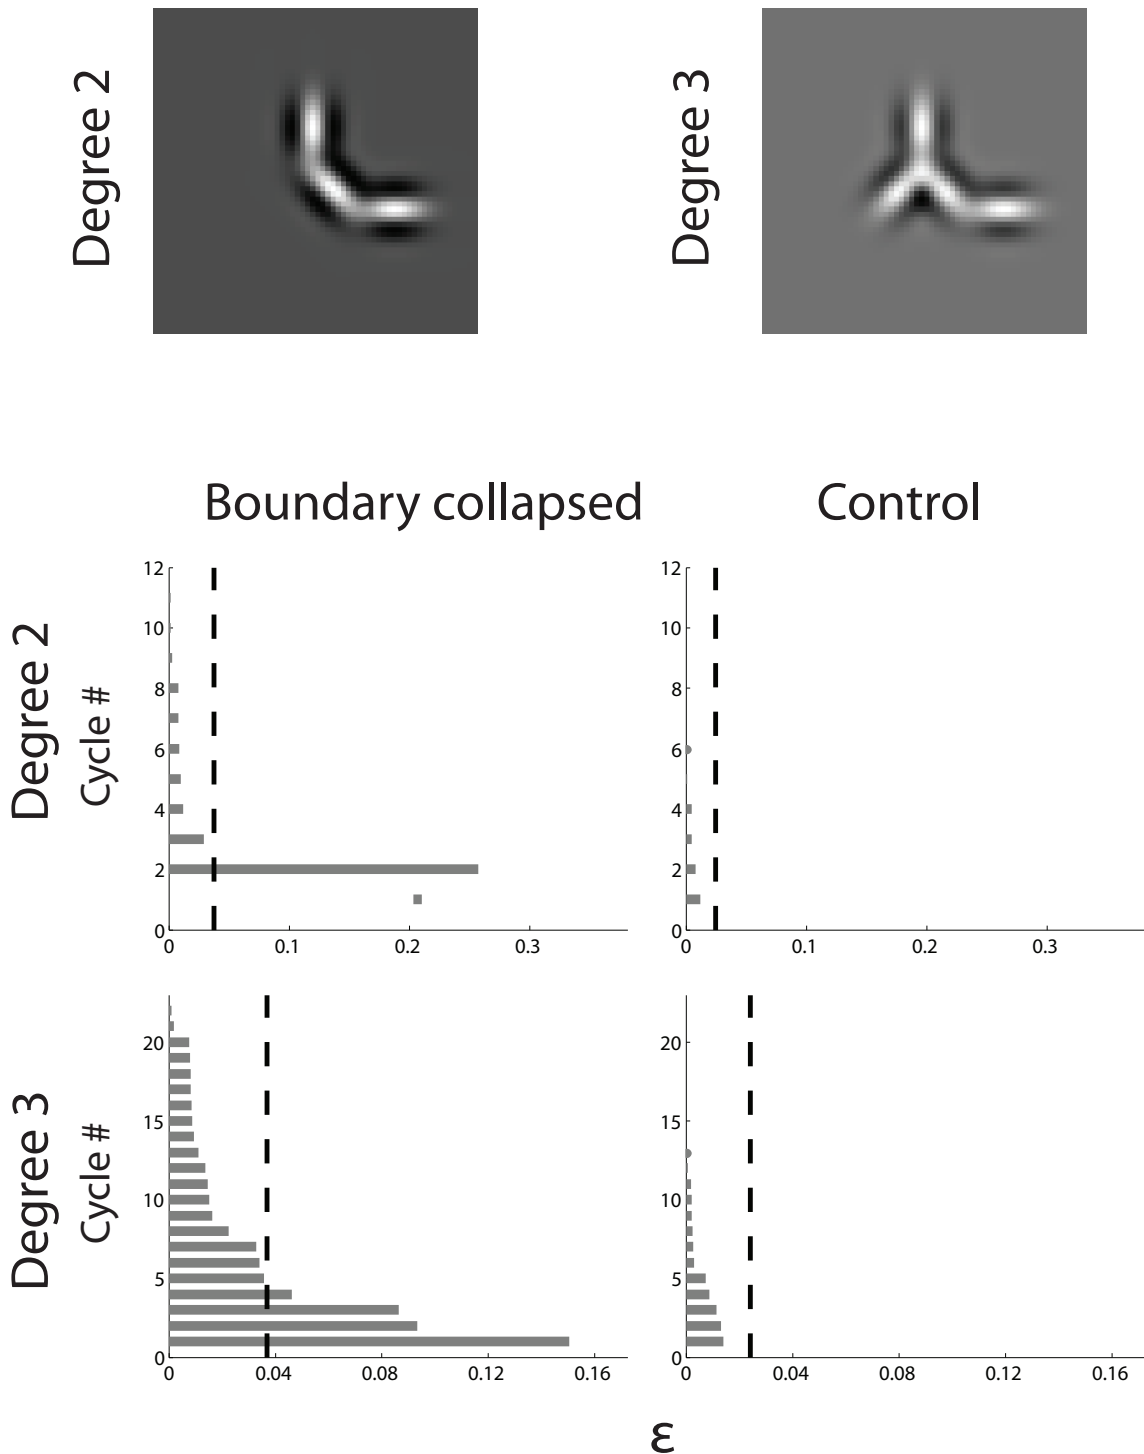


Figure 5: Vertex with degree 2, and its bar code (top) with boundary collapsed to a point, and control. Vertex with degree 3, and its bar code (bottom). All bar codes are for homology group 1. The mean b_1/b_2 (where b_2 is the length of the second-longest bar) ratio over 100 runs for the non-vertex stimulus was 9.41, while for the vertex it was 1.36.

threshold levels for binarization, complicated receptive surfaces will simplify and split into disconnected regions. This kind of filtration is essentially the basis for Morse theory, and so there is a great deal of machinery to assist in analyzing such data. This approach could be used to identify components of a RF selective to multiple features, and, in theory, to say something about features' functional contribution to the RF.

7 Conclusions and future directions

We have demonstrated that it is possible to exploit persistent homology calculations to detect invariances in the receptive fields of neurons. A major strength of the topological approach is that it is not limited to identifying invariances which derive from rigid or even affine transformations. Arbitrarily complicated invariances can be detected using this methodology.

Beyond extension to further stages of the visual cortex and more complicated invariances, there are a variety of possible extensions of this work which are in progress. We assume static spatial response surfaces. However, the method can be adapted to handle temporally varying receptive fields; invariances in object motion, such as might be found in the dorsal stream, can be identified by adding temporal dimensions to the parameter space and employing appropriate windowing. In addition, we expect that this method will be well-suited to "adaptive sampling" techniques for exploring the space of stimuli to which a cell responds. We intend to run experiments in which samples to confirm or falsify hypotheses about candidate invariances detected in persistent

homology, thus allowing us to use the limited time for physiological data acquisition more efficiently.

References

- Carlsson, G., & de Silva, V. (2004). Topological estimation using witness complexes. In *Proceedings of the Eurographics symposium on point-based graphics*.
- Edelsbrunner, H., Letscher, D., & Zomorodian, A. (2002). Topological persistence and simplification. *Discrete and computational geometry*, 28:511–533.
- Pasupathy, A. (2006). Neural basis of shape representation in the primate brain. *Progress in Brain Research*, 154, 293-313.
- Perry, P., de Silva, V., and Zomorodian, A. (2005). PLEX package. Release 2.0.
- Wallis, G., & Rolls, E. T. (1997). Invariant face and object recognition in the visual system. *Progress in neurobiology*, 51:167–194, 1997.
- Zomorodian, A., & Carlsson, G. (2005). Computing persistent homology. *Discrete and computational geometry*, 33(2):247–274.
- Zomorodian, A. & Carlsson, G. (2007). Localized homology. *Shape Modeling International, Lyon, France*.

Functional maps of cat visual area 21a reveal a role in processing spatial form

Abstract

Visual area 21a of the cat is an extrastriate visual area with known physiological properties but uncertain topography and homology to primate visual areas. Using optical imaging methods, we were able to locate and characterize the area in multiple subjects. We found it to be topographically organized for orientation and spatial frequency. Its orientation columns were progressively larger than area 18 and area 17. Its spatial frequency, temporal frequency, and velocity preferences were similar to area 17 and distinct from area 18, suggestive of a close functional connection and a role in spatial vision as opposed to motion perception, similar to primate area V4.

Introduction

The cat has long been one of the standard model systems for studying vision because of its high visual acuity, visually dominated behavior, relatively smooth visual cortex, small size and tractability. These properties have enabled many seminal discoveries about the visual cortex, especially primary visual cortex (Hubel and Wiesel, 2005). Comprehensive atlases of the visual areas of the cat's cortex have been developed (Tusa et al., 1978, 1979; Tusa & Palmer, 1980; Heath & Jones, 1971).

However, progress in understanding the cat's visual brain has been impeded by uncertainty about the integrity and function of the higher cortical areas in the cat visual

pathway and the correspondence between cat and primate visual areas. Unlike the primate, the cat has two primary visual areas, area 17 and area 18, which both receive their major input from the lateral geniculate nucleus of the thalamus. Area 18 appears specialized for motion processing; its neurons have larger receptive fields than equivalent neurons in area 17, and they prefer faster moving stimuli (Payne & Peters, 2002).

The primary visual cortex notwithstanding, recent studies using reversible inactivation of cortex have shown that the functional distinction between dorsal and ventral visual streams, corresponding to spatial form and motion processing, is as clear in the cat (Lomber et al., 1996) as in the primate, where the dissociation was first identified (Ungerleider & Mishkin, 1982). The cat ventral stream includes areas 21a, 21b, 20a, and 20b, while the dorsal stream includes area LS and anterior areas of the middle suprasylvian sulcus (Payne, 1993).

Visual area 21a sits at a crucial junction for spatial form vision. It is the lowest visual area in the ventral stream that receives no afferents from the LGN. Its dominant input is thought to be from X retinal ganglion cells via the supragranular layers of area 17 and, to a lesser degree, area 18 (Dreher et al., 1993, Michalski et al., 1993). Compared to area 17, single neurons in area 21a are reported to be less motion selective, less direction selective, and equally or slightly more orientation selective (Toyama et al., 1994; Dreher et al., 1993). Area 21a's specialization for form vision is consistent with its

representation of the central upper visual hemifield, similar to primate V4 (Tusa & Palmer, 1980).

Understanding the functional architecture of area 21a would allow us to better understand its role in the visual pathway of the cat, and more generally how the cortex transforms visual signals between two levels of a visual hierarchy. However, little has been reported about the functional organization of area 21a or any non-primary visual area in the cat. One study has measured the spacing of orientation columns observed using episodic imaging in area 21a (Huang et al., 2006), but it used few stimuli, did not control for individual variability, and did not characterize the global organization of the orientation map. While some neurons in area 21a have been shown to respond better to one eye than another, the presence of an ocular dominance (OD) map has not been determined. In addition, there is no information available about whether spatial frequency is organized in columns, as it is in areas 17 and 18.

One persistent difficulty with studying area 21a has been the difficulty of locating its boundaries in individual animals. It has traditionally been identified by its histology, although some studies have shown it can be identified with laborious systematic electrophysiological mapping (Toyama et al., 1994). However, techniques such as optical imaging can rapidly define the extent of a visually responsive area on the exposed surface of the cortex, facilitating the characterization of areas with variable location. The speed of optical imaging also allows for comparisons of different areas within single animals, removing a major source of variability in the study of functional architecture.

Methods

Surgery

The surgical preparation and maintenance procedures are similar to those described previously (Dadvand et al., 2006; Kalatsky & Stryker, 2003). All experimental procedures were approved by the UCSF Institutional Animal Care and Use Committee.

Cats were initially sedated with intramuscular ketamine (25 mg/kg) and midazolam (0.07 mg/kg) and then anesthetized with the inhaled anesthetic isoflurane (3%-4% in O₂). After the placement of a femoral catheter, barbiturate anesthesia was substituted (sodium pentobarbital, 2.5 mg/kg bolus). Glycopyrrolate (0.02 mg/kg) was given intramuscularly to reduce tracheal secretions, antibiotic (cefazolin, 22 mg/kg) was given intravenously to prevent infection, and dexamethasone (2 mg/kg) was given intravenously to prevent edema. Core temperature was maintained at 37.5° C using a feedback-regulated heating pad. A tracheotomy was performed, and the animal was placed in a stereotaxic apparatus and ventilated with 67% N₂O, 33% O₂. Eyedrops containing 1% ophthalmic atropine sulfate and 10% phenylephrine hydrochloride were applied to dilate the pupils and retract the nictitating membrane. Contact lenses of the appropriate strength were fitted with the aid of a streak retinoscope.

The level of anesthesia was continually monitored with core temperature, electrocardiogram heart rate, pulse oxygen saturation, expired CO₂, and peak airway

pressure. The absence of reflexes was regularly checked, and pentobarbital was administered as necessary to maintain a surgical plane of anesthesia.

A craniotomy was made over one hemisphere to simultaneously expose areas 17, 18 and 21a. The typical extent was from 5 mm anterior to 7 mm posterior, midline to 15 mm lateral, relative to stereotaxic zero.

To prevent eye movements, neuromuscular blockade was induced by continuous infusion of pancuronium bromide (0.055 mg/kg/hr) mixed with lactated Ringer's solution with 2.5% dextrose (total volume of fluid infused was 5 ml/kg/hr). The animal was maintained on ventilation for the duration of the experiment. For better control over anesthesia, EEG wires were placed between the occipital and temporal edges of the craniotomy. The optic disks were mapped with an ophthalmoscope, and the center of the visual field was determined. Finally the dura mater was reflected, and agarose (3% in saline) at body temperature was placed over the exposed cortex, and sealed with a glass cover slip. The EEG was maintained in slow-wave or mixed slow-wave/fast activity for the duration of the experiment.

Imaging Procedure

Optical images of the cortical intrinsic signal were obtained using a Dalsa 1M30P CCD camera controlled by custom software (Kalatsky & Stryker, 2003). Using a tandem lens "macroscope" (Nikon) in two different configurations (50 mm over 50 mm, 85 mm over 50 mm) we could image areas of 12 mm x 12 mm or 7.2 mm x 7.2 mm. The surface vascular pattern was recorded under green (546 ± 10 nm) illumination, and the intrinsic

signal was recorded under illumination with either broad spectrum light filtered at 610 ± 10 nm or LED illumination (Luxeon LXHL-NH94, 617nm), and a red filter (610 ± 10 nm) was inserted between the camera lenses. The focal plane was 0.6 mm below the vasculature on the cortical surface. During imaging, frames were acquired at 30 fps, 1024 x 1024 pixels and were binned 2x2 spatially and by 4 frames temporally for a final image size of 512 x 512 pixels, final sampling rate of 7.5 Hz, and 16 bit well depth.

Imaging stimuli

Optical imaging stimuli were generated by custom software and displayed on a high refresh rate monitor (Nokia Multigraph 445X; 1024 x 768, 120 Hz) with screen size 40 cm x 30 cm, at a distance of 40cm from the eyes. Stimuli were presented at 100% contrast (gamma-corrected), to the eye contralateral to the recording hemisphere (unless otherwise specified).

The primary stimulus employed was the “drifting rotating grating,” a square- or sine-wave drifting grating that fully rotates around the center of the visual field with a period of 30 s, presented continuously for 6 or 8 minutes. To find responsive visual areas, compute ocular dominance, and orientation and direction maps, we used a square-wave grating with spatial frequency=0.2 cycles/° and temporal frequency=2.0 cycles/s (velocity 10 °/s). Sinusoidal gratings at this spatial frequency produce a strong intrinsic signal response in areas 17 and 18 (Issa et al., 2000), and a square-wave grating includes higher spatial frequencies and should stimulate neurons preferring a wide range of spatial frequencies. For spatial frequency and temporal frequency maps, we used

sinusoidal gratings and held one parameter constant while we varied the other. In experiments using spatial frequencies > 0.4 , artificial pupils were always placed in front of the animal's eyes to better ensure that stimuli were in focus.

Imaging analysis

We fit the response of each individual pixel with a sinusoid of the same period as the stimulus to obtain a robust estimate of the best phase and amplitude (Kalatsky & Stryker, 2003). These phases and amplitudes formed the direction map, and the phases and amplitudes fit using a sinusoid with half the period formed the orientation map. To eliminate artifacts from the orientation map, we centered each phase-amplitude distribution to a vector mean of zero. This corresponds to the assumption that all phases are represented with equal strength over the map, which is accurate for direction and orientation maps of any reasonable size. To reduce noise, maps were usually smoothed, after averaging or combining, with a Gaussian kernel with $\sigma = 100\mu\text{m}$.

Each map was demarcated by a manually stenciled template, with borders based on the orientation phase and amplitude maps, as well as vessels in the surface image. Large vessels with a large solid area of hemodynamic artifact were excluded from the template, but small vessels (which did not disrupt the orientation maps) were not. The border between area 17 and 18 was identified based on the different amplitudes of the two areas' maps at different spatial frequencies; typically area 18 was stronger at 0.05 cycles/ $^{\circ}$, while area 17 was stronger at 0.2 or 0.4 cycles/ $^{\circ}$ (Issa et al., 2000).

Delay correction of imaging maps

If the intrinsic signal response were instantaneous, every pixel's response peak phase would correspond to one drift direction (and orientation), and this would define the orientation map. However, the intrinsic signal has a variable hemodynamic delay, which must be removed or it will bias the orientation estimate.

The standard procedure for correcting the delay is to run experiments rotating the grating in both clockwise and counterclockwise directions, in separate acquisition periods. By adding the two angle estimates together, we obtain an estimate of the doubled delay angle $2d$ (modulo 2π). Previous studies (Kalatsky & Stryker, 2003) have divided by 2 to obtain the delay angle d (modulo π) for each pixel, but this is a noisy computation even after local smoothing and is inherently ambiguous. When the stimulus period is 15 s, for example, the delay angle will be ambiguous ± 7.5 s. If the signal is locking to the late bloodflow response, the natural distribution of delays may be approximately Gaussian, centered at 9 s with an SD of 2 s. Using the older technique, delay angle estimates will frequently be off by π , causing errors of π radians in our estimation of the preferred phase.

Instead, we compute a histogram of all responsive pixels in the visual area, and for each pixel we take d_{target} = whichever of d or $d+\pi$ shows a peak around 9 s, corresponding to the standard hemodynamic response peak. Taking the distribution of d_{target} over the region of interest, we calculate the mean and use it as the delay angle to rotate each of

the clockwise and counterclockwise maps. Finally the two maps are vector averaged together to produce a best estimate.

In some weaker maps, instead of tracking the peak, the sinusoid fit will lock to the initial hemodynamic dip (~ 2 s) instead of the repolarization phase (~ 9 s; Kalatsky, personal communication), and the original delay angle estimate d will be correct. We can detect these maps by orientation-orientation scatterplot comparisons to other experiments in the same area, because instead of having identical orientations, one of the maps will be offset by π overall. Once this is detected, the map's delay angle can be corrected accordingly. This correction is important for analyses that depend exclusively on the signal phase, such as when comparing intrinsic signal imaging data with electrophysiological data; for analyses that depend only on the signal amplitude, such as ocular dominance maps, the correction has no effect.

Map similarity

The similarity of orientation maps was quantified using the canonical correlation (Mardia & Jupp, 1999), a measurement of the correlation between two circular variables based on the correlation of their direction sines and cosines. Because we were interested in the structure of the map rather than its strength, only phases, not amplitudes, were used in this comparison. In simulations, canonical correlation has a nearly inversely linear relationship with the amplitude of vector noise added to a map (not shown), which makes it ideal for comparing continuous intrinsic signal maps.

Ocular Dominance maps

We mapped the ocular dominance of each pixel by comparing its amplitudes when stimulated through the contralateral and ipsilateral eyes. Although the contralateral and ipsilateral maps were acquired in sequential runs, the map amplitude fluctuated from run to run. Making the assumption that high amplitude pixels (e.g. the centers of OD columns) should have similar strengths in contralateral and ipsilateral maps (Cheng et al., 2001) we multiplicatively rescaled amplitude values in each map (excluding the top 5%) to the range 0-1, and set the top 5% to 1. We then computed an ocular dominance index for each point, $ODI = (Amplitude_{Contra} - Amplitude_{Ipsi}) / (Amplitude_{Contra} + Amplitude_{Ipsi})$.

Spatial and Temporal Frequency maps

Since the drifting rotating grating yields strong orientation maps in areas 17, 18 and 21a, we used it to obtain spatial and temporal frequency maps for each area. We ran orientation maps at logarithmically spaced spatial frequencies $s = [0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1.2, 1.6]$ cycles/°, in randomized order. At each pixel, the best-fit phase p_s and amplitude a_s were taken to be vectors v_s , and averaged to find the consensus phase for that pixel c .

From the rectified dot product of the consensus phase and the phase of a single spatial frequency map, we obtained a measurement of map quality for each pixel at that spatial frequency, $w_s = \langle p_s, c \rangle$. This reflects the fact that, while amplitudes were often variable between maps and only somewhat predictive of quality, good maps at near-optimal

spatial frequencies all showed similar best orientations, while maps at suboptimal spatial frequencies were noisy and orientation was more variable. These weights formed a spatial frequency profile, $w = [w_1 \dots w_n]$. This profile can be thought of as a (noisy) spatial frequency tuning curve for the intrinsic signal at a given pixel. To obtain a single best spatial frequency, we took the dot product $d_i = \langle w, f_i \rangle$ with a series of filters f_i , which corresponded to a tuning curve for a different peak spatial frequency. These 40 filters were Gaussians with constant maximum value, linearly spaced over the range of spatial frequencies, with tuning width of $\sigma = 0.6$ octaves. The peak spatial frequency yielding the maximal d_i was taken as the spatial frequency for that pixel.

Temporal frequency maps were constructed similarly to spatial frequency maps. Instead of varying the spatial frequency of the drifting rotating grating, the drift temporal frequency was changed between runs, to values [0.5, 1, 2, 3, 4] cycles/s. Since we had fewer samples in the temporal frequency spectrum than the spatial frequency spectrum, we weighted the maps by their amplitude without utilizing phase consistency; this method minimized error in simulated data (not shown).

Pinwheels

Pinwheels in orientation maps were identified as in Löwel et al. (1998). The polar phase-amplitude vector was transformed into Cartesian coordinates z_x and z_y , and the contour lines $z_x = 0$ and $z_y = 0$ were plotted on top of the orientation map. Points where those lines intersected were visually inspected, and those that showed clear evidence of radially organized orientation domains were included in the analysis. If there were

multiple intersections within 0.2 mm of each other, only the location showing clearest radial organization was analyzed.

Column spacing

We quantified the spacing of orientation columns using a wavelet fitting technique (Kaschube et al., 2000). Maps were scaled from 0 to 1 and filtered to remove wavelengths > 3.0 cycles/mm. At each pixel in the map, wavelets of 8 different wavelet orientations and 8 wavelengths logarithmically spaced between 0.4 and 3.0 cycles/mm were tested, and the best wavelength was taken from the wavelet that produced the maximum response regardless of wavelet orientation.

Because orientation is a circular, continuous variable, we discretized orientation into 8 sub-maps, each of which represented a single orientation preference, and the same wavelet analysis was applied to each map. For example, the 22.5° map contained all points with orientation preference 0° - 45° , cosine weighted by similarity so that points with orientation 22.5° had value 1 and points with orientation 0° had value 0. The best wavelength for all points with orientation preferences 11.25° - 33.75° was set to the wavelength from the 22.5° map, and the same was done for other orientation preferences.

The column spacing for each area was computed as the mean over all pixels in the area. The variability of spacing was estimated by calculating a standard error, σ/\sqrt{N} , using the standard deviation of all pixels and estimating N, the number of independent columnar units in the area, as $(\text{total area} / \pi r^2)$ where r is half the estimated spacing.

Physiological Recording

After imaging, in some experiments we replaced the solid cover slip over the chamber with one with a small (~2 mm) circular aperture. Using a tungsten microelectrode (Microprobe, impedance ~2 M Ω) we targeted random sites within the strongly responsive areas of visual cortex. Penetrations were made at the same angle as the camera, relative to the brain, generally within 10 degrees of the surface normal.

Stimuli were generated in Matlab using the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997). Orientation stimuli were sinusoidal drifting gratings (8 different orientations, spatial frequency 0.2 cycles/ $^{\circ}$, temporal frequency 2.0 cycles/s, 100% contrast), presented to the contralateral eye. At the most responsive multi-unit orientation, a spatial frequency tuning curve was constructed by varying the spatial frequency (0.05 to 1 cycles/ $^{\circ}$ in 8 log-spaced steps).

Extracellular signals were acquired using a System 3 workstation (Tucker-Davis Technologies), filtered from 0.7 to 7 kHz and sampled at 25 kHz. Spiking events were detected on-line by voltage threshold crossing, and a 1 ms waveform sample was acquired around the time of threshold crossing. Spike sorting was performed as in Niell and Stryker (2008), and the orientation and spatial frequency tuning curves for each unit were constructed from the mean firing rate during stimulation. Orientation selectivity Index (OSI) was computed from the orientation tuning response curves as $(r_{\text{preferred}} - r_{\text{null}}) / (r_{\text{preferred}} + r_{\text{null}})$. Direction Selectivity Index (DSI) was computed as $(r_{\text{preferred}} - r_{\text{anti-preferred}}) / (r_{\text{preferred}} + r_{\text{anti-preferred}})$.

Results

Location of area 21a in optical maps

Figure 1a shows a representative example of three smoothed optical maps of the intrinsic signal response to stimulus orientation that were simultaneously acquired in areas 17, 18, and 21a of one animal. Imaging higher visual areas under barbiturate anesthesia is more difficult than imaging early visual areas 17 and 18. However, in animals where we obtained any orientation maps in area 17 or area 18, we were also able to obtain such maps in area 21a in 13 of 14 cases. Area 21a is apparently early enough in the visual pathway that it can still reliably yield optical maps under anesthesia.

In hemispheres where areas 17 and 18 failed to yield orientation maps, we never successfully imaged area 21a. Since these are the input areas to 21a, their good physiological condition appears to be a prerequisite for normal function in 21a. Ideally, we would like to simultaneously image all three areas in the same hemisphere.

However, because of brain curvature, the limited depth of field of the macroscope, and the sensitivity of optical imaging to angle of illumination, it can be difficult to find an imaging window and camera position permitting simultaneous imaging of all three areas, especially in larger brains. We attempted simultaneous imaging in 8 experiments, and succeeded in 6 of those.

We observed that the position of area 21a varied between animals, along both the lateral-medial and the anterior-posterior axes. Of these two axes, lateral-medial variability is more problematic for the experimenter because of the presence of the lateral suprasylvian sulcus; if a 21a map were positioned very laterally it would extend into the sulcus, making it challenging to image and more difficult to map electrophysiologically. We observed at least 3 mm of lateral-medial variability. The mappable area of 21a was always adjacent to the lateral suprasylvian sulcus. In some cases, it covered most of the gyrus, extending 3/4 of the way to the lateral sulcus; in others, it extended less than halfway. In two animals, relatively little of the map was exposed on the gyrus, but there were indications the map extended down into the lateral suprasylvian sulcus. In stereotaxic coordinates, the medial border was L6-L10, and the center of the visible map fell between L9-L12. The lateral border could not be precisely ascertained since it was in the lateral suprasylvian sulcus and the sulcus is located on a steeply slanting part of the brain not well described by stereotaxic coordinates, but the position of the sulcus varied from approximately L12-L14.

Anterior-posterior variability poses a less significant difficulty for imaging and microelectrode recordings than does lateral-medial variability, because area 21a's position along the surface of the gyrus has little effect on its accessibility. Anterior-posterior variability was of the same magnitude as lateral-medial variability. The center of the map was located 4 ± 2 mm anterior to the bend at the posterior end of the lateral suprasylvian sulcus; relative position seemed a more reliable landmark for the center of the area than stereotaxic coordinates.

We did not observe any adjacent maps either medially, on the medial lip of the lateral suprasylvian gyrus, or posteriorly, at the bend of the gyrus. Both of these locations are predicted to be the central visual field representation of area 19 (Tusa et al., 1979). We also did not observe maps to the anterior in parietal area 7, which is usually reported to be unresponsive under anesthesia (Payne, 1993).

Size of area 21a in optical maps

The size of the exposed map varied from 3.5 mm² to 26.6 mm², and the distribution of map sizes was bimodal, with modes at ~8 mm² and ~22 mm². The smaller mode probably resulted from our experimental procedure; when we were unable to see a map high on the surface of the gyrus, we would try to align the camera to focus down into the sulcus, yielding a number of small maps.

Our full-screen stimulus extended to approximately 25° along the horizontal axis of the visual field and 20° along the vertical axis; since area 21a only represents the central visual field, this stimulus encompasses the full extent its visual field according to Tusa and Palmer (1980)'s standardized map. Since area 21a is organized retinotopically, our stimulus should have activated nearly the entire area. The surface area of 21a is 25±2 mm², estimated from Hilgetag and Grant (2000). Thus, the largest maps we found represented almost the entirety of area 21a; the smallest maps invariably bordered (and presumably extended down into) the sulcus, thus representing only a small fraction of area 21a.

Corresponding visual field locations were imaged in areas 17, 18, and 21a, and the position of the monitor was never changed between imaging of the different areas. The imaged region of area 17 represents a narrow strip of the contralateral visual field along the vertical meridian from about 5° above the horizontal meridian to 10° below (Tusa et al., 1978), while the imaged region of area 18 represented the lower contralateral visual field, from 0° to -10°, extending out to 5-10° along the horizontal meridian (Tusa et al., 1979). Similarly, the exposed region of area 21a represents a quadrant-shaped region of the contralateral visual field from 25° above the horizontal meridian to about 5° below, extending out to 20° from the vertical meridian (Tusa & Palmer, 1980). Because the cortical magnification factor enlarges the representation of the central visual field, the majority of the imaged regions in all three areas fell within 5° of the center of the visual field. Functionally, inactivation or stimulation of the exposed area 21a map affects the response in the exposed areas 17 and 18 maps (Huang et al., 2004), confirming that the retinotopic correspondence between exposed areas is quite close.

The exposed portion of area 21a is frequently portrayed as a wedge pointed towards the lateral sulcus, based on the atlas of Tusa and Palmer (1980). We rarely observed this shape; the visible map was generally rectangular. In a minority of maps the anterior and posterior borders of area 21a were angled so as to converge medially. In cases where a large part of the map dove into the sulcus, the overall shape of the area may have been occluded, and some of the map may lie posterior to the visible portion.

Orientation columns in area 21a are larger than in areas 17 and 18.

The spacing of orientation columns appears to increase from area 17 to area 18, and from area 18 to area 21a (Figure 1a). We quantified this trend by measuring the spacing between orientation columns preferring identical orientations, using wavelet analysis to determine the period (local wavelength) of the orientation map. Figure 1b shows the trend of increasing spacing across areas, plotted across all animals for which we have multiple maps. On average, column spacing was 5% larger in area 18 than area 17 and 18% larger in area 21a than area 18. Note that, while all individual animals show an increasing trend from 17 to 18 to 21a, the column spacing varies significantly between animals, by as much as 25% (~0.4 mm).

A paired comparison (Figure 1c) eliminates the effect of inter-animal variability and shows that the column spacing in 21a is larger than in area 17 (N=5, sign test $p < 0.06$), and the spacing in area 21a was not quite significantly larger than area 18 (N=6, $p < 0.11$). The spacing in 18 is larger than in area 17 (n=7, $p < 0.01$), as has been previously observed (e.g. Löwel & Singer, 1987).

In some animals, the column spacing in area 21a seems to be much larger than area 18; in others, it was only somewhat larger. This variability was real, and was not explained by the correlations in column spacing within single animals. Using the wavelet analysis, we estimated the local column spacing at all pixels in all area 21a maps, and analyzed the variability of this measure. Column spacing was normalized so that, in each animal, area 18 corresponded to 1.0. Figure 1d shows that, while the column spacing in area 17

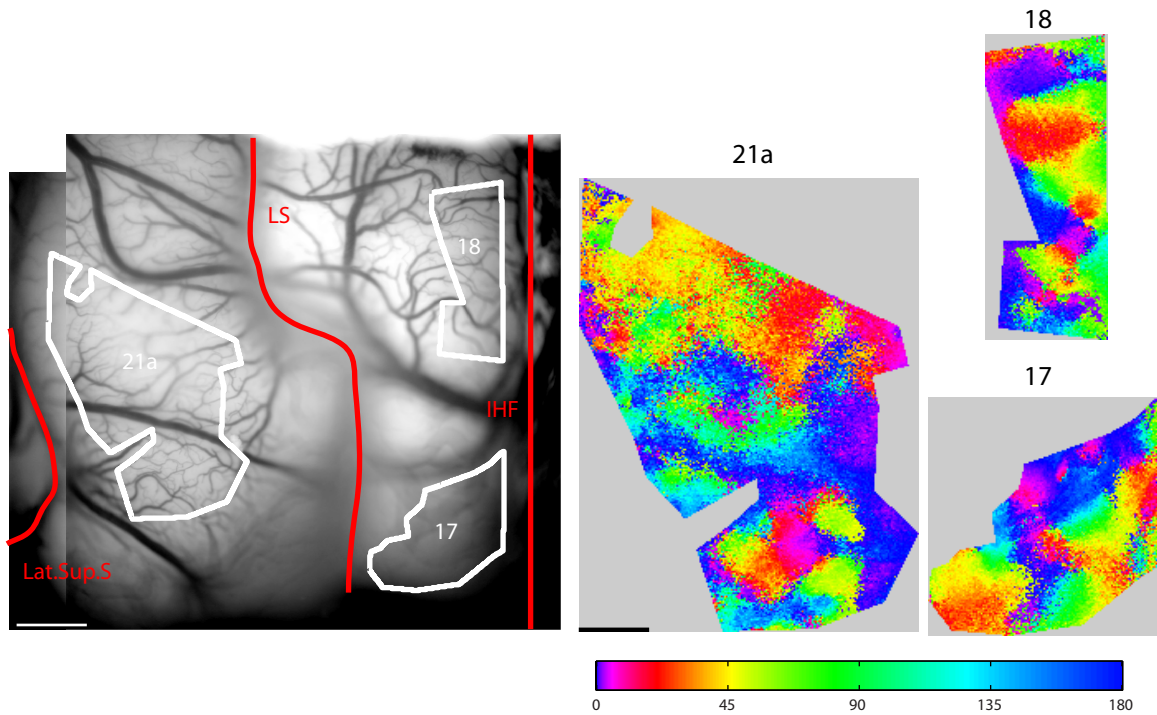


Figure 1a. Orientation maps of area 17, 18, and 21a acquired simultaneously in a single cortical hemisphere.
 (Left) Imaging exposure with sulci labeled and imaged areas marked. Lat. Sup. S = Lateral suprasylvian sulcus; LS = Lateral sulcus; IHF = Interhemispheric fissure.
 (Right) Orientation maps, with best orientation shown by key. Scale bars = 1 mm.




Figure 1b. Trend in orientation column spacing across areas 17, 18, and 21a. For each area, overall spacing is the maximum of a Gaussian fit to the distribution of wavelet periods over all pixels. Single animals are joined by lines (only one hemisphere shown per animal) and horizontally staggered for clarity. Error bars are SEM for $n =$ number of column units in each map.




Figure 1c. Paired comparison of orientation column spacing across areas. Diagonal line indicates equal period. Each animal is shown as a point; error bars are SEM as in Figure 1b.

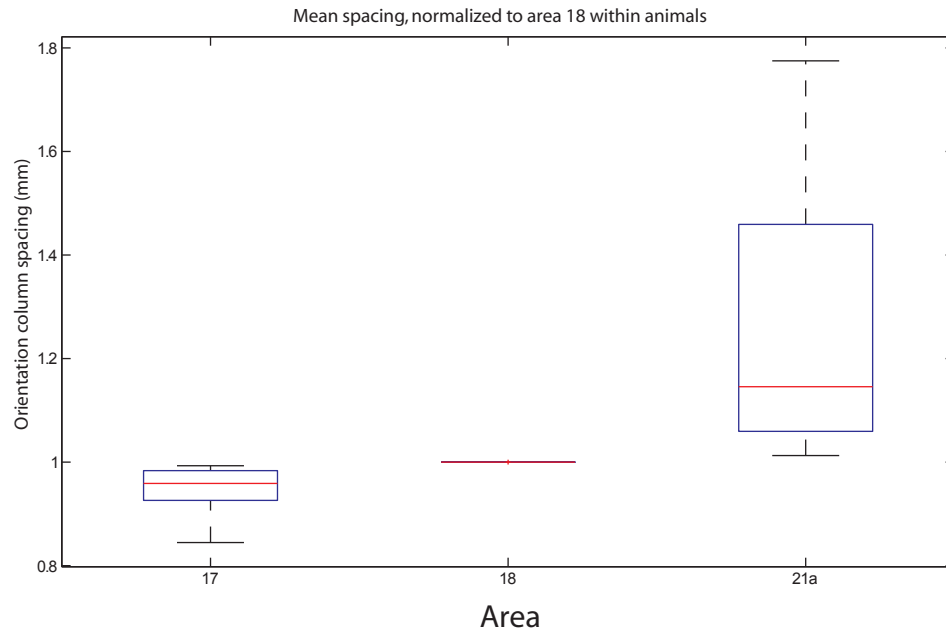


Figure 1d. Orientation column spacing comparison including all image pixels across all animals. The local column spacing at every pixel is normalized by the mean spacing in area 18 in the same animal. The mean normalized spacing is then taken over each area. The distribution shown is this mean normalized spacing, over all animals.

is relatively consistent relative to area 18, column spacing in area 21a is variable relative to area 18, with the two middle quartiles ranging from 1.1 to almost 1.5 times the spacing of columns in area 18.

Within a single visual area of one animal, column spacing was relatively constant across all pixels. The amount of variation did not seem to differ between areas 17, 18, and 21a, as shown by the size of error bars in Figure 1b. There did not appear to be multiple subareas within the maps of 21a, and no gradient of column spacing across the map was evident. This evidence, along with the homogeneity of the other types of map in area 21a, indicates that we are observing a unitary area rather than multiple neighboring areas.

Orientation column spacing in areas 17, 18 and 21a was uncorrelated with animal weight in the normal range of 2-4 kg for ages 6 months-2 years, with the caveat that in two very large (male) animals with weights over 6 kg, the column spacing in area 17 and 18 was larger. Kaschube et al. (2002) have reported a similar lack of correlation between ocular dominance columns in area 17 and animal weight. Column spacing was also not correlated with the sex of the animal.

Area 21a prefers higher spatial frequencies than area 18

The difference between the spatial frequency preferences of areas 17 and 18 can be easily identified from their different orientation map amplitudes at a single spatial frequency; area 18 was visibly stronger at 0.05 cycles/°, while area 17 was stronger at

0.2 and 0.4 cycles/°. This consistent amplitude difference between 17 and 18 was useful for identifying their border, as described by Issa et al. (2000).

It is nearly always possible to directly compare the amplitudes of area 17 and area 18 because the two neighboring areas are usually imaged in the same run, are similarly illuminated because the surface of the lateral gyrus is very flat, and have similar levels of intrinsic signal activity. In contrast, the imaging conditions of areas 17 and 18 and area 21a often differed: they could not always be captured in a single imaging plane, so they were imaged in separate runs with different camera and lighting positions. Since intrinsic signal is dependent on animal state, and the imaging conditions were different, the raw response amplitudes of areas 17 and 18 and area 21a were often not directly comparable.

Thus, to compare area 21a to the other areas, we computed combined spatial frequency maps by averaging multiple orientation maps acquired at different spatial frequencies, as described in Methods. Figure 2a shows a representative set of spatial frequency maps across areas 17, 18, and 21a in one hemisphere. Spatial frequency maps were generally consistent across multiple runs and contrast levels (50% vs. 100%).

The mean preferred spatial frequency across maps of areas 17, 18, and 21a showed a consistent trend: area 21a preferred higher spatial frequencies than area 18, but similar, or slightly higher, spatial frequencies to area 17. Figure 2b shows these trends in mean preferred spatial frequency, plotted across all animals for which we have spatial frequency maps in multiple areas. These trends were robust and did not depend on the

exact shape of the spatial frequency distribution; they were identical for pixel means, geometric means, and medians (not shown). The effect stands out in the paired comparisons (Figure 2c).

Grouping all animals together, the median preferred spatial frequency of pixels in area 21a was 0.48 cycles/°, higher than in area 17 (0.34 cycles/°) and considerably higher than in area 18 (0.23 cycles/°). These differences were all highly significant ($p < 0.0001$). This spatial frequency distribution was similar to the distribution of single unit optimal spatial frequencies recorded by Tardif et al. (1996), although ours was shifted slightly towards lower frequencies, possibly because we always stimulated through the eye contralateral to the hemisphere we were imaging. Single units in area 21a stop responding at a lower spatial frequency when stimulated through their nondominant eye (Tardif et al., 1996) compared to their dominant eye. The observed distribution was not similar to that reported by Morley & Vickery (1997), which peaked at 0.15 cycles/°. However, only 30% of the cells they recorded were responsive to gratings, compared to 85% of Tardif et al.'s, and 83% (54 of 65) of our microelectrode sample, so their sample population may not be representative. When area 21a is activated or silenced, area 17's responses to gratings of spatial frequency 0.5 cycles/° are most affected (Huang et al., 2004); this is the average spatial frequency we observed in area 21a, suggesting that the feedback projection to area 17 may be representative of the general area 21a population.

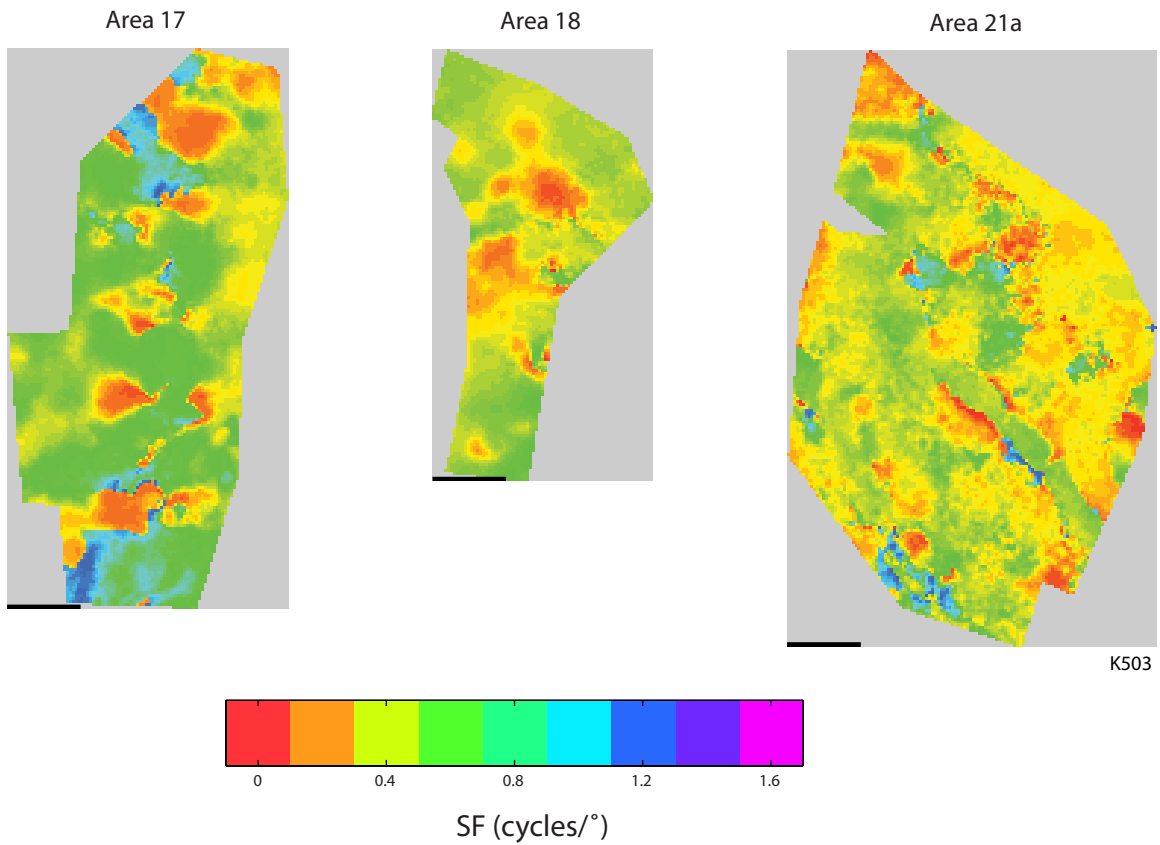


Figure 2a. Spatial frequency maps acquired in a single hemisphere in one animal. Color indicates preferred spatial frequency (as shown on key). Scale bar = 1 mm.




Figure 2b. Trend in mean spatial frequency across areas for each animal. Error bars SEM.




Figure 2c. Paired comparison of mean spatial frequency across areas. Diagonal line indicates equality. Each animal is shown as a point; error bars are SEM.

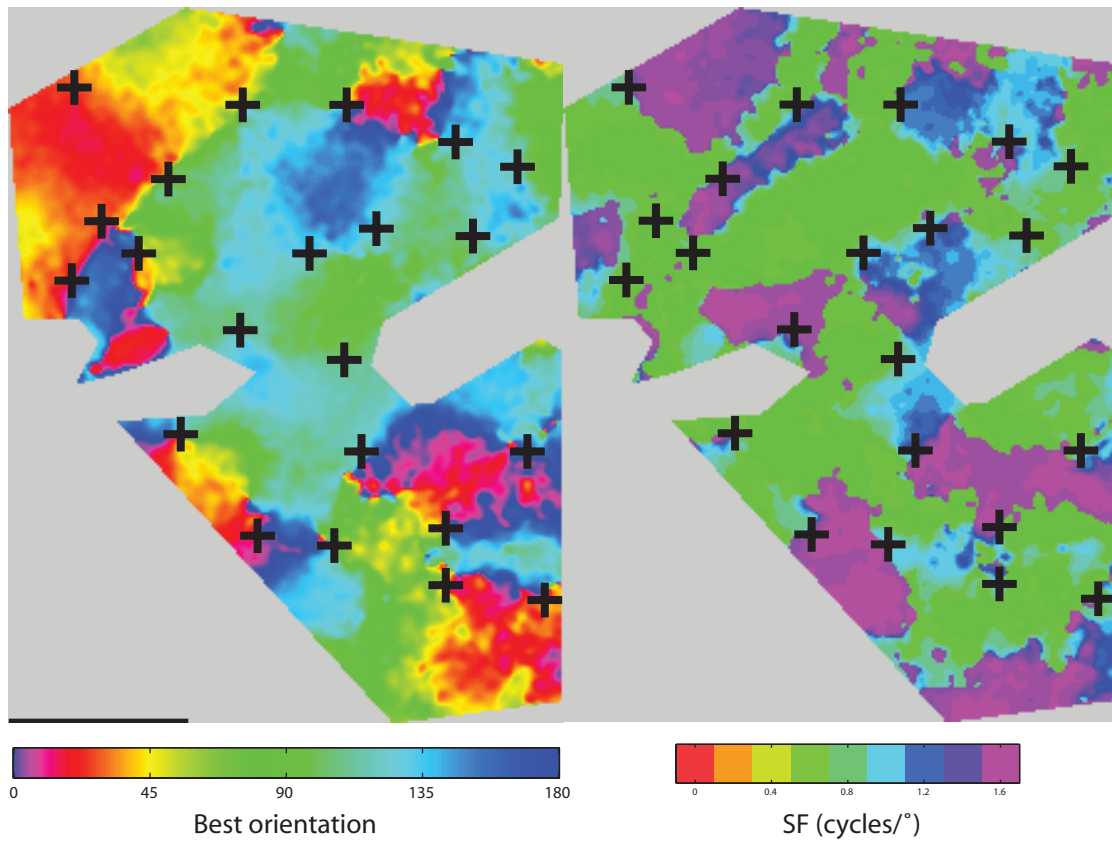


Figure 2d. Orientation map (left) and spatial frequency map (right) of area 21a in one animal. Pinwheels, marked with crosses on both maps, colocalize with the high SF domains (blue and purple areas, at right). Note that, on the orientation map (at left), the marked pinwheels do not exactly coincide with the boundaries of orientation domains when the map is globally normalized (as displayed here), since each pinwheel is marked when the map is locally normalized around it. Scale bar = 1 mm.

Intriguingly, 3 of the 6 spatial frequency maps of area 21a had the largest fractions of high spatial frequencies (> 0.7 cycles/°) we had seen in any of the 16 spatial frequency maps we made in all visual areas. This is likely significant, since there are a number of reasons why the area might seem to prefer spuriously low spatial frequencies (e.g. poor optics or poor responsiveness in early visual areas) but preference for high spatial frequencies is less likely to be artifactual. These maps were normal in strength and organization, and 2 of the 3 were imaged simultaneously with area 17 or 18 maps preferring normal (lower) spatial frequencies, so this result was not caused by a singularly strong imaging run. This observation is consistent with reports of such high peak spatial frequencies in 25% of recorded cells in 21a (Tardif et al., 1996).

Orientation pinwheels in area 21a may correlate with high spatial frequencies

The overall organization of spatial frequency maps in areas 17, 18, and 21a is similar (as shown in Figure 2a). Larger regions of medium spatial frequencies, with slowly changing spatial frequency gradient, are interspersed with patches of higher and sometimes lower spatial frequency, with linear discontinuities at the patches' borders.

Orientations in area 21a are distributed randomly and continuously with no evidence of long range order, even at the visible borders of the area. Therefore, like in areas 17 and 18, orientation pinwheels are by mathematical necessity present in the map (Schwartz, 1977). We observed mean pinwheel densities of 1.71 ± 0.17 , 2.31 ± 0.07 , and 1.50 ± 0.34 pinwheels/mm² in areas 17, 18, and 21a respectively. These were lower than the

densities reported by Huang et al. (2006) of 2.29 ± 0.14 pinwheels/mm² in area 17, and 2.87 ± 0.37 pinwheels/mm² in area 21a, as well as Löwel et al.'s (1998) estimate of 2.7 ± 0.1 pinwheels/mm² in area 17, but higher than Bonhoeffer and Grinvald's (1993) estimate of 1.2 pinwheels/mm² in area 18. However, our measurements were performed over relatively small maps and were subject to border effects.

In area 17 and 18, orientation pinwheels coincide with areas of high spatial frequency preference. Figure 2d shows, for one area 21a map, the orientation map with pinwheels marked and those pinwheels overlaid on a spatial frequency map. Note that the pinwheels seem to fall within or near the borders of the high spatial frequency zones. Pooling the data from areas 17 and 18, we observed that pinwheel density was elevated in the high spatial frequency regions (3.07 ± 1.85 for spatial frequencies $> 0.6+$ cycles/°, compared to 2.06 ± 0.94 for $0.3 - 0.6$ cycles/°; $p < 0.05$). It was also elevated in area 21a, but not significantly (2.27 ± 1.63 for > 0.6 cycles/°, compared to 1.85 ± 0.93 for $0.3 - 0.6$ cycles/°; $p = 0.16$). Issa et al. (2000) have proposed that pinwheels colocalize with the small fraction of high spatial frequency domains in areas 17 and 18 to ensure that all orientations are represented at these relatively rare spatial frequencies; this colocalization may still be important in area 21a.

Area 21a may prefer lower temporal frequencies than area 18

Temporal frequency maps were obtained by combining data from several orientation maps made at different temporal frequencies, similar to the procedure for spatial frequency maps. When these maps were averaged over each area, the trend for mean

preferred temporal frequency was the inverse of the trend for spatial frequency. Area 17 clearly prefers lower temporal frequencies than area 18, as expected; this is evident in the line plot (figure 3a, top) and is significant in the box plot (figure 3a, bottom). Area 21a may prefer lower temporal frequencies than area 18, although this trend is not significant in the data or the paired plots (Figure 3b). There is no significant difference in the temporal frequency preference of area 17 and area 21a.

The temporal frequency distribution we observed in area 21a was consistent with the distribution of low-pass and band-pass single cells reported by Morley and Vickery (1997). It was very different from the surprisingly high distribution, with mean 7.0 ± 4.0 cycles/s, reported by Tardif et al. (1996).

We did not observe reproducible temporal frequency structure in our maps of area 21a. The existence of temporal frequency maps in primary visual cortex is still an open question (Issa et al., 2008). Using optical imaging, Khaytin et al. (2008) find that orientation domains in bush baby V1 remain constant in position as the temporal frequency of a stimulus changes, suggesting that temporal frequency is distributed uniformly across the cortical representation. Any temporal frequency map in primary visual cortex is likely to be weak; DeAngelis et al. (1999) reported that, in pairs of single neurons in area 17 recorded simultaneously with a microelectrode, temporal frequency was more weakly correlated than spatial frequency, and much less correlated than orientation. The selectivity of temporal frequency tuning is even lower in area 21a than

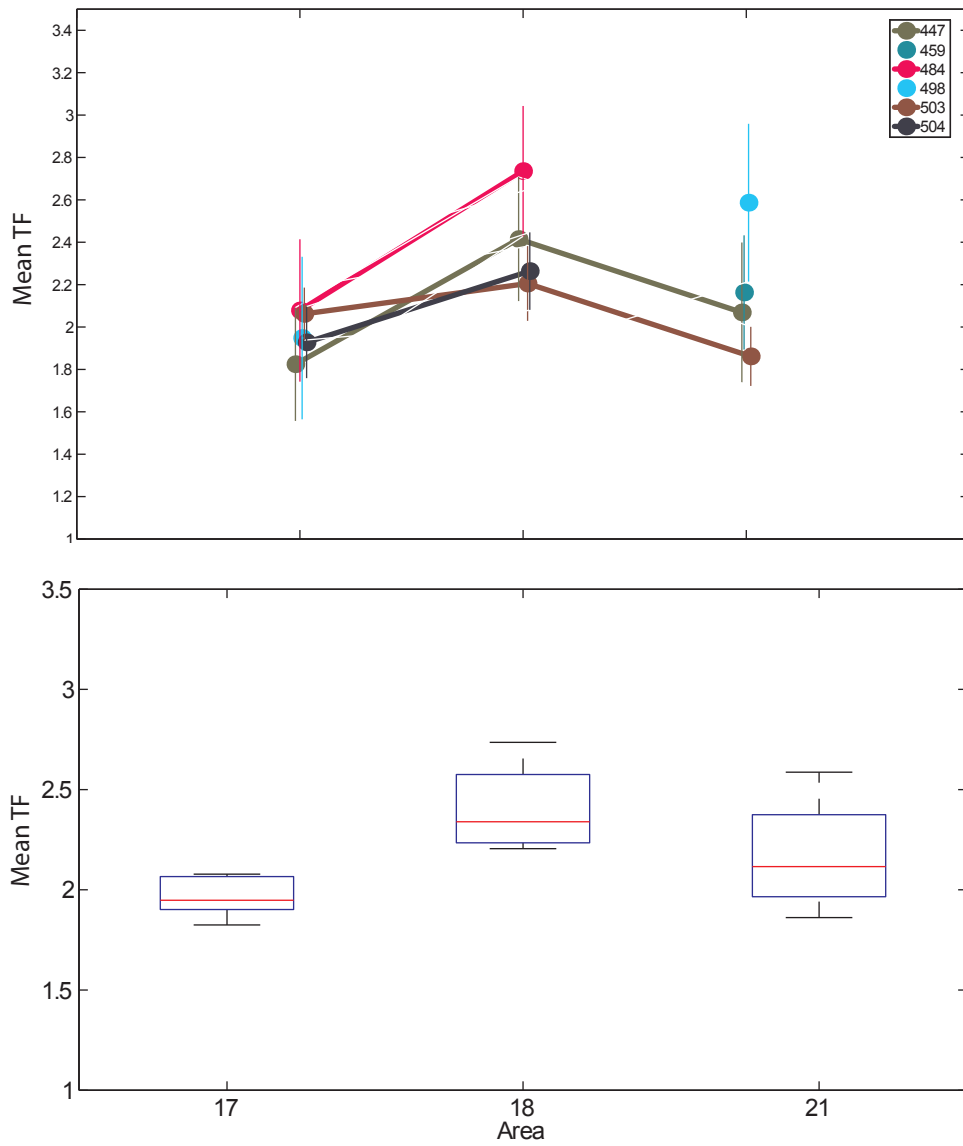


Figure 3a. (top) Trend in mean temporal frequency across areas for each animal. Error bars SEM for each map as in Figure 1b. (bottom) Boxplot showing same data. Error bars SEM across animals.




Figure 3b. Paired comparison of mean temporal frequency across areas. Diagonal line indicates equality. Each animal is shown as a point; error bars are SEM.

in area 17 (Kayser & König, 2006), predicting that any temporal frequency map that exists would be even weaker.

Area 21a's combination of low temporal frequency preference and high spatial frequency preference indicates a strong preference for low grating velocities. Our observation of a mean preferred temporal frequency of 0.5 cycles/s (at spatial frequency = 0.2 cycles/°) is consistent with a mean preferred velocity of 2.5 °/s. Likewise, our observation of a mean preferred spatial frequency of 0.5 cycles/° (at temporal frequency = 2.0 cycles/s) is consistent with a similarly low mean preferred velocity of 4.0 °/s. Of course, we did not explicitly sample the space of velocities, so these should be interpreted as estimates. However, we never observed optical maps in area 21a in response to stimuli moving faster than 20 °/s (temporal frequency=4.0 cycles/s, spatial frequency = 0.2 cycles/°). These measurements are consistent with the reported distribution of single neurons' velocity preferences (Dreher et al., 1993; Wimborne & Henry, 1992). Both the peak velocity estimates and the maximum velocity are consistent with the hypothesis that area 21a is a spatial form vision area which responds poorly to fast-moving stimuli.

Direction maps are weaker in area 21a than in area 18

Using continuous optical imaging of responses to drifting rotating gratings, direction maps for a visual area can be acquired simultaneously with orientation maps (see Methods). In general the preferred direction of a pixel is expected to be perpendicular to the orientation of its optimal bar stimulus. The consistency between orientation and

direction maps can provide independent confirmation of the validity of both orientation and direction maps because they are derived from two independent signals that are independently fit by the continuous intrinsic signal analysis procedure. Figure 4a shows sample pairs of orientation and direction phase maps and their associated amplitude maps.

Direction maps in area 21a were weaker than in areas 17 and 18, in the sense that the consistency between best orientation and direction was lower in area 21a than in 17 and 18. (The absolute amplitude of direction maps is not a very reliable measure; see Discussion.) Figure 4b shows the consistency between orientation and direction preference in the three areas shown in Figure 4a. While consistency was low overall (see Discussion), the median correlation was higher in areas 17 (0.18) and 18 (0.19) than in area 21a (0.09). The difference was significant in the Mann-Whitney test in both cases ($p < 0.05$). Area Areas 17 and 18 have a substantially higher fraction of directionally selective cells than area 21a (55%, 42%, and 17% respectively; Dreher et al., 1993), which probably accounts for the weakness of the overall direction map. We observed that all 21a direction maps were less consistent than all but one of the area 17 and 18 maps, indicating the consistency of this measurement across animals.

Ocular dominance maps are not apparent in area 21a

Orientation maps are very consistent from one eye to another (Figure 5a). In general, the highest amplitude areas of the orientation map are very stable from eye to eye, and run to run; but regions of the map in which responses are weaker amplitude sometimes

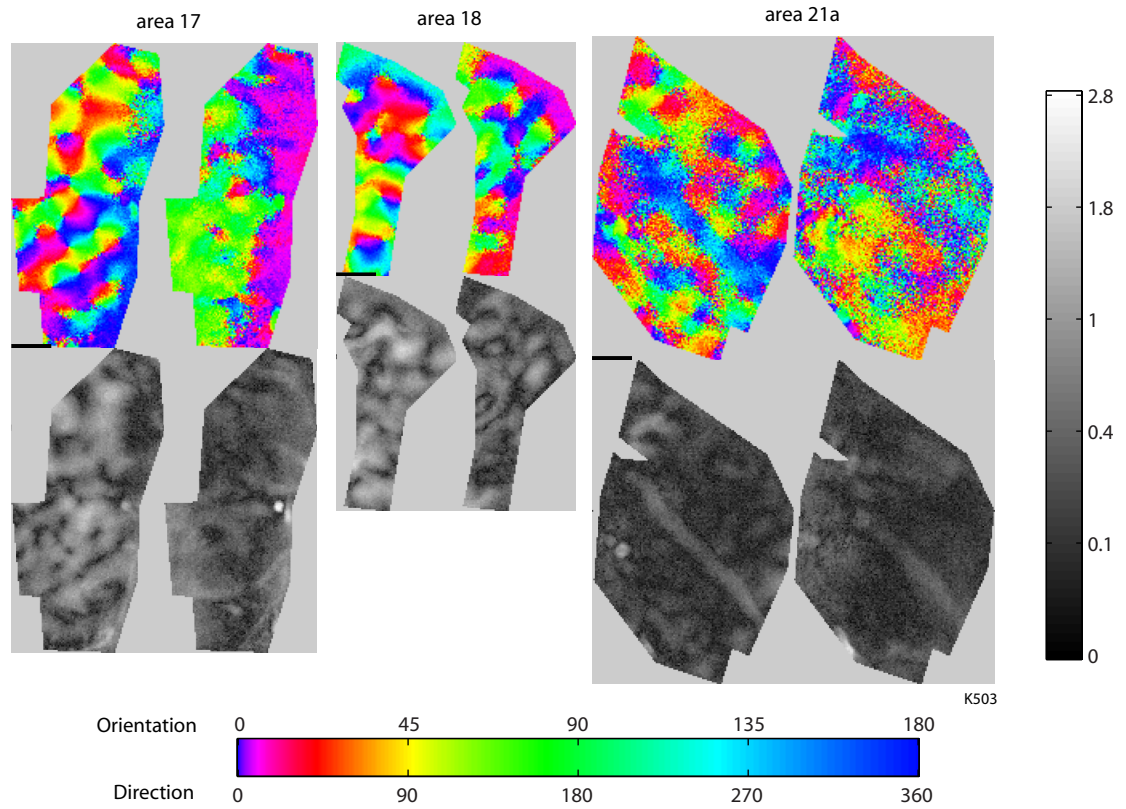


Figure 4a: Orientation vs. direction maps. Orientation maps (phase-amplitude pair, left) and direction maps (same, right) are shown for areas 17, 18, and 21a. The color scales for direction and orientation maps are identical (see color key) but the colors correspond to different preferences. Amplitudes are consistent across maps (see grayscale key). Scale bar = 1 mm.




Figure 4b: Orientation-direction consistency maps for areas 17, 18, and 21a in one cortical hemisphere. Hue indicates best orientation; intensity indicates similarity between best orientation and direction (see color scale). Scale bar = 1 mm.

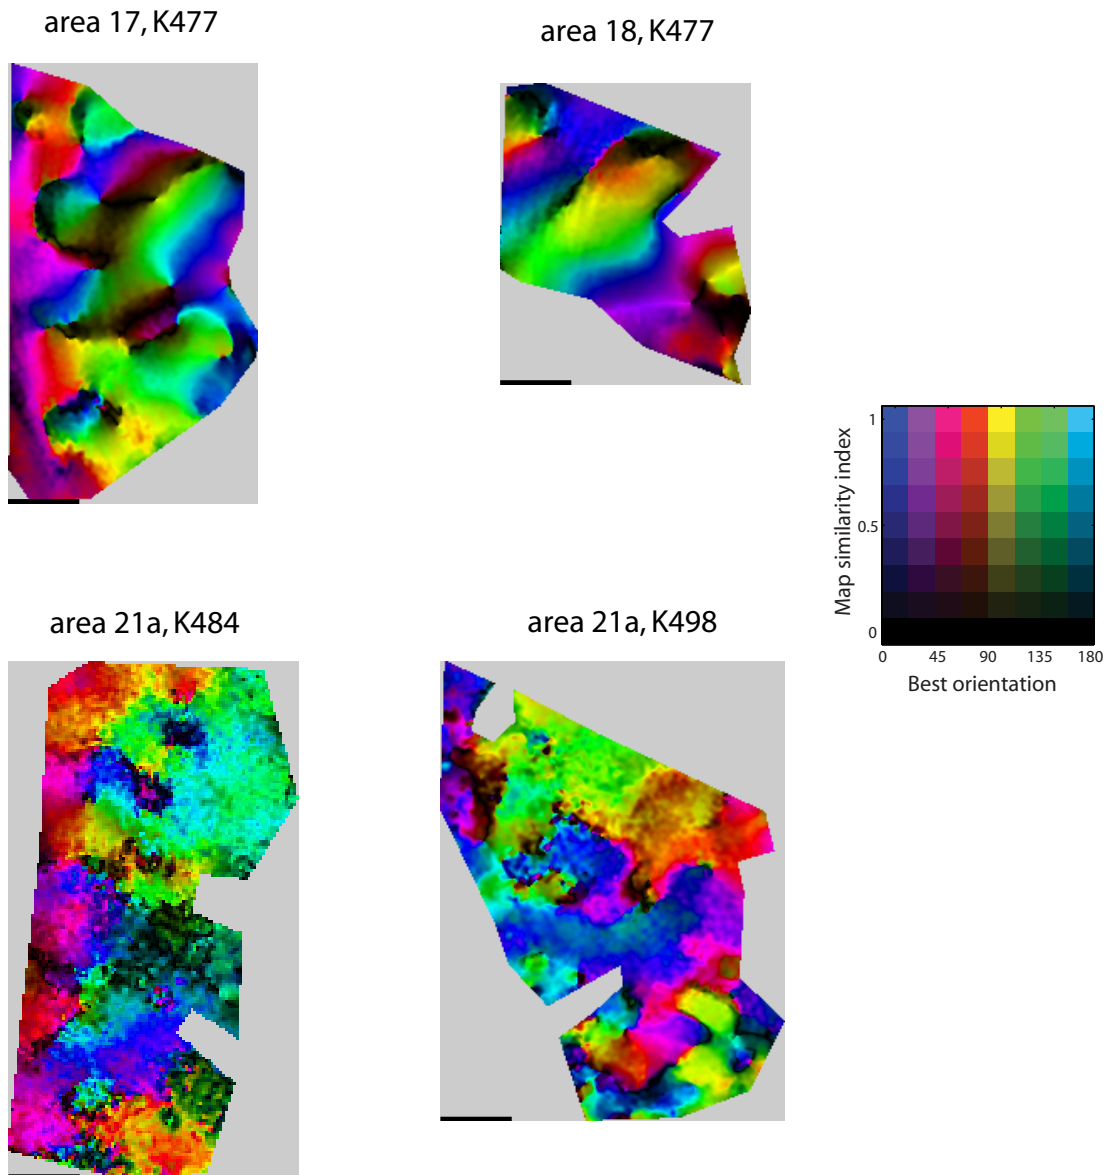


Figure 5a. Contralateral–ipsilateral orientation consistency maps. Color indicates preferred angle in contralateral map; intensity indicates consistency between two independent runs. Scale bar = 1 mm.

shift to the best orientation of a stronger neighboring region. This is a predictable result of looking at the cortical bloodflow signal, which is dominated by the most active nearby group of cells.

Indeed, orientation maps in area 21a of one hemisphere are as consistent in preferred orientation between contralateral and ipsilateral eyes as they are between two separate maps using the contralateral eye. In area 21a, the median smoothed orientation correlation between contralateral and ipsilateral eyes was 0.55; between two contralateral eye runs, it was 0.52. These medians were not significantly different. Correlation values for valid maps ranged from 0.21 to 0.83.

Comparing the amplitudes of orientation maps in areas 17 and 18 derived from the two eyes revealed spatially periodic stripes similar to the ocular dominance columns that have been shown by other means, such as episodic imaging (not shown). Such striped maps in 21a were not consistently observed. The lack of clear ocular dominance maps in area 21a was not the result of poor signal, since both high and low amplitude area 21a maps showed a similar lack of ocular dominance organization.

Microelectrode recordings agree with imaging maps

We used microelectrodes to record responses to different stimulus orientations from neurons in penetrations targeted based on imaging in three experiments, and in another experiment we recorded from stereotaxically defined 21a. We observed a strong correlation between the preferred orientation of neurons recorded with a microelectrode and the local preferred orientation in the optical maps ($r = 0.59$; figure

6a). To test the significance of this observation, we measured the correlation between the preferred orientation estimated from imaging and electrophysiological data and compared it to the correlation after randomly permuting the imaging orientation of cells while holding their electrophysiological preferred orientation constant (10,000 times). The observed data was highly significant ($p \ll 0.0001$; outside the distribution of the bootstrap).

The consistency of orientation tuning along single penetrations is similar to previous reports (Huang et al., 2006; Dreher et al., 1993). The single neurons studied were well tuned for orientation (median OSI = 0.51) but not for direction (median DSI = 0.28), as has been previously reported in area 21a (Toyama et al., 1994).

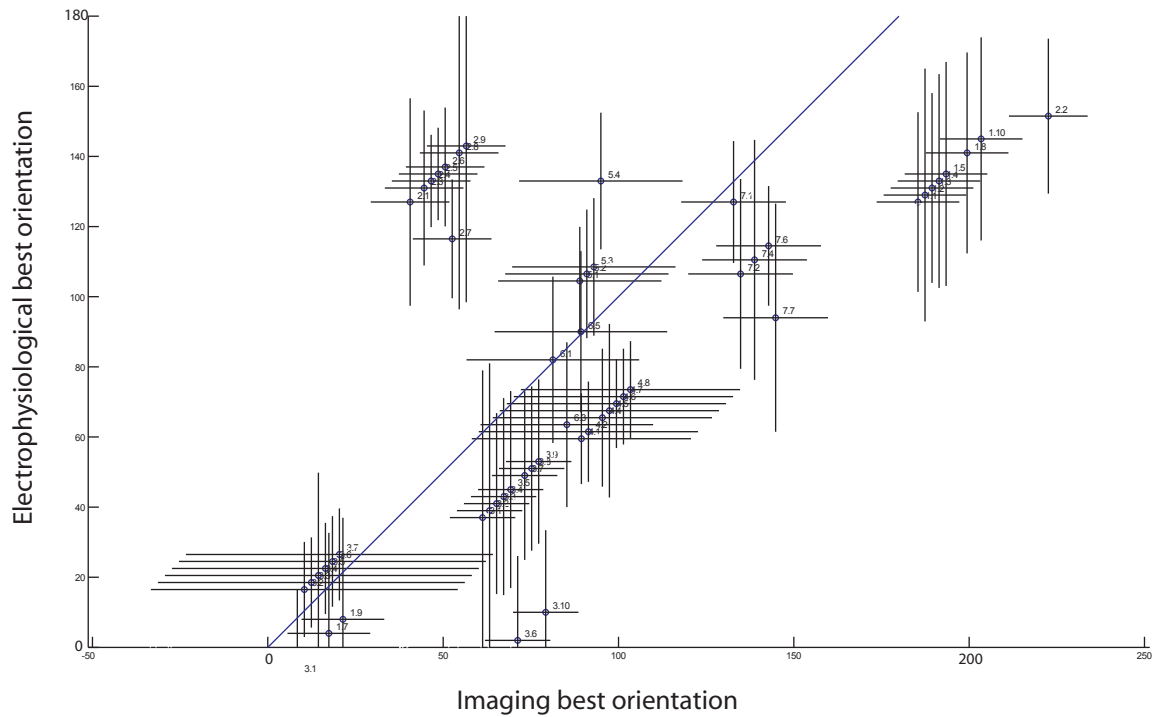


Figure 6a. Scatterplot of orientation preference determined by optical imaging versus orientation preference of single units recorded with a microelectrode. Cells were recorded at several sites along each penetration. They are vertically and horizontally staggered by up to $\pm 5^\circ$ so that error bars are distinct. Cells are labelled as 'penetration number.site number'. Electrophysiological orientation was the peak of the cell's orientation tuning curve, and error bars indicate the full-width at half height of the tuning curve. Imaging orientation was taken from maps smoothed with a $100\mu\text{m}$ kernel. Error bars are the local (circular) standard deviation within a window of radius $100\mu\text{m}$.

Discussion

We have characterized a unitary visual area located on the posterior suprasylvian gyrus of the cat, consistent with the location of histological area 21a. By mapping it with intrinsic signal imaging and verifying our imaging findings with microelectrode physiology, we have confirmed its identity as area 21a. We have shown that the area's preference statistics for spatial and temporal frequency, its lack of ocular dominance organization, and its weak direction selectivity are all consistent with its hypothesized role as the second area in the ventral stream object recognition hierarchy.

Overall, there was substantial positional variability in area 21a. Its position along the anterior-posterior axis was generally consistent with most reports (Dreher et al., 1993; Tusa & Palmer, 1980) although somewhat anterior to the report of Toyama et al. (1994). This finding reemphasizes the need, when studying area 21a, to confirm its location, preferably by imaging, or alternatively by histology or systematic electrophysiological mapping.

We were not successful in obtaining clear intrinsic signal maps of retinotopy despite numerous attempts. These would have been especially helpful in demarcating this visual area and in evaluating the accuracy of the standard retinotopic map provided by Tusa and Palmer (1980). The weakness of the retinotopic map in area 21a reflects a common problem when receptive fields are large and neurons highly selective for another stimulus property such as orientation. Indeed, topographic intrinsic signal maps at the scale of single hypercolumns have not been reported by most authors in area 17.

Our study suggests that area 21a does not have ocular dominance columns like those we observe in areas 17 and 18. Löwel and Singer (1987) observed OD columns in areas 17 and 18 after injection of (^3H)proline, and suggested they were variably present in area 19 and possibly in area 21a. The presence of OD columns would seem to imply the presence of strongly contralateral- and ipsilateral-preferring neurons. Wang and Dreher (1996) and Dreher et al. (1993) have shown that the distribution of OD values in area 21a is peaked in the middle (binocular) category, narrower than the distribution of area 17, and very few cells in 21a respond only to ipsilateral stimulation while a number respond only to simultaneous binocular stimulation. Such weak OD tuning does not seem consistent with the presence of columns. It is possible that weak OD tuning could exist in a weak OD map, but it seems unlikely and our imaging studies do not disclose one.

In all areas, direction maps were weaker and less consistent than orientation maps, even in area 18, where the episodically imaged direction map is stronger than the episodic orientation map (Swindale et al., 2003). This has two major sets of causes; the first is that our imaging protocol is optimized for orientation maps. The period of the orientation stimulus was optimized to be in a low-noise region of the time spectrum; the period of the direction stimulus was fixed relative to it. The noise power in the intrinsic signal falls off at higher frequencies (Kalatsky & Stryker, 2003) so direction maps, being lower frequency, will generally have higher noise. Also, the frequency of the orientation signal is higher, so the number of periods is also higher, which can improve the signal/noise ratio.

The second major cause is the difficulty of fitting direction maps acquired with the continuous intrinsic signal imaging technique. In theory direction maps can be computed from the same mapping stimulus as orientation maps, by using the first harmonic of the stimulus period to produce a direction map and the second harmonic to produce an orientation map. In practice, the situation is not that simple. While for orientation maps the continuous imaging procedure fits two periods of a sinusoid to the generally two-lobed orientation tuning curve of each image pixel, yielding a reasonably unbiased fit, for direction maps it fits a sinusoid to a signal which is at best Gaussian (on a circle, for a pixel with pure direction selective tuning) and at worst nonexistent (for a pixel with pure orientation selective tuning). This fitting decreases the amplitude of direction maps, which may be responsible for the overall weakness of these maps in all areas. While Swindale et al. (2003) provide a useful solution for episodic imaging by fitting Gaussian tuning curves to the response maps, applying their method to continuous imaging is nontrivial.

Orientation column spacing

Despite the relatively small size of area 21a, its orientation domains are relatively large.

The surface areas of areas 17, 18, and 21a are approximately 284 ± 7 , 125 ± 10 , and 25 ± 2 mm² (after 10-15% shrinkage from histological processing; Hilgetag & Grant, 2000).

Thus, while area 17 contains ~50 complete hypercolumns (representing all orientations), and area 18 contains ~15, area 21a contains as few as 3. Since all combinations of retinotopy and orientation must be represented within the map, the constraint of

imposing a smooth retinotopic map on a very coarse-grained orientation map may constrain the complex retinotopy of area 21a (Tusa & Palmer, 1980).

The column spacing we observed was larger than reported by Huang et al. (2006).

Stimulating area 21a with 4 episodic gratings (0°, 45°, 90°, 135°) and cocktail normalizing the resulting maps (i.e. dividing by the overall average map), they found maps with a patchy, striated appearance similar to ocular dominance maps. They reported patch widths of 0.97 ± 0.12 mm in area 21a, smaller than in area 17 (1.04 ± 0.19 mm).

However, with only four stimuli it is difficult to be certain about the true orientation preference of every pixel. This patch width measurement encompasses an orientation selective region the size of which depends on the number of stimuli, normalization, and the image enhancement used (histogram equalization). In contrast, our continuous measurement of orientation allows for any orientation preference. Likewise, our wavelet column spacing measurement measures orientation from the center of one orientation column to the next column with the same orientation preference and is not affected by the exact stimulus or image processing parameters.

Primate homologues of area 21a

The assignment of area 21a to object recognition has been historically controversial.

Sherk and Mulligan (1993) argued that 21a might be a subpart of area LS (a motion vision area similar to primate MT) based on connectivity and the complexity of retinotopy in areas LS and 21a. However, more recent work shows that the two areas' thalamic afferents are distinct (Dreher et al., 1996a) and their cortical afferents originate

from different layers (Conway et al., 2000). A controlled physiological experiment comparing areas LS and 21a at a single eccentricity showed a clear difference in RF size, ocular dominance, firing rate, direction tuning, velocity tuning, and response to visual noise (Dreher et al., 1996b). While area LS is located inside the sulcus and cannot be easily imaged, all of our results further confirm that area 21a's properties are those of an object vision area, not a motion area.

Primate area V4 is area 21a's natural homologue (in the sense of Payne, 1993). V4 is also an intermediate visual area with no direct input from the LGN. Its orientation, direction, spatial frequency, and temporal frequency response properties are similar to area 21a (Payne, 1993; Lomber, 2001). However, there are intriguing differences between the functional architecture of area V4 and area 21a. Orientation maps have been observed by optical imaging in parafoveal V4, but not reliably at eccentricities larger than 3° (Ghose & Ts'o, 1997). This difference may reflect foveal specialization or a specialized temporal pathway as the authors suggest. However, although the cat lacks a specialized fovea, the representation of its area centralis in area 21a is highly magnified, so our observation that the entirety of 21a has orientation maps suggests a difference from V4. The trend of increasing column size along the visual pathway is also found in primates; orientation domains in V4 are three times larger than in V1, although they similar in size to V2. The relatively larger size of domains in V2 may reflect the complexity of its organization relative to cat area 18 (e.g. cytochrome oxidase stripes).

Differences between continuous intrinsic imaging of cats and rodents

Our laboratory has recently published a number of papers using the continuous imaging technique to make highly replicable quantitative measurements of ocular dominance in juvenile rodents, in particular contralateral and ipsilateral response amplitude in mouse V1 near the time of the critical period (Cang et al., 2005). Under some anesthetic regimens, the amplitudes of mouse retinotopic maps are very consistent between runs and between experimental sessions. However, our response amplitude measurements in the cat are not nearly so consistent and vary substantially over intervals of just 15 minutes.

While some of this difference may be attributable to age and anesthetic factors, other aspects may be specific to larger animals like the cat. In comparison to the rodent, many more large sources of physiological noise are present at periods closer to that of the stimulus, including ventilation and heartbeat. Controlling the plane of anesthesia is also more difficult in the cat, especially over the course of long experiments. Pulsation of the cortex and the resulting changes in reflectivity are a much greater problem in the exposed stabilized cat brain than in the trans-cranial mouse preparation. Even if these sources of noise could be controlled, the simplicity of the mouse's visual topography – a uniform, un-fractured retinotopic map – gives it signal/noise advantages compared to non-uniform and fractured cat functional maps.

Future directions

To further elucidate the function of the different areas of the cat's extrastriate visual cortex, it would be interesting to study functional maps in areas adjacent to 21a. Ideally, one would want to compare area 21a, part of the ventral stream, with its adjacent area LS, which is part of the dorsal stream and would be predicted to have direction and orientation maps similar to primate area MT (Malonek et al., 1994). However, almost the entirety of area LS falls within the lateral suprasylvian sulcus (Sherk & Mulligan, 1993), and cannot be optically imaged using standard methods. As an alternative, one could try to image area 19, an intermediate visual area with dorsal stream-like response properties, suggested homology to primate V3, and extensive exposure along the posterior suprasylvian gyrus, representing the upper visual quadrant. While we did not observe maps in the part of area 19 bordering area 21a, stimulus parameters, camera positioning, and the visual field position of the stimulus were not optimized for area 19. After optimizing these parameters, one might observe orientation domains or luminance-modulation domains similar to those observed in owl monkey V3 (Kaskan et al., 2009).

An extremely interesting direction would be to advance further down the ventral stream and study maps in visual areas 21b, 20a, and 20b of the cat using intrinsic signal imaging. These areas have retinotopic maps which are as well organized as 21a (Tusa & Palmer, 1980). However, it is not known whether they contain complex feature detecting cells such as face cells (although it seems likely given the observation of face cells in the temporal lobe of the sheep; Kendrick & Baldwin, 1987). It may be possible to

address the problem of complex feature selectivity in a targeted manner, by using optical imaging to identify face-selective regions to target with microelectrodes; in the primate, using fMRI to target electrophysiological recordings vastly enriches the population of face cells (Tsao et al., 2006). Much of area 20a is exposed on the cortical surface near the tail of the lateral sulcus ventral to area 19, although being on the underside of the brain it would have to be approached differently than earlier visual areas. It is possible to obtain optical maps using stimulus objects and faces in the inferotemporal cortex of anesthetized primates (Tsunoda et al., 2001), and these stimuli may also be effective in higher ventral stream areas of the cat.

References

- Bonhoeffer, T., & Grinvald, A. (1993). The layout of iso-orientation domains in area 18 of cat visual cortex: Optical imaging reveals a pinwheel-like organization. *Journal of Neuroscience*, *13*(10), 4157-4180.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433-436.
- Cang, J., Kalatsky, V. A., Löwel, S., & Stryker, M. P. (2005). Optical imaging of the intrinsic signal as a measure of cortical plasticity in the mouse. *Visual Neuroscience*, *22*(5), 685-691.
- Cheng, K., Waggoner, R. A., & Tanaka, K. (2001). Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron*, *32*(2), 359-374.
- Conway, B., Boyd, J. D., Stewart, T. H., & Matsubara, J. A. (2000). The projection from V1 to extrastriate area 21a: A second patchy efferent pathway colocalizes with the CO blob columns in cat visual cortex. *Cerebral Cortex*, *10*(2), 149-159.
- Dadvand, L., Stryker, M. P., & Frank, M. G. (2006). Sleep does not enhance the recovery of deprived eye responses in developing visual cortex. *Neuroscience*, *143*(3), 815-826.
- DeAngelis, G. C., Ghose, G. M., Ohzawa, I., & Freeman, R. D. (1999). Functional micro-organization of primary visual cortex: Receptive field analysis of nearby neurons. *Journal of Neuroscience*, *19*(10), 4046-4064.
- Dreher, B., Djavadian, R. L., Turlejski, K. J., & Wang, C. (1996a). Areas PMLS and 21a of cat visual cortex are not only functionally but also hodologically distinct. *Progress in Brain Research*, *112*, 251-276.
- Dreher, B., Michalski, A., Ho, R. H., Lee, C. W., & Burke, W. (1993). Processing of form and motion in area 21a of cat visual cortex. *Visual Neuroscience*, *10*(1), 93-115.
- Dreher, B., Wang, C., Turlejski, K. J., Djavadian, R. L., & Burke, W. (1996b). Areas PMLS and 21a of cat visual cortex: Two functionally distinct areas. *Cerebral Cortex*, *6*(4), 585-599.
- Ghose, G. M., & Ts'o, D. Y. (1997). Form processing modules in primate area V4. *Journal of Neurophysiology*, *77*(4), 2191-2196.
- Heath C.J., J. E. G. (1971). *The anatomical organization of the suprasylvian gyrus of the cat*. Berlin: Springer-Verlag.
- Hilgetag, C. C., & Grant, S. (2000). Uniformity, specificity and variability of corticocortical connectivity. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *355*(1393), 7-20.

- Huang, L., Chen, X., & Shou, T. (2004). Spatial frequency-dependent feedback of visual cortical area 21a modulating functional orientation column maps in areas 17 and 18 of the cat. *Brain Research*, 998(2), 194-201.
- Huang, L., Shou, T., Chen, X., Yu, H., Sun, C., & Liang, Z. (2006). Slab-like functional architecture of higher order cortical area 21a showing oblique effect of orientation preference in the cat. *NeuroImage*, 32(3), 1365-1374.
- Hubel, D. H., & Wiesel, T. N. (2004). *Brain and visual perception : The story of a 25-year collaboration*. New York: Oxford University Press.
- Issa, N. P., Trepel, C., & Stryker, M. P. (2000). Spatial frequency maps in cat visual cortex. *Journal of Neuroscience*, 20(22), 8504-8514.
- Issa, N. P., Rosenberg, A., & Husson, T. R. (2008). Models and measurements of functional maps in V1. *Journal of Neurophysiology*, 99(6), 2745-2754.
- Kalatsky, V. A., & Stryker, M. P. (2003). New paradigm for optical imaging: Temporally encoded maps of intrinsic signal. *Neuron*, 38(4), 529-545.
- Kaschube, M., Wolf, F., Geisel, T., & Löwel, S. (2002). Genetic influence on quantitative features of neocortical architecture. *Journal of Neuroscience*, 22(16), 7206-7217.
- Kaschube, M., Wolf, F., Geisel, T., & Löwel, S. (2000). Quantifying the variability of patterns of orientation domains in the visual cortex of cats. *Neurocomputing*, , 32 415-423.
- Kaskan, P. M., Lu, H. D., Dillenburger, B. C., Kaas, J. H., & Roe, A. W. (2009). The organization of orientation-selective, luminance-change and binocular-preference domains in the second (V2) and third (V3) visual areas of new world owl monkeys as revealed by intrinsic signal optical imaging. *Cerebral Cortex*, 19(6), 1394-1407.
- Kayser, C., & Konig, P. (2006). Feature selectivity in area 21a of the cat. *Neuroreport*, 17(8), 809-812.
- Kendrick, K. M., & Baldwin, B. A. (1987). Cells in temporal cortex of conscious sheep can respond preferentially to the sight of faces. *Science*, 236(4800), 448-450.
- Khaytin, I., Chen, X., Royal, D. W., Ruiz, O., Jermakowicz, W. J., Siegel, R. M., & Casagrande, V. A. (2008). Functional organization of temporal frequency selectivity in primate visual cortex. *Cerebral Cortex*, 18(8), 1828-1842.
- Lomber, S. G. (2001). Behavioral cartography of visual functions in cat parietal cortex: Areal and laminar dissociations. *Progress in Brain Research*, 134, 265-284.
- Lomber, S. G., Payne, B. R., Cornwell, P., & Long, K. D. (1996). Perceptual and cognitive visual functions of parietal and temporal cortices in the cat. *Cerebral Cortex*, 6(5), 673-695.
- Löwel, S., Schmidt, K. E., Kim, D. S., Wolf, F., Hoffsummer, F., Singer, W., & Bonhoeffer, T. (1998). The layout of orientation and ocular dominance domains in area 17 of strabismic cats. *The European Journal of Neuroscience*, 10(8), 2629-2643.

- Löwel, S., & Singer, W. (1987). The pattern of ocular dominance columns in flat-mounts of the cat visual cortex. *Experimental Brain Research*, 68(3), 661-666.
- Malonek, D., Tootell, R. B., & Grinvald, A. (1994). Optical imaging reveals the functional architecture of neurons processing shape and motion in owl monkey area MT. *Proceedings of the Royal Society of London B*, 258(1352), 109-119.
- Mardia, K. V., & Jupp, P. E. (1999). *Directional statistics* (2nd ed.). New York: Wiley.
- Michalski, A., Wimborne, B. M., & Henry, G. H. (1993). The effect of reversible cooling of cat's primary visual cortex on the responses of area 21a neurons. *Journal of Physiology*, 466, 133-156.
- Morley, J. W., & Vickery, R. M. (1997). Spatial and temporal frequency selectivity of cells in area 21a of the cat. *Journal of Physiology*, 501 (Pt 2), 405-413.
- Niell, C. M., & Stryker, M. P. (2008). Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience*, 28(30), 7520-7536.
- Payne, B. R. (1993). Evidence for visual cortical area homologs in cat and macaque monkey. *Cerebral Cortex*, 3(1), 1-25.
- Payne, B. R., & Peters, A. (2002). The concept of cat primary visual cortex. In B. R. Payne, & A. Peters (Eds.), *The cat primary visual cortex* (pp. 1-129). San Diego: Academic Press.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437-442.
- Schwartz, E. L. (1977). Afferent geometry in the primate visual cortex and the generation of neuronal trigger features. *Biological Cybernetics*, 28(1), 1-14.
- Sherk, H., & Mulligan, K. A. (1993). A reassessment of the lower visual field map in striate-recipient lateral suprasylvian cortex. *Visual Neuroscience*, 10(1), 131-158.
- Swindale, N. V., Grinvald, A., & Shmuel, A. (2003). The spatial pattern of response magnitude and selectivity for orientation and direction in cat visual cortex. *Cerebral Cortex*, 13(3), 225-238.
- Tardif, E., Bergeron, A., Lepore, F., & Guillemot, J. P. (1996). Spatial and temporal frequency tuning and contrast sensitivity of single neurons in area 21a of the cat. *Brain Research*, 716(1-2), 219-223.
- Toyama, K., Mizobe, K., Akase, E., & Kaihara, T. (1994). Neuronal responsiveness in areas 19 and 21a, and the posteromedial lateral suprasylvian cortex of the cat. *Experimental Brain Research*, 99(2), 289-301.
- Tsao, D. Y., Freiwald, W. A., Tootell, R. B., & Livingstone, M. S. (2006). A cortical region consisting entirely of face-selective cells. *Science*, 311(5761), 670-674.
- Tsunoda, K., Yamane, Y., Nishizaki, M., & Tanifuji, M. (2001). Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. *Nature Neuroscience*, 4(8), 832-838.

- Tusa, R. J., & Palmer, L. A. (1980). Retinotopic organization of areas 20 and 21 in the cat. *Journal of Comparative Neurology*, 193(1), 147-164.
- Tusa, R. J., Palmer, L. A., & Rosenquist, A. C. (1978). The retinotopic organization of area 17 (striate cortex) in the cat. *Journal of Comparative Neurology*, 177(2), 213-235.
- Tusa, R. J., Rosenquist, A. C., & Palmer, L. A. (1979). Retinotopic organization of areas 18 and 19 in the cat. *Journal of Comparative Neurology*, 185(4), 657-678.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549-586). Cambridge, MA: MIT Press.
- Wang, C., & Dreher, B. (1996). Binocular interactions and disparity coding in area 21a of cat extrastriate visual cortex. *Experimental Brain Research*, 108(2), 257-272.
- Wimborne, B. M., & Henry, G. H. (1992). Response characteristics of the cells of cortical area 21a of the cat with special reference to orientation specificity. *Journal of Physiology*, 449, 457-478.

Conclusion: experimental approaches to the visual ventral stream

The visual cortex of the human brain allows us to perceive and recognize millions of objects, seemingly effortlessly. How this happens is not just a profoundly interesting computational problem, but is more widely important, because object vision serves as a primary model for understanding both the functional circuitry of the entire cortex and the function of hierarchies of cortical areas.

Visual processing in the cortex of many species involves two divergent networks: a dorsal stream that processes an object's location and motion, and a ventral stream that processes and recognizes an object's spatial form. The dorsal stream has been studied extensively at all levels (Milner & Goodale, 2006), but it has not yielded much information about the specific function of the cortical circuit; the study of the ventral stream has lagged behind. While we know a great deal about the response of primary visual cortex (V1) to simple stimuli and a considerable amount about inferotemporal cortex (IT)'s response to faces, we know much less about intermediate visual areas or intermediate stimuli.

A rational strategy for understanding object recognition is to work forward from V1, discovering how successive ventral stream cortical areas represent and transform visual stimuli. Since V1 is so well characterized, the natural next step is to investigate the transformation of stimulus information that takes place between V1 and successive ventral stream areas.

The function of ventral stream areas

A key property of ventral stream areas is invariance: neurons that respond to a certain object will tolerate and continue to respond to that object in a variety of positions and three-dimensional orientations. Invariance allows neurons to recognize an object regardless of its location, view, and illumination, and allows objects to be compared and categorized. In primary visual cortex, simple cells prefer one position of an object edge, while complex cells respond to the same edge in several positions. More general invariances occur in higher visual cortex (reviewed in Wallis & Rolls, 1997). How do these invariances arise?

Most models of how the ventral stream might work are extensions of pre-existing models of V1. Riesenhuber and Poggio (2004) propose a “standard model” of object recognition: a hierarchy of alternating simple-like and complex-like cells, in which the simple-like cells detect feature conjunctions and the complex-like cells group inputs from simple-like cells to permit invariance and flexibility in recognition. While the implicit analogy to the “standard model” of particle physics is hyperbolic, this model does capture the intuitive understanding of many vision scientists.

In this model, the role of an intermediate visual area such as primate V4 or cat 21a is to capture simple features of an object’s outline, such as contours, angles, and corners. These features are easy to build up conjunctively from edge detectors such as V1 neurons; they are invariant to scale, and are highly informative about overall shape, forming a distributed “structural representation” of the parts of an object (Connor,

2004). This model is consistent both with the observation that cells in V4 and 21a are still quite responsive to the simple edges that maximally stimulate cells in V1, and with the complex contour selectivity found in some studies of area V4 (Pasupathy, 2006).

The dimensionality problem in the ventral stream

As experimenters, we would like to test the standard model and understand the function of the hierarchy of ventral stream visual areas, but there is a major obstacle: classical sensory physiology experiments will fail in higher ventral stream areas because of the high dimensionality of stimulus space.

Consider first a V1 simple cell, with a receptive field well-fit by a Gabor wavelet. For this cell, we find the optimal stimulus by exhaustively exploring the seven-dimensional parameter space defined by x position, y position, orientation, length, width, period, and phase. In theory, we could exhaustively characterize the cell by showing it a few hours of pre-computed stimuli, but in practice even this is too cumbersome. Instead, one fits a Gabor wavelet to a receptive field found by reverse correlation, and assumes (usually without testing) that this fit yields the optimal wavelet stimulus.

Now, consider that neurons in higher visual areas integrate numerous V1 inputs. For these neurons, exhaustive exploration of a parameter space is likely to prove impossible; if we tried to map a trivial V4 receptive field (with only two V1 inputs) by fitting multiple wavelets, we would need to exhaustively explore a space with 14

parameters, which would take an unfeasibly long time. Exhaustive exploration methods do not scale with increasing receptive field intricacy.

The problem dimensions: stimulus parameterization and search strategy

The experimenter's usual challenge is to find the stimulus to which a ventral stream neuron is maximally responsive. While a variety of response functions could be maximized (firing rate, latency of spiking, mutual information about stimuli, even synchrony between pairs of neurons), we will assume that in general the experimenter is trying to maximize the neuron's firing rate. (This is not always the most efficient strategy for characterizing a neuron; see Paninski, 2005.) This problem can be described as "active learning" – we observe the neuron's responses, build a model, test it, observe the test results, then (sometimes) repeat the process. One classical example of active learning is an experimenter recording in V1 who first finds the best orientation of a neuron and then uses this knowledge about its orientation to characterize its spatial frequency tuning. For our purposes it is irrelevant whether the experiment is human- or computer-controlled; only the algorithm matters.

There are two major dimensions to the experimental problem. The first describes the extent to which the stimulus is parametric or not. A wavelet can be described with a small number of parameters. An image of the natural world, however, can only be represented with a large number of pixels – forming a stimulus space so vast (10^6 pixels for a 1000x1000 image) that it is not useful to think in terms of its dimensions. (Note

that when using “natural scenes” stimuli, the dimensionality of the stimulus space is not proportional to the number of pixels; it is proportional to the large, but less vast, dimensionality of the manifold defined by natural scenes within that pixel space.)

The second dimension is search strategy. In any optimization problem there is a tradeoff between exploration and exploitation (Thrun et al., 2005). A strategy can be thought of as either searching for new maxima even if it decreases the response (exploration), or as improving on or maintaining existing maxima (exploitation). A strategy can also do both sequentially. Since experimental time is always limited by how long a single cell can be recorded from, there will always be a tradeoff between strategies. (Once all data has been acquired, analyzing or fitting it with a model can be revealing, but it does not count as either exploration or exploitation because there is no way to test whether the model can extrapolate to new data.)

As an illustration of this classification scheme, consider the ancient example of an experimenter qualitatively mapping a cat visual neuron with a light gun projected onto a tangent screen (Hubel & Wiesel, 2004). This is a parametric stimulus with five parameters: x position, y position, orientation, length and width. As he moves the gun across the screen looking for the neuron, he is exploring; when he finds a responsive region and adjusts the orientation, length, and width, he is exploiting a maximum (the responsive subfield he has identified). When he checks to see if other subfields are present, he is exploring again; when he finds one and optimizes its parameters, he is again exploiting.

Classification of previous studies

We will describe and classify several attempts to study intermediate ventral stream areas. This classification is not normative, and the ideal exploration-exploitation tradeoff will depend on the nature of the study: an exhaustive characterization of a single neuron will require mostly exploitation, while a multielectrode study of coding in a neural population will necessarily rely mostly on exploration. Even so, we believe many studies would benefit from a more balanced combination of exploration and exploitation, or at least more explicit contemplation of the tradeoff.

Since the computerization of neurophysiology in the 1970s, the standard quantitative approach has become systematic exploration, with little or no exploitation, within a parametric space. At one extreme is the work of Gallant and Van Essen (1996), who explored a parametric space of hyperbolic gratings, finding that some neurons seemed to prefer them to linear gratings. However, these preferences proved difficult to interpret, because the stimulus space was not very closely linked to shape perception. Pasupathy (2006) explored a parametric space of curved line segments and shapes, inspired by psychological theories of shape perception, while recording from primate V4. She found a number of highly responsive cells, some of which appeared to respond to stimuli with specific curvature in a specific position.

The risk of the exploration-only approach is that it is possible to miss an optimal stimulus without realizing it; the possibility of experimental serendipity is foreclosed, and at the experiment's end, the experimenter is faced with a data set suggesting

interesting directions that cannot be pursued because the neuron is gone. For an exploration-only approach to be successful, the neuron's responses must be sufficiently invariant that some of the small set of predetermined samples will produce a significant response. The odds can be improved if exploration is seeded with stimuli that are highly likely to produce a response – either because they are experimentally overtrained (see below), or because they are part of the animal's natural visual environment.

One example of a primarily exploitation-based approach was pioneered by Tanaka (reviewed in Tanaka, 1996). This “decomposition” approach first uses a moderately large, fixed set of nonparametric “natural” exploration stimuli (e.g. faces, body parts, fruits). However, once a neuron is found that responds to some of these stimuli, its preferred stimuli are decomposed into quasi-parametric geometric primitives by the experimenter (using image manipulation software) in order to find minimal “critical features” that excite the cell.

“Optimization” or “active learning” approaches (reviewed in Benda et al., 2007) are typically explicitly designed to balance exploration and exploitation. These methods attempt to explore until a good region is found; then it switches to exploit this region's maximum. One parametric optimization approach was partially described in Chapter 2, and tested in recordings in cat area 21a (Caywood & Stryker, 2005). Using a parametric configuration of wavelets, it is possible to explore large regions of stimulus space, and once a response is found, to adjust the wavelets locally to maximize the response. Topological invariance detection, as described in Chapter 2, can be integrated into the

approach as one way to exploit response maxima. Along these lines, Yamane et al. (2008) have optimized the response of IT neurons by perturbing random solid shapes with bumps, dents and ridges (formally, non-uniform rational b-splines).

A risk of a parametric approach is the assumption that the parameter space is the right one for the neuron. The experimenter's understanding of the cortical area's function may be incorrect, and so the stimulus space may be inappropriate for its neurons. Or a seemingly reasonable stimulus space may be subtly misleading (as described for color-opponent gratings, in chapter 1, figure 7). Still, if an experiment with parametric stimuli produces an optimal stimulus with a high firing rate, and the neural response changes systematically as the stimulus is changed from the optimum, one can conclude that the parameter space is reasonably informative about the neuron.

Non-parametric approaches to the problem have thus far been primarily exploration-based. While pixel space is a vast parameter space, even compared to a high-dimensional configuration space, restricting stimuli to the space defined by the Fourier statistics of natural scenes has been productive in early visual areas and some intermediate visual areas (reviewed in Wu et al., 2006), probably because stimuli within that subspace are effective in driving V1, which makes monosynaptically connected areas such as V2 and V4 reasonably responsive. The natural scenes approach, thus far, has not been productive in higher visual areas, which respond to more specific image correlations (and thus stimuli in a smaller subspace).

Since non-parametric data are difficult to interpret as visual features, many of these studies attempt to fit a parametric model to the response data (David et al., 2006). The maximally informative dimension method (Sharpee et al., 2004) works along similar lines to fit a complex, pixel-based model with a large number of parameters that explains the spiking response, although it has not yet been applied to higher visual areas.

Földiák (2001) describes a non-parametric exploitation-based approach. He optimizes neurons in primary visual cortex beginning with a random starting stimulus, then adding granular noise and measuring whether the response increases. His particular implementation of the approach is susceptible to nonstationarity or adaptation in the neuron's response, and only yields a single maximum. However, it could be turned into a more general approach by performing his optimization technique several times from different starting points – a limited form of exploration.

Evading dimensionality with training

Another approach is to opt out of the dimensionality problem entirely. One popular technique is to pretrain an animal on an object recognition task (reviewed in Freedman & Miller, 2007). When an animal is rewarded for performing a certain visual task, neurons in the temporal lobe will begin to represent the class of stimuli used in that task, and will develop selectivity for individual stimuli. This has been observed for wire-frame synthetic “paperclip” stimuli (Logothetis et al., 1995) as well as non-naturalistic stimuli resembling Rorschach blots (Sakai et al., 1994). While training significantly

enriches the fraction of cells in the higher ventral stream visual areas, it has not been shown to affect the statistics of intermediate areas such as V4. There is also a significant risk that what is learned will not generalize, because the task becomes “overtrained” – in other words, it takes the cortex from its normal operating mode to an impoverished state. (See Behrmann et al., 2005 for a fascinating observation of disruptive overtraining of a prosopagnosic human subject on an object recognition task.)

Using functional maps to specify the parameter space

Another approach that can reduce the effective dimensionality of the stimulus space is to use functional map information to target recordings. Using fMRI data from the primate, this has been used to successfully record from a highly enriched (97%) population of face-selective cells in the “middle face patch” of primate IT (Tsao et al., 2006). This patch is probably an independent visual area, since it is seemingly too large (16 mm²) to be part of the columnar functional architecture of IT cortex. Knowing that nearly all neurons in this area prefer face-like stimuli allows an experimenter to study them using a relatively low-dimensional parameter space optimized for coding faces. A similar fMRI imaging approach has recently been used to reveal color selective areas in posterior IT (Conway et al., 2007). However, it is still an open question how well this approach generalizes; highly face and color selective areas represent a minority of the ventral stream, and the majority of IT cortex has no known specialization.

Using functional maps to constrain dimensionality

In Chapter 3, we showed that functional maps of orientation and spatial frequency can be imaged in the cat's intermediate visual area 21a, homologous to primate area V4.

Retinotopic maps have also been recorded in area 21a (Tusa & Palmer, 1980). In primate V4, several parameters are reported to be functionally mapped, including retinotopy, orientation and surround suppression (Ghose & Ts'o, 1997). In V3, orientation, luminance and disparity maps have been found (Kaskan et al., 2009). What do all these mapped variables imply about the function of neurons in these areas?

It is tempting to argue that understanding the parameters of an area's map will be the key that unlocks the responses of its neurons. But knowing that an intermediate visual area like V4 has a strong orientation map does not necessarily mean that its primary function is edge detection. Likewise, knowing that it has a retinotopic map seems relatively unhelpful, since the ventral stream appears to eliminate positional information fairly quickly, subject to the limited spread of anatomical connections between cortical areas.

Indeed, it has recently been questioned whether maps have any functional consequence at all (Horton & Adams, 2005). This question is particularly applicable to ocular dominance maps, since ocular dominance domains are not preserved at higher levels of the visual system than V1 (or cat areas 17, 18 and 19). Ocular dominance maps may therefore simply recapitulate the structure of eye-specific layers in the LGN. Even orientation maps seem to be unnecessary for highly selective orientation tuning, at least

in animals with small visual cortices such as squirrels (Van Hooser et al., 2005) and mice (Ohki et al., 2005; Niell & Stryker, 2008). Perhaps orientation domains in cat area 21a merely recapitulate the structure of its major source of input, area 17.

A different approach views functional maps as the inevitable outcome of self-organizing Hebbian developmental processes operating in a sufficiently large, cortical area with strong, competitive lateral connections (von der Malsburg, 1973; Miller, 1996). In this view, the presence of a cortical map for a certain parameter implies that certain correlations are present in the responses of neurons in that area, causing them to wire together into a map. Likewise, its absence implies that such correlations are not present. For a map to self-organize with distinct domains for a certain parameter, the range of values for that parameter must be large relative to the selectivity of the neurons.

It should be noted that self-organization rules are feedback systems, and thus tend to produce all-or-nothing results – i.e., either a strong map or none at all. It is apparently rare for the relevant correlations to be so borderline that a “capricious” map with incomplete expression within or across animals is produced. By itself, the presence of a map for a certain parameter will not be a sensitive indicator of the absolute strength of that parameter’s correlations; in the presence of feedback, similar correlation values can produce different results. However, the relative strength and interleaving of maps of different parameters may be more informative; this may be an interesting question for future theoretical research.

As an example of how functional maps can constrain single-neuron studies, consider Chapter 3's data showing cat area 21a prefers a temporal frequency of 2.0 cycles/s, responds to temporal frequencies from 0.5 to 4.0 cycles/s, and yet lacks a temporal frequency map. Based on this, one could (retroactively) predict the temporal frequency tuning of neurons in area 21a will be broad, probably including low-pass cells, and that even neurons with band-pass tuning will overlap. This is exactly what has been observed (Morley & Vickery, 1997). Thus, an experimenter could conclude that temporal frequency is unlikely to be an important parameter when maximizing the responses of neurons in area 21a.

As another example, consider the observation that in IT cortex, certain stimuli (e.g. human heads) are represented so that as the view changes, the representation progressively moves along the cortical surface. This suggests that view position may be part of the functional architecture in these higher visual areas (Wang et al., 1998; Tsunoda et al., 2001; but see Sato et al., 2009). A recent model has proposed that such functional maps can be learned from the spatiotemporal correlations of natural objects (Michler et al., 2009). However, this body of research seems to introduce a paradox. View-invariant object recognition is theorized to be one of the fundamental operations of the ventral stream (Logothetis et al., 1995; although see Andresen et al., 2009), so why is view the only parameter known to be functionally organized into columns with smoothly varying tuning? Perhaps this map just recapitulates the structure of its less view-invariant input area. However, face and eye position are very important social and behavioral cues (Langton et al., 2000), and perhaps view is actually an important tuning

parameter of single neuron responses in IT cortex. It could be interesting to study the functional connections of these face columns with brain areas involved in social behavior, e.g. the limbic system.

There are several other uses for data about the functional architecture of a higher cortical visual area beyond supporting inferences about the role of that area in the computations underlying vision, such as in developmental neurobiology, neuroanatomy, and studies of neural plasticity. However, these inferences, even if imperfect, can provide useful information about cortical function in areas that lack a clear basis set of optimal stimuli. Functional architecture, as studied here, remains a powerful clue to function itself.

References

- Andresen, D. R., Vinberg, J., & Grill-Spector, K. (2009). The representation of object viewpoint in human visual cortex. *NeuroImage*, *45*(2), 522-536.
- Behrmann, M., Marotta, J., Gauthier, I., Tarr, M. J., & McKeeff, T. J. (2005). Behavioral change and its neural correlates in visual agnosia after expertise training. *Journal of Cognitive Neuroscience*, *17*(4), 554-568.
- Benda, J., Gollisch, T., Machens, C. K., & Herz, A. V. (2007). From response to stimulus: Adaptive sampling in sensory physiology. *Current Opinion in Neurobiology*, *17*(4), 430-436.
- Caywood, M. S., & Stryker, M. P. (2005). Functional maps of cat visual area 21a reveal role in processing spatial form. Society for Neuroscience abstract 508.17.
- Connor, C. E. (2004). Shape dimensions and object primitives. In L. M. Chalupa, & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1080-1089). Cambridge, MA: MIT Press.
- Conway, B. R., Moeller, S., & Tsao, D. Y. (2007). Specialized color modules in macaque extrastriate cortex. *Neuron*, *56*(3), 560-573.
- David, S. V., Hayden, B. Y., & Gallant, J. L. (2006). Spectral receptive field properties explain shape selectivity in area V4. *Journal of Neurophysiology*, *96*(6), 3492-3505.
- Földiák, P. (2001). Stimulus optimization in primary visual cortex. *Neurocomputing*, *38-40*, 1217-1222.
- Freedman, D. J., & Miller, E. K. (2008). Neural mechanisms of visual categorization: Insights from neurophysiology. *Neuroscience and Biobehavioral Reviews*, *32*(2), 311-329.
- Gallant, J. L., Connor, C. E., Rakshit, S., Lewis, J. W., & Van Essen, D. C. (1996). Neural responses to polar, hyperbolic, and cartesian gratings in area V4 of the macaque monkey. *Journal of Neurophysiology*, *76*(4), 2718-2739.
- Ghose, G. M., & Ts'o, D. Y. (1997). Form processing modules in primate area V4. *Journal of Neurophysiology*, *77*(4), 2191-2196.
- Horton, J. C., & Adams, D. L. (2005). The cortical column: A structure without a function. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *360*(1456), 837-862.

- Hubel, D. H., & Wiesel, T. N. (2004). *Brain and visual perception: The story of a 25-year collaboration*. New York: Oxford University Press.
- Kaskan, P. M., Lu, H. D., Dillenburger, B. C., Kaas, J. H., & Roe, A. W. (2009). The organization of orientation-selective, luminance-change and binocular-preference domains in the second (V2) and third (V3) visual areas of new world owl monkeys as revealed by intrinsic signal optical imaging. *Cerebral Cortex*, *19*(6), 1394-1407.
- Langton, S. R., Watt, R. J., & Bruce, I. I. (2000). Do the eyes have it? Cues to the direction of social attention. *Trends in Cognitive Sciences*, *4*(2), 50-59.
- Logothetis, N. K., Pauls, J., & Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, *5*(5), 552-563.
- Michler, F., Eckhorn, R., & Wachtler, T. (2009). Using spatiotemporal correlations to learn topographic maps for invariant object recognition. *Journal of Neurophysiology*, *102*(2), 953-964.
- Miller, K. D. (1996). *Receptive fields and maps in the visual cortex: Models of ocular dominance and orientation columns*. In E. Domany, J. L. van Hemmen & K. Schulten (Eds.), *Models of neural networks III* (pp. 55-78) Springer-Verlag.
- Milner, A. D., & Goodale, M. A. (2006). *The visual brain in action* (2nd ed.). Oxford: Oxford University Press.
- Morley, J. W., & Vickery, R. M. (1997). Spatial and temporal frequency selectivity of cells in area 21a of the cat. *Journal of Physiology*, *501* (Pt 2), 405-413.
- Niell, C. M., & Stryker, M. P. (2008). Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience*, *28*(30), 7520-7536.
- Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P., & Reid, R. C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature*, *433*(7026), 597-603.
- Paninski, L. (2005). Asymptotic theory of information-theoretic experimental design. *Neural Computation*, *17*(7), 1480-1507.
- Pasupathy, A. (2006). Neural basis of shape representation in the primate brain. *Progress in Brain Research*, *154*, 293-313.
- Riesenhuber, M., & Poggio, T. (2004). How the visual cortex recognizes objects: The tale of the standard model. In L. M. Chalupa, & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1640-1653). Cambridge, MA: MIT Press.

- Sakai, K., Naya, Y., & Miyashita, Y. (1994). Neuronal tuning and associative mechanisms in form representation. *Learning & Memory*, 1(2), 83-105.
- Sato, T., Uchida, G., & Tanifuji, M. (2009). Cortical columnar organization is reconsidered in inferior temporal cortex. *Cerebral Cortex*, 19(8), 1870-1888.
- Sharpee, T., Rust, N. C., & Bialek, W. (2004). Analyzing neural responses to natural signals: Maximally informative dimensions. *Neural Computation*, 16(2), 223-250.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience*, 19, 109-139.
- Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic robotics*. Cambridge, MA: MIT Press.
- Tsao, D. Y., Freiwald, W. A., Tootell, R. B., & Livingstone, M. S. (2006). A cortical region consisting entirely of face-selective cells. *Science*, 311(5761), 670-674.
- Tsunoda, K., Yamane, Y., Nishizaki, M., & Tanifuji, M. (2001). Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. *Nature Neuroscience*, 4(8), 832-838.
- Tusa, R. J., & Palmer, L. A. (1980). Retinotopic organization of areas 20 and 21 in the cat. *Journal of Comparative Neurology*, 193(1), 147-164.
- Van Hooser, S. D., Heimel, J. A., Chung, S., Nelson, S. B., & Toth, L. J. (2005). Orientation selectivity without orientation maps in visual cortex of a highly visual mammal. *Journal of Neuroscience*, 25(1), 19-28.
- von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, 14(2), 85-100.
- Wallis, G., & Rolls, E. T. (1997). Invariant face and object recognition in the visual system. *Progress in Neurobiology*, 51(2), 167-194.
- Wang, G., Tanifuji, M., & Tanaka, K. (1998). Functional architecture in monkey inferotemporal cortex revealed by in vivo optical imaging. *Neuroscience Research*, 32(1), 33-46.
- Wu, M. C., David, S. V., & Gallant, J. L. (2006). Complete functional characterization of sensory neurons by system identification. *Annual Review of Neuroscience*, 29, 477-505.

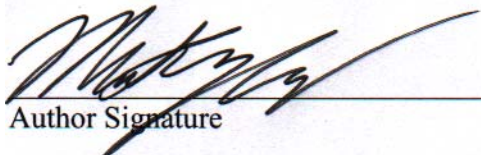
Yamane, Y., Carlson, E. T., Bowman, K. C., Wang, Z., & Connor, C. E. (2008). A neural code for three-dimensional object shape in macaque inferotemporal cortex. *Nature Neuroscience*, *11*(11), 1352-1360.

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses, dissertations, and manuscripts. Copies of all UCSF theses, dissertations, and manuscripts will be routed to the library via the Graduate Division. The library will make all theses, dissertations, and manuscripts accessible to the public and will preserve these to the best of their abilities, in perpetuity.

Please sign the following statement:

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis, dissertation, or manuscript to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.



Author Signature

9/7/2009

Date