

Selectivity and sparseness in the responses of striate complex cells

Sidney R. Lehky^{a,*}, Terrence J. Sejnowski^{b,c}, Robert Desimone^d

^a *Cognitive Brain Mapping Laboratory, RIKEN Brain Science Institute, Hirosawa 2-1, Wako-shi, Saitama 351-0198, Japan*

^b *Howard Hughes Medical Institute, The Salk Institute, 10010 North Torrey Pines Road, La Jolla, CA 92037, USA*

^c *Department of Biology, University of California, San Diego, La Jolla, CA 92093, USA*

^d *Laboratory of Neuropsychology, National Institute of Mental Health, Bethesda, MD 20892, USA*

Received 19 February 2004; received in revised form 23 June 2004

Abstract

Probability distributions of macaque complex cell responses to a large set of images were determined. Measures of selectivity were based on the overall shape of the response probability distribution, as quantified by either kurtosis or entropy. We call this non-parametric selectivity, in contrast to parametric selectivity, which measures tuning curve bandwidths. To examine how receptive field properties affected non-parametric selectivity, two models of complex cells were created. One was a standard Gabor energy model, and the other a slight variant constructed from a Gabor function and its Hilbert transform. Functionally, these models differed primarily in the size of their DC responses. The Hilbert model produced higher selectivities than the Gabor model, with the two models bracketing the data from above and below. Thus we see that tiny changes in the receptive field profiles can lead to major changes in selectivity. While selectivity looks at the response distribution of a single neuron across a set of stimuli, sparseness looks at the response distribution of a population of neurons to a single stimulus. In the model, we found that on average the sparseness of a population was equal to the selectivity of cells comprising that population, a property we call ergodicity. We raise the possibility that high sparseness is the result of distortions in the shape of response distributions caused by non-linear, information-losing transforms, unrelated to information theoretic issues of efficient coding.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Vision; Macaque monkey; Modeling; Information theory; Sparse coding

1. Introduction

How the visual system encodes images includes questions of why individual visual units have particular receptive field organizations, and how populations of units act together. One approach to systematically characterizing receptive fields is to quantify their stimulus selectivities. Here we are concerned with determining selectivities of striate complex cells to “complicated” stimuli, including natural images. This will involve con-

sideration primarily of non-linear models of complex cells, although some single-unit data will be analyzed as well. Complex cells comprise a large percentage of the units in V1, with estimates ranging from around 40–90% of the total (De Valois, Albrecht, & Thorell, 1982; Hubel & Wiesel, 1968; Schiller, Finlay, & Volman, 1976).

Related to selectivity is the idea of sparseness, which attempts to explain receptive fields in terms of information theoretic notions of efficient coding (Atick, 1992; Barlow, Kaushal, & Mitchison, 1989; Field, 1994; Field, 1999; see Simoncelli & Olshausen (2001) and Simoncelli (2003), for reviews). Under sparse coding, a small subset within a neural population will respond strongly to a stimulus, while most will respond poorly. Sparse codes

* Corresponding author. Address: Computational Neuroscience Laboratory, The Salk Institute, 10010 North Torrey Pines Road, La Jolla, CA 92037, USA.

E-mail address: sidney@salk.edu (S.R. Lehky).

can be efficient from an information theoretic standpoint when they arise from sets of receptive fields that are matched to statistical regularities in the environment. Efficiency in this case means reduced redundancy (decorrelation) among responses of units within a population while seeking to preserve image information as far as possible. This decorrelation is produced by linear operations on the stimulus input, in the form of convolution with a set of receptive fields. (Some recent sparseness models go beyond decorrelation, and attempt to pick out higher order visual structures by including an additional non-linear layer; for example see Karklin & Lewicki (2003)).

In this study, we are concerned with what we call “non-parametric” selectivity. Non-parametric selectivity is determined by the shape of the probability density function (pdf) of response magnitudes to a large set of stimuli (Fig. 1B). The shape is quantified by measures such as kurtosis or entropy, which will be explained more fully in Section 2. These measures seek to pick out pdf’s that are “peakier” than a Gaussian, and with heavier tails. Such distributions indicate responses that are close to spontaneous levels for most stimuli but occasionally are much larger, with intermediate responses being less common than under a Gaussian distribution. The Gaussian distribution serves as a reference distribution indicative of low selectivity. In contrast to this non-parametric selectivity, the more commonly used parametric selectivity does not depend on response probability distributions, but rather measures the bandwidth of response tuning curves to some parameter, such as orientation or spatial frequency (Fig. 1A).

The use of non-parametric selectivity is appropriate when dealing with “complicated” stimuli, such as natural images, that are not ordered by some metric. One might want to use such stimuli to characterize the sys-

tem under more ecological conditions, or perhaps because receptive fields are so complicated that we do not understand along what parameters they parse their input.

Selectivity is defined here in terms of the probability distribution of responses of a single unit to a population of stimuli. Sparseness, on the other hand, is in some sense the converse of selectivity. Sparseness is determined by the distribution of responses of a population of units to a single stimulus. In both the cases, although the probability distributions are measuring different things, the measurement of the distribution shapes can be calculated using the same methods.

To compare our terminology with that used in some previous studies, our “non-parametric selectivity” is equivalent to the “lifetime sparseness” used by Willmore and Tolhurst (2001). What Vinje and Gallant (2000, 2002) call the sparseness of complex cells is, under our terminology, not sparseness but non-parametric selectivity.

Although selectivity and sparseness measure different things, they are nonetheless related quantities. In particular, the average sparseness of a population over multiple stimulus inputs must equal the average selectivity of the neurons within the population (Földiák, 2002), provided responses of units are uncorrelated. Therefore populations that exhibit sparse coding will be composed of neurons showing high selectivity. We label systems exhibiting this equivalence between selectivity and sparseness as “ergodic”. (This is a term taken from statistical mechanics, where the average of a single system across time is compared with the average of an ensemble of systems at one time.) The relationship between sparseness and selectivity will be further discussed in the section on ergodicity.

Algorithms that generate sparse codes have been applied to natural images (Bell & Sejnowski, 1997;

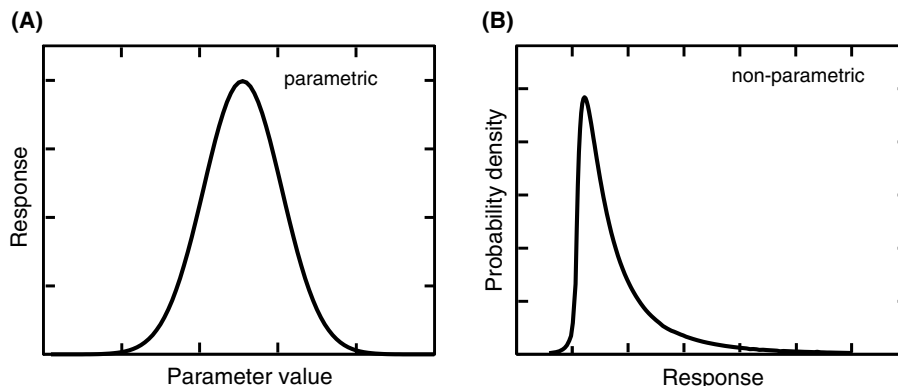


Fig. 1. Two approaches to defining selectivity. (A) Parametric selectivity: Selectivity is defined in terms of the bandwidth of the tuning curve for a stimulus parameter (orientation, spatial frequency, etc.). (B) Non-parametric selectivity: selectivity is defined by a statistic, such as kurtosis or entropy, describing the shape of the probability distribution of response magnitudes. This measure is appropriate when a stimulus set, such as natural images, is not ordered by some parameter.

Olshausen & Field, 1996, 1997; van Hateren & van der Schaaf, 1998). The resulting receptive fields resemble Gabor functions, similar to striate simple cells (Jones & Palmer, 1987), although not incorporating the rectifying non-linearity of actual simple cells. These results have been used to support the hypothesis that V1 implements statistically efficient coding through sparseness.

Moving to V1 complex cells, Vinje and Gallant (2000, 2002) have also interpreted observation of sparseness in their data as supporting the efficient coding hypothesis. While both our data and modeling confirm that complex cell responses are indeed sparse, we shall question whether a high sparseness index is indicative of efficient coding in this case. Although efficient coding can be a useful basis for understanding receptive field organization in the early, linear stages of visual processing, other conceptual frameworks may be necessary for dealing with non-linear, information-losing transforms such as occur in complex cells and beyond in the visual pathways. After documenting some response properties of complex cells and modeling them, we shall present a more general discussion of the limitations of information theoretic modeling in the context of such a system.

2. Materials and methods

2.1. Data acquisition

Twenty-four V1 units were recorded from an anesthetized macaque monkey (*M. fascicularis*). Receptive fields were within 5° of fixation. All were known to be complex cells from the equality of responses to both white and black bars against a gray background. Stimuli consisted of 157 synthetic patterns of two types, 78 random textures and 79 shaded paraboloid figures, displayed within a circular aperture of 1.5° (examples shown in Fig. 2). Patterns were presented in random order, and interspersed with more traditional stimuli such as bars and gratings, although responses to those simpler stimuli are not considered here. Presentation of each pattern was repeated 30 times. Stimulus duration was 200 ms, with a 250 ms blank period between stimuli, a fast presentation pace in order to accumulate data for many patterns. Details of the physiological techniques have been described previously (Lehky, Sejnowski, & Desimone, 1992), and the data presented here represent a subset of data from that publication.

2.2. Data analysis

The peristimulus time histogram (PSTH) for each image was determined from the spike train running with a 50 ms lag relative to stimulus presentation, to take V1




Fig. 2. Examples of synthetic stimulus images, and resulting responses. Responses are taken from experimental data for a single complex cell. By presenting a large set of images and tabulating the probability distribution of the resulting responses, the non-parametric selectivity of a unit can be measured.

latency into account. Spike rates were estimated using a two-stage adaptive kernel technique (Silverman, 1986). In the first stage, each spike was convolved with a Gaussian having $\sigma_{\text{init}} = 20$ ms and the resulting curves summed to give a preliminary PSTH. In the second stage, going back to the raw spike train, each spike was convolved with a Gaussian whose σ was inversely related to the local spike rate estimated in the preliminary PSTH, with $\sigma = \sigma_{\text{init}} \sqrt{m/r}$, where m is the geometric mean of the spike rate over the stimulus duration, and r is the spike rate at a given instant. Data from all repetitions of each stimulus were pooled prior to forming the kernel estimate.

For each unit, 157 PSTH's were determined, one for each stimulus image. Each PSTH was then reduced to a single summary number. We use mean response as the summary statistic in the presentation below, but using peak responses instead would not have made much difference. From these 157 numbers a probability distribution of response magnitudes was compiled. As the observed probability distributions for the 24 units were positively skewed, for descriptive purposes they were fit with gamma distributions:

$$f(r | a, b) = \frac{1}{b^a \Gamma(a)} r^{a-1} e^{-r/b} \quad (1)$$

2.3. Measuring selectivity

Non-parametric selectivity measures should depend solely on the response distribution shape, and not its scale (variance) or location (mean). A number of measures are available, although all have drawbacks of various sorts and there is room for the development of new methods for dealing with this issue.

2.3.1. Kurtosis

The most common selectivity index is the reduced kurtosis of the response probability distribution. This is the normalized fourth moment of the distribution, reduced by subtracting three:

$$S_K = \frac{\langle (r_i - \bar{r})^4 \rangle}{\sigma^4} - 3 \quad (2)$$

where r_i is the unit's response to the i th image, \bar{r} is the mean response over all images, σ is the standard deviation of the responses, and $\langle \cdot \rangle$ is the mean value operator. Subtracting three normalizes the measure so that a Gaussian has a kurtosis of zero. Larger values of S_K correspond to greater selectivity. A practical problem with using kurtosis is that it involves raising data to the fourth power, which makes estimates of this measure highly sensitive to noise in the data.

2.3.2. Activity fraction

Another selectivity statistic in the literature is:

$$A = \left(\sum_{i=1}^n r_i/n \right)^2 / \sum_{i=1}^n (r_i^2)/n \quad (3)$$

(Rolls & Tovee, 1995), where n is the number of stimulus images. It measures the fraction of units active on average over a set of inputs, generalized for continuous-valued units rather than binary ones. Small values of A indicate high selectivity. The activity fraction measure was slightly modified by Vinje and Gallant (2000):

$$S_A = \frac{1 - A}{1 - 1/n} \quad (4)$$

thereby inverting and rescaling the index.

A problem with activity fraction is that, in addition to being a measure of probability distribution shape, it is also sensitive to the distribution's mean and variance. The close relation between activity fraction and variance can be seen by comparing Eq. (3) with a common formula for variance:

$$\sum_{i=1}^n (r_i^2)/n - \left(\sum_{i=1}^n (r_i/n) \right)^2 \quad (5)$$

The activity fraction index takes the two terms in the variance formula and divides rather than subtracts them, leading to a measure that is still proportional to variance.

To properly use this measure, the data should be standardized for mean and variance beforehand. We will not be using activity fraction, in preference for the entropy measure described next.

2.3.3. Entropy

Introduced here is a measure of selectivity based on the entropy of the response probability distribution. This has the advantage of connecting more directly to information theoretic definitions of sparseness than does kurtosis. In this measure, selectivity is quantified as the decrease in entropy relative to a Gaussian distribution, which has maximum entropy for a fixed variance:

$$S_E = H_G - H(r) \quad (6)$$

where H_G is the entropy of a Gaussian, and $H(r)$ is the entropy of a unit's response distribution taken from the data. Equating high selectivity with low entropy captures the characteristic of a highly selective unit having little response to most stimuli and a large response to a few. The value of S_E has a minimum of zero, and increases with no upper bound as selectivity increases (because the entropy of the data $H(r)$ can have negative values, as it involves a continuous-valued variable, response magnitude). Since the variance of a distribution affects its entropy, all calculations are done after rescaling the data to unit variance.

The entropy of the Gaussian reference distribution is given by Rieke, Warland, van Steveninck, and Bialek (1997):

$$H_G = \frac{1}{2} \log_2(2\pi e \sigma^2) = 2.074 \text{ bits} \quad (7)$$

with variance normalized to one. Entropy of the probability distribution of the data is:

$$H(r) = - \int p(r) \log_2(p(r)) dr \quad (8)$$

where $p(r)$ is the response probability density function (pdf). For practical calculations Eq. (8) is discretized to:

$$H(r) = - \sum_{j=1}^M p(r_j) \log_2(p(r_j)) \Delta r \quad (9)$$

where responses from n stimulus images have been placed in M bins.

The value of $H(r)$ will depend on the bin size Δr , which in turn depends on the number of bins into which the response range has been divided. The number of bins therefore needs to be standardized, and this is done by defining the number to be $M = \sqrt{n}$, where n is the number of images in the stimulus set. Expressing the selectivity index S_E (Eq. (6)) in terms of Eqs. (7) and (9), we get:

$$S_E = 2.074 + \sum_{j=1}^M p(r_j) \log_2(p(r_j)) \Delta r \quad (10)$$

If selectivity of a neuron is defined in terms of its response entropy, then selectivity can be related to the mutual information between stimulus and response:

$$\begin{aligned} I(r, s) &= H(r) - H(r|s) \\ &= -S_E + (H_G - H(r|s)) \end{aligned} \quad (11)$$

$H(r|s)$ is the conditional entropy of the response given the stimulus, and is essentially entropy due to noise in the system. Given fixed noise entropy $H(r|s)$ (as well as H_G , fixed by definition), increasing selectivity decreases mutual information. In other words, there is a conflict between maximizing selectivity of a unit and maximizing information transfer.

A drawback to the entropy measure is that each cell needs to be tested with a large number of stimulus images in order for entropy to be accurately estimated. This is demonstrated in Fig. 3, which shows the entropy of different-sized samples of a Gaussian distributed random variable. Entropy asymptotically approaches the theoretical value from below as stimulus set size increases. In general, it is not possible to accurately determine the shape of a probability distribution from a small data sample, particularly the tails, and that will affect calculations of entropy.

A second potential problem with S_E is that there are situations where, unlike in Fig. 3, the entropy does not converge as the size of the stimulus set increases. However, that will not be an issue for any probability distributions encountered in this study.

A general point about non-parametric selectivity is that its value depends not only on the receptive field organization of a cell, but also on the particular stimulus set used. Non-parametric selectivity is always relative to the stimulus set, not absolute. This is not true of para-




Fig. 3. Entropy of a Gaussian distribution as a function of the size of the stimulus set. Entropy increases with the size of the stimulus set up to an asymptote of 2.07, which is the theoretical value for a Gaussian with unit variance. This demonstrates that under the entropy measure of selectivity (Eq. (10)), a probability distribution will have an artifactually low entropy if stimulus sample size is small. A stimulus set of at least several dozen images is required to get a reasonable estimate of this parameter.

metric selectivity, where the parameter of interest by its nature defines the stimulus set.

2.4. Modeling

In addition to data from striate complex cells, we looked at selectivities of model units. This allowed us to expand the stimulus set beyond what was used experimentally, and allows examination of the effects of varying receptive field properties in a well-defined manner.

Two models of complex cells will be presented, with rather subtle differences in their receptive fields that lead to large differences in their selectivities. The first is based on Gabor functions. For this model, the complex cell responses are defined as the quadrature pair summation of two subunits with Gabor receptive fields (Emerson, Korenberg, & Citron, 1992; Szulborski & Palmer, 1990), at sine and cosine phase respectively:

$$C = \sqrt{(G_{\sin} * S)^2 + (G_{\cos} * S)^2} \quad (12)$$

where C is the complex cell response, and $G * S$ is the result of convolving a Gabor subunit with the stimulus image. This class of energy model for complex cells is widely used (Adelson & Bergen, 1985; Heeger, 1992; Ohzawa, DeAngelis, & Freeman, 1990; Pollen & Ronner, 1983; Spitzer & Hochstein, 1988). The Gabor subunits resemble striate simple cells, except that they are not half-wave rectified. Each Gabor subunit can be viewed, therefore, as representing the pooled output of a pair of simple cells with opposite contrast polarities.

The Gabor functions are sinusoidal plane waves with spatial frequency f , orientation θ , and phase ϕ , under a Gaussian envelope:

$$G_{\phi} = e^{-\frac{1}{2} \left[\left(\frac{x'}{\sigma_x} \right)^2 + \left(\frac{y'}{\sigma_y} \right)^2 \right]} \sin(2\pi f x' + \phi) \quad (13)$$

where x' and y' are within the rotated coordinate system:

$$\begin{aligned} x' &= x \cos(\theta) + y \sin(\theta) \\ y' &= -x \sin(\theta) + y \cos(\theta) \end{aligned} \quad (14)$$

Twenty-four model complex cells were created, having four spatial frequency tuning curves and six orientations per spatial frequency. Peak spatial frequencies were $f = [0.031, 0.062, 0.125, \text{ and } 0.250]$ cycles/pixel, which given 128×128 input images translates to $[4, 8, 16, 32]$ cycles/picture. Orientations were $[0^\circ, 30^\circ, 60^\circ, 90^\circ, 120^\circ, 150^\circ]$. With respect to the Gaussian envelope, σ_x defines the spatial frequency bandwidth, and was set to $0.65/f$. The ratio σ_y/σ_x defines the orientation tuning bandwidth, and was set to 1.7. These appear to be physiologically realistic values (De Valois et al., 1982; Kulikowski & Vidyasagar, 1986), producing a spatial frequency bandwidth of 1.6 octaves and orientation bandwidth of 33° , full width at half-maximum.

Sine and cosine Gabor functions are only approximate quadrature pairs. This reflects the fact that while the sine Gabor integrates to zero, the cosine Gabor does not, so that the two functions have different mean values or different zero spatial frequency amplitudes. A quadrature pair should be identical other than a 90° phase shift, and the Gabor pair does not fulfill that condition. In observational terms, this means that complex cells constructed from Gabor pairs are only approximately phase invariant, showing a ripple in their responses as a grating is drifted across their receptive fields.

In addition to the above Gabor model units, a second set of 24 model complex cells was constructed that were perfectly phase invariant. We call this the Hilbert model, because synthesizing these units required the use of the Hilbert transform. In this model, the two subunits were, first, a sine Gabor the same as before, and second, instead of a cosine Gabor, the Hilbert transform of a sine Gabor. Given this pair of subunits, they were combined in the same manner as previously:

$$C = \sqrt{(G_{\sin} * S)^2 + (H(G_{\sin}) * S)^2} \quad (15)$$

where $H(\cdot)$ represents the Hilbert transform. The Hilbert transform in essence takes the Fourier transform of a waveform, shifts the phase by 90° , and then does the inverse Fourier transform, producing a waveform that is identical to the original other than the phase shift. We used the “hilbert” command in the signal processing toolbox of Matlab (www.mathworks.com), and mathematical details of implementing the transform are given in their documentation. Morrone and Burr (1988) have previously used the Hilbert transform for modeling receptive fields in the context of biological vision.

The Hilbert transform of a sine Gabor is almost identical to a cosine Gabor (Fig. 4), except that it integrates to zero. The difference in the receptive field profiles is so small it would be almost impossible to detect through mapping experiments. The difference does show up clearly, however, in their spatial frequency amplitude spectra (Fig. 5) at low frequencies.

So, to summarize, we had two sets of 24 model complex cells. One was based on the Gabor model and the other on the Hilbert model, and the only difference was a tiny change in the receptive field profile for one of their two subunits.

The model units were tested with the same set of 157 images that had been presented to the actual complex cells (example images shown in Fig. 2). During testing, each image was centered on the model receptive field, in the same manner as was done during the actual neurophysiological experiments. Statistics were collected for response probability distributions, and once it had been verified that the model units had similar properties to those seen in the data, they were presented with an expanded stimulus set of 500 natural images (Fig. 6). Each

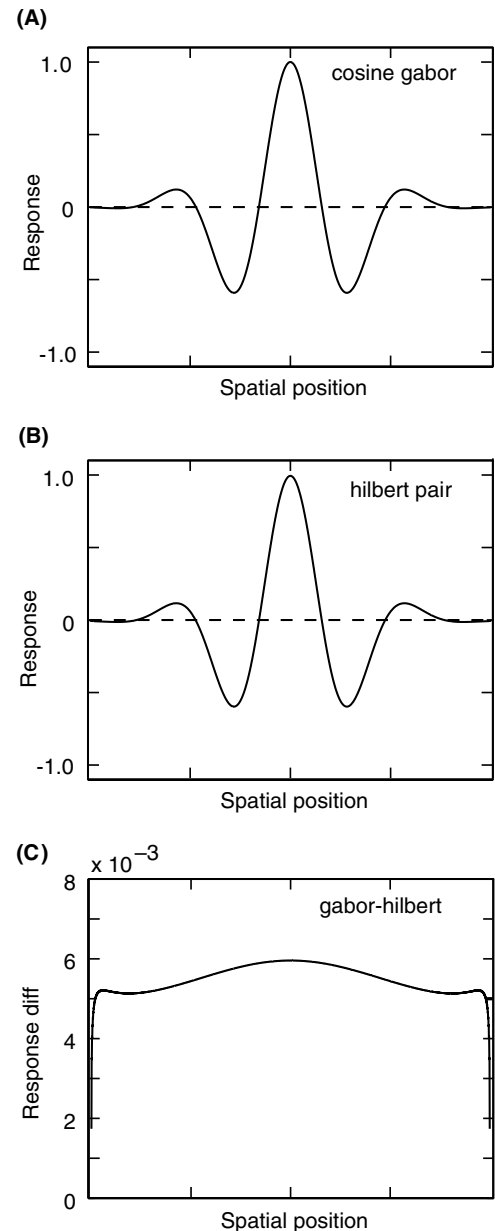


Fig. 4. Subunit receptive field profiles for Gabor and Hilbert models of complex cells. Although differing only slightly, they lead to very different statistics in their responses to natural images. (A) Gabor model: the two subunits had receptive fields consisting of a sine Gabor function, and a cosine Gabor (pictured here). (B) Hilbert model: the two subunits had receptive fields consisting of a sine Gabor function, the same as before, and the Hilbert transform of a sine Gabor (pictured here). This Hilbert transform receptive field resembles a cosine Gabor, except that it integrates to zero while the Gabor does not. (C) Difference between the cosine Gabor and the Hilbert transform receptive fields (note scale).

of the natural images was sampled multiple times by each model unit, as the receptive field was shifted about to different patches within the image. The total number of sampled image patches ranged from 500 at the lowest spatial frequency up to 24,000 at the highest frequency.

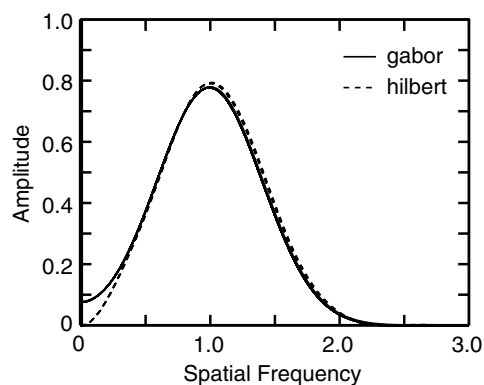


Fig. 5. Spatial frequency spectra of a cosine Gabor and the Hilbert transform of a sine Gabor. Differences in response statistics to natural images for the two models arise because of these different spatial frequency spectra. Spatial profiles of the two receptive fields are pictured in Fig. 4. Spatial frequency is on a relative scale, with curve peak set to 1.0.



Fig. 6. Examples from among the 500 natural images presented to model complex cells.

High frequency units had smaller receptive field diameters, allowing a greater number of image samples.

3. Results

Responses of an example striate complex cell to several images are shown in Fig. 2. Clearly the response magnitudes differ substantially from image to image, and it is straightforward to examine the statistical distribution of these responses over a set of inputs. In addition to differences in response magnitude there are differences in response temporal waveforms, which will not be considered here.

Response probability distributions for three example neurons from the data are shown in Fig. 7. They are positively skewed and are well fit by gamma distribu-

tions (Eq. (1)), also shown in Fig. 7, although no theoretical significance is placed on that description. These three examples illustrate distributions with selectivities near the maximum, minimum and median over all recorded units, with both the kurtosis S_K (Eq. (2)) and entropy S_E (Eq. (10)) selectivity indices indicated in each case.

The distributions of the selectivity indices for the 24 neuron in the data are shown in Fig. 8. Also given are their median values, $S_K = 0.84$ and $S_E = 0.23$, indicating that selectivities of units are typically substantially




Fig. 7. Histograms of response probability distributions for three macaque complex cells. Cells were presented with a set of 157 synthetic stimulus images. Also shown is the best-fit gamma function in each case, and spontaneous activity (vertical dashed line). Both the kurtosis (S_K) and entropy (S_E) measures of selectivity are given for each unit. These examples show response distributions for units with high and low selectivities, plus a distribution close to the median.

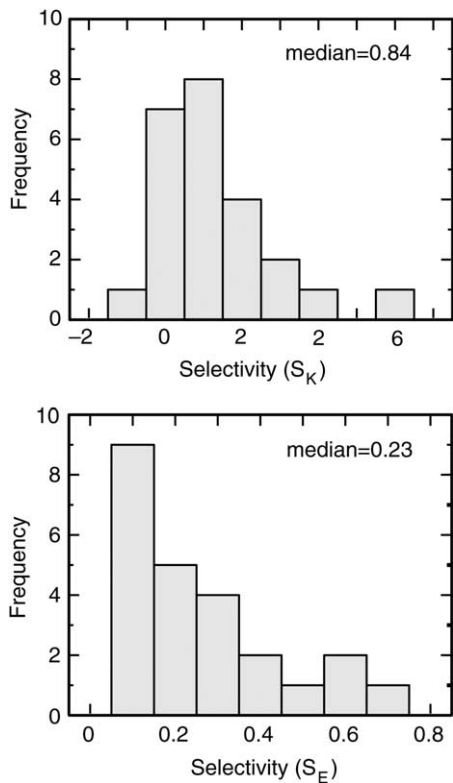


Fig. 8. Distributions of kurtosis and entropy measures of selectivity for 24 macaque complex units presented with synthetic images.




Fig. 9. Negative correlation in the data between selectivity and the median stimulus response of the 24 complex cells. This is shown for both the kurtosis (S_K) and entropy (S_E) measures of selectivity.

greater than for the Gaussian reference distribution (which would have $S_K = S_E = 0$).

There was a negative correlation between the median activity of a neuron and selectivity (Fig. 9). The least active neurons were the most selective. However, the level of activity in itself does not determine selectivity. For example, a Gaussian distribution with a high mean or a low mean both indicate exactly the same selectivity. There must be a change in the shape of the response distribution that correlates with activity, and in this case the distributions change shape by becoming more skewed as mean activity drops. Highly skewed distributions register as more selective under our measures. Increased skewness as the mean decreases is a general property of random variables whose variances are large relative to their means, and which are also constrained to have only non-negative values (as is the case with firing rates). The relationship shown in Fig. 9 is an indication that high selectivity measures can arise from distortions in the shape of response probability distributions due to nonlinearities in the system, rather than the sort of linear transforms discussed by Field (1994), as will be discussed further below.

3.1. Selectivity in model complex cells

The model complex cells had positively skewed response probability distributions, and negative correlations between response magnitude and selectivity, both features of the data. Fig. 10 shows the distributions, for synthetic images, of the kurtosis and entropy selectivity indices under the Gabor model (Eq. (13)), and the Hilbert model (Eq. (15)). The median selectivities for the two models and for the data, again for synthetic images, are summarized in Table 1. The table shows that the selectivity of the data is bracketed by the two models, being greater than the Gabor model and less than the Hilbert model, but closer to the Gabor model. Being intermediate to the two models in this manner suggests that actual complex cells have subunits that do not integrate to zero, with the discrepancy from zero being somewhat smaller than for Gabor model units. An implication of this is that actual complex cells should not exhibit perfect phase invariance but show a ripple response to drifting gratings. This is in fact an observed feature of complex cells (reviewed by Spitzer & Hochstein, 1988). Higher selectivity when receptive fields integrate to zero, as shown in Table 1, has previously been noted by Baddeley (1996a) in modeling of linear units resembling retinal ganglion cells.

Having verified that the modeling provides a reasonable match to the data for synthetic stimulus images, we can now look at model responses to natural images. These are also shown in Table 1. Selectivities of model complex cells are higher for natural images than synthetic ones (with the exception of one condition).

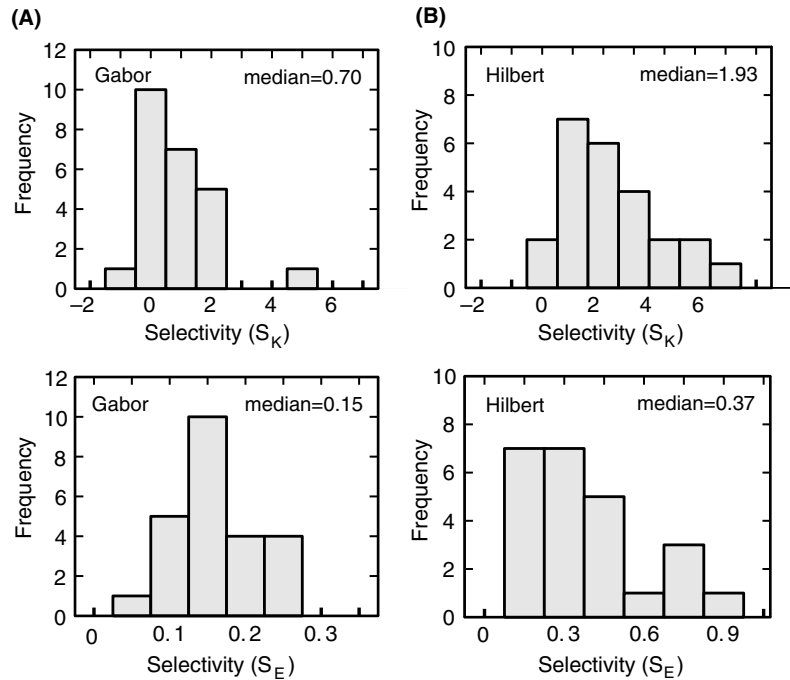


Fig. 10. Distributions of kurtosis and entropy measures of selectivity for Gabor and Hilbert model complex cells presented with synthetic images. These are analogous to selectivity measures for the data shown in Fig. 8. (A) Gabor model. (B) Hilbert model.

Table 1

Selectivity measures for complex cell experimental data and the two models of complex cells

	Kurtosis		Entropy	
	Synthetic	Natural	Synthetic	Natural
Gabor model	0.7	2.2	0.15	0.08
Data	0.8	{2.8}	0.23	{0.24}
Hilbert model	1.9	7.1	0.37	0.53

Results for both synthetic and natural images are included. Values in brackets are estimated by linear interpolation between the Gabor and Hilbert models, weighed in accord with the results for synthetic images.

Selectivities of biological complex cells to natural images were estimated by performing a linear interpolation between the Gabor and Hilbert models, weighed by the synthetic image results. The estimated selectivities are given within brackets in Table 1. The complex cell kurtosis estimated here for natural images, 2.8, is not far from to the value of 4.1 reported by Vinje and Gallant (2000, 2002) for macaque complex cells.

3.2. Spatial frequency dependence of selectivity

Selectivities are higher for the Hilbert model than the Gabor model, regardless of which measure is used (Table 1), and the question arises why that is. To examine this issue, we start by plotting the response probability distributions separately for model complex cells tuned to different spatial frequencies. At each spatial frequency, responses for units tuned to all orientations

were pooled, as we did not note interesting orientation specific effects. This is done for the Gabor model in Fig. 11A, and the Hilbert model in Fig. 11B, in both the cases using natural images as inputs.

Examination of Fig. 11 shows that there is a large difference between the models in their response probability distributions for different spatial frequencies. Response distributions of the Gabor model are almost independent of spatial frequency tuning, while those of the Hilbert model show distributions whose skewness (and selectivities) increase sharply for units tuned to higher spatial frequencies.

A possible explanation for this difference arises when one examines the spatial frequency amplitude spectra of the stimuli (Fig. 12), and the spectra of the Gabor and Hilbert model subunits (Fig. 5). The image amplitude spectra (both natural and synthetic) exhibit a $1/f$ frequency dependence, as has been widely reported (for example, Baddeley, 1996a; Field, 1987; Ruderman & Bialek, 1994). The stimulus amplitudes at low spatial frequencies are enormous compared to those at high frequencies. Although it is the phase spectrum and the alignment of phases to produce localized forms that lead to the important structures in natural images, the amplitude spectrum can be used as an index indicating stimulus intensity at different spatial scales, when averaging over a large set of image samples.

Continuing development of the argument here, the Gabor model includes subunits with significant sensitivity to low spatial frequencies, including zero frequency,

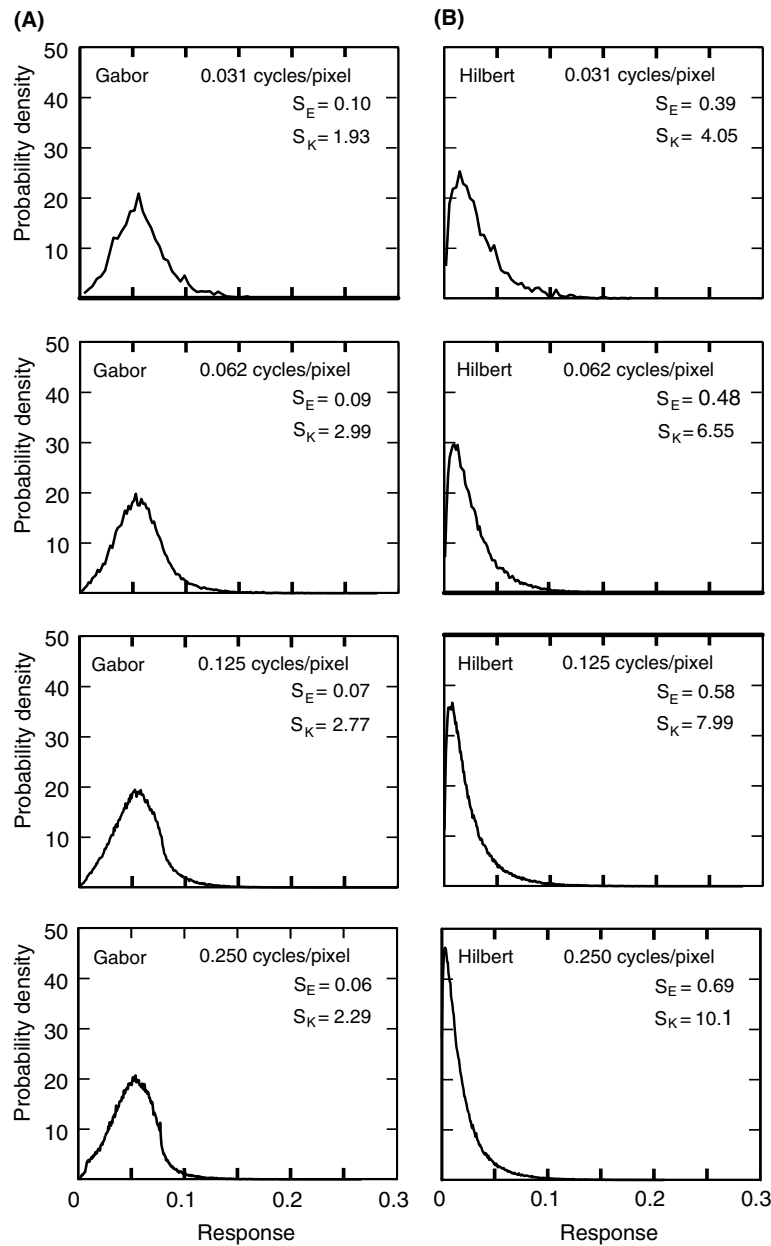


Fig. 11. Response distributions of model complex cells with different spatial frequency tunings, when presented with natural images. Peak frequencies of the spatial tuning curves are indicated, as are the selectivity measures. These results show that the selectivities of Gabor units are independent of spatial frequency tuning, while selectivities of Hilbert units increase for high frequency units. Response is on a relative scale with the optimal stimulus producing a value of 1.0. (A) Gabor model. (B) Hilbert model.

and this low frequency sensitivity is largely independent of the position of the peak. Given a $1/f$ input stimulus, the response distributions of such units will be dominated by the strong, non-specific signal at the left tail of their frequency tuning, rather than the weak signal at their peak. We therefore see Gabor model responses in Fig. 11A that are independent of the tuning curve peak, and with relatively low selectivity.

The Hilbert model, on the other hand, incorporates subunits that have reduced sensitivities to low spatial frequencies and are completely insensitive to zero fre-

quency (Fig. 5). Without sensitivity to the strong signal at frequencies near zero, Hilbert response distributions are more dependent on the peak location of their spatial frequency tuning curves than Gabor distributions are. As spatial frequency tuning increases, localized structures at those spatial scales become increasingly rare, and responses of Hilbert units show increases selectivity.

As actual complex units have properties intermediate between the Gabor and Hilbert models (Table 1), we predict that their selectivities will show moderately spa-

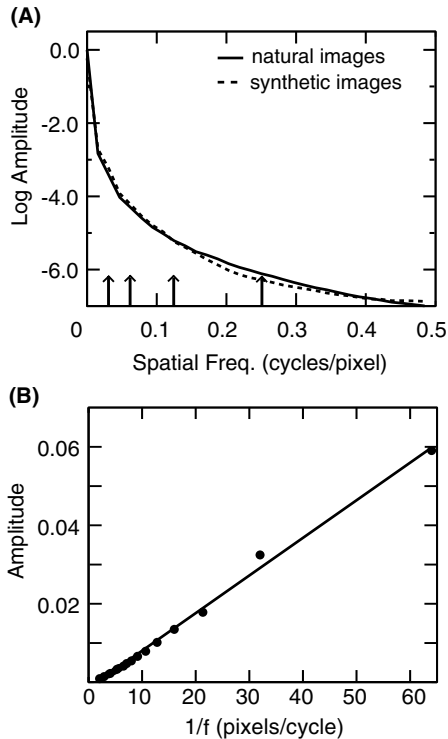


Fig. 12. Spatial frequency spectra of the stimulus images. (A) Amplitude spectra of synthetic and natural stimulus images, averaged over all images in each set. The four vertical lines at the bottom indicate the spatial tuning peaks of the four sets of model units whose response distributions are graphed in Fig. 11. (B) Amplitude spectra plotted as a function of inverse frequency, showing that the stimulus images had approximately $1/f$ spectra.

tial frequency dependence, in between that which appears in Fig. 11A and B.

3.3. Response distribution tails

Probability distributions associated with high selectivity or sparseness are often described as being “heavy-tailed”, so we shall examine tail properties here. Events falling on the tails are by definition rare, so it is difficult to characterize them with the small samples typically available from the neurophysiological data. However with a model, such as we developed for complex cells, this limitation is bypassed. Heavy-tailed distributions have been used to model a variety of systems in economics, communications engineering, and physics (Adler, Feldman, & Taqqu, 1998), following the seminal work of Mandelbrot (1963), for situations in which there is an occasional large extremal event mixed in with the usual small events.

The tail of a distribution is defined as the complement of the cumulative distribution function $\bar{F}_r(r) = 1 - F(r)$ (Bryson, 1983). A “heavy-tail” is one that decreases more slowly than some reference distribution. An exponentially decaying tail is commonly used as the dividing




Fig. 13. Tail of the response distribution of a high spatial frequency Hilbert unit. Linearity of the semi-log plot indicates that the tail follows an exponential decay. Dashed line shows linear best fit.

line between light-tailed and heavy-tailed distributions. This would classify the Gaussian distribution (whose tail is the square of an exponential) as light-tailed, and distributions with power law tails (the Cauchy distribution for example) as heavy-tailed because they decay much more slowly.

We examined the tail of the response distribution showing the highest selectivity, that coming from a high spatial frequency Hilbert model complex cell (Fig. 11B, bottom panel). A semi-log plot of its right tail yields a straight line (Fig. 13), indicating that it does not follow a power law but rather is exponentially decaying. This is consistent with previous reports of exponential tails for striate simple cells and inferotemporal units (Baddeley et al., 1997; Treves, Panzeri, Rolls, Booth, & Waksman, 1999). The response distribution of the model complex cell is therefore not heavy-tailed, despite being leptokurtotic. Although high kurtosis arises when a distribution is heavy-tailed, it can also arise if the distribution is thin-tailed and skewed. Skewness and not heavy-tailedness appears to be the source of the high measures of selectivity seen in complex cells, judging from these modeling results.

3.4. Ergodicity: the relationship between selectivity and sparseness

Another issue is the relationship between the distribution of responses of single units across time when presented with a set of images (selectivity) and the distribution of responses within a neural population measured simultaneously (sparseness). If neural responses are presented on a matrix in which each column represents a different neuron and each row represents a different stimulus image, the question is how do response distributions compare if they are measured along columns (selectivity) or along rows (sparseness). Why be concerned with this issue? As a technical matter, it's far easier to isolate units one at a time while presenting each

one with many stimuli, rather than record an entire population simultaneously, even though the population sparseness may be the quantity of theoretical interest. Therefore it is useful to understand the relationship between selectivity and sparseness. For convenience in the following discussion, we shall use the terms “selectivity” and “sparseness” to refer to the response distributions themselves, and not just summary statistics on those distributions such as entropy or kurtosis.

If the selectivities for individual units are the same as the population sparseness, we call the neural system *ergodic*, by analogy with the concept from statistical mechanics. To refine this a bit more, if the selectivity of each individual unit is the same as average population sparseness, then the system is *strongly ergodic*. Obviously under strong ergodicity all units have the same selectivity. It should be noted that units can respond to quite different sets of stimuli and yet still have identical selectivities (response probability distributions), as we are defining the term. If individual units have different selectivities, but the average selectivity is the same as the average population sparseness, then the system is *weakly ergodic*. Weak ergodicity necessarily occurs if responses of units are uncorrelated.

When examining ergodicity in simulations here, the population size of model units was expanded from the previous 24 to 126 by increasing the number of different spatial frequency and orientation tunings included in the population. Fig. 14 shows response probability distributions for Gabor model complex cells plotted both ways, by individual units (selectivity) and by population (sparseness). Fig. 15 shows selectivity and sparseness for Hilbert model units.

The top panel of Fig. 14 shows response distributions for individual units (selectivities). The distributions are displayed with histograms having 10 bars. Each color represents the responses of a different unit. By following a color across the histogram, one can see the response distribution of a particular unit when presented with the entire stimulus set. Low spatial frequency units are at the blue end of the spectrum, and high frequency units are at the red end. The black outline around each bar indicates the average distribution over all units.

The bottom panel shows response distributions across the population (sparseness) for individual stimuli. The population always remains the same, and each color represents the response distribution of that population to a different stimulus image. The black outline bars show the average distribution of responses for the population over the stimulus set (average sparseness). There is high variability for the different colors within each histogram bar, as the response distribution of the population jumps about with each individual image.

The black outline bars in the two panels are practically identical, which indicates that the average selectivity of individual units is the same as the average

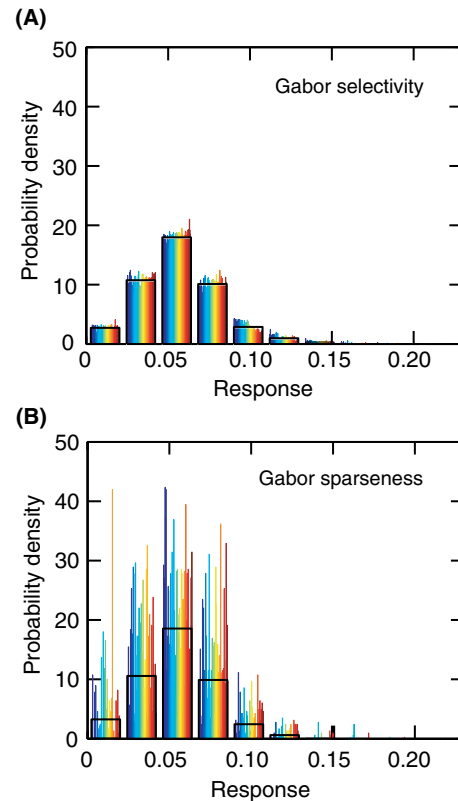


Fig. 14. Comparison of selectivity and sparseness in Gabor model units. Responses are from a population of 126 model units with various orientation and spatial frequency tunings, when presented with natural images. (A) Histogram with 10 bars showing response distribution of individual units presented with a series of images (selectivity). Each color indicates a different unit, and the black outline bars show the average distribution over all units. (B) Histogram showing distribution of responses within an entire population measured simultaneously (sparseness). Each color corresponds to a different stimulus image presented to the population. The black outline bars show the average distribution of the population over the entire set of stimulus images. Response is on a relative scale with the optimal stimulus producing a value of 1.0. (The crowding of colored bars causes them to bleed into white spaces between them, making the solid colored areas appear taller than they actually are.)

sparseness of the population. That satisfies the condition for the system to be weakly ergodic. Looking more closely at the response distributions for the individual Gabor complex cells in Fig. 14A, they all appear to be very similar regardless of spatial frequency tuning. That is indicated by the flatness within each histogram bar of all the different colors. Therefore, the selectivity of each individual unit matches the average sparseness of the population. Thus, Gabor model complex cells have response distributions that are strongly ergodic.

Moving from Gabor units to Hilbert units, Fig. 15 is analogous to Fig. 14. The near identity of the black outline histogram bars in the top and bottom panels of Fig. 15 show that Hilbert units satisfy weak ergodicity, as did Gabor units. In other words, average selectivity is the same as average sparseness. However, unlike Gabor

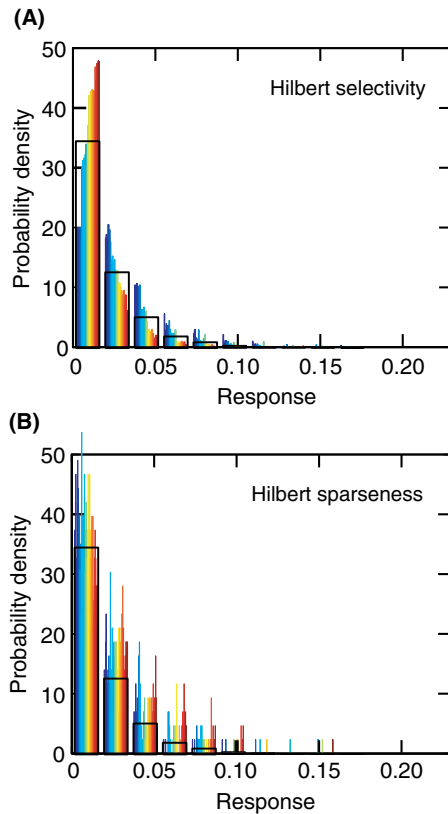


Fig. 15. Comparison of selectivity and sparseness in Hilbert model units. Histograms show responses for (A) individual units across time (selectivity) and (B) across the population simultaneously (sparseness). This is analogous to Fig. 14.

units, Hilbert units each have different response distributions as a function of spatial frequency, indicated by the fact that histogram bars in Fig. 15A are not flat for different colors. Therefore, Hilbert units are not strongly ergodic.

Actual complex cells are expected to have properties intermediate between the Gabor and Hilbert models (again referring to Table 1). Therefore, we predict they will be weakly ergodic but not strongly ergodic. This means that the average selectivity of complex cells will be equal to average sparseness across the population, but that the selectivity of each individual complex cell will not match average sparseness.

4. Discussion

Non-parametric selectivity measures were calculated from both striate complex cell data and model complex cells. These indicated average selectivities of individual units that were moderately greater than that of a Gaussian reference distribution (Table 1).

The numerical equivalence between selectivity (response distributions of individual units across time when presented with a sequence of stimulus images) and

sparseness (response distributions across all units in a population measured simultaneously), which we call ergodicity, was shown in simulations of model cells (Fig. 14 and 15). This model prediction of ergodicity remains to be confirmed by experimental studies.

If one accepts the existence of ergodic equivalence, then the selectivity measures of individual units in Table 1 can also be considered as average sparseness measures of the population as a whole. That is, populations of striate complex cells would be expected to show a moderate degree of elevated sparseness relative to a Gaussian reference distribution. However, we shall argue below that such sparseness would not necessarily indicate efficient coding.

4.1. Comparison with previous models

Our simulations indicating the equivalence of selectivity and sparseness are in complete disagreement with the simulations of Willmore and Tolhurst (2001), who found no relation between the two. It is not clear why this difference exists, particularly as there are conditions under which selectivity and sparseness are mathematically required to be identical, as Földiák (2002) has pointed out. One difference between our methods is that they measured selectivity/sparseness indices for each individual response distribution and then averaged the indices, whereas in Figs. 14 and 15 we are averaging the probability distributions on which sparseness measurements are based rather than averaging the sparseness measurements themselves.

Three factors were found to influence non-parametric selectivity in our models of complex cells. The first is whether the linear subunits integrate exactly to zero or not (or equivalently, whether they have a DC response or not). Complex cells with subunits that integrate to zero have higher selectivity. The second is location of the peaks of spatial frequency tuning curves of the complex cells. Units tuned to higher spatial frequencies are more selective, again provided receptive fields integrate to zero. The third is average response of the complex cells over the entire stimulus set. Units with low average responses are more skewed and produce higher selectivity measures.

Baddeley (1996a) has previously reported, from simulations, that linear, circularly-symmetric units with inhibitory surround, similar to retinal ganglion cells, produce higher selectivity if the receptive fields integrate to zero. We can confirm that this property still holds true for oriented, non-linear, phase-independent complex cells. However, we have offered a different explanation for the origin of this effect. Baddeley (1996a) ascribed it to non-stationarity in the image statistics, whereas our explanation focuses on the $1/f$ nature of the image spatial frequency spectrum (in conjunction with alignments in its phase spectrum). Baddeley (1996a)

also reported that for his circular linear units, higher spatial frequency tuning led to higher selectivity, which we also see in complex cells. Again, however, the explanations for these effects differ, with Baddeley ascribing it to image non-stationarity while we offer an explanation in terms of the $1/f$ image spectrum.

Although the entropy or kurtosis values we observed in model complex cells indicates a moderate degree of sparseness, the modeling indicates that sparseness measures are very sensitive to slight changes in subunit receptive field profiles, with higher sparseness for receptive fields that integrate to zero. We know from data in the literature (Spitzer & Hochstein, 1988) that receptive fields do not integrate to zero, based on observations of a residual ripple in complex cell responses to drifting gratings which indicate imperfect phase invariance. This suggests that evolutionary pressure to create very high sparseness in these units is weak, or constrained by other objectives.

Furthermore, the exponential tails seen in the response distributions of model complex cells (Fig. 13) do not lead to high sparseness values compared to other possible shapes, such as a power law tail. Exponential tails have also been noted in V1 simple cells and inferotemporal cells (Baddeley et al., 1997; Treves et al., 1999), indicating that this property is commonplace and perhaps ubiquitous in visual cells. On the other hand, power law tails, leading to much higher sparseness, have never been reported. Rather than being associated with high sparseness and information efficiency, exponential tails can be associated with energy (metabolic) efficiency (Baddeley, 1996b; Baddeley et al., 1997), as they maximize output entropy for a fixed firing rate. See Laughlin and Sejnowski (2003) for a more general discussion of metabolic efficiency as a constraint in brain organization.

The fact that minor changes in the receptive field organization underlying complex cells can lead to large changes in sparseness/selectivity raises the possibility that those parameters may be under dynamic control. For example, it is possible that attention can cause slight adjustments in the receptive field structure leading to a change in selectivity, although such an effect has not been reported.

4.2. Comparisons with previous experimental data

Vinje and Gallant (2000, 2002) have previously reported data on sparseness in macaque striate complex cells. On our terminology, they were actually measuring non-parametric selectivity rather than sparseness. However, by the ergodicity principle that our modeling indicates, selectivity is equivalent to sparseness. Our data leads to similar estimates of selectivity in complex cells to those of Vinje and Gallant.

Vinje and Gallant (2000, 2002) further reported that selectivity in macaque striate complex cells increases

when the stimulus diameter is expanded to include the non-classical receptive field surround. Based on the present study we can identify two possible mechanisms that may underlie this effect. The first relates to the negative correlation between selectivity and the average activity level (Fig. 9). The observed increase in selectivity for broad stimuli may simply be secondary to non-specific inhibition from the non-classical surround, which would reduce the average activity level of the unit. We know from Gallant, Connor, and Van Essen (1998) that the non-classical surround does have an inhibitory effect on activity in these units. The second possibility is that the surround affects the receptive field profiles of the complex cell subunits so that they more closely integrate to zero. As we have seen, changing the receptive fields in this manner greatly increases selectivity.

The first mechanism, lateral inhibition, is so non-specific and ubiquitous that increased sparseness resulting from it could easily be an epiphenomenal side effect rather than a deliberate means of increasing coding efficiency. On the other hand, if the second mechanism, fine tuning of receptive field profiles, were shown to occur, that could more convincingly be interpreted as a purpose-built mechanism for increasing efficiency.

A fundamental issue in which we would disagree with Vinje and Gallant is the idea that high sparseness measures calculated from data are necessarily an indicator of statistically efficient coding in the system. Baddeley (1996a) and Treves et al. (1999) have already established that high sparseness measures can arise for reasons unrelated to coding efficiency, and this is a point we would like to expand on.

On one explanatory level, excess sparseness in complex cells arises because they have response distributions that are skewed (but not heavy-tailed). The skewness in turn arises from the squaring non-linearity by which subunit responses of complex cells are combined (Eq. (12)). (To give a simple example, if a set of normally distributed random numbers are squared, the resulting probability density function will be skewed and exhibit large excess kurtosis.) Thus, we see high sparseness measures arising from non-linear distortions in the neural response distributions, rather than the linear transforms that underlie information theoretic explanations of the origins of sparseness (Field, 1994). It is not self-evident that sparseness measures arising from such non-linear probability distortions need to be interpreted in terms of efficient coding. Rather, the non-linearities seen in complex cells could be involved in the implementation of image processing algorithms, and unrelated to issues in statistical efficiency.

4.3. Implications for the efficient coding hypothesis

It may be objected that perhaps the non-linear distortions are a *means* of implementing the high sparseness

required for efficient coding. However, a feature of the squaring non-linearity underlying complex cells is that it is an information-losing transform. Information theory deals with the reliability and efficiency of information transmission in the presence of noise, and the significance of any deterministic information-losing transforms within the system is orthogonal to the concerns of the theory, other than as arbitrary external constraints that an information theoretic analysis must deal with. Information theory will not predict nor explain the nature of information-losing transforms within the visual system, for example why complex cells implement a particular non-linearity and not some other. Yet it seems likely that understanding the functional role of information-losing transforms will prove central to understanding vision, as the organism creates representations that highlight aspects of the environment that are of ecological significance.

Simoncelli and Olshausen (2001) appear to recognize this problem, presenting a weakened form of the efficient coding hypothesis in which information is not preserved. Efficient coding is in that case only relative to whatever information is spared at each stage. “The hypothesis states only that information must be represented efficiently; it does not say what information should be represented...”, as they say. Obviously, identifying what information is being represented must come prior to determining if that information is efficiently represented. This reinforces the point we are making, that it is premature to interpret high sparseness measures calculated from complex cell data as indicative of efficient coding, as Vinje and Gallant do, without understanding the requirements of the visual algorithms being implemented.

If one accepts that high sparseness measures can occur for reasons unrelated to information efficient coding, it still remains possible to retain some suggested benefits of sparse coding without invoking information-theoretic optimality arguments for its origin. For example, the ability of sparse codes to increase the storage capacity of associative memory under some models (Baum, Moody, & Wilzek, 1988; Palm, 1980; Treves & Rolls, 1991) remains whether the sparseness arises from information-efficient transforms or from information-losing system non-linearities. Sparseness also has value in reducing energy costs regardless of the other benefits.

In more general terms, the criticism here of visual models that center on efficient coding and redundancy reduction is that they attempt to explain properties of receptive fields purely in terms of the statistical properties of the input stimulus, without considering the goals of the organism for which the visual apparatus was constructed. To understand why V1 has particular receptive fields, it may be necessary to look not only at the structure of the stimulus, but also at the structure of the higher visual areas into which V1 feeds, analyze what these

higher areas are trying to accomplish, and determine what kind of inputs best serve those ends (Lehky & Sejnowski, 1999). Including considerations of these higher areas may produce visual representations that not only reflect image statistics but also incorporate the requirements for visuomotor coordination and other behaviors the organism needs in order to survive.

The idea that a theory of sensory processing can be developed purely by examining the internal structure of the stimuli without any reference to the organism as an integrated sensorimotor system has been criticized in particular by those espousing an “embodied” viewpoint (for example, see Churchland, Ramachandran, & Sejnowski, 1994; Clark, 1997; Lakoff, 1987; Merleau-Ponty, 1945; Varela, Thompson, & Rosch, 1991; Winograd & Flores, 1987). Perhaps a broader way of saying the same thing is that sparse coding is motivated by issues in information theory, and as with all information theoretic models it is fundamentally concerned with the syntax of the signal rather than semantics or “meaning”. Ultimately it may not be possible to have a satisfactory theory of the brain without confronting the constellation of issues relating to meaning (of which the problem of categorization is a part). The basic point here is that by focusing on issues arising from information theory (such as sparse coding), one is led to ask fundamentally the wrong kinds of questions concerning visual processing at higher levels.

It is important to emphasize that what is being challenged here is the use of information theoretic optimality principles as a core explanation of why sensory systems are structured the way they are, and not the application of information theory as a tool for data analysis per se. Information theoretic analyses of data (for example, Optican & Richmond, 1987; Rolls, 2003) can lead to interesting insights on sensory processing without making the claim that the system is optimized along information theoretic principles.

In view of the above criticisms, explanations of receptive field structure in terms of information theoretic measures of efficient coding are most convincing when confined to peripheral parts of sensory pathways, such as the retina and lateral geniculate nucleus, which involve linear signal transforms and reduced contamination by cognitive (non-stimulus) feedback. The presence of information-losing non-linearities at higher levels, such as in the complex cells studied here, indicate other factors in addition to information theory that should be taken into account in order to understand receptive fields.

Acknowledgments

The experimental phase of this project was supported by the McDonnell-Pew Foundation while SRL was at

NIMH. SRL thanks Dr. Keiji Tanaka for his support during the modeling phase of the project at RIKEN.

References

- Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America, A, 2*, 284–299.
- Adler, R. J., Feldman, R., & Taquq, M. S. (Eds.). (1998). *A practical guide to heavy tails: statistical techniques and applications*. Boston: Birkhäuser.
- Atick, J. J. (1992). Could information theory provide an ecological theory of sensory processing? *Nature, 3*, 213–251.
- Baddeley, R. (1996a). Searching for filters with “interesting” output distributions: an uninteresting direction to explore? *Network, 7*, 409–421.
- Baddeley, R. (1996b). An efficient code in V1? *Nature, 381*, 560–561.
- Baddeley, R., Abbot, L. F., Booth, M. C. A., Sengpiel, F., Freeman, T., Wakeman, E. A., & Rolls, E. T. (1997). Responses of neurons in primary and inferior cortices to natural scenes. *Proceedings of the Royal Society of London, B, 264*, 1775–1783.
- Barlow, H. B., Kaushal, T. P., & Mitchison, G. J. (1989). Finding minimum entropy codes. *Neural Computation, 1*, 412–423.
- Baum, E. B., Moody, J., & Wilzek, F. (1988). Internal representations for associative memory. *Biological Cybernetics, 59*, 217–228.
- Bell, A. J., & Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision Research, 37*, 3327–3338.
- Bryson, M. C. (1983). Heavy-tailed distributions. In S. Kotz, C. B. Read, & N. L. Johnson (Eds.), *Encyclopedia of statistical sciences* (vol. 3, pp. 598–601). New York: John Wiley.
- Churchland, P. S., Ramachandran, V. S., & Sejnowski, T. J. (1994). A critique of pure vision. In C. Koch & J. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 23–60). Cambridge, MA: MIT Press.
- Clark, A. (1997). *Being there: putting brain, body and world together again*. Cambridge, MA: MIT Press.
- De Valois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research, 22*, 545–559.
- Emerson, R. C., Korenberg, M. J., & Citron, M. C. (1992). Identification of complex-cell intensive nonlinearities in a cascade model of cat visual cortex. *Biological Cybernetics, 66*, 291–300.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America, A, 4*, 2379–2394.
- Field, D. J. (1994). What is the goal of sensory coding? *Neural Computation, 6*, 559–601.
- Field, D. J. (1999). Wavelets, vision, and the statistics of natural scenes. *Philosophical Transactions of the Royal Society of London A, 357*, 2527–2542.
- Földiak, P. (2002). Sparse coding in the primate cortex. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (2nd ed., pp. 895–898). Cambridge, MA: MIT Press.
- Gallant, J. L., Connor, C. E., & Van Essen, D. C. (1998). Neural activity in areas V1, V2, and V4 during free viewing of natural stimuli compared to controlled viewing. *NeuroReport, 9*, 2153–2158.
- Heeger, D. J. (1992). Half-squaring in responses of cat striate cells. *Visual Neuroscience, 9*, 427–443.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of the monkey striate cortex. *Journal of Physiology (London), 195*, 215–243.
- Jones, J. P., & Palmer, L. A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology, 58*, 1233–1258.
- Karklin, Y., & Lewicki, M. S. (2003). Learning higher order structure in natural images. *Network, 14*, 483–499.
- Kulikowski, J. J., & Vidyasagar, T. R. (1986). Space and spatial frequency: analysis and representation in the macaque striate cortex. *Experimental Brain Research, 64*, 5–18.
- Lakoff, G. (1987). *Women, fire, and dangerous things: what categories reveal about the mind*. Chicago: University of Chicago Press.
- Laughlin, S. B., & Sejnowski, T. J. (2003). Communication in neuronal networks. *Science, 301*, 1870–1874.
- Lehky, S. R., & Sejnowski, T. J. (1999). Seeing white: qualia in the context of decoding population codes. *Neural Computation, 11*, 1261–1280.
- Lehky, S. R., Sejnowski, T. J., & Desimone, R. (1992). Predicting responses of nonlinear neurons in monkey striate cortex to complex patterns. *Journal of Neuroscience, 12*, 3568–3581.
- Mandelbrot, B. B. (1963). The variation of certain speculative prices. *Journal of Business, 36*, 394–419.
- Merleau-Ponty, M. (1945). *Phenomenologie de la perception (translated 1962)*. London: Routledge & Kegan Paul.
- Morrone, M. C., & Burr, D. C. (1988). Feature energy in human vision: a phase dependent energy model. *Proceedings of the Royal Society of London, B, 235*, 221–245.
- Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1990). Stereoscopic depth discrimination in the visual cortex: neurons ideally suited as disparity detectors. *Science, 249*, 1037–1041.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature, 381*, 607–609.
- Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Research, 37*, 3311–3325.
- Optican, L. M., & Richmond, B. J. (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *Journal of Neurophysiology, 57*, 162–178.
- Palm, G. (1980). On associative memory. *Biological Cybernetics, 36*, 19–31.
- Pollen, D. A., & Ronner, S. (1983). Visual cortical neurons as localized spatial frequency filters. *IEEE Transactions on Systems, Man, and Cybernetics, 13*, 907–916.
- Rieke, F., Warland, D., van Steveninck, R. D. R., & Bialek, W. (1997). *Spikes: exploring the neural code*. Cambridge, MA: MIT Press.
- Rolls, E. T., & Tovee, M. J. (1995). Sparseness of the neuronal representation of stimuli in the primate temporal cortex. *Journal of Neurophysiology, 73*, 713–726.
- Rolls, E. T. (2003). An information theoretic approach to the contribution of the firing rates and the correlations between firing rates. *Journal of Neurophysiology, 89*, 2810–2822.
- Ruderman, D. L., & Bialek, W. (1994). Statistics of natural images: scaling in the woods. *Physical Review Letters, 73*, 814–817.
- Schiller, P. H., Finlay, B. L., & Volman, S. F. (1976). Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *Journal of Neurophysiology, 39*, 1288–1319.
- Silverman, B. W. (1986). *Density estimation for statistics and data analysis. Monographs on statistics and applied probability 26*. London: Chapman & Hall.
- Simoncelli, E. P. (2003). Vision and statistics of the visual environment. *Current Opinion in Neurobiology, 13*, 144–149.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience, 24*, 1193–1216.
- Spitzer, H., & Hochstein, S. (1988). Complex-cell receptive field models. *Progress in Neurobiology, 31*, 285–309.
- Szulborski, R. G., & Palmer, L. A. (1990). The two-dimensional spatial structure of nonlinear subunits in the receptive fields of complex cells. *Vision Research, 30*, 249–254.

- Treves, A., Panzeri, S., Rolls, E. T., Booth, M., & Waksman, E. A. (1999). Firing rate distributions and efficiency of information transmission of inferior temporal neurons to natural stimuli. *Neural Computation*, *11*, 601–632.
- Treves, A., & Rolls, E. T. (1991). What determines the capacity of autoassociative memories in the brain? *Network*, *2*, 371–397.
- van Hateren, J. H., & van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society of London, B*, *265*, 359–366.
- Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind*. Cambridge, MA: MIT Press.
- Vinje, W. E., & Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, *287*, 1273–1276.
- Vinje, W. E., & Gallant, J. L. (2002). Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. *Journal of Neuroscience*, *22*, 2904–2915.
- Willmore, B., & Tolhurst, D. J. (2001). Characterizing the sparseness of neural codes. *Network*, *12*, 255–270.
- Winograd, T., & Flores, F. (1987). *Understanding computers and cognition: a new foundation for design*. Reading, MA: Addison Wesley.