



Color opponency is an efficient representation of spectral properties in natural scenes

Te-Won Lee^{a,b,*}, Thomas Wachtler^c, Terrence J. Sejnowski^{a,b,d}

^a *Institute for Neural Computation, University of California, San Diego, 9500 Gilman, La Jolla, CA 92093-0523, USA*

^b *Computational Neurobiology Laboratory, The Salk Institute, La Jolla, CA 92037, USA*

^c *Universität Freiburg, Biologie III/Neurobiologie und Biophysik, 79104 Freiburg, Germany*

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

Received 23 October 2001; received in revised form 13 May 2002

Abstract

The human visual system encodes the chromatic signals conveyed by the three types of retinal cone photoreceptors in an opponent fashion. This opponency is thought to reduce redundant information by decorrelating the photoreceptor signals. Correlations in the receptor signals are caused by the substantial overlap of the spectral sensitivities of the receptors, but it is not clear to what extent the properties of natural spectra contribute to the correlations. To investigate the influences of natural spectra and photoreceptor spectral sensitivities, we attempted to find linear codes with minimal redundancy for trichromatic images assuming human cone spectral sensitivities, or hypothetical non-overlapping cone sensitivities, respectively. The resulting properties of basis functions are similar in both cases. They are non-orthogonal, show strong opponency along an achromatic direction (luminance edges) and along chromatic directions, and they achieve a highly efficient encoding of natural chromatic signals. Thus, color opponency arises for the encoding of human cone signals, i.e. with strongly overlapping spectral sensitivities, but also under the assumption of non-overlapping spectral sensitivities. Our results suggest that color opponency may in part be a result of the properties of natural spectra and not solely a consequence of the cone spectral sensitivities.

© 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Color opponency; Independent component analysis; Receptive fields; Efficient coding; Natural images; Human cone sensitivities

1. Introduction

The human visual system encodes the chromatic signals conveyed by the three types of retinal cone photoreceptors in an opponent fashion. This color opponency is often interpreted as an attempt to remove correlations in the signals of different cone types that are introduced by the strong overlap of the cone spectral sensitivities (Buchsbaum & Gottschalk, 1983). This explanation is in accordance with Barlow's hypothesis (Barlow, 1961) that the goal of sensory information processing is to transform the input signals such that the redundancy between the inputs is reduced. Buchsbaum and Gottschalk (1983) studied an orthogonal decorrelation

scheme to remove the correlations introduced by the strong spectral overlap of the human photoreceptor sensitivities. The result was a cone opponency that matched qualitatively the opponent characteristics of parvocellular neurons in the retina and lateral geniculate nucleus (LGN) of primates. However, as the authors noted, the overlaps of the cone spectral sensitivities are not the only source of correlations across wavelengths. Naturally occurring spectra are known to be fairly smooth (e.g. Maloney, 1986; Marimont & Wandell, 1992; Stiles, Wyszecki, & Ohta, 1977) and therefore may contribute substantially to redundancies in the cone signals (Brill & Benzsawel, 1985). Ruderman, Cronin, and Chiao (1998) applied principal component analysis (PCA) to analyze cone signals in response to natural spectra. They found small differences between results obtained with natural spectra and results obtained with random spectra that were uncorrelated across wavelengths. A shortcoming of orthogonal decorrelation as obtained by PCA in the analysis of chromatic signals is

* Corresponding author. Address: Institute for Neural Computation, University of California, San Diego, 9500 Gilman, La Jolla, CA 92093-0523, USA. Tel.: +1-858-534-9662; fax: +1-858-587-0417.

E-mail addresses: tewon@salk.edu (T.-W. Lee), thomas@salk.edu (T. Wachtler), terry@salk.edu (T.J. Sejnowski).

that it will lead inevitably to color opponency (Buchsbau & Gottschalk, 1983). There is, however, no reason why the visual system should be restricted to orthogonal decorrelation. Therefore, a more general analysis method should be used in order to investigate possible causes of color opponency. In a previous study (Wachtler, Lee, & Sejnowski, 2001), we found that non-orthogonal color-opponent codes can indeed be highly efficient. Here, we show that color opponency constitutes a sparse code for natural images, and that color opponency is efficient not only because of the correlations introduced by the overlaps of the cone spectral sensitivities, but also because there are substantial correlations across wavebands in natural spectra.

We investigated samples of spectral images (Párraga, Brelstaff, Troscianko, & Moorehead, 1998) of natural scenes. Our goal was to find efficient representations of the chromatic sensory information for overlapping as well as non-overlapping spectral sensitivities such that spatial and chromatic redundancy is reduced significantly. The method we used for finding statistically efficient representations is independent component analysis (ICA). ICA is a way of finding a linear non-orthogonal coordinate system in multivariate data that minimizes mutual information among the axial projections of the data. The directions of the axes of this coordinate system (basis functions) are determined by both second and higher-order statistics of the original data. This contrasts to PCA which uses solely second order statistics and has orthogonal basis functions. The ICA algorithm used here differs from results obtained in our earlier studies (Wachtler et al., 2001) since the nonlinear function is flexible and adapts to the density of the data, i.e. it does not enforce sparseness and focuses only on the independence assumption.

2. Methods

2.1. Data

We analyzed a set of 256×256 pixels hyperspectral images (Párraga et al., 1998). Each pixel is represented by radiance values for 31 wavebands of 10 nm width, sampled in 10 nm steps between 400 and 700 nm. The pixel size corresponds to $0.056^\circ \times 0.056^\circ$ of visual angle. The images were recorded around Bristol, UK, either outdoors, or inside the glass houses of Bristol Botanical Gardens. We chose eight of these images which had been obtained outdoors under apparently different illumination conditions. Each image contains a gray reflectance standard mounted on a tripod. Image samples taken for analysis were restricted to the image regions outside the areas occupied by these objects. The vector of 31 spectral radiance values of each pixel was converted to a vector of three cone excitation values whose components

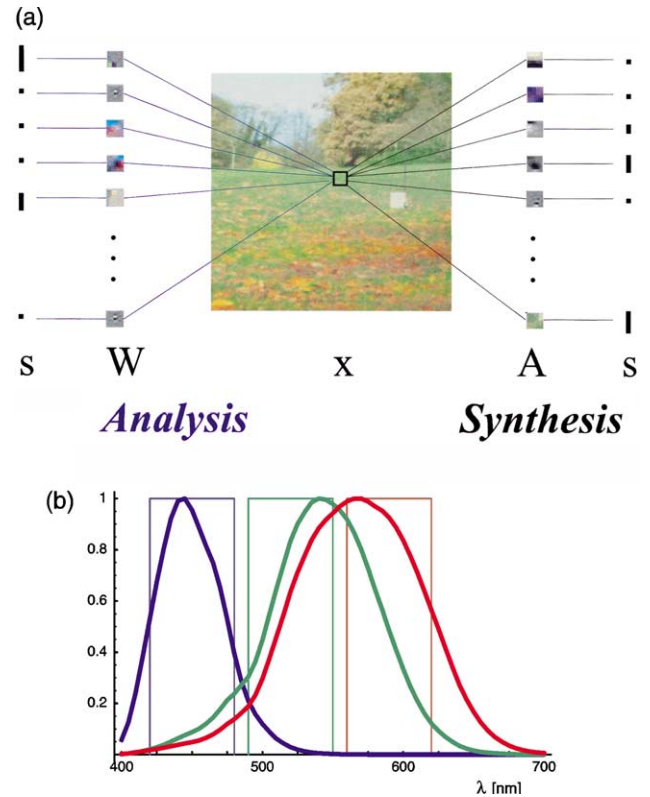


Fig. 1. Analysis of natural scenes. (a) Synthesis and analysis model. Synthesis: a small image patch x is generated linearly using a set of basis functions A and its corresponding source coefficients or activations s . Analysis: the image patch can be linearly decomposed into ICA filters W and their corresponding independent source activations s . The small bars for the activations indicate the strength of the activation. Sparse sources have mostly weak or zero activations and only a few strong activations. (b) For the main analysis, we used estimates of the human cone sensitivities (Stockman, MacLeod, & Johnson, 1993) (thick lines). In an additional analysis, we used hypothetical rectangular, non-overlapping sensitivities (thin lines).

were the inner products of the radiance vector with the cone sensitivity vectors (Fig. 1(b)). The logarithms of these values were used in the analysis (Ruderman et al., 1998). As cone spectral sensitivity functions, we used the estimates of Stockman et al. (1993), and, in addition, hypothetical non-overlapping sensitivities. The latter were rectangular functions with absorptions between 420 and 480 nm (S'), 490 and 550 nm (M'), and 560 and 620 nm (L'), respectively. For additional analyses, narrower non-overlapping functions were used (see Section 3). From the image data, 7×7 pixel image patches were chosen randomly, yielding $7 \times 7 \times 3 = 147$ dimensional vectors.

2.2. Analysis

The goal of ICA is to perform a linear transform which makes the resulting source outputs as statistically independent from each other as possible (Bell & Sejnowski, 1995; Comon, 1994; Jutten & Herault, 1991;

Lee, 1998). ICA assumes an unknown source vector \mathbf{s} with mutually independent components s_i . The observed data \mathbf{x} can be represented as a linear combination of source components s_i such that

$$\mathbf{x} = \mathbf{A}\mathbf{s} = \sum_{i=1}^N \mathbf{a}_i s_i, \quad (1)$$

where \mathbf{A} is a scalar square matrix and the columns of \mathbf{A} (\mathbf{a}_i) are the basis functions. In a generative image model (Bell & Sejnowski, 1997; Olshausen & Field, 1996) \mathbf{A} represents the basis functions generating the observed image pattern in the real world whereas $\mathbf{W} = \mathbf{A}^{-1}$ refers to the ICA filters that transform the image pattern into activations or source coefficients $\mathbf{s} = \mathbf{W}\mathbf{x}$. Since \mathbf{A} and \mathbf{s} are unknown, the goal of ICA is to adapt the basis functions by estimating \mathbf{s} so that the individual components s_i are statistically independent and this adaptation process minimizes the mutual information between the components s_i . A learning algorithm can be derived using the information maximization principle (Bell & Sejnowski, 1995) or the maximum likelihood estimation method (Pearlmutter & Parra, 1996) which can be shown to be equivalent in this case (Cardoso, 1997). In our experiments, we used the infomax learning rule with natural gradient extension and the learning algorithm for updating the basis functions is

$$\Delta \mathbf{A} \propto \mathbf{A}[\mathbf{I} - \varphi(\mathbf{s})\mathbf{s}^T], \quad (2)$$

where \mathbf{I} is the identity matrix, $\varphi(\mathbf{s}) = -(\partial p(\mathbf{s})/\partial \mathbf{s})/p(\mathbf{s})$ and \mathbf{s}^T denotes the matrix transpose of \mathbf{s} . $\Delta \mathbf{A}$ is the change of the basis functions that is added to \mathbf{A} . The change in $\Delta \mathbf{A}$ will converge to zero once the adaptation process is complete. Note that $\varphi(\mathbf{s})$ requires a density model for $p(s_i)$. We assumed a parametric exponential power density $p(s_i) \propto \exp(-|s_i|^{q_i})$ and updated the parameters q_i during the adaptation process to match the distribution of the estimated sources (Lee & Lewicki, 2000). This was accomplished by finding the maximum posteriori value of q_i given the observed data. The ICA algorithm can thus characterize a wide class of statistical distributions including uniform, Gaussian, Laplacian, and other so-called sub- and super-Gaussian densities. In other words, our experiments do not constrain the coefficients to have a sparse distribution, unlike some previous methods (Bell & Sejnowski, 1997; Olshausen & Field, 1996).

For adapting the basis functions we used Eq. (2). Training was done in 500 training steps, each using a set of spectra of 40,000 image patches, 5000 chosen randomly from each of the eight images, excluding the image regions corresponding to the reflectance standard. The parameters q_i of the exponential power density were initialized as Gaussian densities and updated after each training step.

2.2.1. Displaying the results

To visualize the learned basis functions (Fig. 2), we used the method by Ruderman et al. (1998) and plotted for each basis function a 7×7 pixel matrix, with the color of each pixel indicating the combination of L , M , and S cone responses as follows. The values for each patch were normalized to values between 0 and 255, with 0 cone excitation corresponding to a value of 128. Thus, the R, G, and B components of each pixel represent the relative excitations of L , M , and S cones, respectively. This yields a pseudo-color representation which, however, is qualitatively similar to a true color rendering of the respective cone values. To further illustrate the chromatic properties of the basis functions, we convert the L , M , S vector of each pixel to its coordinates in a cone-opponent color space (Derrington, Krauskopf, & Lennie, 1984; MacLeod & Boynton, 1979). The x -axis corresponds to the difference between L and M cone stimulation, the y -axis corresponds to the difference between S -cone stimulation and the sum of L and M cone stimulation, and the z -axis corresponds to the sum of L and M cone stimulation. Here, we plot x versus y coordinates, i.e. the projection of the color coordinates onto the isoluminant plane. For each pixel of the basis functions, a point is plotted at its corresponding location in the x - y plane of that color space (Fig. 2(b) and (d)). The color of the points are the same as used for the pixels in the left part of the figure. Thus, although only the projection onto the isoluminant plane is shown, the third dimension (i.e. luminance) can be inferred by the brightness of the points.

3. Results

3.1. Color opponency

Fig. 2(a) shows the learned ICA basis functions in pseudo-color representation (see Section 2) for overlapping cone spectra. In Fig. 2(b), the chromaticities of the pixels in each basis function are shown in a cone-opponent color space. The basis functions are shown in order of decreasing L_2 -norm (definition of L_2 -norm: square root of the sum of the vector elements squared). There are three main types of basis functions, namely homogeneous chromatic, oriented achromatic (luminance edges), and color-opponent (color edges) basis functions. These results are qualitatively similar to those obtained in our earlier analysis (Wachtler et al., 2001). Luminance edges tended to have higher L_2 -norm than color edge basis functions. This reflects the fact that in the natural environment, luminance variations are generally larger than chromatic variations (Párraga et al., 1998). The achromatic basis functions were localized and oriented, similar to those found in the analysis of grayscale natural images (Bell & Sejnowski, 1997;

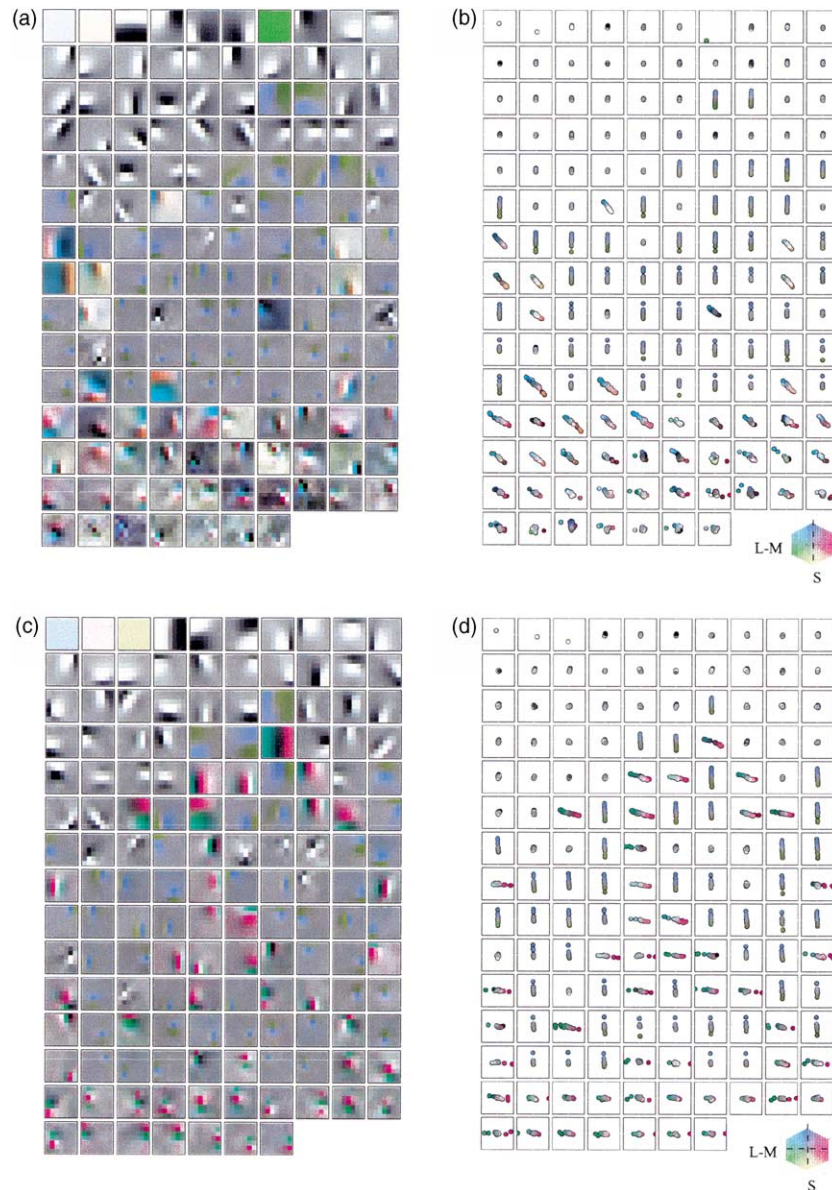


Fig. 2. Basis functions derived from natural scenes. 147 ICA basis functions of trichromatic image patches for overlapping cone spectra (a,b), and non-overlapping spectra (c,d). Subplots (a) and (c) show the spatiochromatic structure of the basis functions. Each patch represents a 7×7 pixel basis function. Red, green, and blue pixel values correspond to L , M , and S cone stimulation, respectively. Subplots (b) and (d) show the chromaticities of basis function pixels projected onto the isoluminant plane of cone-opponent color space. Horizontal axis: $L-M$; vertical axis: $S-(L+M)$. See text for further details.

Olshausen & Field, 1996). Most of the chromatic basis functions, particularly those with strong contributions, were color opponent, i.e. the chromaticities of their pixels lay roughly along a line through the origin of our color space. Most chromatic basis functions with relatively high contributions were modulated between light blue and dark yellow, in the plane defined by luminance and S -cone modulation. Those with lower L_2 -norm were highly localized, but still were mostly oriented. Other chromatic basis functions showed opponency axes corresponding to blue versus orange colors. The chromaticities of these basis functions occupied mainly the

second and fourth quadrant. The basis functions with lowest contributions were less strictly aligned in color space, but still tended to be color opponent, mostly along a bluish-green/orange direction. The opponency axes of both blue–yellow and red–green basis functions were tilted with respect to the color space axes. This most likely reflects the distribution of the chromaticities in our images.

In natural images, $L-M$ and S coordinates in our color space are negatively correlated (Webster & Mollon, 1997). ICA finds the directions that correspond to maximally independent signals, i.e. extracts statistical

structure of the inputs. PCA did not yield basis functions in these directions (Ruderman et al., 1998; Wachtler et al., 2001), probably because it is limited by the orthogonality constraint.

3.2. Non-overlapping spectral sensitivities

In order to determine the contribution of natural spectra to the statistical dependencies in the receptor signals (Buchsbaum & Gottschalk, 1983), we repeated the analysis, but with hypothetical receptor sensitivities (L' , M' , S') that have no overlap, while sampling in roughly the same spectral regions as the L , M , and S cones (see Section 2 and Fig. 2(c)). The resulting basis functions were strongly color opponent, as for overlapping cone sensitivities. The chromatic axis of opponency of the red–green basis functions was slightly closer to the L – M color space axis than in the case of overlapping sensitivities. This indicates that the axis of opponency reflects the differences in correlations of L and M cone, and of M - and S -cone signals. Our finding of strong opponency with non-overlapping sensitivities suggests that in natural spectra the correlations of radiance values between different wavelengths are sufficiently high to require a color-opponent code in order to represent the chromatic structure efficiently. To further investigate this point, we calculated the correlations between L and M , and between L' and M' signals, respectively, for different input ensembles. For spectra drawn from the hyperspectral image dataset, the correlation between L and M signals was 0.99, the correlation between L' and M' was 0.94. For random spectra that were uncorrelated across wavelengths, the correlation between L and M was 0.91, the correlation between L' and M' was, by definition, zero. Thus, the correlations in cones with non-overlapping sensitivities due to natural spectra were even slightly larger than correlations due to the overlaps of human cone sensitivities alone. Although ICA takes also higher-order statistics into account, these results support the conclusion that color opponency is in part a result of the properties of natural spectra.

To study in more detail the effect of the properties of the non-overlapping cone sensitivities itself we have carried out experiments for other cone sensitivity characteristics. Sampling in regions that are narrower and further apart than the human cone sensitivities might not lead to cone opponency. However, when we used sensitivities of 30 nm width, either centered on the peaks of the human cone sensitivities, or further apart in the visible spectrum (at 420, 530, and 640 nm), all cases led to opponent basis functions. Opponency emerged even with narrower sensitivities, although for 20 and 10 nm, measurement noise in the data led to very slow convergence of the algorithm and less pronounced opponency. Our experiments suggest that color opponency

can be achieved with many other non-overlapping cone sensitivity functions.

3.3. Coding efficiency and sparseness

The densities for the coefficients s_i in Eq. (1) are highly sparse (Fig. 3) with an average kurtosis value of 18.6. This shows that the independence criterion alone is sufficient to learn sparse image codes. Although no sparseness constraint was imposed, most of the obtained ICA coefficients were extremely sparse, i.e. the data \mathbf{x} were encoded in the sources \mathbf{s} in such a way that the coefficients of \mathbf{s} were mostly around zero; there is only a small percentage of informative values (non-zero coefficients). From an information coding perspective this implies that we can encode and decode each chromatic image patch with only a small percentage of the basis functions. In contrast, Gaussian densities as assumed in PCA were not sparsely distributed and a large portion of the basis functions was required to represent the chromatic images. In this sense, the basis functions that produce sparse distributions are statistically efficient codes.

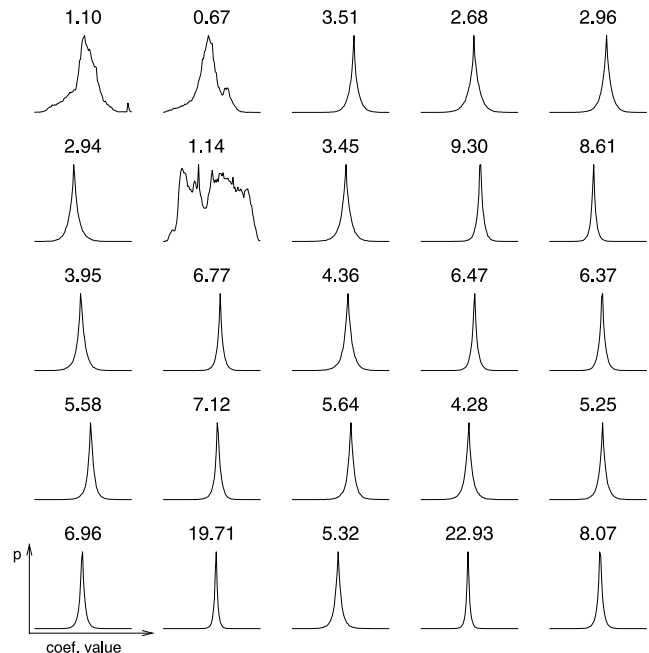


Fig. 3. ICA coefficients are sparsely distributed. The distributions of the coefficients for ICA in Fig. 2(a) for the first 25 basis functions. The coefficients were computed from 40,000 randomly selected image patches using the ICA basis functions. The normalized kurtosis value is shown on top of each distribution. This value indicates how sparse the distribution is compared to a Gaussian distribution. The normalized kurtosis for the Gaussian density is zero and the sparser the density the greater its value. Negative values indicate sub-Gaussian densities such as the uniform density. Averaged over all distributions for the ICA coefficients the normalized kurtosis was 18.6, suggesting highly sparse source density functions producing efficient codes.

To quantitatively measure the encoding difference we compared the coding efficiency between ICA and PCA using Shannon's theorem to obtain a lower bound on the number of bits required to encode a spatiochromatic pattern (Lewicki & Olshausen, 1999; Lewicki & Sejnowski, 2000). The average number of bits required to encode 40,000 patches randomly selected from the eight images in Fig. 1 with a fixed noise coding precision of $\sigma_x = 0.059$ was 1.83 bits for ICA and 4.46 bits for PCA. Note that the encoding difference for achromatic image patches using ICA and PCA was about 20% in favor of ICA (Lewicki & Olshausen, 1999). The encoding difference in the chromatic case was significantly higher (>100%) and suggests that there is a large amount of chromatic redundancy in the natural scenes. We believe that this redundancy is mostly due to non-orthogonal projections in the data that is captured by ICA using the higher-order statistical information. To verify our findings, we computed the average pairwise mutual information I in the original data ($I_x = 0.1522$), the PCA representation ($I_{\text{PCA}} = 0.0123$) and the ICA representation ($I_{\text{ICA}} = 0.0096$). ICA was able to further reduce the redundancy between its components, and its basis functions therefore represent more efficient codes. Although those codes were allowed to produce source coefficients regardless of their density, sparse coefficients with average normalized kurtosis values of 18.6 for ICA and 6.6 for PCA. Interestingly, the basis functions in Fig. 2(a) produced only sparse coefficients except for basis function 7 (green basis function), which has a nearly uniform distribution (Fig. 3), suggesting that this basis function is active almost all the time. The reason is that the green color component is present in a large number of image patches of the natural scenes. For the non-overlapping sensitivities, the coding efficiency was 2.23 bits per pixel, the average mutual information was 0.0102 and the average kurtosis was 17.8. In comparison to the ICA results obtained by overlapping spectral sensitivities the small differences are not significant (Table 1) and this control experiment again supports the argument that color opponency may be a consequence of efficient encoding of spectral information.

We repeated the experiments and used other nonlinearities $\varphi(\mathbf{s})$ in the learning rule in Eq. (2) and obtained similar results. Basis functions obtained with the exponential power distributions or the simple Laplacian

prior (Wachtler et al., 2001) were statistically most efficient. The results using the JADE algorithm (Cardoso, 1999) and the fixed-point ICA algorithm (Hyvaerinen & Oja, 1997) resulted in qualitatively similar basis functions with slightly higher entropy than the exponential power function ICA algorithm, which is due to a more accurate density estimation using a more flexible ICA algorithm.

In comparison of ICA to other decorrelation algorithms, there are many ways to decorrelate the cone signals. However, many decorrelation algorithms have constraints in the basis functions such as orthogonality in PCA and zero-phase in symmetric whitening filters. ICA has no constraint in the structure of the basis functions and it presents the most efficient coding scheme compared to all other decorrelation algorithms because it not only removes second order redundancies but other higher-order redundant information. There is no analytical way to achieve the same. Certain decorrelation algorithms such as PCA may give a very limited number of color-opponent basis functions but this could have been a result of the constraints (orthogonality in PCA or symmetry in ZCA) and not a result of finding an efficient coding scheme for the input signals.

4. Discussion

In studies where ICA had been applied to achromatic images of natural scenes (Bell & Sejnowski, 1997; Lewicki & Olshausen, 1999; Olshausen & Field, 1996; van Hateren & van der Schaaf, 1998), the resulting basis functions were localized and orientation sensitive. In other words, the basis functions represented luminance edges useful for representing image features. Those basis functions are non-orthogonal and adapted to the structure inherent in the data. We analyzed chromatic images and obtained, in addition to achromatic edges, basis functions that correspond to color edges. These basis functions were likewise non-orthogonal and their coefficients had sparse distributions, indicating a highly efficient encoding of the data. Recently, Caywood, Willmore, and Tolhurst (2001) analyzed the same hyperspectral data set that we used and found similar results. Hoyer and Hyvaerinen (2000) had used uncompressed RGB images of natural scenes to adapt basis functions. Our results using L , M , S cone inputs derived from spectra of natural scenes yielded qualitatively similar results, thus supporting the hypothesis of these authors that nonlinearities in the generation of RGB images do not strongly influence the qualitative structure of the results. A quantitative comparison, in particular with respect to the opponency axes, is not possible, since calibration data for the acquisition of the RGB images are not available. Results obtained from images encoded with the lossy JPEG compression

Table 1
Coding efficiencies for raw data, PCA and ICA

	Kurtosis	Mutual information	Entropy (bits/pixel)
Raw	0.4	0.1522	7.50
PCA	6.6	0.0123	4.46
ICA	18.6	0.0096	1.83
ICA (no-overlap)	17.8	0.0102	2.23

(Tailor, Finkel, & Buchsbaum, 2000) would be even less comparable. In an attempt to accurately model the retinal properties Doi and Inui (2000) applied natural scenes (RGB images) to simulated trichromatic cone mosaic outputs and obtained similar results.

The results presented here extend earlier findings (Wachtler et al., 2001) by showing that the sparseness arises due to the statistical structure of the data. Furthermore, we show that color opponency may serve not only to remove spectral correlations introduced by overlapping photoreceptor sensitivities, but also correlations inherent in the natural spectra. Indeed, our results indicate that those two influences are about equally strong. Thus, encoding signals of cones with non-overlapping spectral sensitivities leads to equally strong opponency and similar coding efficiency as with overlapping sensitivities. In other words, given the properties of natural spectra, the differences of having cone sensitivities with less overlap than the human cones would be minor in terms of coding efficiency.

Regarding interpretations of the results obtained by ICA, one issue to consider is whether one should focus on the basis functions \mathbf{A} , or rather on the corresponding ICA filters \mathbf{W} . For achromatic images of natural scenes, there are small differences (Bell & Sejnowski, 1997; Olshausen & Field, 1996) between the characteristics of the basis functions and ICA filters. For the chromatic case, the differences are not negligible. In our analysis, we adapted the basis functions, corresponding to the image features that are represented. The filters \mathbf{W} were derived by inverting the basis function matrix \mathbf{A} and therefore are constrained by the inversion process. In particular, a filter will give no response to any but its corresponding basis function. Furthermore, the basis function corresponding to a filter is not necessarily the stimulus that will yield the largest response. The filter for which a particular basis function would give the largest response would essentially look like this basis function, thus achieving the largest dot product of basis function and filter. The basis functions correspond to what the algorithm estimates as independent causes of the images, as well as to filters for which these causes would be ‘best’ stimuli. Therefore, we focus on the basis functions rather than the filters.

It has been pointed out that the achromatic ICA basis functions resemble receptive-field characteristics of neurons in the visual cortex (Bell & Sejnowski, 1997; Lewicki & Olshausen, 1999; Olshausen & Field, 1996; van Hateren & van der Schaaf, 1998), leading to the conclusion that the receptive fields of visual neurons reflect an efficient code for the visual signals. If the visual system encodes chromatic signals efficiently, one could expect to find similar correspondences in the chromatic domain. The visual system of trichromatic primates encodes retinal color signals in a color-opponent fashion. In the LGN of monkeys, color selective cells re-

spond best for chromatic stimuli along the coordinate axes of a cone-opponent color space (Derrington et al., 1984), indicating a coding scheme that decorrelates the signals from different cone classes (Buchsbaum & Gottschalk, 1983). In the visual cortex, the chromatic tuning of color selective cells is not restricted to the color space axes (De Valois, Cottaris, Elfar, Mahon, & Wilson, 2000; Hanazawa, Komatsu, & Murakami, 2000; Lennie, Krauskopf, & Sclar, 1990; Wachtler, Sejnowski, & Albright, 1999). Owing to our method of deriving cone responses from natural spectra, we can make some quantitative statements about efficient encoding of retinal signals under natural conditions. Our results indicate that the chromatic axes for efficient opponent encoding do not necessarily correspond to the ‘cardinal’ directions of color opponency as found in LGN cells. While the results of PCA show chromatic organization along these axes, ICA basis functions show opponency along intermediate directions. Comparing the properties of these coding schemes with the properties of visual neurons, one could speculate that, in retina and LGN, the visual system may take advantage of PCA-like encoding to encode as much signal as possible with the limited number of fibers in the optic nerve, while recoding the chromatic signals in the cortex to achieve a sparse code that better reflects the statistical structure of the environment (Field, 1994).

Often, double-opponent cells (Daw, 1967) have been associated with the cortical coding of color (Hubel & Wiesel, 1968; Livingstone & Hubel, 1984). Our chromatic ICA basis functions do not show the typical double-opponent organization. Usually, the contributions of different cone types in our basis functions have both positive and negative regions, but these regions are out of phase, such that maximal stimulation (positive or negative) of one cone type coincides with close-to-zero stimulation of another cone type. This is more similar to single opponency with anisotropic center-surround organization than to classical double opponency.

Two further properties of our ICA basis functions do not match the classical double-opponent receptive field. First, the opponency axis of red–green basis functions does not correspond to pure L versus M signals, but indicates contributions from S cones. Interestingly, a recent study (Conway, 2001) investigating cone inputs to V1 neurons reported double-opponent receptive fields where L input had opposite sign to both M and S input. Thus, the opponency axis did not coincide with the L – M axis, but rather was tilted in the same way as the red–green basis functions of our results. Second, the spatial organization of the chromatic basis functions does not show circular symmetry like the classical double-opponent cell’s receptive field. The spatial structure of the chromatic basis functions is similar to the achromatic ones, which have been interpreted as luminance edges. Likewise, the chromatic basis functions may represent

chromatic edges. Commonly, color coding neurons are assumed to be non-oriented, despite the undisputed opinion that chromatic contrast is an important feature in color vision. However, oriented chromatically selective cells in primary visual cortex have been described (Johnson, Hawken, & Shapley, 2001; Michael, 1978; Ts'o & Gilbert, 1988). In view of our results, this may suggest that indeed efficient encoding may be a factor influencing the properties of cortical receptive fields.

Although the analysis of natural signals can lead to important insights in the encoding strategies of the visual system, one has to be aware of the limits of such comparisons. The assumption of a linear encoding of the visual signals is an idealization. Furthermore, the relevance of the visual signals, e.g. in signalling food, should be taken into account (Mollon, 1989; Osorio & Vorobyev, 1996). However, even without these additional criteria, our results show that non-orthogonal basis functions provide an efficient encoding of chromatic natural images, and that color opponency is not a mere consequence of overlapping cone spectral sensitivities. This supports the hypothesis that cortical neurons represent the intrinsic spatiochromatic structure of the natural environment in a statistically efficient manner.

Acknowledgements

We are grateful to C. Parraga, G. Brelstaff, T. Troscianko, and I. Moorehead for making the hyperspectral image data set available. We would like to thank Michael Lewicki for many discussions and for providing us with the exponential power density function algorithm. We thank Bruno Olshausen, David Field and Eizaburo Doi for fruitful discussions. We are grateful to the reviewers for critical comments that improved the presentation of the results. T.-W. Lee was supported by the Swartz Foundation. T.J. Sejnowski was supported by the Howard Hughes Medical Institute.

References

- Barlow, H. (1961). Possible principles underlying the transformation of sensory messages. In *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- Bell, A. J., & Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7, 1129–1159.
- Bell, A. J., & Sejnowski, T. J. (1997). The 'independent components' of natural scenes are edge filters. *Vision Research*, 37(23), 3327–3338.
- Brill, M. H., & Benzsawel, T. (1985). Remarks on signal-processing explanations of the trichromacy of vision. *Journal of the Optical Society of America A*, 2, 1794–1796.
- Buchsbaum, G., & Gottschalk, A. (1983). Trichromacy, opponent colours coding and optimum colour information transmission in the retina. *Proceedings of the Royal Society London B*, 220, 89–113.
- Cardoso, J. (1999). High-order contrasts for independent component analysis. *Neural Computation*, 11, 157–192.
- Cardoso, J.-F. (1997). Infomax and maximum likelihood for blind source separation. *IEEE Signal Processing Letters*, 4(4), 112–114.
- Caywood, M., Willmore, B., & Tolhurst, D. J. (2001). The color tuning of independent components of natural scenes matches V1 simple cells. Vision Sciences Society First Annual Meeting.
- Comon, P. (1994). Independent component analysis—a new concept? *Signal Processing*, 36(3), 287–314.
- Conway, B. R. (2001). Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1). *Journal of Neuroscience*, 21, 2768–2783.
- Daw, N. W. (1967). Goldfish retina: Organization for simultaneous color contrast. *Science*, 158, 942–944.
- Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology*, 357, 241–265.
- De Valois, R. L., Cottaris, N. P., Elfar, S. D., Mahon, L. E., & Wilson, J. A. (2000). Some transformations of color information from lateral geniculate nucleus to striate cortex. *Proceedings of the National Academy of Science, USA*, 97, 4997–5002.
- Doi, E., & Inui, T. (2000). Whitening and independent component filters of simulated cone mosaic output have spatio-chromatic properties of pLGN and V1. In *7th international conference on neural information processing* (pp. 500–505).
- Field, D. J. (1994). What is the goal of sensory coding? *Neural Computation*, 6, 559–601.
- Hanazawa, A., Komatsu, H., & Murakami, I. (2000). Neural selectivity for hue and saturation of colour in the primary visual cortex of the monkey. *European Journal of Neuroscience*, 12, 1753–1763.
- Hoyer, P., & Hyvaerinen, A. (2000). Independent component analysis applied to feature extraction from colour and stereo images. *Network: Computation in Neural Systems*, 11(3), 191–210.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195, 215–244.
- Hyvaerinen, A., & Oja, E. (1997). A fast fixed-point algorithm for independent component analysis. *Neural Computation*, 9, 1483–1492.
- Johnson, E. N., Hawken, M. J., & Shapley, R. (2001). The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience*, 4, 409–416.
- Jutten, C., & Herault, J. (1991). Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture. *Signal Processing*, 24, 1–10.
- Lee, T.-W. (1998). *Independent component analysis: Theory and applications*. Boston, MA: Kluwer Academic Publishers, ISBN: 0 7923 8261 7.
- Lee, T.-W., & Lewicki, M. S. (2000). The generalized Gaussian mixture model using ICA. *International workshop on ICA* (Vols. 239–244).
- Lennie, P., Krauskopf, J., & Sclar, G. (1990). Chromatic mechanisms in striate cortex of macaque. *Journal of Neuroscience*, 10, 649–669.
- Lewicki, M. S., & Olshausen, B. A. (1999). A probabilistic framework for the adaptation and comparison of image codes. *Journal of the Optical Society of America A: Optics, Image Science, and Vision*, 16(7), 1587–1601.
- Lewicki, M. S., & Sejnowski, T. J. (2000). Learning overcomplete representations. *Neural Computation*, 12(2), 337–365.
- Livingstone, M. S., & Hubel, D. H. (1984). Anatomy and physiology of a color system in the primate visual cortex. *Journal of Neuroscience*, 4, 309–356.
- MacLeod, D. I. A., & Boynton, R. M. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America*, 69, 1183–1186.
- Maloney, L. T. (1986). Evaluation of linear models of surface spectral reflectance with small number of parameters. *Journal of the Optical Society of America A*, 3, 1673–1683.

- Marimont, D. H., & Wandell, B. A. (1992). Linear models of surface and illuminant spectra. *Journal of the Optical Society of America A*, 9, 1905–1913.
- Michael, C. R. (1978). Color vision mechanisms in monkey striate cortex: Simple cells with dual opponent-color receptive fields. *Journal of Neurophysiology*, 41, 1233–1249.
- Mollon, J. D. (1989). “Tho’ she kneel’d in that place where they grew...” the uses and origins of primate colour vision. *Journal of Experimental Biology*, 146, 21–38.
- Olshausen, B., & Field, D. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607–609.
- Osorio, D., & Vorobyev, M. (1996). Colour vision as an adaptation to frugivory in primates. *Proceedings of the Royal Society London B*, 263, 593–599.
- Párraga, C. A., Brelstaff, G., Troscianko, T., & Moorehead, I. (1998). Color and luminance information in natural scenes. *Journal of the Optical Society of America A*, 15, 563–569. Available: <http://www.crs4.it/~gjb/ftpJOSA.html>.
- Pearlmutter, B., & Parra, L. (1996). A context-sensitive generalization of ICA. In *International conference on neural information processing* (pp. 151–157).
- Ruderman, D. L., Cronin, T. W., & Chiao, C.-C. (1998). Statistics of cone responses to natural images: Implications for visual coding. *Journal of the Optical Society of America A*, 15, 2036–2045.
- Stiles, W. S., Wyszecki, G., & Ohta, N. (1977). Counting metameric object-color stimuli using frequency-limited spectral reflectance functions. *JOSA*, 67, 779–784.
- Stockman, A., MacLeod, D. I. A., & Johnson, N. E. (1993). Spectral sensitivities of the human cones. *Journal of the Optical Society of America A*, 10, 2491–2521. Available: <http://www-cvrl.org>.
- Taylor, D. R., Finkel, L. H., & Buchsbaum, G. (2000). Color-opponent receptive fields derived from independent component analysis of natural images. *Vision Research*, 40, 2671–2676.
- Ts’o, D. Y., & Gilbert, C. D. (1988). The organization of chromatic and spatial interactions in the primate striate cortex. *Journal of Neuroscience*, 8, 1712–1727.
- 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 London B*, 265, 359–366.
- Wachtler, T., Lee, T.-W., & Sejnowski, T. (2001). Chromatic structure of natural scenes. *Journal of the Optical Society of America A*, 18, 65–77.
- Wachtler, T., Sejnowski, T. J., & Albright, T. D. (1999). Interactions between stimulus and background chromaticities in macaque primary visual cortex. *Investigative Ophthalmology & Visual Science*, 40, S641.
- Webster, M. A., & Mollon, J. D. (1997). Adaptation and the color statistics of natural images. *Vision Research*, 37, 3283–3298.