

Chromatic structure of natural scenes

Thomas Wachtler

Howard Hughes Medical Institute, Computational Neurobiology Laboratory, The Salk Institute, La Jolla, California 92037, and Sloan Center for Theoretical Neurobiology, The Salk Institute, La Jolla, California 92037

Te-Won Lee

Howard Hughes Medical Institute, Computational Neurobiology Laboratory, The Salk Institute, La Jolla, California 92037, and Institute for Neural Computation, University of California, San Diego, La Jolla, California 92093

Terrence J. Sejnowski

Howard Hughes Medical Institute, Computational Neurobiology Laboratory, The Salk Institute, La Jolla, California 92037, Institute for Neural Computation, University of California, San Diego, La Jolla, California 92093, and Department of Biology, University of California, San Diego, La Jolla, California 92093

Received December 1, 1999; revised manuscript received July 27, 2000; accepted August 2, 2000

We applied independent component analysis (ICA) to hyperspectral images in order to learn an efficient representation of color in natural scenes. In the spectra of single pixels, the algorithm found basis functions that had broadband spectra and basis functions that were similar to natural reflectance spectra. When applied to small image patches, the algorithm found some basis functions that were achromatic and others with overall chromatic variation along lines in color space, indicating color opponency. The directions of opponency were not strictly orthogonal. Comparison with principal-component analysis on the basis of statistical measures such as average mutual information, kurtosis, and entropy, shows that the ICA transformation results in much sparser coefficients and gives higher coding efficiency. Our findings suggest that nonorthogonal opponent encoding of photoreceptor signals leads to higher coding efficiency and that ICA may be used to reveal the underlying statistical properties of color information in natural scenes. © 2001 Optical Society of America
OCIS codes: 330.1690, 330.1720, 100.2960, 330.7310.

1. INTRODUCTION

The efficient encoding of visual sensory information is an important task for image processing systems, and its study may provide insights into coding principles of biological visual systems. Barlow¹ proposed that the goal of sensory information processing is to transform the input signals such that the redundancy between the inputs is reduced. Recently several methods have been proposed for finding linear basis functions for achromatic images of natural scenes that reduce the redundancy in the images. Olshausen and Field² used a sparseness criterion and found basis functions that were localized and oriented like receptive fields in the mammalian visual cortex. Similar results were obtained with the infomax independent component analysis (ICA) algorithm³ and a Bayesian approach.⁴

Although luminance dominates the visual structure of the natural environment, many species have developed color vision to make use of the information in the variations in spectral composition of the light that reaches the eye. It would therefore be of interest to find efficient codes for the chromatic structure of natural images.

Analyses of color coding have focused mostly on coding efficiency with respect to the postreceptoral signals.^{5,6} Buchsbaum and Gottschalk⁵ found opponent coding to be the most efficient way to encode human photoreceptor sig-

nals. In an analysis of cone responses to natural scenes using principal-component analysis (PCA), Ruderman *et al.*⁶ found principal components close to the opponent functions of Buchsbaum and Gottschalk. As the study by Ruderman *et al.* showed, the decorrelating functions depend mainly on the spectral properties of the photoreceptors.⁶ These properties, however, may not be determined solely by the spectral statistics in the environment. Physical, biological, and ecological factors may impose constraints on the number and spectral position of photoreceptor sensitivities.^{7,8} Therefore the PCA-like opponent coding of photoreceptor signals, while optimal for transmitting chromatic information through the optic nerve, may not necessarily reflect the chromatic structure of natural scenes and may be not the ultimate goal of the visual system.⁹ It is known that, while neurons in the lateral geniculate nucleus of trichromatic primates show responses along orthogonal axes (cardinal directions) of cone-opponent color space,¹⁰ cortical cells are not restricted to responses along these directions.^{11,12} This suggests that a nonorthogonal coding scheme may be more appropriate for encoding the chromatic structure of natural images.

Our goal in this paper is to find efficient representations of chromatic sensory information such that its redundancy is reduced significantly. To this end, we use

ICA to adapt or *learn* the spectral and spatial representations of natural chromatic images. In Section 2 we describe the motivation for using ICA for analyzing spectral images. In Section 3 we illustrate the spectral decomposition using ICA to find basis functions that efficiently encode natural spectra. In Section 4 we describe the spatial properties of natural color images, which are found by decomposing small image patches into spatial basis functions that encode the image in an efficient manner. In Section 5 we discuss our results and compare them with other methods.

2. LEARNING CODES FOR COLOR IMAGES WITH INDEPENDENT COMPONENT ANALYSIS

A. Independent Component Analysis

The ICA algorithm was initially proposed to solve the blind source separation problem; i.e., given only mixtures of a set of underlying sources, the task is to separate the mixed signals and retrieve their original sources^{13–16} without knowing how the signals were mixed nor the distribution of the sources. In contrast to correlation-based transformations, such as PCA, ICA not only decorrelates the signals (using second-order statistics) but also reduces higher-order statistical dependencies, thus reducing the mutual information between the output signals, making them as statistically independent as possible. Furthermore, the ICA bases are not restricted to being orthogonal.

ICA assumes a source vector \mathbf{s} with components s_i that are mutually independent. The sources are not observed directly, but linear combinations of the sources \mathbf{x} are given such that

$$\mathbf{x} = \mathbf{A}\mathbf{s}, \quad (1)$$

where \mathbf{A} is a $N \times M$ scalar matrix. The columns of \mathbf{A} are called the basis functions. We assume that there is no additive sensor noise. The goal of ICA is to find the basis functions by adaptation or learning given only the observed data \mathbf{x} . During the adaptation process a cost function such as the mutual information function is minimized by using an adaptation rule. Once the minimum is achieved, the sources s_i will be as linearly independent as possible.

The many variants of the ICA adaptation algorithm can be loosely categorized into parametric^{3,4,17} and nonparametric^{18,19} algorithms. The main difference between these two classes of algorithms is that the parametric approach assumes an adaptive or nonadaptive prior distribution on the source densities $p(\mathbf{s})$, whereas the nonparametric approach tries to approximate the underlying statistics using cumulants up to 4th order. Our primary interest is to learn efficient codes; i.e., the distribution of $p(\mathbf{s})$ should be as sparse as possible. The sparseness constraint has been used to learn efficient codes for natural images.^{2,4} Sparseness in this case assumes that the data \mathbf{x} are encoded in the sources \mathbf{s} in such a way that the coefficients of \mathbf{s} are mostly close to zero.

Figure 1 shows an example of a Laplacian density function [$p(\mathbf{s}) \propto \exp(-|\mathbf{s}|)$] in comparison with a Gaussian density function. Both distributions are zero mean and

have unit variance. The Laplacian prior is sharply peaked around zero and has heavy tails; there is only a small percentage of informative values (nonzero coefficients) in the tails of the Laplacian distribution, and most of the data values are close to zero; i.e., the data is sparsely distributed. From a coding perspective this assumes that we can encode and decode the data with only a small percentage of the coefficients \mathbf{s} . For Gaussian densities however, the data is not sparsely distributed, and a larger portion of the coefficients \mathbf{s} is required to encode the data. In this sense, the basis functions (codes) that produce sparse distributions are statistically efficient codes.

In our experiments, we use the infomax learning rule with natural gradient extension as derived in Appendix A. Furthermore, we impose the Laplacian prior on $p(\mathbf{s})$. The adaptation of the basis functions is then

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

where \mathbf{I} is the identity matrix, sign is the sign function, \mathbf{s}^T denotes the matrix transpose of \mathbf{s} , and $\Delta \mathbf{A}$ is the change of the basis functions that is added to \mathbf{A} . We also assume that \mathbf{A} is a square matrix; i.e., the number of sources is equal to the number of sensors. Once the adaptation process is complete, $\Delta \mathbf{A}$ will be a zero matrix.

The goal in our analysis is to find the basis functions \mathbf{A} such that the data can be linearly decomposed into extremely sparse distributions. The adapted basis functions are efficient codes for representing spectral data.

B. Hyperspectral Images of Natural Scenes

We analyzed a set of eight hyperspectral images of natural scenes. A detailed description of the images is given by Párraga *et al.*²⁰ Briefly, the full data set consists of 29 images with a size of 256×256 pixels. Each pixel is represented by radiance values for 31 wave bands of 10-nm width, sampled in 10-nm steps between 400 and 700 nm. The pixel size corresponds to 0.056×0.056 deg of visual angle. The images were recorded around Bristol, either outdoors or inside the glass houses of Bristol Botanical Gardens. We chose eight of these images that had been obtained outdoors under apparently different illumination conditions (Fig. 2). Note that in each image a gray plate mounted on a tripod can be seen. This plate, a Kodak GrayCard, was intended as a reference reflectance for calibration. For each image, we determined the area in the image occupied by the reflectance standard and tripod. Image samples taken for analysis were restricted to the image regions outside these areas.




Fig. 1. Schematic comparison between a Laplacian density (left) and a Gaussian density (right).

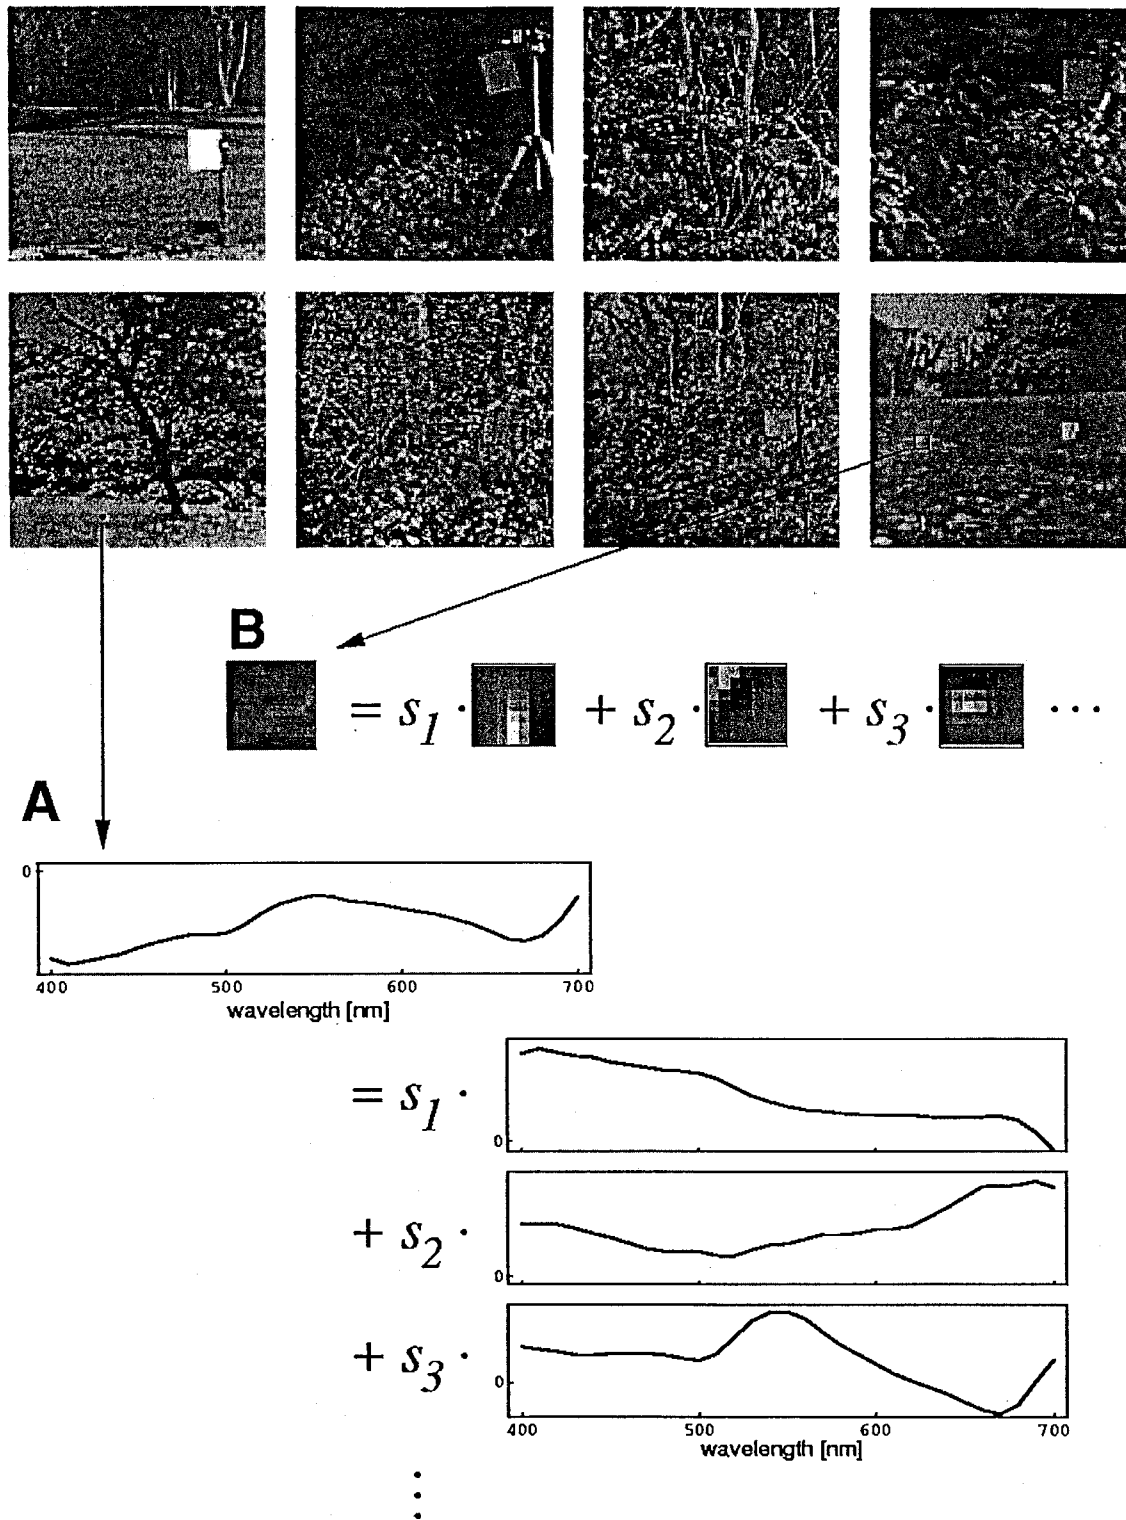


Fig. 2. Eight hyperspectral color images of natural scenes. A, Schematic of decomposition of the observed spectrum into spectral basis functions. B, Schematic of decomposition of the observed image patch into spatiochromatic basis functions.

Figure 2 illustrates two experiments using ICA. In Fig. 2A the spectral data of image pixels are used to learn basis functions so that the observed spectra can be represented as a linear superposition of scaled basis functions. Figure 2B illustrates the linear decomposition of the observed image patch into spatial basis functions. The Sections 3 and 4 describe the learned basis functions for the two experiments.

3. INDEPENDENT COMPONENTS OF NATURAL SPECTRA

In this experiment, we randomly sampled spectra of image pixels from the eight images in Fig. 2. As in the studies by Field⁹ and Ruderman *et al.*,⁶ the logarithms of the radiance values were used. Besides the issues concerning logarithms discussed by Ruderman *et al.*, we see

another possible benefit of using logarithmic values when investigating color signals. Important processes that affect spectral distributions, such as reflection and transmission, are of a multiplicative nature. When converted to logarithms, these processes become additive operations and thus may be more suitable for ICA, which assumes linear combinations of sources. Unlike in the study by Ruderman *et al.*, no attempt was made to correct for differences in illumination in the different images. Thus our data contained variations caused by illuminants and by object reflectances within images as well as between images. In preliminary analyses²¹ we had found that speed of convergence of the algorithm was slow and that small peaks of one sample (10-nm) width occurred in otherwise fairly smooth basis functions. These peaks changed when other combinations of images were used, while the basis functions themselves were stable. We attributed these small peaks to artifacts introduced by changes in the scenes during image acquisition (e.g., by movements due to wind or by changes in the illumination). Since natural photoreceptors have much wider spectral sensitivities, we consider variations on a spectral range of 10 nm not as relevant for color vision. Therefore we smoothed the spectra by convolution with a triangular kernel of 10-nm half-width at half-height. This increased the speed of convergence and reduced the small peaks in the basis functions. Training was done in 1000 training or adaptation steps. For each training step,

5000 spectra of each of the eight images were chosen randomly from the regions outside the images of the reflectance standard. For each such set of 40,000 spectra, the mean spectrum vector of the ensemble was subtracted, and the data were scaled to unit variance for each ensemble. Learning of the ICA basis functions was done on sequential blocks of 100 spectra from the data, with Eq. (2). The learned basis functions for the pixel spectra are shown in Fig. 3, plotted in order of decreasing L_2 norm. The lengths of the ICA basis functions (L_2 norm) are proportional to the power along the components and therefore measure the contributions of the basis functions to the data, as do the eigenvalues for PCA. The inset of Fig. 3 shows the distribution of L_2 -norm values, i.e., the relative contributions of the basis functions to the pixel spectra in the eight images.

The first four basis functions are broadband spectra. Basis function 1 is relatively uniform, with a slight bias toward long wavelengths. Basis functions 2 and 3 contribute mainly in the short- and long-wavelength part of the visible spectrum, respectively. Basis function 4 is slightly biased toward short wavelengths. Basis function 5 has a single peak near 540–550 nm and a dip at higher wavelengths. Basis function 6 has a peak that is broader and is shifted to longer wavelengths around 570 nm. Basis function 7 has a peak at the location of the peak of basis function 5 but lacks the long-wavelength dip. Note that we show the basis functions obtained from the loga-




Fig. 3. Analysis of pixel spectra using ICA: the learned basis functions of pixel spectra, ordered by decreasing L_2 norm. The number in the upper left corner of each plot denotes the basis function number as referred to in the text. The inset shows the histogram of the L_2 norms of the basis functions.




Fig. 4. Analysis of pixel spectra using PCA: PCA basis functions of pixel spectra, ordered by decreasing eigenvalues. The inset shows the histogram of the eigenvalues.

rhythmic data; but since the spectral contrast is relatively low within each basis function, the corresponding functions in linear values have similar shapes.

A. Comparison with Principal Component Analysis

For comparison, we calculated the PCA basis functions for the same image spectra. Figure 4 shows the resulting basis functions, ordered by decreasing eigenvalue (inset). In contrast to the ICA results, many basis functions show multiple zero crossings, and there are many basis functions of high spectral frequency. The third basis function has a peak in the region of ICA basis function 5, which confirms that there is variation with such spectral characteristics in natural spectra. The eigenvalue spectrum decreases strongly, with the first three components explaining more than 90% of the variance in the images.

B. Comparing Coding Efficiency

We have 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 pixel spectrum^{4,22}.

$$\# \text{bits} \geq -\log_2 P(\mathbf{x}_t | \mathbf{A}) - N \log_2(\sigma_x), \quad (3)$$

where N is the dimensionality of the input spectrum \mathbf{x}_t , σ_x is the coding precision (standard deviation of the noise introduced by errors in encoding), and $P(\mathbf{x}_t | \mathbf{A})$ is the likelihood of the data given the bases.

The average number of bits required to encode 40,000 pixel spectra randomly selected from the eight images in

Table 1. Coding Efficiencies of ICA and PCA for Pixel Spectra and Image Patches

Data	Method	Entropy (bits/pixel)	Mutual Information	Average Kurtosis
Pixel spectra	ICA	2.03	0.0194	11.8
	PCA	4.79	0.0237	6.7
Image patches	ICA	1.73	0.0093	19.7
	PCA	4.46	0.0123	6.6

Fig. 2 with a fixed-noise coding precision of $\sigma_x = 0.059$ was 2.03 bits for ICA and 4.79 bits for PCA, thus indicating that ICA basis functions encode natural spectra more efficiently than PCA components.

To compare the form of the resulting coefficient densities for PCA and ICA, we computed the normalized kurtosis of the source coefficients for both PCA- and ICA-transformed data. The average normalized kurtosis was 11.8 for the resulting ICA source coefficients and 6.7 for those of the PCA. Note that the normalized kurtosis for a Gaussian density is zero. This suggests that the ICA basis functions produce a very sparse representation for the spectra.

To illustrate how ICA representation minimizes the mutual information between its components, we computed the average pairwise mutual information in the spectra data (see Table 1). The average mutual informa-

tion was 0.2932 in the original data, 0.0237 in PCA-transformed data, and 0.0194 in the ICA-transformed data.

4. INDEPENDENT COMPONENTS OF COLOR IMAGES

The goal of the experiment illustrated in Fig. 2B is to analyze the spatial properties of chromatic information of the data set. To reduce the computational effort and the number of resulting basis functions, we converted the spectra of each pixel to a vector of three cone excitation values. Thus we attempted to find an efficient postreceptoral code for natural images. Conversion to human cone-excitation values was done by multiplying the radiance value for each wavelength by the corresponding values for the human cone sensitivities²³ and summing over the resulting values. Thus, the dimensionality of each image was reduced from $256 \times 256 \times 31$ dimensions to $256 \times 256 \times 3$ dimensions. From these data, 7×7 pixel image patches were chosen randomly, yielding $7 \times 7 \times 3 = 147$ dimensional vectors. As in the case of the spectra, the data were transformed to logarithmic values. To adapt 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, again excluding the image regions corresponding to the reflectance standard.

To visualize the resulting components, we used the method by Ruderman *et al.*⁶ and plotted for each basis function a 7×7 pixel matrix, with the color of each pixel indicating the combination of long-, middle-, and short-wavelength sensitive (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 red, green, and blue (R, G, and B) components of each pixel represent the relative excitations of L, M, and S cones, respectively. In Fig. 5, the 147 components are ordered by decreasing L_2 norm. A histogram of the L_2 norms is shown in Fig. 6. To further illustrate the chromatic properties of the basis functions, we convert the L, M, S vector of each pixel to its projection onto the isoluminant plane of a cone-opponent color space. We used a space similar to the color spaces used by MacLeod and Boynton²⁴ and Derrington *et al.*¹⁰ with coordinate axes defined by L-M, S, and L+M+S modulation, respectively. In our plots we show the projections on the isoluminant plane, with L-M on the horizontal axis and S on the vertical axis. For each pixel of the basis functions, a point is plotted at its corresponding location in that color space (Fig. 5, bottom). The colors of the points are the same as those used for the pixels in the top 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.

Figure 5 shows homogeneous and oriented achromatic and chromatic basis functions. When ordered by decreasing L_2 norm, achromatic basis functions tend to appear before chromatic basis functions. This reflects the fact that in the natural environment, luminance variations are generally larger than chromatic variations.²⁰

The achromatic basis functions are localized and oriented, similar to those found in the analysis of gray-scale natural images.^{2,3} Most of the chromatic basis functions, particularly those with strong contributions, are color opponent; i.e., the chromaticities of their pixels lie roughly along a line through the origin of our color space. Basis functions 4 and 5 code green and purple colors, respectively (see Fig. 5, bottom). These are the only basis functions with strong contributions in the first and third quadrants. Most chromatic basis functions with relatively high contributions are modulated between light blue and dark yellow in the plane defined by luminance and S-cone modulation. Those with lower L_2 norm are highly localized but still are mostly oriented. There are other chromatic basis functions with tilted orientations, corresponding to blue versus orange colors. The chromaticities of these basis functions occupy mainly the second and fourth quadrants. The basis functions with the lowest contributions are less strictly aligned in color space but still tend to be color opponent, mostly along a bluish-green-orange direction. There are no basis functions with chromaticities along the horizontal axis, corresponding to pure L versus M cone opponency, as with PCA basis functions.⁶ In Fig. 5 we show some of the filters (column marked "W") corresponding to the basis functions. The filters likewise are localized and oriented, and their directions in color space are nonorthogonal. Note that the chromatic directions in the filters are enhanced in comparison with the achromatic direction, since in the basis functions they have smaller contributions. This leads to chromatic variations in the achromatic filters. For better visualization, we chose to scale the data along the chromatic directions to approximately equal variance by scaling with factors of 3 (S) and 8 (L-M), respectively, before computing the inverse.

Figure 6A shows a histogram of the relative contributions of the first 50 basis functions. In Fig. 6B we plot the distributions of the coefficients of the first 49 basis functions across the images. Most distributions have very high kurtosis. The mirror-image pair of basis functions 4 and 5 have relatively low kurtosis, indicating contribution to many pixels. The contribution of basis function 5 to one of the images is shown in Fig. 7. The 256×256 pixel image was first divided into 7×7 -pixel image patches, and the contribution was calculated by computing the source coefficients s_5 . The s_5 coefficients show the low contribution in the sky area and its high contribution in the tree scene, indicating that this basis function encodes green and dark areas. The gray plate in the image contains very low green contribution as does the sky region.

A. Comparison with Principal Component Analysis

We compared the ICA results with principal components of the data. For single pixels, we found the same principal components as Ruderman *et al.*⁶ The three principal axes were $(1/\sqrt{3})(1.00l + 1.00m + 0.99s)$, $(1/\sqrt{6})(1.02l + 0.96m - 2.01s)$, and $(1/\sqrt{2})(0.99l - 1.01m + 0.02s)$, with l , m , and s the unit direction vectors in logarithmic L-, M-, and S-cone space. The values for the ICA basis functions for single pixels were $(0.32l + 0.29m + 0.66s)$, $(0.38l + 0.53m + 0.33s)$, and $(0.57l + 0.46m$

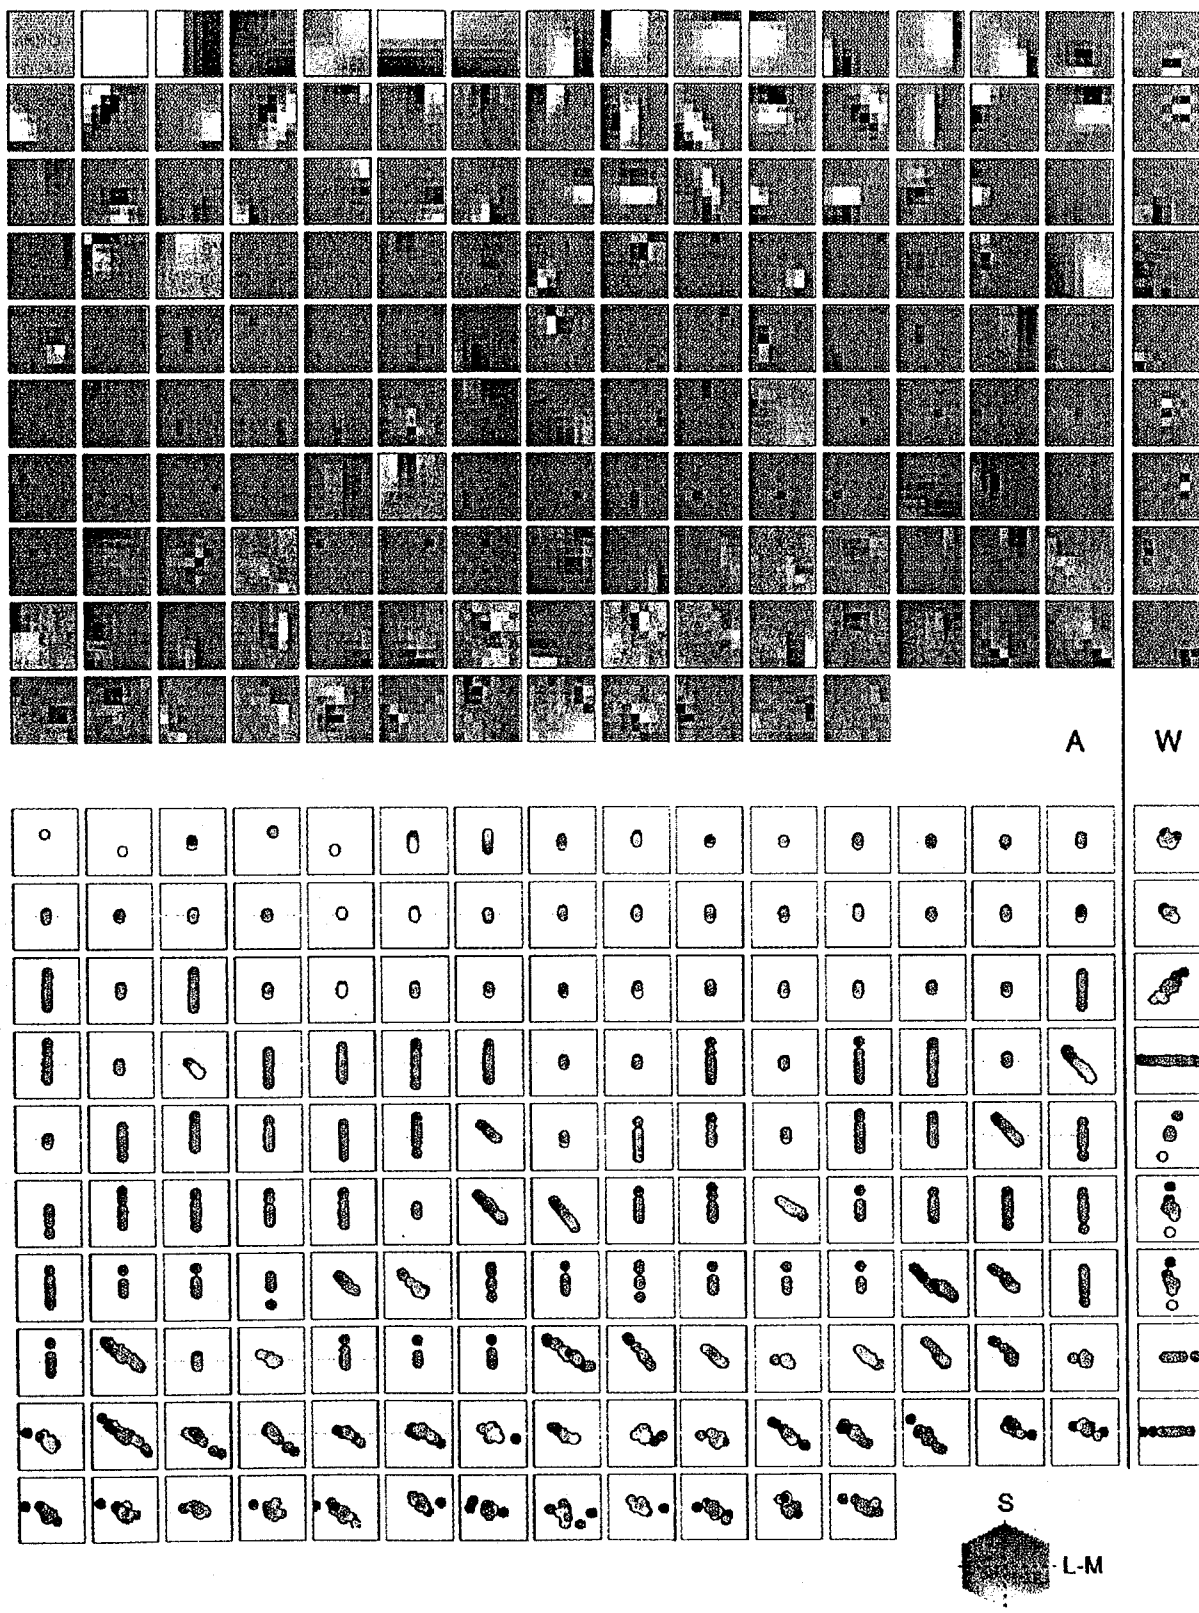


Fig. 5. Analysis of image patches using ICA. Top, spatiochromatic structure of the 147 learned ICA basis functions of image patches (7×7 pixels, three chromatic dimensions). The R, G, and B values of the color of each pixel correspond to the relative excitation of L, M, and S cones, respectively. The functions are in order of decreasing L_2 norm from left to right and top to bottom. The additional column marked W on the right shows the filters for the rightmost column of basis functions (marked A). Bottom, chromaticities of the basis functions, plotted in cone-opponent color-space coordinates. Horizontal axes, L- versus M-cone variation. Vertical axes, S-cone variation. Each dot represents the coordinate of a pixel of the respective basis function, projected onto the isoluminant plane. Luminance can be inferred from the brightness of the dot. Note that for basis functions that vary mainly in luminance, the dots tend to lie on top of one another.

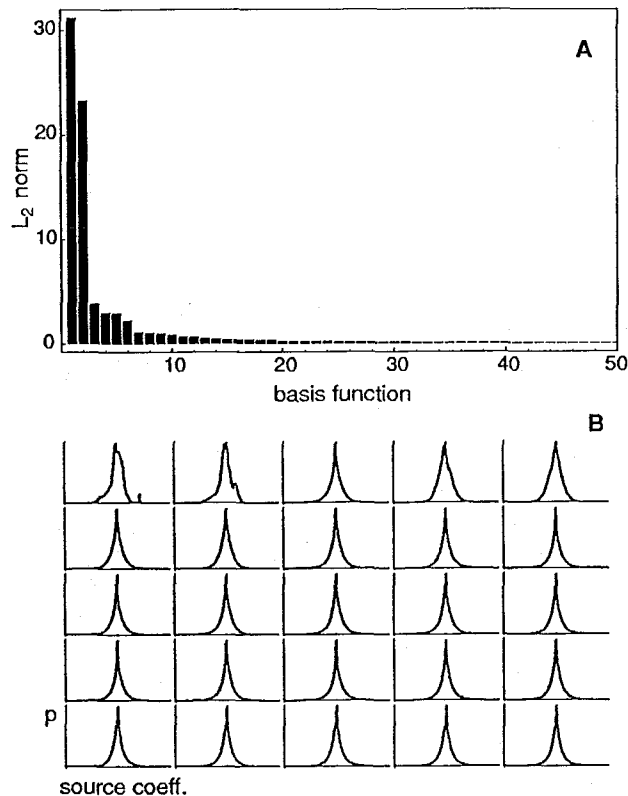


Fig. 6. Relative contributions of ICA basis functions. A, Histogram of the L_2 norms of the first 50 ICA basis functions in Fig. 5. B, Histograms of coefficient values for the first 25 ICA basis functions in Fig. 5. Most distributions have high kurtosis.




Fig. 7. Example of contribution levels for ICA-patch basis function 5 (green image patch) in one of the images. Left, gray-scale rendering of original image. Right, contribution level. Dark regions correspond to low contributions of basis function 5.

+ 0.26s). The ICA result reflects the fact that most natural colors are not highly saturated; i.e., usually all cone types respond to a given color.

Analyzing 3×3 image patches, Ruderman *et al.*⁶ found a separation of chromatic and spatial dimensions such that in each patch, color varied along one chromatic axis, and the same spatial patterns were found for each of the three chromatic axes. Our analysis of 7×7 image patches (Figs. 8 and 9) shows a similar trend, but there are clear deviations from this rule. A number of patches show pixels with colors other than the principal colors. As a consequence, there are no pure L-M components with low spatial frequency. Thus the separation of chromatic and spatial dimensions is not perfect, which indicates that S and L-M directions are correlated to some degree in the data.²⁵

B. Comparing Coding Efficiency

We compared the coding efficiency of ICA and PCA for the image patches by using the same method as described for the pixel spectra in Section 3. The average number of bits required to encode 40,000 image patches randomly selected from the eight images in Fig. 2 with a fixed noise coding precision of $\sigma_x = 0.059$ was 1.73 bits for ICA and 4.46 bits for PCA. The difference between ICA and PCA was significant and supports the argument that basis functions for efficient coding are not orthogonal.

To compare the forms of the resulting coefficient densities for PCA and ICA, we again computed the normalized kurtosis of the source coefficients for both PCA- and ICA-transformed data. The average normalized kurtosis for the resulting source coefficients was 19.7 for ICA and 6.6 for PCA. This again suggests that the ICA basis functions produce a very sparse representation for the image patches.

In case of the chromatic image patches, the average pairwise mutual information was 0.1522 in the original data, 0.0123 in PCA-transformed data, and 0.0093 in the ICA-transformed data. Thus the patches were not independent in the basis functions found by ICA, but they were more independent than in the images or in the PCA basis.

5. DISCUSSION

We applied ICA to a set of hyperspectral images of natural scenes. Our goal was to find efficient codes for the chromatic structure in the natural environment, assuming sparse probability density functions for the source coefficients. Our data set consisted of images of terrestrial natural scenes, containing mainly plants (grass, trees, etc.) and rocks. No strong shadows, no specular highlights, no fruit, and only few flowers were in the images.

A. Basis Functions of Spectra

For single-pixel spectra, ICA found basis functions that were in most cases relatively broadband and smooth, in contrast to PCA basis functions obtained from the same data. Spectra of naturally occurring illumination are typically broadband and can be described with a small number of parameters.²⁶⁻²⁸ We can ask whether there is a relation between the ICA basis functions and natural spectra. In a recent study, Chiao *et al.*²⁹ investigated illuminant and reflectance spectra in natural environments. Our first four ICA basis functions are fairly broadband with slight biases in the short- or the long-wavelength region of the visible spectrum. Their shapes are qualitatively similar to some of the illumination spectra of forest scenes [Fig. 2(a) of Chiao *et al.*,²⁹ illuminants 2, 4, and 5]. ICA basis function 7 with its peak near 550 nm has a shape similar to light from forest shade (illuminants 1 and 3 of Chiao *et al.*²⁹). Chiao *et al.* determined the first three principal components of reflectance spectra of forest scenes [Fig. 2(a) of Chiao *et al.*²⁹]. Our basis function 5, with its peak at medium wavelengths and its dip at short wavelengths, matches qualitatively the varia-

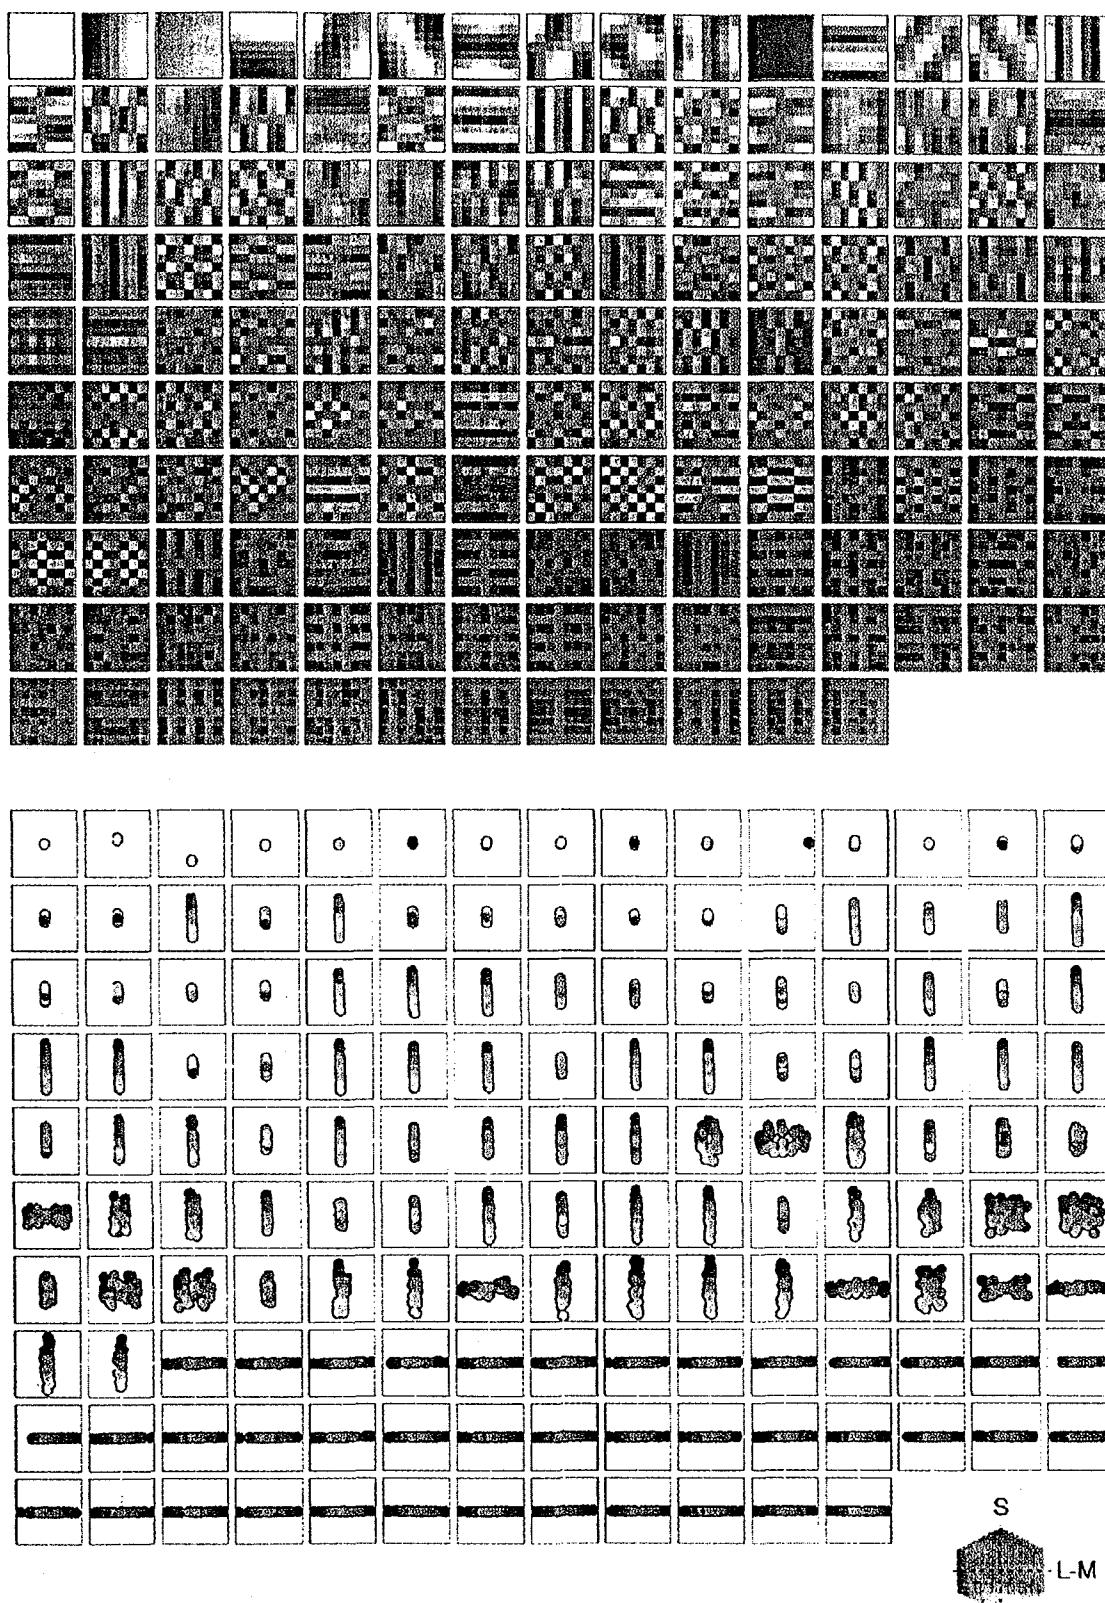


Fig. 8. Analysis of image patches using PCA. Top, spatiochromatic structure of the 147 PCA basis functions of image patches. The components are ordered by decreasing eigenvalues. Bottom, chromaticities of the basis functions in cone-opponent color-space coordinates. The data are plotted in the same way as the ICA basis functions in Fig. 5.

tions of reflectance spectra as represented by the second and third principal components shown by Chiao *et al.* For comparison, in Fig. 10 we show basis function 5 together with reflectance spectra of leaves as determined by Vrhel

*et al.*³⁰ The peak of basis function 5 matches closely the peaks in the reflectance spectra. The similarity of some basis functions with natural spectra raises the question whether ICA can be used to separate the chromatic con-

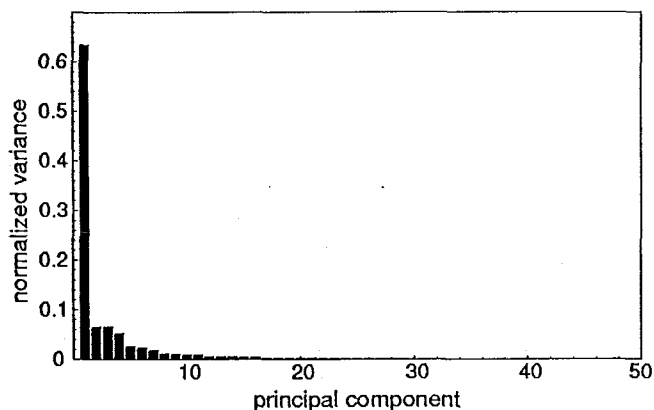


Fig. 9. Eigenspectrum of the first 50 principal components in Fig. 8.




Fig. 10. Comparison of ICA basis function 5 (dashed curve) with reflectance spectra of leaves after Vrhel *et al.*³⁰ (solid curves). All spectra have been scaled such that their peaks near 550 nm have a value of 1.

tributions of illuminants and reflectances, which vary independently in natural scenes. This will be investigated in the future.

Interestingly, owing to the spectral position of its peak near 550 nm, basis function 5 leads to stronger stimulation of the M cones than the other basis functions, when normalized to equal total radiance (data not shown). Likewise, basis function 6, with its peak near 570 nm, is more effective in stimulating the L cones than the other basis functions. The question of the relation between primate cone spectra and natural spectra has been discussed in several studies. It has been proposed that L- and M-cone sensitivities evolved in frugivorous monkeys for the purpose of finding food.^{31,32} The images of our data set did not contain fruit, but the search for an efficient code nevertheless yielded functions that strongly contribute in the spectral region of the L- and M-cone sensitivities.

B. Basis Functions of Images

When applied to small image patches, ICA found homogeneous basis functions and edges as well as achromatic and chromatic basis functions. Most of the chromatic basis functions showed pronounced opponency; i.e., the chromaticities of their components fell along lines through the origin of color space. However, the directions of these

lines did not always coincide with the color-space axes. Whereas achromatic and blue–yellow basis functions lie along the luminance and S-cone axes, respectively, the chromaticities of red–green basis functions form lines tilted with respect to the L–M-cone axis. 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.²⁵ ICA finds the directions that correspond to maximally independent signals; i.e., it extracts the statistical structure of the inputs. For the corresponding ICA filters, we likewise found non-orthogonal directions. PCA did not yield basis functions in these directions, probably because it is limited by the orthogonality constraint. Although it is known that chromatic properties of neurons in the lateral geniculate nucleus of primates correspond to variations along the axes of cone opponency,¹⁰ cortical neurons show sensitivities for intermediate directions.¹¹ Since the results of PCA and ICA differ in a qualitatively similar way, we suspect that opponent coding along the cardinal directions of cone opponency is used by the visual system to reliably transmit visual information to the cortex, where the information is recoded in order to better reflect the statistical structure of the environment.⁹

We had reduced the spectral dimensions by converting the spectra to the three-dimensional human cone space. This was done not only to reduce the computational requirements but also to enable us to display and interpret the results. We could have chosen any lower-dimensional basis, but since we are interested in visual information processing, the photoreceptor responses seem a natural and sensible choice. A recent study by Taylor *et al.*³³ analyzed JPEG images with the extended infomax ICA algorithm,¹⁷ resulting in basis functions defined in RGB space. Our basis functions are qualitatively similar to these in the sense that there are achromatic basis functions and also chromatic basis functions mainly along two directions in the respective color space. The results are not comparable on a more detailed level, though, since the data were qualitatively different, and different ICA algorithms were used, which leads to codes of different efficiencies (see below).

C. Independent Component Analysis Methods

To ensure that our results are not sensitive to the choice of prior and ICA method, we applied other ICA methods, including the methods of Hyvärinen and Oja,¹⁸ Cardoso,¹⁹ and Lewicki,³⁴ to the same data set and found that they produced similar results. The resulting basis functions looked similar, the coefficients were sparsely distributed, and the encoding differences between the different ICA methods were less than 10%. Note that these methods do not use a fixed distribution on the coefficients of the independent components. We initialized the algorithms with Gaussian sources, but the algorithms converged to a solution where the sources were sparsely distributed, similar to Laplacian densities. This indicates that the results presented here are insensitive to a specific ICA method. The results obtained with the Laplacian prior were most similar (encoding difference was less than 1%) to the ICA method by Lewicki.³⁴ The reasoning behind using a Laplacian prior is that this distribution

models a sparse density on the coefficients, and sparseness is essential for efficient encoding of signals.² Other sparse-source-density models could have been used for this purpose, presumably resulting in similar basis functions. However, not all super-Gaussian densities give sparse codes. It has been shown by Kreutz-Delgado *et al.*³⁵ that super-Gaussianness is a necessary but not sufficient condition for sparse signal coding when a *maximum a posteriori* coding method is found. Therefore the ICA method used by Bell and Sejnowski,³ by Lee *et al.*,¹⁷ and also by Tailor *et al.*³³ results in basis functions that give less-sparse distributions of \mathbf{s} than the method by Lewicki and Olshausen⁴ and Lee *et al.*³⁶ when applied to gray-scale images.

6. CONCLUSIONS

Independent-component analysis of natural spectra yielded basis functions that were broadband or had relatively broad peaks. Their shapes show similarities to spectra of natural illuminants and to natural reflectance spectra, thus indicating that ICA may be used to investigate the chromatic structure of natural scenes. The analysis of trichromatic image patches revealed achromatic and color-opponent basis functions with non-orthogonal opponency directions. The comparison with principal-component analysis suggests that nonorthogonal basis functions lead to more efficient encoding of chromatic information.

APPENDIX A: INDEPENDENT COMPONENT ANALYSIS

ICA is a way of finding a linear nonorthogonal coordinate system in any multivariate data that minimizes mutual information among the axial projections of the input data. The directions of the axes of this coordinate system are determined by both second-order and higher-order statistics of the original data. The goal of ICA is to perform a linear transform that makes the resulting source outputs as statistically independent of each other as possible.¹³⁻¹⁶

Assume that there is an M -dimensional zero-mean vector $\mathbf{s} = [s_1, \dots, s_M]^T$, whose components are mutually independent. The vector \mathbf{s} corresponds to M independent scalar-valued source signals s_i . We can write the multivariate probability density function of the vector as the product of marginal independent distributions:

$$p(\mathbf{s}) = \prod_{i=1}^M p_i(s_i). \quad (\text{A1})$$

A data vector $\mathbf{x} = [x_1, \dots, x_N]^T$ is observed such that

$$\mathbf{x} = \mathbf{A}\mathbf{s}, \quad (\text{A2})$$

where \mathbf{A} is an $N \times M$ scalar matrix. The columns of \mathbf{A} are called the basis functions. The mixing is assumed to be instantaneous, so there is no time delay of the source i mixing into channel j . In case of $N > M$, the rank of the data correlation matrix $\langle \mathbf{x}\mathbf{x}^T \rangle$ can be used to estimate the number of sources. In case of $N < M$, an overcomplete

set of basis functions²² can be used to adapt \mathbf{A} . For simplicity, we assume a square mixing matrix \mathbf{A} so that $N = M$.

As the components of the observed vectors are no longer independent, the multivariate probability density function will not satisfy the product equality in Eq. (A1). The mutual information $I(\mathbf{x})$ of the observed vector is given by the Kullback-Leibler divergence $D(\cdot, \cdot)$ of the multivariate density from the density written in product form:

$$I(\mathbf{x}) = \int p(\mathbf{x}) \log \frac{p(\mathbf{x})}{\prod_{i=1}^N p_i(x_i)} d\mathbf{x} = D \left(p(\mathbf{x}) \left\| \prod_{i=1}^N p_i(x_i) \right. \right). \quad (\text{A3})$$

The mutual information is positive and is equal to zero only when the components x_i are independent.³⁷

The goal of ICA is to find a linear transformation \mathbf{W} of the dependent sensor signals \mathbf{x} that makes the outputs \mathbf{u} as independent as possible,

$$\mathbf{u} = \mathbf{W}\mathbf{x} = \mathbf{W}\mathbf{A}\mathbf{s}, \quad (\text{A4})$$

so that \mathbf{u} is an estimate of the sources. The sources are exactly recovered when \mathbf{W} is the inverse of \mathbf{A} up to a permutation and scale change.

A learning algorithm for adapting the basis functions can be derived by using the information-maximization principle¹⁵ or the maximum-likelihood estimation (MLE) method.³⁸⁻⁴⁰ In the MLE approach the data likelihood is derived by marginalizing over the sources:

$$p(\mathbf{x}|\mathbf{A}) = \int p(\mathbf{x}|\mathbf{s}, \mathbf{A})p(\mathbf{s})d\mathbf{s}. \quad (\text{A5})$$

Note that the information-maximization approach to ICA and the MLE approach are equivalent in this case.⁴⁰ Since no additive noise is assumed and \mathbf{A} is a square matrix in Eq. (A2), there is a unique expression for the data, and the conditional-likelihood function becomes a delta function:

$$p(\mathbf{x}|\mathbf{s}, \mathbf{A}) = \delta(\mathbf{x} - \mathbf{A}\mathbf{s}). \quad (\text{A6})$$

Note that the integral in Eq. (A5) for scalars is $\int \delta(x - as)f(s)ds = (1/a)f(x/a)$. The data likelihood can be expressed as³⁸

$$p(\mathbf{x}|\mathbf{A}) = \frac{p(\mathbf{s})}{|\det(\mathbf{A})|}. \quad (\text{A7})$$

The log-likelihood function is

$$\log p(\mathbf{x}|\mathbf{A}) = \sum_{i=1}^N \log p(s_i) - \log |\det(\mathbf{A})|. \quad (\text{A8})$$

The learning rule is obtained by maximizing the log likelihood with respect to \mathbf{A} and using the natural gradient extension $\mathbf{A}\mathbf{A}^T$ (Ref. 41):

$$\Delta \mathbf{A} \propto \mathbf{A}\mathbf{A}^T \frac{\partial}{\partial \mathbf{A}} \log p(\mathbf{x}|\mathbf{A}) = -\mathbf{A}[\mathbf{I} - \varphi(\mathbf{s})\mathbf{x}^T], \quad (\text{A9})$$

where $\varphi(\mathbf{s}) = -[\partial p(\mathbf{s})/\partial \mathbf{s}]/[p(\mathbf{s})]$. Equation (A9) is the general ICA learning rule and involves the prior knowledge on $p(\mathbf{s})$ that is necessary to implement the application specific ICA learning rule.

Our primary interest is to learn efficient codes; i.e., the distribution of $p(\mathbf{s})$ should be as sparse as possible. Sparseness in this case assumes that the data information of \mathbf{x} is encoded in \mathbf{s} in such a way that the coefficients \mathbf{s} are mostly near zero. A Laplacian prior [$p(\mathbf{s}) \propto \exp(-|\mathbf{s}|)$], for example, is sharply peaked near zero and has heavy tails. The sparseness constraint was used in other studies²⁻⁴ to learn efficient codes for natural images. The Laplacian prior on $p(\mathbf{s})$ reduces to a simple sign function for $\varphi(\mathbf{s}) = \text{sign}(\mathbf{s})$, and the learning rule for our analysis is

$$\Delta \mathbf{A} \propto \mathbf{A}[\mathbf{I} - \text{sign}(\mathbf{s})\mathbf{s}^T]. \quad (\text{A10})$$

This is the learning rule we used for the experiments in Figs. 2A and 2B to find the basis functions.

ACKNOWLEDGMENTS

We thank Michael Lewicki and David Field for fruitful discussions and C. Párraga, G. Brelstaff, T. Troscianko, and I. Moorehead for making the hyperspectral image data set available. We also thank two anonymous reviewers for valuable comments. T.-W. Lee was supported by the Swartz Foundation.

Address correspondence to Thomas Wachtler at his current address: Universität Freiburg, Biologie III/Neurobiologie und Biophysik, Schänzlestrasse 1, 79104 Freiburg, Germany; e-mail, wachtler@biologie.uni-freiburg.de.

REFERENCES

- H. B. Barlow, "Possible principles underlying the transformation of sensory messages," in *Sensory Communication*, W. A. Rosenblith, ed. (MIT Press, Cambridge, Mass., 1961), pp. 217-234.
- B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature* **381**, 607-609 (1996).
- A. J. Bell and T. J. Sejnowski, "The 'independent components' of natural scenes are edge filters," *Vision Res.* **37**, 3327-3338 (1997).
- M. S. Lewicki and B. A. Olshausen, "A probabilistic framework for the adaptation and comparison of image codes," *J. Opt. Soc. Am. A* **16**, 1587-1601 (1999).
- G. Buchsbaum and A. Gottschalk, "Trichromacy, opponent colours coding and optimum colour information transmission in the retina," *Proc. R. Soc. London Ser. B* **220**, 89-113 (1983).
- D. L. Ruderman, T. W. Cronin, and C.-C. Chiao, "Statistics of cone responses to natural images: implications for visual coding," *J. Opt. Soc. Am. A* **15**, 2036-2045 (1998).
- J. N. Lythgoe and J. C. Partridge, "Visual pigments and the acquisition of visual information," *J. Exp. Biol.* **146**, 1-20 (1989).
- D. R. Williams, N. Sekiguchi, W. Haake, D. Brainard, and O. Packer, "The cost of trichromacy for spatial vision," in *From Pigments to Perception*, A. Valberg and B. B. Lee, eds., NATO ASI Series A, Life Sciences, Vol. 203 (Plenum, New York, 1991), pp. 11-22.
- D. J. Field, "What is the goal of sensory coding?" *Neural Comput.* **6**, 559-601 (1994).
- A. M. Derrington, J. Krauskopf, and P. Lennie, "Chromatic mechanisms in lateral geniculate nucleus of macaque," *J. Physiol.* **357**, 241-265 (1984).
- P. Lennie, J. Krauskopf, and G. Sclar, "Chromatic mechanisms in striate cortex of macaque," *J. Neurosci.* **10**, 649-669 (1990).
- T. Wachtler, T. J. Sejnowski, and T. D. Albright, "Interactions between stimulus and background chromaticities in macaque primary visual cortex," *Invest. Ophthalmol. Visual Sci. Suppl.* **40**, S641 (1999).
- C. Jutten and J. Herault, "Blind separation of sources, part I: an adaptive algorithm based on neuromimetic architecture," *Signal Process.* **24**, 1-10 (1991).
- P. Comon, "Independent component analysis—a new concept?" *Signal Process.* **36**, 287-314 (1994).
- A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural Comput.* **7**, 1129-1159 (1995).
- T.-W. Lee, *Independent Component Analysis: Theory and Applications* (Kluwer Academic, Boston, Mass., 1998).
- T.-W. Lee, M. Girolami, and T. J. Sejnowski, "Independent component analysis using an extended infomax algorithm for mixed sub-gaussian and super-gaussian sources," *Neural Comput.* **11**, 409-433 (1999).
- A. Hyvärinen and E. Oja, "A fast fixed-point algorithm for independent component analysis," *Neural Comput.* **9**, 1483-1492 (1997).
- J. F. Cardoso, "High-order contrasts for independent component analysis," *Neural Comput.* **11**, 157-192 (1999).
- C. A. Párraga, G. Brelstaff, and T. Troscianko, "Color and luminance information in natural scenes," *J. Opt. Soc. Am. A* **15**, 563-569 (1998). (<http://www.crs4.it/~gjb/fpJOSA.html>)
- T.-W. Lee, T. Wachtler, and T. Sejnowski, "The spectral independent components of natural scenes," *Tech. Rep. INC-9901* (Institute for Neural Computation, University of California, San Diego, San Diego, Calif., 1999).
- M. S. Lewicki and T. J. Sejnowski, "Learning overcomplete representations," *Neural Comput.* **12**, 337-365 (2000).
- A. Stockman, D. I. A. MacLeod, and N. E. Johnson, "Spectral sensitivities of the human cones," *J. Opt. Soc. Am. A* **10**, 2491-2521 (1993). (<http://www-cvrl.ucsd.edu>)
- D. I. A. MacLeod and R. M. Boynton, "Chromaticity diagram showing cone excitation by stimuli of equal luminance," *J. Opt. Soc. Am.* **69**, 1183-1186 (1979).
- M. A. Webster and J. D. Mollon, "Adaptation and the color statistics of natural images," *Vision Res.* **37**, 3283-3298 (1997).
- G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd ed. (Wiley, New York, 1982).
- L. T. Maloney, "Evaluation of linear models of surface spectral reflectance with small number of parameters," *J. Opt. Soc. Am. A* **3**, 1673-1683 (1986).
- J. Romero, A. Garcia-Beltrán, and J. Hernández-Andrés, "Linear basis for representation of natural and artificial illuminants," *J. Opt. Soc. Am. A* **14**, 1007-1014 (1997).
- C.-C. Chiao, T. W. Cronin, and D. Osorio, "Color signals in natural scenes: characteristics of reflectance spectra and effects of natural illuminants," *J. Opt. Soc. Am. A* **17**, 218-224 (2000).
- M. J. Vrhel, R. Gershon, and L. S. Iwan, "Measurement and analysis of object reflectance spectra," *Color Res. Appl.* **19**, 4-9 (1994).
- J. D. Mollon, "Tho' she kneel'd in that place where they grew...: the uses and origins of primate colour vision," *J. Exp. Biol.* **146**, 21-38 (1989).
- D. Osorio and M. Vorobyev, "Colour vision as an adaptation to frugivory in primates," *Proc. R. Soc. London Ser. B* **263**, 593-599 (1996).
- D. R. Taylor, L. H. Finkel, and G. Buchsbaum, "Spatiochromatic independent component filters of natural scenes," *Soc. Neurosci. Abstr.* **25**, 1935 (1999).
- M. S. Lewicki, "A flexible prior for independent component analysis," *Neural Comput.* (to be published).
- K. Kreutz-Delgado, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski, "Convex/schurconvex (csc) log-priors and sparse coding," in *6th Joint Symposium on Neural Computation* (Institute for Neural Computation, University of California, San Diego, San Diego, Calif., 1999), Vol. 9, pp. 65-71.
- T.-W. Lee, M. S. Lewicki, and T. J. Sejnowski, "Unsuper-

- vised classification with nongaussian mixture models using ica," in *Advances in Neural Information Processing Systems 11* (MIT Press, Cambridge, Mass., 1999), pp. 508–514.
37. T. Cover and J. Thomas, eds, *Elements of Information Theory* (Wiley, New York, 1991), Vol. 1.
 38. D. MacKay, "Maximum likelihood and covariant algorithms for independent component analysis," Report (University of Cambridge, Cavendish Lab, Cambridge, UK, August 1996).
 39. B. Pearlmutter and L. Parra, "A context-sensitive generalization of ICA," in *Proceedings of the International Conference on Neural Information Processing* (Springer, Singapore, 1996), pp. 151–157.
 40. J-F. Cardoso, "Infomax and maximum likelihood for blind source separation," *IEEE Signal Process. Lett.* **4**(4), 112–114 (1997).
 41. S. Amari, "Natural gradient works efficiently in learning," *Neural Comput.* **10**, 251–276 (1998).