

Emergence of simple-cell receptive field properties by learning a sparse code for natural images

Bruno A. Olshausen* & David J. Field

Department of Psychology, Uris Hall, Cornell University, Ithaca, New York 14853, USA

THE receptive fields of simple cells in mammalian primary visual cortex can be characterized as being spatially localized, oriented¹⁻⁴ and bandpass (selective to structure at different spatial scales), comparable to the basis functions of wavelet transforms^{5,6}. One approach to understanding such response properties of visual neurons has been to consider their relationship to the statistical structure of natural images in terms of efficient coding⁷⁻¹². Along these lines, a number of studies have attempted to train unsupervised learning algorithms on natural images in the hope of developing receptive fields with similar properties¹³⁻¹⁸, but none has succeeded in producing a full set that spans the image space and contains all three of the above properties. Here we investigate the proposal^{8,12} that a coding strategy that maximizes sparseness is sufficient to account for these properties. We show that a learning algorithm that attempts to find sparse linear codes for natural scenes will develop a complete family of localized, oriented, bandpass receptive fields, similar to those found in the primary visual cortex. The resulting sparse image code provides a more efficient representation for later stages of processing because it possesses a higher degree of statistical independence among its outputs.

We start with the basic assumption that an image, $I(x, y)$, can be represented in terms of a linear superposition of (not necessarily orthogonal) basis functions, $\phi_i(x, y)$:

$$I(x, y) = \sum_i a_i \phi_i(x, y) \quad (1)$$

The image code is determined by the choice of basis functions, ϕ_i . The coefficients, a_i , are dynamic variables that change from one image to the next. The goal of efficient coding is to find a set of ϕ_i that forms a complete code (that is, spans the image space) and results in the coefficient values being as statistically independent as possible over an ensemble of natural images. The reasons for desiring statistical independence have been elaborated elsewhere^{9,12,19}, but can be summarized briefly as providing a strategy for extracting the intrinsic structure in sensory signals.

One line of approach to this problem is based on principal-components analysis^{14,15,20}, in which the goal is to find a set of mutually orthogonal basis functions that capture the directions of maximum variance in the data and for which the coefficients are pairwise decorrelated, $\langle a_i a_j \rangle = \langle a_i \rangle \langle a_j \rangle$. The receptive fields that result from this process are not localized, however, and the vast majority do not at all resemble any known cortical receptive fields (Fig. 1). Principal components analysis is appropriate for capturing the structure of data that are well described by a gaussian cloud, or in which the linear pairwise correlations are the most important form of statistical dependence in the data. But natural scenes contain many higher-order forms of statistical structure, and there is good reason to believe they form an extremely non-gaussian distribution that is not at all well captured by orthogonal components¹². Lines and edges, especially curved and fractal-like edges, cannot be characterized by linear pairwise statistics^{6,21} and so a method is needed for evaluating the representation that can




FIG. 1 Principal components calculated on 8×8 image patches extracted from natural scenes by using Sanger's rule¹⁴. The full set of 64 components is shown, ordered by their variance (by columns, then by rows). The oriented structure of the first few principal components does not arise as a result of the oriented structures in natural images, but rather because these functions are composed of a small number of low-frequency components (the lowest spatial frequencies account for the greatest part of the variance in natural scenes⁸). Reconstructions based solely on the first row of functions will merely yield blurry images. Identical-looking components are obtained for images with the same amplitude spectrum as natural images but with randomized phases (that is, $1/f$ noise).

take into account higher-order statistical dependences in the data.

The existence of any statistical dependences among a set of variables may be discerned whenever the joint entropy is less than the sum of individual entropies, $H(a_1, a_2, \dots, a_n) < \sum_i H(a_i)$, otherwise the two quantities will equal. Assuming that we have some way of ensuring that information in the image (joint entropy) is preserved, then a possible strategy for reducing statistical dependences is to lower the individual entropies, $H(a_i)$, as much as possible. In Barlow's terms¹⁹, we seek a minimum-entropy code. We conjecture that natural images have 'sparse structure'—that is, any given image can be represented in terms of a small number of descriptors out of a large set^{8,12}—and so we shall seek a specific form of low-entropy code in which the probability distribution of each coefficient's activity is unimodal and peaked around zero.

The search for a sparse code can be formulated as an optimization problem by constructing the following cost function to be minimized:

$$E = -[\text{preserve information}] - \lambda[\text{sparseness of } a_i] \quad (2)$$

where λ is a positive constant that determines the importance of the second term relative to the first. The first term measures how well the code describes the image, and we choose this to be the mean square of the error between the actual image and the reconstructed image:

$$[\text{preserve information}] = - \sum_{x,y} \left[I(x, y) - \sum_i a_i \phi_i(x, y) \right]^2 \quad (3)$$

The second term assesses the sparseness of the code for a given image by assigning a cost depending on how activity is distributed among the coefficients: those representations in which activity is spread over many coefficients should incur a higher cost than those in which only a few coefficients carry the load. The cost function we have constructed to meet this criterion takes the sum

* Present address: Center for Neuroscience, UC Davis, Davis, California 95616, USA.

of each coefficient's activity passed through a nonlinear function $S(x)$:

$$[\text{sparseness of } a_i] = - \sum_j S\left(\frac{a_j}{\sigma}\right) \quad (4)$$

where σ is a scaling constant. The choices for $S(x)$ that we have experimented with include $-e^{-x^2}$, $\log(1+x^2)$ and $|x|$, and all yield qualitatively similar results (described below). The reasoning behind these choices is that they will favour among activity states with equal variance those with the fewest number of non-zero coefficients. This is illustrated in geometric terms in Fig. 2.

Learning is accomplished by minimizing the total cost functional, E (equation (2)). For each image presentation, E is minimized with respect to the a_i . The ϕ_i then evolve by gradient descent on E averaged over many image presentations. Thus for a given image, the a_i are determined from the equilibrium solution to the differential equation:

$$\dot{a}_i = b_i - \sum_j C_{ij} a_j - \frac{\lambda}{\sigma} S'\left(\frac{a_i}{\sigma}\right) \quad (5)$$

where $b_i = \sum_{x,y} \phi_i(x,y) I(x,y)$ and $C_{ij} = \sum_{x,y} \phi_i(x,y) \phi_j(x,y)$. The learning rule for updating the ϕ is then:

$$\Delta \phi_i(x_m, y_n) = \eta \left\langle a_i \left[I(x_m, y_n) - \hat{I}(x_m, y_n) \right] \right\rangle \quad (6)$$

where \hat{I} is the reconstructed image, $\hat{I}(x_m, y_n) = \sum_i a_i \phi_i(x_m, y_n)$, and η is the learning rate. One can see from inspection of equations (5) and (6) that the dynamics of the a_i , as well as the learning rule for the ϕ_i , have a local network implementation. An intuitive way of understanding the algorithm is that it is seeking a set of ϕ_i for which the a_i can tolerate 'sparsification' with minimum reconstruction error. Importantly, the algorithm allows for the basis functions to be overcomplete (that is, more basis functions than meaningful dimensions in the input) and non-orthogonal⁵, without reducing the degree of sparseness in the representation. This is because the sparseness cost function, S , forces the system to choose, in the case of overlaps, which basis functions are most effective for describing a given structure in the image.

The learning rule (equation (6)) was tested on several artificial datasets containing controlled forms of sparse structure, and the




FIG. 2 The cost function for sparseness, plotted as a function of the joint activity of two coefficients, a_i and a_j . In this example, $S(x) = \log(1+x^2)$. An activity vector that points towards a corner, where activity is distributed equally between coefficients, will incur a higher cost than a vector with the same length that lies along one of the axes, where the total activity is loaded onto one coefficient. The gradient tends to 'sparsify' activity by differentially reducing the value of low-activity coefficients more than high-activity coefficients. Alternatively, the sparseness cost function may be interpreted as the negative logarithm of the prior probability of the a_i (ref. 23), assuming statistical independence among the a_i (that is, a factorial distribution), and with the shape of the distribution specified by S (in this case a Cauchy distribution).

results of these tests (Fig. 3) confirm that the algorithm is indeed capable of discovering sparse structure in input data, even when the sparse components are non-orthogonal. The result of training the system on 16×16 image patches extracted from natural scenes is shown in Fig. 4a. The vast majority of basis functions are well localized within each array (with the exception of the low-frequency functions). Moreover, the functions are oriented and selective to different spatial scales. This result should not come as a surprise, because it simply reflects the fact that natural images contain localized, oriented structures with limited phase alignment across spatial frequency⁶. The functions ϕ_i shown are the feedforward weights that, in addition to other terms, contribute to the value of each a_i (refer to term b_i in equation (5)). To establish the correspondence to physiologically measured receptive fields, we mapped out the response of each a_i to spots at every position: the results of this analysis show that the receptive fields are very similar in form to the basis functions (Fig. 4b). The entire set of basis functions forms a complete image code that spans the joint space of spatial position, orientation and scale (Fig. 4c) in a manner similar to wavelet codes, which have previously been shown to form sparse representations of natural images^{8,12,22}. The average spatial-frequency bandwidth is 1.1 octaves (s.d., 0.5) with an average aspect ratio (length/width) of 1.3 (s.d., 0.5), which are characteristics reasonably similar to those of simple-cell receptive fields (~ 1.5 octaves, length/width ~ 2)⁵. The resulting histograms have sparse distributions (Fig. 4d), decreased entropy

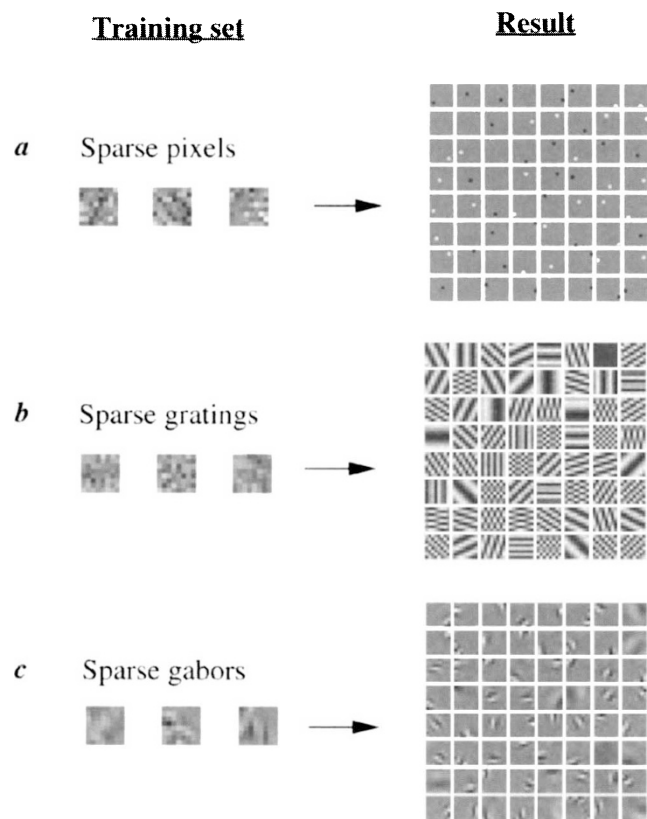
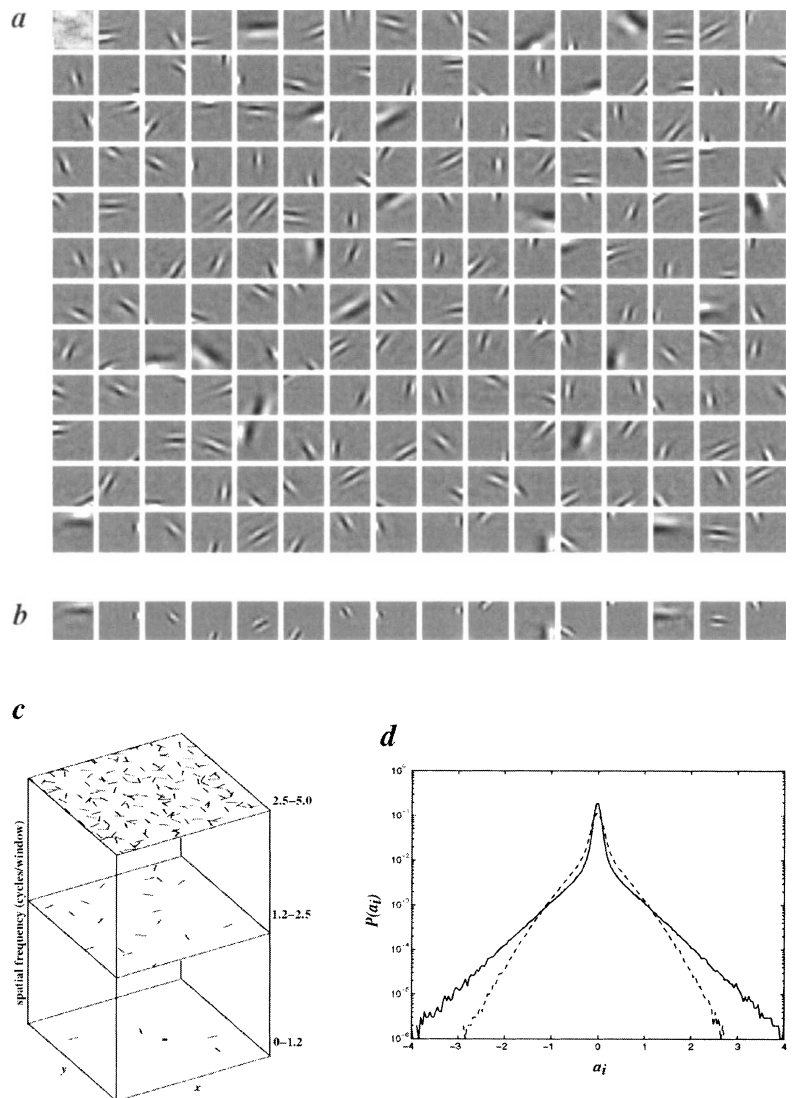


FIG. 3 Test cases. Representative training images are shown at the left and the resulting basis functions that were learned from these examples are shown at the right. In a, images were composed of sparse pixels: each pixel was activated independently according to an exponential distribution, $P(x) = e^{-|x|}/Z$. In b, images were composed similarly to a, except with gratings instead of pixels (that is, 'sparse pixels' in the Fourier domain). In c, images were composed of sparse, non-orthogonal Gabor functions with the method described by Field¹². In all cases, the basis functions were initialized to random initial conditions. The learned basis functions successfully recover the sparse components from which the images were composed. The form of the sparseness cost function was $S(x) = -e^{-x^2}$, but other choices (see text) yield the same results.

FIG. 4 Results from training a system of 192 basis functions on 16×16 -pixel image patches extracted from natural scenes. The scenes were ten 512×512 images of natural surroundings in the American northwest, preprocessed by filtering with the zero-phase whitening/lowpass filter $R(f) = fe^{-f/f_0}$, $f_0 = 200$ cycles/picture (see also ref. 9). Whitening counteracts the fact that the mean-square error (or m.s.e.) preferentially weights low frequencies for natural scenes, whereas the attenuation at high spatial-frequencies eliminates artefacts of rectangular sampling. The a_i were computed by the conjugate gradient method, halting when the change in E was less than 1%. The ϕ_i were initialized to random values and were updated every 100 image presentations. The vector length (gain) of each basis function, ϕ_i , was adapted over time so as to maintain equal variance on each coefficient. A stable solution was arrived at after $\sim 4,000$ updates ($\sim 400,000$ image presentations). The parameter λ was set so that $\lambda/\sigma = 0.14$, with σ^2 set to the variance of the images. The form of the sparseness cost function was $S(x) = \log(1 + x^2)$. **a**, The learned basis functions, scaled in magnitude so that each function fills the grey scale, but with zero always represented by the same grey level (black is negative, white is positive). **b**, The receptive fields corresponding to the last row of basis functions in **a**, obtained by mapping with spots (single pixels preprocessed identically with the images). The principal difference may be accounted for by the fact that sparsifying of activity makes units more selective in which aspects of the stimulus they respond to. **c**, The distribution of the learned basis functions in space, orientation and scale. The functions were subdivided into high-, medium- and low-spatial-frequency bands (in octaves), according to the peak frequency in their power spectra, and their spatial location was plotted within the corresponding plane. Orientation preference is denoted by line orientation. **d**, Activity histograms averaged over all coefficients for the learned basis functions (solid line) and for random initial conditions (broken line). In both cases, $\lambda/\sigma = 0.14$, showing that the learned basis functions can accommodate a higher degree of sparsification. Note that even the random basis functions have positive kurtosis due to sparsification. The width of each bin used in calculating the entropy was 0.04.



(4.0 bits compared with 4.6 bits before training), and increased kurtosis (20 compared with 7.0) for a mean-square reconstruction error that is 10% of the image variance.

These results demonstrate that localized, oriented, bandpass receptive fields emerge when only two global objectives are placed on a linear coding of natural images: that information be preserved, and that the representation be sparse. These two objectives alone are sufficient to account for the principal spatial properties of simple-cell receptive fields. A number of unsupervised learning algorithms based on similar principles have been

proposed for finding efficient representations of data²³⁻³⁰, all of which seem to have the potential to arrive at results like these. What remains as a challenge for these algorithms, and also for ours, is to provide an account of other response properties of simple cells (for example, direction selectivity), as well as the complex response properties of neurons at later stages of the visual pathway, which are noted for being highly nonlinear. An important question, then, is whether these higher-order properties can be understood by considering the remaining forms of statistical dependence that exist in natural images. □

Received 10 November 1995; accepted 25 April 1996.

- Hubel, D. H. & Wiesel, T. N. *J. Physiol., Lond.* **195**, 215–244 (1968).
- De Valois, R. L., Albrecht, D. G. & Thorell, L. G. *Vision Res.* **22**, 545–559 (1982).
- Jones, J. P. & Palmer, L. A. *J. Neurophysiol.* **58**, 1233–1258 (1987).
- Parker, A. J. & Hawken, M. J. *J. opt. Soc. Am.* **A5**, 598–605 (1988).
- Daugman, J. G. *Computational Neuroscience* (ed. Schwartz, E.) 403–423 (MIT Press, Cambridge, MA, 1990).
- Field, D. J. in *Wavelets, Fractals, and Fourier Transforms* (eds Farge, M., Hunt, J. & Vascillicos, C.) 151–193 (Oxford Univ. Press, 1993).
- Srinivasan, M. V., Laughlin, S. B. & Dubs, A. *Proc. R. Soc. Lond.* **B216**, 427–459 (1982).
- Field, D. J. *J. opt. Soc. Am.* **A4**, 2379–2394 (1987).
- Atick, J. J. *Network* **3**, 213–251 (1992).
- van Hateren, J. H. *Nature* **360**, 68–70 (1992).
- Ruderman, D. L. *Network* **5**, 517–548 (1994).
- Field, D. J. *Neur. Comput.* **6**, 559–601 (1994).
- Barrow, H. G. in *IEEE First Int. Conf. on Neural Networks* Vol. 4, (eds Caudill, M. & Butler, C.) 115–121 (Institute of Electrical and Electronics Engineers, 1994).
- Sanger, T. D. in *Advances in Neural Information Processing Systems* Vol. 1 (ed. Touretzky, D.) 11–19 (Morgan-Kaufmann, 1989).
- Hancock, P. J. B., Baddeley, R. J. & Smith, L. S. *Network* **3**, 61–72 (1992).
- Law, C. C. & Cooper, L. N. *Proc. natn. Acad. Sci. U.S.A.* **91**, 7797–7801 (1994).
- Fyfe, C. & Baddeley, R. *Network* **6**, 333–344 (1995).

- Schmidhuber, J., Eldracher, M. & Foltin, B. *Neur. Comput.* **8** 773–786 (1996).
- Barlow, H. B. *Neur. Comput.* **1**, 295–311 (1989).
- Linsker, R. *Computer* 105–117 (March, 1988).
- Olshausen, B. A. & Field, D. J. *Network* **7**, 333–339 (1996).
- Daugman, J. G. *IEEE Trans. biomed. Engng.* **36**, 107–114 (1989).
- Harpur, G. F. & Prager, R. W. *Network* **7**, 277–284 (1996).
- Foldiak, P. *Biol. Cybernet.* **64**, 165–170 (1990).
- Zemel, R. S. thesis, Univ. Toronto (1993).
- Intrator, N. *Neur. Comput.* **4**, 98–107 (1992).
- Bell, A. J. & Sejnowski, T. J. *Neur. Comput.* **7**, 1129–1159 (1995).
- Saund, E. *Neur. Comput.* **7**, 51–71 (1995).
- Hinton, G. E., Dayan, P., Frey, B. J. & Neal, R. M. *Science* **268**, 1158–1161 (1995).
- Lu, Z. L., Chubb, C. & Sperling, G. Technical Report MBS 96-15 (Institute for Mathematical Behavioral Sciences, University of California at Irvine, 1996).

ACKNOWLEDGEMENTS. We thank M. Lewicki for helpful discussions at the inception of this work, and C. Lee, C. Brody, G. Harpur, F. Girosi and M. Riesenhuber for useful input. This work was supported by grants from NIMH to both authors. Part of this work was carried out at the Center for Biological and Computational Learning at the Massachusetts Institute of Technology.

CORRESPONDENCE and requests for materials to be addressed to B.A.O. (e-mail: bruno@ai.mit.edu). The program for running the simulation, as well as the images used in training, are available via <http://redwood.psych.comell.edu/sparsenet/sparsenet.html>.