

INDEPENDENT COMPONENT FILTERS OF NATURAL IMAGES COMPARED WITH SIMPLE CELLS IN PRIMARY VISUAL CORTEX

J.H. van Hateren and A. van der Schaaf

Department of Biophysics, University of Groningen, Nijenborgh 4,
NL-9747 AG Groningen, The Netherlands

(Preprint November 21, 1997; reference: Proc.R.Soc.Lond.B 265:359-366, 1998)

SUMMARY

Properties of the receptive fields of simple cells in macaque cortex were compared with properties of independent component filters generated by independent component analysis (ICA) on a large set of natural images. Histograms of spatial frequency bandwidth, orientation tuning bandwidth, length and aspect ratio of the receptive fields match well. This indicates that simple cells are well tuned to the expected statistics of natural stimuli. There is no match, however, in calculated and measured distributions for the peak of the spatial frequency response: the filters produced by ICA do not vary their spatial scale as much as simple cells do, but are fixed to scales close to the finest ones allowed by the sampling lattice. Possible ways to resolve this discrepancy are discussed.

INTRODUCTION

In this article we investigate to what extent the statistical properties of natural images can be used to understand the variation of receptive field properties of simple cells in the mammalian primary visual cortex. The receptive fields of simple cells have been studied extensively (e.g., Hubel & Wiesel 1968, DeValois et al. 1982a, DeAngelis et al. 1993): they are localised in space and time, have band-pass characteristics in the spatial and temporal frequency domains, are oriented, and are often sensitive to the direction of motion of a stimulus. Here we will concentrate on the spatial properties of simple cells.

Several hypotheses as to the function of these cells have been proposed. As the cells preferentially respond to oriented edges or lines, they can be viewed as edge or line detectors. Their joint localisation in both the spatial domain and the spatial frequency domain has led to the suggestion that they mimic Gabor filters, minimising uncertainty in both domains (Daugman 1980, Marcelja 1980). More recently, the match between the operations performed by simple cells and the wavelet transform has attracted attention (e.g., Field 1993). The approaches based on Gabor filters and wavelets basically consider processing by the visual cortex as a general image processing strategy, relatively independent of detailed assumptions about image statistics. On the other hand, the edge and line detector hypothesis is based on the intuitive notion that edges and lines are both abundant and important in images. This theme of relating simple cell properties with the statistics of natural images was explored extensively by Field (1987, 1994). He proposed that the cells are optimized specifically for coding natural images. He argued that one possibility for such a code, sparse coding, simplifies further processing in the visual system, because it produces a representation of the stimulus that helps detection of coincidences (Barlow 1972, 1994). Indeed, Olshausen & Field (1996) showed that imposing sparseness on the output of receptive fields being trained on natural images produced receptive fields similar to those of simple cells (see also Harpur 1997).

This result was very recently put into the context of independent component analysis (ICA) by Bell and Sejnowski (1997a, 1997b) and Hurri et al. (1996); for a discussion of the connection between the various algorithms see Olshausen & Field (1997). For (linear) ICA one considers an ensemble of signals, each produced by an unknown linear superposition of unknown independent (elementary) signals. By presenting the ICA algorithm with a large number of examples of such signals, it is able to reconstruct the elementary signals, at least if the elementary signals have non-gaussian probability densities (i.e., the distribution of the strengths with which each elementary signal is present in a set of images is not a Gaussian). ICA on natural

images again produces receptive fields like those of simple cells (Bell & Sejnowski 1997a,b; Hurri et al. 1996, Hurri 1997). Although the components produced by ICA on natural images are not completely independent, they are as independent as possible by a linear transformation.

It should be noted here that the independent component model of the primary visual cortex should not be regarded as a full model of simple cells in primary cortex. As it is a linear, non-adaptive model, many aspects of simple cells are ignored, such as contrast adaptation (e.g., Sclar et al. 1989), contrast normalisation (Heeger 1992), nonlinearities involved in orientation tuning (Volgushev et al. 1996), adaptation to various stimulus statistics (Zipser et al. 1996), and so on. Nevertheless, the model has a clear information theoretic interpretation, and is strictly based on stimulus statistics. One may hope that nonlinearities and adaptation can be added in a process of stepwise refinement, with the linear model as a solid basis. This appears to be a successful strategy in earlier stages of vision (vertebrate retina and LGN, fly visual system), where linear theories are quite successful (e.g., Srinivasan et al. 1982, Atick 1992, van Hateren 1992a,b, Linsker 1993, Dong & Atick 1995, Dan et al. 1996), whereas it is known that the early visual system contains (e.g., Laughlin 1981, van Hateren 1997) and needs (e.g., Ruderman and Bialek 1994) stimulus related nonlinearities that modify the linear default.

If decomposing an image into independent components is indeed one of the main functions of simple cells, it is expected that the distribution of their properties, such as spatial frequency bandwidth and orientation tuning bandwidth, is determined by the statistics of the visual environment. Olshausen & Field (1996) report spatial frequency bandwidths and aspect ratios (ratio of length and width of the receptive field) close to those measured in simple cells. Here we extend this result by performing ICA on a large set of calibrated images, and comparing a series of properties of the resulting receptive fields with those of receptive fields measured in simple cells. We find that there is a good correspondence between the distributions for spatial frequency bandwidth, orientation tuning bandwidth, aspect ratio, and receptive field length. For the peak of the spatial frequency sensitivity, however, the results deviate strongly: whereas ICA yields peaks close to the maximum spatial frequency allowed by the sampling lattice, measurements in simple cells show a much broader distribution. The implications and possible resolution of this discrepancy are discussed.

METHODS

Images

The image set consisted of 4212 images obtained with a Kodak DCS420 digital camera (with a 28 mm camera lens). For the intensity this camera uses 12-bit sampling internally, which is then reduced to and stored as 8-bit data via a nonlinear scale table. As this table is recorded for each image, it can be used afterwards to expand the 8-bit data to a linear scale. Although the latter scale is, strictly speaking, not genuinely 12 bit deep, it is effectively close to it: details in the shadows are retained, whereas high peaks of intensity (see e.g. van Hateren 1997) are more roughly quantised, but without the clipping that an 8-bit system would necessitate. The linearity of the camera was checked with a set of calibrated neutral density filters, and found to be satisfactory. The (slight) blur in the images caused by the optical system of the camera was measured and corrected through the following procedure (performed separately for each diaphragm of the camera lens used). Spatial point spread functions of the camera were measured for a small point source presented at a large number of random positions. The resulting images (sampled point spread functions) were Fourier transformed, resulting in slightly different amplitude spectra depending on the exact position of the image of the point source on the CCD sampling grid. Of these spectra, the group with the shallowest high frequency fall-off was selected for inverse filtering the images of natural scenes. Only the images of point sources corresponding to those spectra would thus be reconstructed as a sharp point source. Point sources at other positions with steeper spectra would remain slightly blurred in a reconstruction. As a result of this procedure, the deblurring does not cause spurious structures (such as fringes at the edges) in the images of natural scenes. After inverse filtering, the images had a resolution of 1536x1024. This was subsequently reduced to 768x512 by block averaging, before extracting image patches for the ICA. This procedure reduces the risk of any remaining calibration insufficiencies influencing the ICA. Noise in these

images is negligible. The final images had an angular resolution of approximately 2 minutes of arc per pixel. The images were taken in various environments (wood, open landscapes, urban areas). In a particular environment a typical series of 100-200 consecutive images was taken. Therefore, when going through the series of 4212 images, the images statistics may change regularly. Some of this variability was incorporated in the results of the ICA by using samples from consecutive images rather than random ones taken from the entire set for a particular ICA run (see below).

A second set of images that was used (12x12 video in Fig. 3) consisted of frames grabbed from television broadcasts. This set was uncalibrated, but more diverse than the calibrated set. Images were taken from programs on a wide range of subjects, including wild life (ranging from arctic to tropical), sports, and movies.

Independent Component Analysis

Algorithms performing (linear) ICA decompose each signal of an ensemble into components (also called 'basis vectors') that are as independent as possible by a linear transformation of the signals. The amplitude of a particular component is extracted by a corresponding weight vector (also called a 'filter', see Bell & Sejnowski 1997). Of the various algorithms for performing ICA, the recent algorithm presented by Hyvärinen & Oja (1997a,b) and Hyvärinen (1997) was used. This algorithm implements ICA by finding filters that produce extrema of the kurtosis (the kurtosis is a measure of how 'peaked' a distribution is; high kurtosis means a high central peak and long tails in the distribution). Note that either maxima or minima in the kurtosis can be found (making it more general than maximising sparseness), although we found in practice only maxima for our data set. The idea of the method is that the independent components must have an extremum of the kurtosis, as any impure component (i.e., a linear superposition of two or more pure independent components) would result in a change of the kurtosis towards zero (in the limit a kurtosis of zero, as the central limit theorem states that a linear summation of a large number of independent variables, each with finite variance, will produce a Gaussian distribution, which has zero kurtosis). The Hyvärinen algorithm (the fixed-point algorithm, with a serial deflation scheme as in Hyvärinen & Oja 1997b, using function g_2 as in Hyvärinen 1997) was implemented as a parallel program on a Cray J932 at the Centre for High-Performance Computing of the University of Groningen. A typical run was performed on 120000 different samples of 18x18 image patches taken evenly spaced from 100-120 consecutive images of the image data set; such a run took about half an hour. Histograms were compiled from 33-160 of such runs, based on different subsets of the image set.

Before processing, the logarithm of the intensities was taken, firstly, because this incorporates the contrast invariance of natural scenes, secondly, because it leads to better behaved first-order statistics of natural images (e.g., Ruderman 1994), and, thirdly, because it is similar to the operations performed by the first stages of visual systems (e.g., van Hateren 1997). Note, however, that for the low contrast stimuli as typically used for measuring simple cell properties (e.g., the histograms in Fig. 4) taking the logarithm hardly makes any difference. The results of the ICA algorithm were obtained in PCA whitened space (see below), and subsequently transformed back to the original image space (with intensities represented on a logarithmic scale). All results presented below are in this space, and the filters [e.g. Fig. 1(b)] are thus meant to work on the (logarithmic) original images.

Early calculations were performed without dimensionality reduction in the PCA (the principal component analysis which is part of the data whitening used before applying the Hyvärinen algorithm). This leads to about 70% oriented IC filters, and 30% non-oriented. The latter only extend over a few pixels, and have about equal power in all four corners of the power spectrum. We consider these filters as an aliasing artefact, and they were not included in the analysis (by selecting only the filters with oriented power spectra). In later calculations, the dimensionality of the data was reduced by about 25% (i.e., from $18 \times 18 = 324$ degrees of freedom to 240 degrees of freedom, by selecting the most significant principal components), which leads only to oriented filters; these were all included in the analysis. Both of the above procedures led to very similar histograms for all parameters investigated. Also reducing the dimensionality further to 160, 80, or 40 degrees of freedom led to similar results, only differing in the spatial scale, but not in the relative distribution of the parameters.




Figure 1. Independent component analysis on natural images (18×18 patches from 103 consecutive images, dimension reduced via PCA to 240 of 324). (a) IC basis vectors, and (b) corresponding IC filters (filtering an image with an IC filter yields the strength of the corresponding basis vector in the image). Signs of basis vectors and filters are arbitrary. (c) Amplitude spectra of the filters of (b), with darker grey values coding larger amplitudes. Zero spatial frequency is at the centre of each patch.

RESULTS

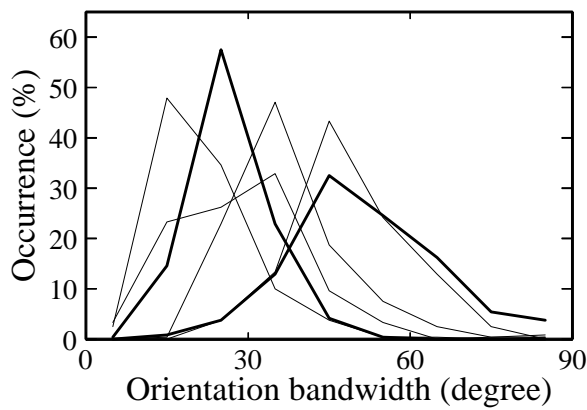
Figure 1 shows an example of independent components [Fig. 1(a)], i.e., the basis vectors, and the corresponding filters required to extract the strength of each component from an image [Fig. 1(b)], i.e., the weight vectors (or transform coefficients). Both basis vectors and filters are shown here (and below) as they appear in image space (after the logarithm of the intensity was taken). Whereas the independent components (ICs) can be considered as the constituting elements of images (see Bell & Sejnowski 1997a,b; Olshausen & Field 1996), the IC filters needed to extract their strengths are analogous to the receptive fields of neurons analysing a scene. Thus the receptive fields of cortical neurons should be compared to the IC filters, and not to the ICs. As discussed by Bell & Sejnowski (1997a,b), the ICs resemble short edges, and the IC filters are similar to simple cell receptive fields, i.e., resembling Gabor filters or wavelets. The filters are usually low-pass in one direction, and band-pass in the orthogonal direction. Figure 1(c) shows the amplitude spectra of the filters, with zero spatial frequency in the centre of each patch. From sets of filters as in Fig. 1(b) and (c) it is possible to extract several descriptive parameters. Following work analysing cortical receptive fields (in particular DeValois et al. 1982a,b; Parker & Hawken 1988), we will investigate here the following properties:

- spatial frequency bandwidth; this is defined here as the full width at half maximum (FWHM) of each filter along the orientation of the peak in the amplitude spectrum; it is expressed in octaves (factors of two in frequency).
- orientation tuning bandwidth; defined as the FWHM along a circle (with its origin at zero spatial frequency) through the peak in the amplitude spectrum.
- peak spatial frequency and peak orientation: spatial frequency and orientation of the peak in the amplitude spectrum.
- length and aspect ratio of the receptive fields; first the (2D) Hilbert transform of the filter was calculated (the quadrature phase filter), which then yields, together with the original filter, the filter's envelope by subsequently squaring both filters, adding the results, and taking the square root (see Field & Tolhurst

1986, DeAngelis et al. 1993). The length is then defined as the FWHM of the envelope along the orientation into which the filter is low-pass, and the width as the FWHM along the orientation into which the filter is band-pass. The aspect ratio is defined here as the ratio of length and width.

We found that the distribution of each parameter depends on several factors. First, it depends on the particular set of images used for the ICA. Figure 2(a) shows several examples of the distribution of orientation tuning bandwidths for different subsets of images. Each of the curves resulted from ICA on a different set of 103 consecutive images taken from the entire set. All 40 sets thus analysed had histograms similar to or scattering between the ones shown in the figure. To further illustrate the variability of both the IC basis vectors and the IC filters, Fig. 2(b)-(e) shows samples of these (representative of the complete basis set) for the two subsets corresponding to the thick lines in Fig. 2(a). The left thick line corresponds to basis vectors in Fig. 2(b) and filters in (c), the right thick line to (d) and (e). The first image set consists mostly of images of a waterside landscape with grass and reed, the second image set consists mostly of images inside a wood, with trees, leaves, and foliage. Observed superficially, the two sets do not appear very much different,

(a)



(b)

(c)

(d)

(e)




Figure 2. Variation of ICs for different ensembles of images. (a) Distribution of orientation tuning bandwidth for 6 different ensembles. (b) Examples of basis vectors for the left thick curve in (a), and (c) corresponding IC filters. (d) Examples of basis vectors for the right thick curve in (a), and (e) corresponding IC filters. Signs of basis vectors and filters are arbitrary. See text for further details.

but as Fig. 2(b)-(e) shows they lead to rather different sets of ICs. This suggests that the number of independent components needed for a large ensemble of images is larger than the number of degrees of freedom of single image patches (see also the Discussion).

Variability of the results of ICA is further investigated in Fig. 3 for the spatial frequency bandwidths of the IC filters. Here we first varied the way of extracting image patches from the database: either patches taken from consecutive image sets (18×18 consecutive as above), or patches drawn randomly from the entire database (18×18 random). The latter procedure still yielded a wide range of bandwidths, but somewhat more restricted than for the consecutive image sets. Decreasing the size of the consecutive image sets from 103 images per set to 26 per set (18×18 short) did not noticeably broaden the distribution further. Second, we varied the type of preprocessing (PCA reduced, with the dimensionality of the patches reduced from 324 to 240), and the size of the patches considered (to 12×12). This produced similar histograms as the 18×18 consecutive results. Finally, an entire different image database was used (12×12 video images generated from video frames grabbed from television broadcasts). Again, the results are not dramatically different.

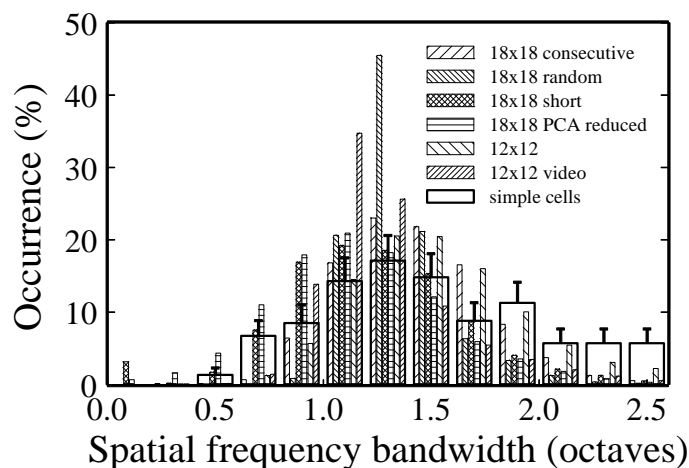


Figure 3. A comparison of different ways to perform the ICA with each other and with measurements in macaque simple cells. Data from DeValois et al. 1982b, their figure 5, foveal simple cells. See text for further details.

Figure 3 also shows measurements of spatial frequency bandwidths (FWHM) in simple cells in macaque cortex (histogram bars, data from DeValois et al. 1982b, cells recorded from the foveal area). Comparing these measurements with the ICA results, we see that the distributions roughly match, despite the variability in ICA results. This variability should probably not be considered as just a random estimation error, but rather as an inherent property associated with systematic changes in the statistics of different ensembles of natural images. Therefore, we consider the standard deviation of the ICA histograms (rather than the standard error of the mean) as a rough but reasonable measure for the reliability of the estimate. Below we will use this, together with the mean of the ICA histograms, for comparing the ICA results with various properties measured in simple cells.

Figure 4(a) shows this for the same data as in Fig. 3. The error for each bin of the histogram of simple cells was estimated by taking the square root of the number of cells in each bin. Figure 4(a) shows that the calculated and measured distributions are roughly similar; a chi-square test yields $p > 0.26$, which means that the hypothesis that the calculated and measured histograms are identical can not be rejected. The mean and the standard deviation of the spatial frequency bandwidths shown here are similar to those reported by Olshausen & Field (1996), although a direct comparison is not fully appropriate because they analyse basis functions rather than filters.

Figure 4(b) shows that calculated and measured orientation tuning bandwidths are similarly distributed (data from both DeValois et al. 1982a and Parker & Hawken 1988). Although the calculated curve appears to peak at a slightly smaller orientation tuning bandwidth (20°-30°) than the measurements in simple cells (30°-40°), the errors in the curves are such that they are not significantly different ($p > 0.18$).

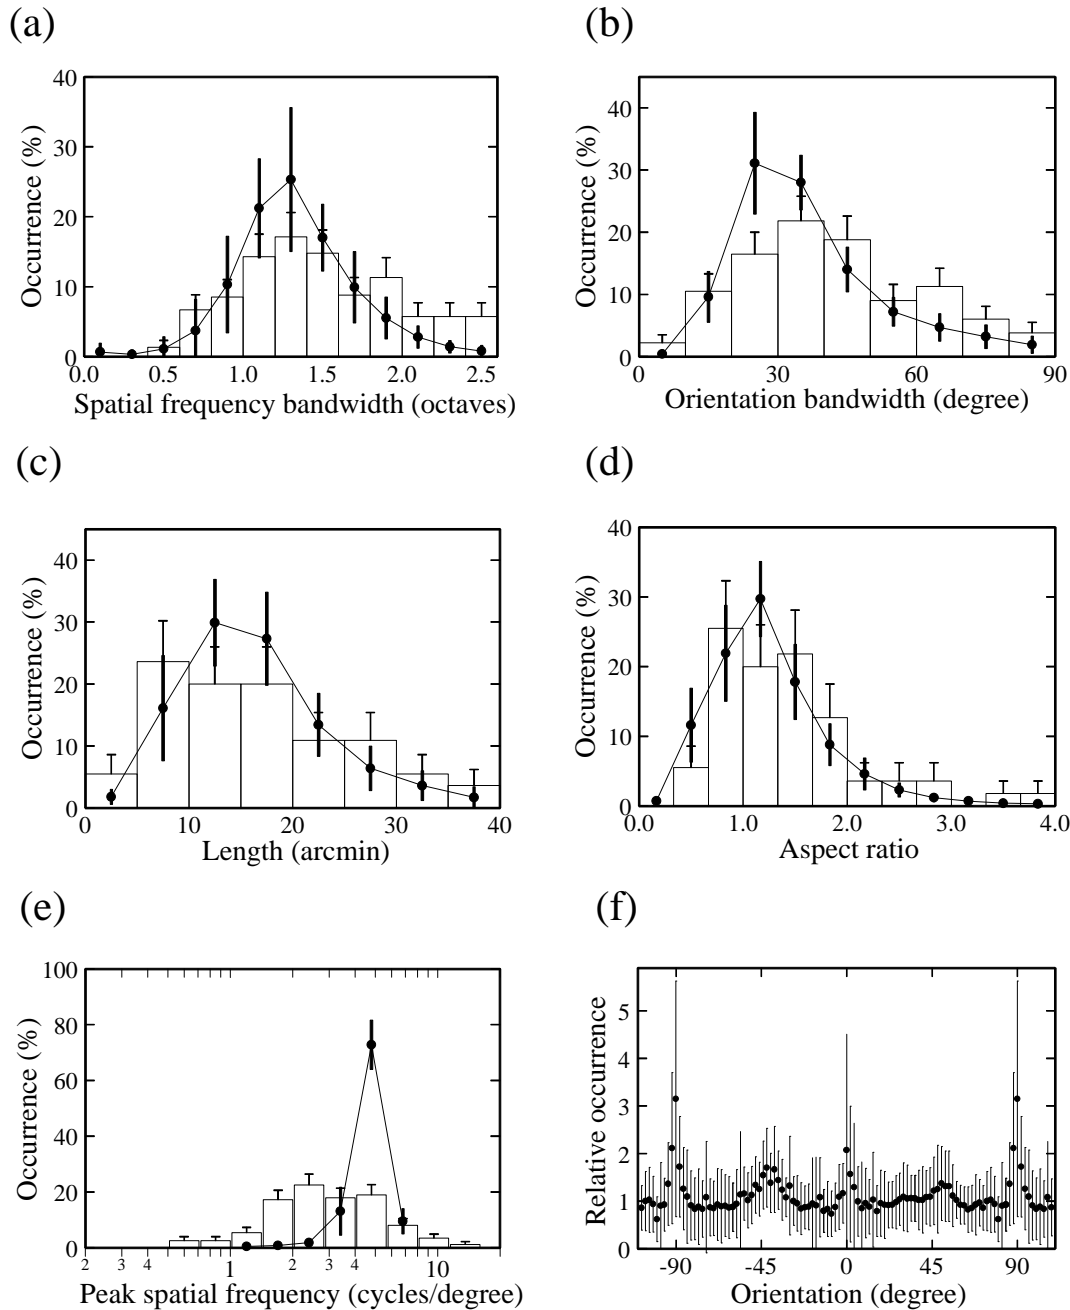


Figure 4. Comparison of IC filters (dots with error bars show mean and s.d. of the calculated histograms as in Fig. 3) to measurements of properties of simple cells in the foveal area of macaque primary visual cortex (histograms). Data from DeValois et al. 1982a [(b), their fig. 3A; merged with data from Parker & Hawken, see below], DeValois et al. 1982b [(a), fig. 5; (e), fig. 6], and Parker & Hawken 1988 [(b), fig. 3a; (c), fig. 4a; (d), fig. 6a]. Scaling in (c) for calculated data: 1.2 pixel/bin. Filters sensitive to horizontal structures correspond to 0° in (f). See text for further details.

Figure 4(c) shows distributions of the length of the calculated filters and measured receptive fields ('height' as measured by Parker & Hawken 1988). Here the abscissa gives the length (in minutes of arc of visual angle) for the histogram of simple cell measurements. The calculated curve was scaled along this axis such that it provided a good match with the measurements. Thus only the shape of the distributions should be compared, not the position of the peaks. This procedure was necessary, because the calculations do not yield absolute visual angles that can be compared directly with those of the primate visual system: firstly, because the spatial resolution of the first data set used (taken with the digital camera) is not as fine as in the

primate fovea (2' between pixels for the image patches, versus approximately 30" for the primate fovea); secondly, because, even if they had been identical, the true sample base as used by the primate cortex may be different from that of the cones; and thirdly, because the spatial calibration of the second data set (images grabbed from television broadcasts) is not known and variable because of diverse camera lenses. If images are approximately scale invariant over the relatively small range of scales involved here, the distribution of lengths should be scale invariant as well. Figure 4(c) shows that the shape of the calculated and measured distributions of receptive field lengths is similar ($p > 0.68$).

The aspect ratio for the calculated filters is defined here as the ratio of length and width of the envelope of the filter (see above). From measurements by Parker & Hawken (1988) of length and width of the central lobe of the simple cell receptive field, the aspect ratio was estimated by assuming that the width of the envelope is about three times the width of the central lobe. Figure 4(d) shows the resulting distributions: again a close correspondence between ICA filters and receptive fields in simple cells ($p > 0.47$). Similar aspect ratios were reported by Olshausen & Field (1996) for their basis functions.

Figure 4(e) shows the distribution of spatial frequencies giving the maximum response for ICA filters and for simple cells (data from DeValois et al. 1982b). Contrary to the results presented above, the two distributions now deviate strongly ($p < 0.001$). Again there is no well-defined spatial scale for the calculated curve, which was arbitrarily positioned. As will be clear from the figure, any other positioning (e.g., aligning the peaks of the distributions) would also give a strong mismatch. The main reason for the mismatch is the tendency for the ICA to produce filters at a scale as close as possible to the sampling grid of the images (see the Discussion).

This property of ICA also causes an excess of filters aligned with the sampling lattice, i.e., horizontally and vertically. Thus a histogram of peak orientations yields a broad distribution with large and sharp peaks at horizontal and vertical orientations. These peaks can not be attributed directly to horizontal and vertical structures in the images, because they also occur when each image patch of the set of 120000 used for the ICA is rotated to a random orientation. In that case no inherent orientation can be present in the set of image patches, whereas there are still peaks in the results at horizontal and vertical orientations, be it somewhat smaller than before. In order to study possibly inherent orientational biases in the image set, despite the tendency of the ICA to align with the sampling lattice, we used the following procedure. We performed a series of ICA runs on 1) image patches all rotated over a fixed angle (30° for the results shown here, similar results were obtained with 45° and 22.5° rotations), and 2) image patches each rotated over a random angle. Set 1 is expected to yield peaks at horizontal and vertical orientations (the artefact), plus, e.g., a peak shifted by 30° from the horizontal if the image set has inherently more horizontal structures. Set 2 is expected just to yield the artefacts at horizontal and vertical orientations. By taking the ratio of the histograms of orientations for set 1 and set 2, the excess at horizontal and vertical orientations in the ICA filters is expected to cancel, at least approximately. Any remaining anisotropy at other orientations must be caused by true anisotropy in the images.

The result of this operation shows peaks at 30° and -60°, i.e., not aligned with the square sampling lattice. Rotating all orientations by -30°, back to the original upright position of the images, produces Fig. 4(f). Dots and error bars show the mean and standard deviation of 40 different image sets. As can be seen, the IC filters are relatively more often oriented horizontally (0°) and vertically (90°) than in other directions, but the effect is so small compared to the variability between subensembles of images that one could easily fit a horizontal line to the data. The width of the horizontal and vertical peaks is remarkably small; this is apparently caused during photography, by habitually aligning the camera frame with dominant horizontal and vertical structures in the image. Thus it depends on how precisely the human eye aligns with such structures during natural vision whether these peaks have any functional significance for human vision. The small peaks at about $\pm 45^\circ$ may also arise from a photographer's bias: one possible cause is that straight lines in the landscape (such as a road or path, or the fringe of a wood) are often framed diagonally, because this appears to produce pleasing, balanced pictures.

DISCUSSION

The independent component model is at present probably one of the most sophisticated ecologically inspired models for understanding the image representation in the array of simple cells in the human primary cortex. Not only does it produce quantitative predictions of receptive fields which compare reasonably well with those measured, it also gives a functional interpretation of the visual processing performed in terms of information theory and natural image statistics. In this article we have put this model to the test by applying it on a large set of images, and comparing the distribution of predicted properties to those measured in simple cells. We showed that most of these distributions are similarly shaped. This applies to spatial frequency bandwidth, orientation tuning bandwidth, length of the receptive fields, and aspect ratio. This result implies two conclusions. Firstly, it strengthens the hypothesis that cortical simple cells strive to produce a representation of natural images with independent variables, each having a highly kurtotic amplitude distribution (i.e., with long tails, leading to sparse coding). Secondly, it suggests that the apparent randomness of simple cell properties may not be the sign of a sloppy design, nor of random variability in development, but may in fact be a deliberate attempt to match the requirements of processing natural images.

A notable exception to the correspondence between IC filters and simple cells is the distribution of the peak of the spatial frequency response. Whereas simple cells have receptive fields acting on different spatial scales (i.e., they show spatial scaling), the IC filters show much less variability [Fig. 4(e)]. This discrepancy may be resolved in several ways. One possibility is that spatial scaling should be imposed as an extra constraint (see e.g. Koenderink 1984, Li & Atick 1994), as it can be argued that spatial scaling is a useful property for higher visual processing (e.g., object recognition relatively independent of viewing distance). Another possibility is that focussing on high spatial frequencies is a property of the particular ICA algorithm used (or even any algorithm maximising kurtosis, see Baddeley 1996). However, ICA performed by Bell and Sejnowski (1997a,b), with a different algorithm, yielded similar results as reported here. The basis vectors calculated by Olshausen & Field (1996, 1997) show a somewhat higher occurrence of low spatial frequency ones, but still considerably less than measured in simple cells [cf. Fig. 4(e)]. A further possibility is that extensions of the model to nonlinear ICA, possibly with overcomplete bases (i.e., more basis vectors than degrees of freedom), may resolve this discrepancy purely in the spatial domain. More likely, however, it will be necessary to include the time domain in the analysis. A preliminary analysis of the results of ICA performed on video sequences shows that spatiotemporal IC filters with peaks at low spatial frequencies are abundant, and are associated with higher speeds in the visual scenes.

Linear ICA yields the same number of basis vectors and filters as the number of degrees of freedom of the input. This is far less than the number of cells in the cortex per independent viewing direction as coded by the optic nerve. Part of this divergence may arise in ICA as well when including the time domain, leading to cells with different temporal frequency and velocity tuning. That there is also a need for divergence purely in the spatial domain is suggested by Fig. 2(b)-(e): different ensembles of images need different sets of IC filters. Current work on overcomplete bases (Olshausen & Field 1997, Lewicki & Olshausen 1998) may lead to similar variability when performed on a sufficiently large and diverse image database.

REFERENCES

- Atick, J.J. 1992 Could information theory provide an ecological theory of sensory processing? *Network* 3, 213-251.
- Baddeley, R.J. 1996 Searching for filters with “interesting” output distributions: an uninteresting direction to explore? *Network* 7, 409-421.
- Barlow, H.B. 1972 Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1, 371-394.
- Barlow, H.B. 1994 What is the computational goal of the neocortex? In *Large scale neuronal theories of the brain* (ed. Koch, C.), pp. 1-22. Cambridge: MIT Press.
- Bell, A.J. & Sejnowski, T.J. 1997a Edges are the ‘independent components’ of natural scenes. In *Advances in neural information processing systems* 9 (eds. Mozer, M.C., Jordan, M.J. & Petsche, T.), pp.831-837. Cambridge: MIT Press.

- Bell, A.J. & Sejnowski, T.J. 1997b The 'independent components' of natural scenes are edge filters. *Vision Res.* 37, 3327-3338.
- Dan, Y., Atick, J.J. & Reid, R.C. 1996 Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J.Neurosci.* 16, 3351-3362.
- Daugman, J.G. 1980 Two-dimensional spectral analysis of cortical receptive field profiles. *Vision Res.* 20, 847-856.
- DeAngelis, G.C., Ohzawa, I. & Freeman, R.D. 1993 Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *J.Neurophysiol.* 69, 1091-1117.
- DeValois, R.L., Yund, E.W & Hepler, N. 1982a The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res.* 22, 531-544.
- DeValois, R.L., Albrecht, D.G. & Thorell, L.G. 1982b Spatial frequency selectivity of cells in macaque visual cortex. *Vision Res.* 22, 545-559.
- Dong, D.W. & Atick, J.J. 1995 Temporal decorrelation: a theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network* 6, 159-178.
- Field, D.J. & Tolhurst, D.J. 1986 The structure and symmetry of simple-cell receptive-field profiles in the cat's visual cortex. *Proc.R.Soc.Lond. B* 228, 379-400.
- Field, D.J. 1987 Relations between the statistics of natural images and the response properties of cortical cells. *J.Opt.Soc.Am. A* 4, 2379-2394.
- Field, D.J. 1993 Scale-invariance and self-similar 'wavelet' transforms: an analysis of natural scenes and mammalian visual systems. In *Wavelets, fractals, and fourier transforms* (eds. Farge, M., Hunt, J.C.R. & Vassilicos, J.C.), pp. 151-193. Oxford: Clarendon Press.
- Field, D.J. 1994 What is the goal of sensory coding? *Neural Computation* 6, 559-601.
- Harpur, G.F. 1997 Low entropy coding with unsupervised neural networks. Thesis, University of Cambridge.
- Hateren, J.H. van 1992a Theoretical predictions of spatiotemporal receptive fields of fly LMCs, and experimental validation. *J.Comp.Physiol. A* 171, 157-170.
- Hateren, J.H. van 1992b Real and optimal neural images in early vision. *Nature* 360, 68-70.
- Hateren, J.H. van 1997 Processing of natural time series of intensities by the visual system of the blowfly. *Vision Res.* 37, 3407-3416.
- Heeger, D.J. 1992 Normalization of cell responses in cat striate cortex. *Visual Neuroscience* 9, 181-197.
- Hubel, D.H. & Wiesel, T.N. 1968 Receptive fields and functional architecture of monkey striate cortex. *J.Physiol.(Lond.)* 195, 215-243.
- Hurri, J. 1997 Independent component analysis of image data. Master thesis, Helsinki University of Technology, Laboratory of Computer and Information Science.
- Hurri, J., Hyvärinen, A., Karhunen, J. & Oja, E. 1996 Image feature extraction using independent component analysis. In *Proc. NORSIG'96*, Espoo, Finland
- Hyvärinen A. 1997 A family of fixed-point algorithms for independent component analysis. In *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'97)*, pp. 3917-3920. Munich, Germany.
- Hyvärinen A. & Oja, E. 1997a One-unit learning rules for independent component analysis. In *Advances in neural information processing systems 9* (eds. Mozer, M.C., Jordan, M.J. & Petsche, T.), pp. 480-486. Cambridge: MIT Press.

- Hyvärinen, A. & Oja, E. 1997b A fast fixed-point algorithm for independent component analysis. *Neural Computation* 9, 1483-1492.
- Koenderink, J.J. 1984 The structure of images. *Biol.Cybern.* 50, 363-370.
- Laughlin, S.B. 1981 A simple coding procedure enhances a neuron's information capacity. *Z.Naturforsch.* 36c, 910-912.
- Lewicki, M.S. & Olshausen, B.A. 1998 Inferring sparse, overcomplete image codes using an efficient coding framework. In *Advances in neural information processing systems 10*, in press.
- Li, Z. & Atick, J.J. 1994 Towards a theory of striate cortex. *Neural Computation* 6, 127-146.
- Linsker, R. 1993 Deriving receptive fields using an optimal encoding criterion. In *Advances in neural information processing systems 5* (eds. Hanson S.J., Cowan, J. & Giles, C.L.), pp. 953-960. San Mateo: Morgan Kaufmann.
- Marcelja, S. 1980 Mathematical description of the responses of simple cortical cells. *J.Opt.Soc.Am.* 70, 1297-1300.
- Olshausen, B.A. & Field, D.J. 1996 Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607-609.
- Olshausen, B.A. & Field, D.J. 1997 Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Res.* 37, 3311-3325.
- Parker, A.J. & Hawken, M.J. 1988 Two-dimensional spatial structure of receptive fields in monkey striate cortex. *J.Opt.Soc.Am. A* 5, 598-605.
- Ruderman, D.L. 1994 The statistics of natural images. *Network* 5, 517-548.
- Ruderman, D.L. & Bialek, W. 1994 Statistics of natural images: scaling in the woods. *Phys.Rev.Lett.* 73, 814-817.
- Sclar, G., Lennie, P. & DePriest, D.D. 1989 Contrast adaptation in striate cortex of macaque. *Vision Res.* 29, 747-755.
- Srinivasan, M.V., Laughlin, S.B. & Dubs, A. 1982 Predictive coding: a fresh view of inhibition in the retina. *Proc.R.Soc.Lond. B* 216, 427-459.
- Volgushev, M., Vidyasagar, T.R. & Pei, X. 1996 A linear model fails to predict orientation selectivity of cells in the cat visual cortex. *J.Physiol.(Lond.)* 496, 597-606.
- Zipser, K., Lamme, V.A.F. & Schiller, P.H. 1996 Contextual modulation in primary visual cortex. *J.Neurosci.* 26, 7376-7389.

Acknowledgements – We thank Aapo Hyvärinen, Bruno Olshausen, Dan Ruderman, and Herman Snippe for comments. This research was supported by the Netherlands Organization for Scientific Research (NWO).