

ICA FEATURES OF IMAGE DATA IN ONE, TWO AND THREE DIMENSIONS

Mika Inki

Neural Networks Research Centre
Helsinki University of Technology
P.O. Box 5400, FIN-02015 HUT, Finland

ABSTRACT

In this paper we study image features obtained by independent component analysis (ICA) in one, two and three dimensions. As three-dimensional data we use an image of a human head obtained with magnetic resonance imaging (MRI). We also study the features when cutting one or two dimensions from the data, and make comparisons to ordinary optical natural image data. Additionally, we look at feature distributions and scale invariance in both data types.

1. INTRODUCTION

Independent component analysis can be considered to be a fundamental generative model for low-level features of many types of natural data, e.g. natural image data. In ICA the observed data is expressed as a linear transformation of latent variables that are nongaussian and mutually independent. We may express the model as

$$\mathbf{x} = \mathbf{A}\mathbf{s} = \sum_i \mathbf{a}_i s_i, \quad (1)$$

where $\mathbf{x} = (x_1, x_2, \dots, x_m)$ is the vector of observed random variables, $\mathbf{s} = (s_1, s_2, \dots, s_n)$ is the vector of the latent variables called the independent components or source signals, and \mathbf{A} is an unknown constant matrix, called the mixing matrix. The columns of \mathbf{A} are often called features or basis vectors. Exact conditions for the identifiability of the model were given in [2], and several methods for estimation of the classic ICA model have been proposed in the literature, see [6] for a review. We consider here only the classic case $m = n$.

Of course, natural image data has not been generated according to the ICA model. Hardly any two features are really independent in images. In order to create natural images, you would have to have a model as complex as nature itself (or at least the subsystem depicted in the images along with the photographer and camera). However, the features found by ICA have many nice properties. First, ICA finds the most nongaussian decomposition of the data, and as the com-

ponents of this decomposition have a supergaussian or sparse distributions, they are in a way optimal for image compression. Second, as ICA features are tuned to the data, their statistics tell much more than mere pixel statistics. And third, ICA features of image data can be considered to be edge detectors [1], thus offering a very natural interpretation. These properties have also connections to biological vision [8].

ICA features of image data have often been described as Gabor filters [6], a claim which we will be looking at more closely here. Gabor functions are originally 1D entities that were generalized into 2D form by Daugman [3]. Intuitively, they can be generalized to higher dimensions in the same way:

$$g(\mathbf{u}) \propto \exp\left[-\sum_{i=1}^d \frac{u_i^2}{2b_i^2} + i(2\pi f u_1 + \theta)\right]. \quad (2)$$

Here $\mathbf{u} = (u_1, \dots, u_d)$ is the d -dimensional position vector, $\mathbf{b} = (b_1, \dots, b_d)$ gives the standard deviations of the gaussian envelope along the coordinate axes, f is the frequency of the function and θ is a phase shift. Any d -dimensional Gabor can now be obtained by coordinate rotation and translation. The real component of this is usually used as the Gabor function in image data, and we will be using it in this paper:

$$r(\mathbf{u}) \propto \exp\left(-\sum_{i=1}^d \frac{u_i^2}{2b_i^2}\right) \cos(2\pi f u_1 + \theta). \quad (3)$$

The ICA features obtained from image data are often such that the fitted Gabor function has a θ that is roughly $\pm\pi/2$, so the function has a zero crossing at the centre, and f is very small, so that only two lobes, one of them positive and the other negative, are visible. Note that it can be easily shown that the real component of a Gabor with $\theta = \pm\pi/2$ approaches the first derivative w.r.t. u_1 of a gaussian when f approaches zero, when multiplicative constants are ignored. Some ICA features in image data are best fit by a Gabor with a θ that is roughly 0 or π , but are in fact actually more like the second derivative of gaussian (with

only three lobes visible), which is not in the closure of Gabor functions.

In this paper we study ICA features of image data in one, two and three dimensions. Some studies have been made of video data [9], where the time dimension can be considered to be a third dimension, but here we study MRI data, which has three similar dimensions. We also qualitatively look at how well Gabor functions describe the features. Additionally, we look at the distributions of the features, and scale invariance.

2. IMAGE DATA

First one should note that most three-dimensional images of the natural world available are medical, and most medical images have restricted access. This means that obtaining 3D data may prove to be difficult, and there is not a great variety in the subjects represented in the images. Therefore some properties normally associated with natural images might not be found in 3D data. One of these properties is scale invariance. If most three-dimensional images are of human bodies, certain features in these images are always of almost the same proportional size. Also, with ordinary natural image data, you observe the statistical properties (e.g. scale invariance) better when you have more images, but with the three-dimensional data available, there is only a small advantage of using more images.

The Zubal Phantom site provides a $256 \times 256 \times 128$ volume element (voxel) MR image of a human head at the web address <http://noodle.med.yale.edu/zubal/>. This is an updated, dedicated MRI head phantom made available by techniques described in [10]. The spacing between the voxels is 1.1 millimeters in x and y directions, and 1.4 millimeters in the z direction (a slice with a fixed value of z -coordinate is a transaxial slice). The image has been cut so that it contains the portions of the head above the jaw. Three transaxial slices of the image are shown in figure 1.

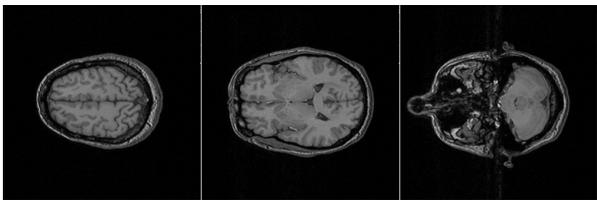


Figure 1: Three transaxial slices of the MR image.

The MR image has some nongaussian pixel noise, which can affect the results. By $2 \times 2 \times 2$ block averaging (and downsampling) we reduce the variance of this

noise and also make it more gaussian by the central limit theorem. We take windows of different sizes from this head. Most of the image is ‘empty’, i.e. 74% of the voxels are outside the head, but as a segmented version of this image is also available, we can easily take only those windows which are inside the head. The skull also produces very noticeable edges (leading and falling) in the image, so we also take only the windows that are completely inside the skull, unless otherwise mentioned.

We compare the features found in MRI data to features found in ordinary optical natural image data. We use 13 grayscale images of landscapes, animals and plants. These images can be found at the web address <http://www.cis.hut.fi/projects/ica/data/images/>. We also use randomly positioned windows of varying sizes to sample the data.

One should note the differences between ordinary natural image data and MRI data. First, the intensity values in MR images relate mostly to hydrogen density, not optical properties. Second, there is no perspective and thus no occlusion. Third, unsegmented MR images can be called natural in the sense that they depict directly (without further processing) objects occurring in the natural world, but the optical natural images we use here can be called natural even in the sense that they are somewhat close to what the human (animal) visual system has adapted to process during evolution. Apparently, no animal sees three-dimensionally and therefore no three-dimensional image could be called natural in the latter sense. In this paper we will call the our optical images ordinary natural images.

Unless otherwise mentioned, we first remove the mean from the windows. We do not further normalize the patches, e.g. normalize the norm. We use FastICA [5] to perform ICA on the data, using the hyperbolic tangent (tanh) nonlinearity, relating to a supergaussian source distribution, $p(s) \propto 1/\cosh(s)$. Note that at no point do we remove PCA components with nonzero eigenvalues from the data. Removing the smallest PCA components, which correspond to the highest frequencies, may increase the ‘waviness’ of the features, making Gabor functions better fits. We show the feature basis and also study feature distributions. Additionally, we look at the two-point correlations in the data.

3. RESULTS

3.1. Feature basis

First in figure 2 we illustrate the results when taking $8 \times 8 \times 8$ windows of the MRI data. Eight consecutive slices represent one feature. Each slice has a fixed value

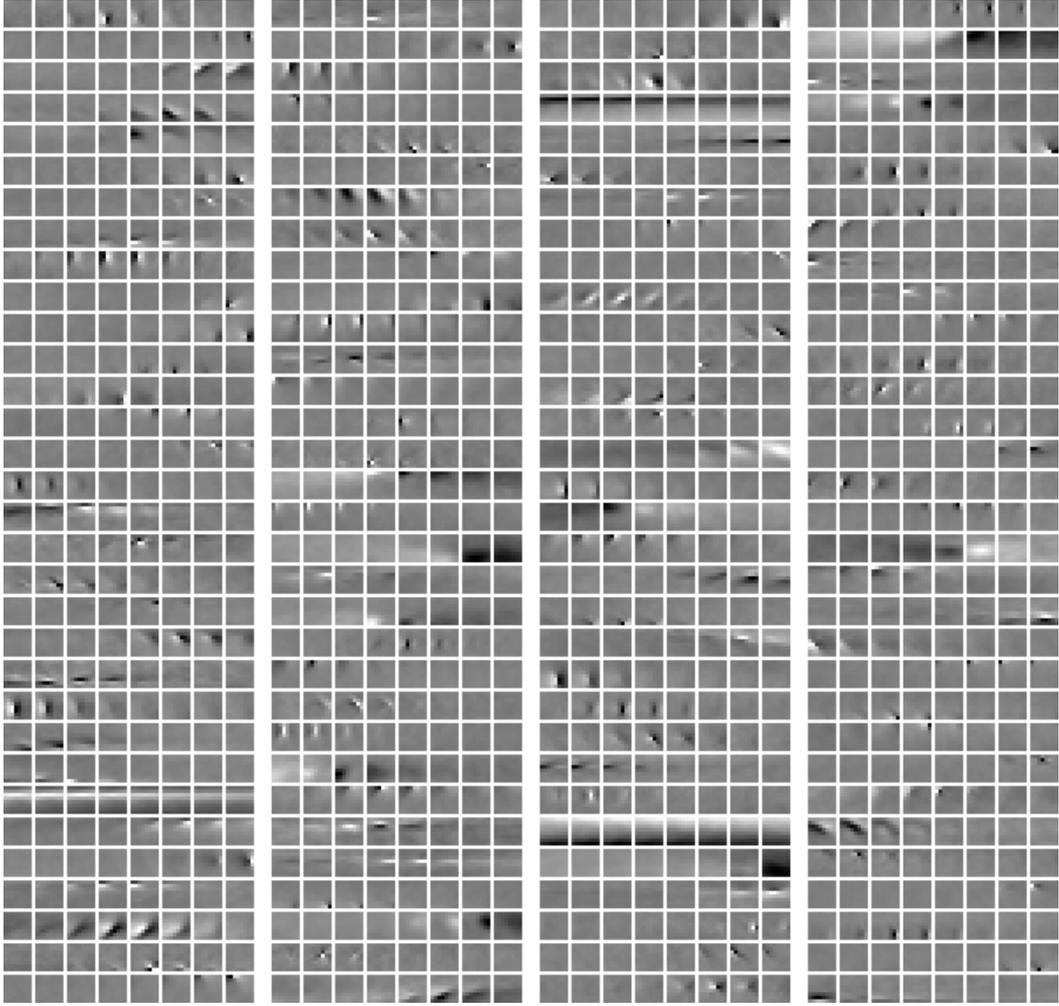


Figure 2: A randomly selected quarter of the $8 \times 8 \times 8$ features found in the MRI head. Eight consecutive slices represent one feature. Intensity values have been normalized for each feature.

of the z -coordinate, increasing from left to right. As the number of slices would be $511 * 8$, i.e. more than four thousand, we only show a randomly selected quarter of these features. A good portion of the features appear to have aligned with the axes (i.e. the direction u_1 in the best fitting Gabor in equation (3) is aligned with one of the axes) and some appear to be approximately diagonal or have other inclinations.

However, Gabor functions do not appear to give very good explanations of the features. In several features the positive and negative sides of the function appear to be quite ‘flat’ (which would suggest a Gabor with a large b_1 and a low frequency f), but the transition between them is rapid (which would suggest a Gabor with a high frequency or a small envelope). The MRI head has many areas with rather ‘flat’ intensity

values and rapid transitions between these areas. This partly explains why one gets such features. When we sampled the skull in the windows, the features were even more step-function like.

One should note that independent components (ICs) are not affected by a reduction in dimension, i.e. the new features should be simply the projections into the lower dimensional space. Therefore $8 \times 8 \times 1$ ICs should be similar to the $8 \times 8 \times 8$ feature slices. On the center of the top row in figure 3 we show these features. Indeed, the features are quite similar, although one difference is apparent. There are no $8 \times 8 \times 1$ features with a static value, as the mean value has been removed from the windows. Also there appear to be no low-frequency features.

On the left side of the top row in figure 3 we have the

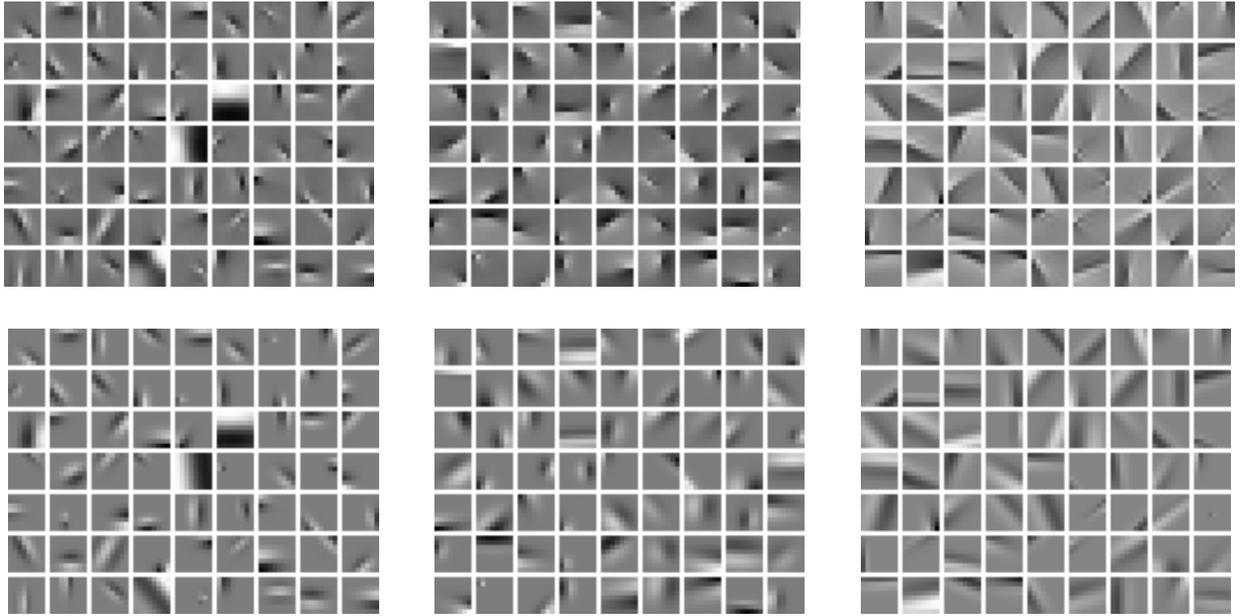


Figure 3: On the left side of the top row we have the features found in 8×8 ordinary natural image data, the center image on the top row represents the $8 \times 8 \times 1$ features found in MRI data when only sampling inside the skull, and on the right side of the top row are the features found when sampling the entire MR image. On the bottom row we have the corresponding fitted Gabor functions.

features found in ordinary natural image data, when using 8×8 windows. When compared to the MRI features, the features in ordinary natural image data are more compact. The MRI features have as fast or even faster transitions between the positive and negative sides, but the nonzero area of the feature extends further. Therefore Gabor functions make better fits for the features in ordinary natural image data. On the right side of figure 3, we have the features found in MRI data when sampling the skull also. These are clearly more step-function like than with either ordinary natural image data or MRI data without the skull. Therefore it can be said that when the data has clear boundaries between areas with flat intensity values, the ICA basis functions become more step-function like.

On the bottom row in figure 3 we have fitted Gabor functions to all the bases on the top row. The Gabor functions are able represent the ordinary natural image features well, but are unable to satisfactorily represent the MRI features. The fitted Gabors are unable to capture the sharp edge (and the overshoot) between the positive (light) and negative (dark) area. Gabor functions are even less able to represent ICA features of MRI data with the skull.

This property is even more visible, when looking at one-dimensional features. In figure 4 we present the

features found on a single line of length 24 in natural image data, $24 \times 1 \times 1$ features of MRI data when only sampling inside the skull, and MRI features when sampling the whole head. Again ordinary natural image data has features that are the most Gabor-like, and the whole MRI head has features the most step-function like.

3.2. Feature distributions

In figure 5 we show for $8 \times 8 \times 8$ MRI data (sampled inside the skull) the average IC distribution, i.e. a distribution obtained by concatenating all the ICs (all of which have a variance of one), the average distribution in a random direction in the whitened data space, the average distribution in PCA eigenvector directions and the average symmetrized voxel distribution. The latter three are much less supergaussian than the average distribution in ICA directions. The ICA components all had rather similar distributions. Note, however, that by the central limit theorem, the distribution in a random direction would be very close to gaussian if the sources found by ICA were really independent. The supergaussianity of the ICA components (and the data in general) justifies using tanh as the nonlinearity in FastICA. The average distribution in a random direction might be qualitatively closer to $p(s) \propto 1/\cosh(s)$

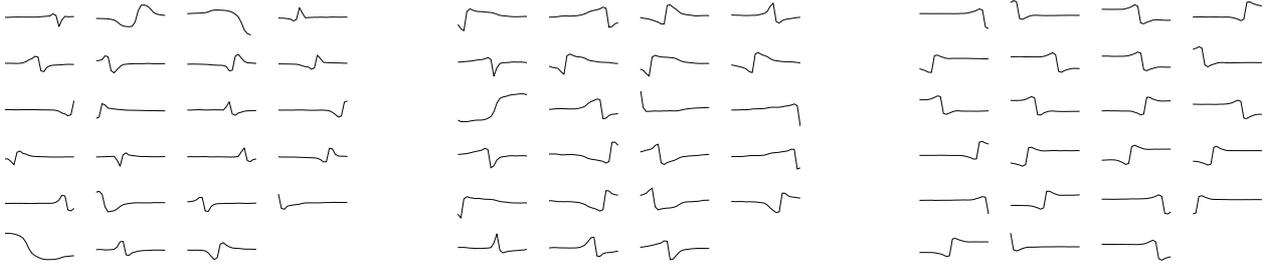


Figure 4: On the left side of we have the features found in 24×1 natural image data, the center image represents the $24 \times 1 \times 1$ features found in MRI data when only sampling inside the skull, and on the right side are the features found when using all sample windows that are not completely outside the head.

than the average IC distribution, but this is not a problem with maximum likelihood estimation.

On the right side of figure 5 we have the same distributions for ordinary natural image data. Note that these distributions are also supergaussian or sparse, and again the average IC distribution is clearly the most supergaussian of all. We do not compare the distributions of MRI and ordinary natural image data any further. MRI data distribution is affected by what is sampled in the windows, and comparing the specifics of two distributions of different things is somewhat difficult.

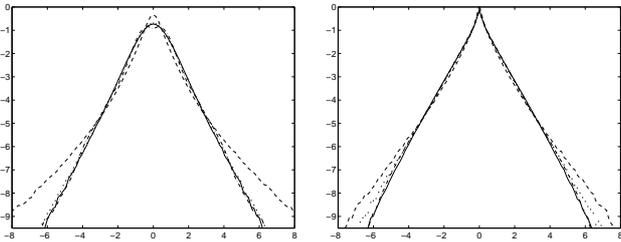


Figure 5: On the left side the logarithms of the average distributions in a random direction (solid line), the average distribution in PCA directions (dotted line), the average symmetrized voxel distribution (dashdotted) and the average distribution in ICA directions (dashed line) with $8 \times 8 \times 8$ MR image data. On the right side the same for 8×8 ordinary natural image data.

In figure 6, we have the average distributions of $8 \times 8 \times 8$, $4 \times 4 \times 4$, $8 \times 8 \times 1$ and $8 \times 1 \times 1$ features. The $4 \times 4 \times 4$, $8 \times 8 \times 1$ and $8 \times 1 \times 1$ windows were sampled from the $8 \times 8 \times 8$ windows. Note that as we only sample inside the skull, these windows might sample some parts of the head differently if this was not done. Also, we didn't separately reduce the mean from these new windows. The distributions for $8 \times 8 \times 8$, $4 \times 4 \times 4$

and $8 \times 8 \times 1$ windows are practically identical. The $8 \times 8 \times 8$ distribution is by a small amount the most supergaussian. The $8 \times 1 \times 1$ features are clearly less supergaussian than the other distributions.

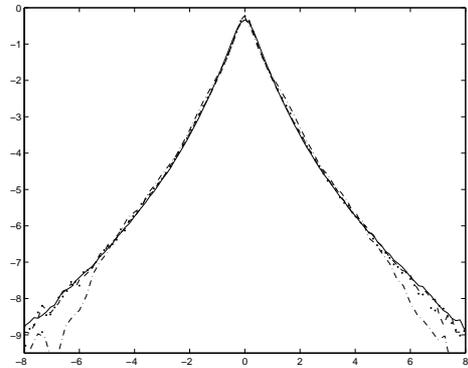


Figure 6: The distributions of $8 \times 8 \times 8$ MRI features (solid line), $4 \times 4 \times 4$ features (dashed line), $8 \times 8 \times 1$ features (dotted line) and $8 \times 1 \times 1$ features (dashdotted).

3.3. Correlation structure

One way of quantifying the scale invariance of an image is to look at the two-point correlation function, i.e.

$$C(r)_e = E\left\{I(\mathbf{u})I\left(\mathbf{u} + \frac{r\mathbf{e}}{\|\mathbf{e}\|}\right)\right\}, \quad (4)$$

where r is a shift and $E\{\cdot\}$ is the expectation taken here over all images I (with zero mean and unit variance) and positions \mathbf{u} . This form assumes translational invariance. If one assumes rotational symmetry, as we do here, the correlation is the same for all $\mathbf{e} \neq \mathbf{0}$, and if one assumes scale invariance, this has a logarithmic form $C(r) = c_1 \log(r) + c_2$. Note that the logarithmic form suggests divergence at zero, but pixel values essentially

correspond to the average intensity value of the image that falls inside the pixel boundaries. Therefore pixel variance is bounded even with nontrivial correlations over nonzero distances. Also, as distance approaches infinity, the covariance diverges. But as images are finite sized, and one may assume that their means have somehow been normalized, this is not a problem. A more thorough analysis of scale invariance properties is beyond the scope of this paper, see for example [4, 7] for more information.

We have calculated the voxel correlation for MRI data (while compensating for the bigger spacing in the z -direction) and pixel correlation for ordinary natural image data. We randomly selected pairs of pixels, calculated the average correlations inside one pixel wide bins (which correspond to distances from one to sixteen pixels), and average distance for the pixel pairs within the bins. With MRI data, we have taken only the voxels inside the skull into consideration. The results are in figure 7. Note that on the x -axis we have the logarithm of distance. The dashed lines in the figures are straight lines fitted to the correlation functions. These lines fit very well. So, although MRI data is not scale invariant in the larger sense, the voxel correlation inside the skull exhibits scale invariance as does pixel correlation in ordinary natural image data.

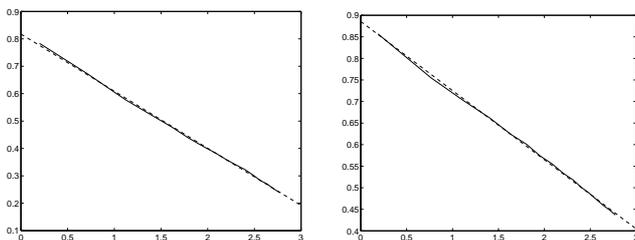


Figure 7: The voxel correlation in MRI data and the pixel correlation in ordinary natural image data as functions of the logarithm of distance on the left and right sides respectively. Note that pixel/voxel variance has been normalized to unity. Dashed line: straight fitted to the data.

4. CONCLUSIONS

In this paper we looked at the ICA features of three-dimensional image data, as which we used an MR image of a human head. The features of two- and one-dimensional restrictions of this image were found to be rather similar to those of ordinary natural image data, which might be surprising considering the physical differences in the imaging process and the differences between the subjects depicted in the images. However,

the ICA features of these restrictions were not as well described as Gabor functions, which may be due to the fact that the MR image contains many flat intensity value areas with rapid transitions, thus producing features more similar to step functions. The three-dimensional features could be similarly described as ‘something like Gabor functions’, with more step-function like properties. They can, however, still be described as edge detectors. We also found that the ICA features of MRI data have supergaussian distributions and that MRI data displays scale invariance, at least on the voxel correlation level.

5. ACKNOWLEDGEMENTS

The author was funded by the Helsinki Graduate School of Computer Science and Engineering, Emil Aaltonen Foundation and The Finnish Foundation of Technology. Dr. Aapo Hyvärinen assisted with the manuscript and Mr. Mika Pollari assisted with MRI data.

6. REFERENCES

- [1] A.J. Bell and T.J. Sejnowski. The ‘independent components’ of natural scenes are edge filters. *Vision Research*, 37:3327–3338, 1997.
- [2] P. Comon. Independent component analysis—a new concept? *Signal Processing*, 36:287–314, 1994.
- [3] J. G. Daugman. Uncertainty relations for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Journal of the Optical Society of America A*, 2:1160–1169, 1985.
- [4] J. Huang and D. Mumford. Statistics of natural images and models. In *In Proc. Conf. Computer Vision and Pattern Recognition*, pages 541–547, 1999.
- [5] A. Hyvärinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Trans. on Neural Networks*, 10(3):626–634, 1999.
- [6] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. Wiley Interscience, 2001.
- [7] D. Mumford and B. Gidas. Stochastic models for generic images. *Quarterly of Applied Mathematics*, LIX(1):85–111, 2001.
- [8] B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, 37:3311–3325, 1997.
- [9] J. H. van Hateren and D. L. Ruderman. Independent component analysis of natural image sequences yields spatiotemporal filters similar to simple cells in primary visual cortex. *Proc. Royal Society, Ser. B*, 265:2315–2320, 1998.
- [10] I.G. Zubal, C.R. Harrell, E.O. Smith, Z. Rattner, G.R. Gindi, and P.B. Hoffer. Computerized three-dimensional segmented human anatomy. *Medical Physics*, 21(2):299–302, 1994.