From: Heinrich H. Bülthoff and Seong-Whan Lee and Tomaso A. Poggio and Christian Wallraven (eds.), Biologically Motivated Computer Vision, pp. 117–126, Springer, 2002.

Image Reconstruction from Gabor Magnitudes ⋆

Ingo J. Wundrich[1], Christoph von der Malsburg[1,2], and Rolf P. Würtz[1]

[1] Institut für Neuroinformatik, Ruhr-Universität Bochum
D-44780 Bochum, Germany
[2] Laboratory for Computational and Biological Vision
University of Southern California, Los Angeles, USA

Abstract. We present an analysis of the representation of images as the *magnitudes* of their transform with complex-valued Gabor wavelets. Such a representation is very useful for image understanding purposes and serves as a model for an early stage of human visual processing. We show that if the sampling of the wavelet transform is appropriate then the reconstruction from the magnitudes is unique up to the sign for almost all images. We also present an iterative reconstruction algorithm derived from the ideas of the proof, which yields very good reconstruction up to the sign and minor numerical errors in the very low frequencies.

Keywords: Gabor wavelets, reconstruction, feature extraction, phase retrieval, Gabor magnitudes, visual cortex.

1 Introduction

In the visual system of humans or monkeys an early stage is concerned with processing image information by convolution with some point spread function and ensuing non-linear operations that enhance edges and lead to some contrast normalization. A recognized problem [13] with edge-enhancing convolutions is their sensitivity to image shift, whereas one of the most important features of human vision is the positional invariance with which objects can be recognized. The two-dimensional power spectrum can represent images in a shift-invariant way, but it has the great drawback of being non-local, each component being influenced by all image pixels. On the other extreme, pixels as image features are maximally localized but achieve notoriously little in terms of analyzing image contents.

We focus here on Gabor functions as an adjustable compromise between pixel representation and Fourier components. They seem to be implemented in the first stages of processing in the visual cortex of higher vertebrates, as the receptive fields of the so-called *simple cells* can be described to some accuracy as Gabor functions [1, 7]. There is also evidence that the *magnitudes* of the Gabor filter responses (short *Gabor magnitudes*) are calculated by another set of cells called *complex cells* [12].

The simplest model for these findings is that simple cell responses are calculated from the image intensities by a feedforward neural net, and that complex cells build on their information by another feedforward net. The complex cells, in turn, can be combined to more complicated feature detectors such as corner detectors [21]. They have also proven useful for higher image understanding tasks such as texture classification [3], recognition of faces [9, 20, 2], vehicles [18], and hand gestures [14]. The deeper reason for this is that the magnitude operation introduces some local shift invariance in the sense that under small shifts in the image the Gabor magnitudes are more robust than the full complex valued responses, because they are much smoother. This robustness is crucial for recognition systems, which have

⋆ This work has been supported by grants from BMBF, ONR and ARO.

to cope with small local deformations. As a practical consequence, similarity landscapes between local features are smoother if magnitudes are used, which makes matching faster and less prone to local maxima [9, 20, 17].

If the Gabor functions are arranged into a wavelet transform and the sampling is dense enough then the original image can be recovered from the transform values with arbitrary quality (except for the DC-value). Given the useful properties of the magnitudes of the Gabor transform an important theoretical question is how much image information can be recovered from that.

There are many results on the reconstruction of images from localized phase [15], and they show convincingly that localized phase is more useful than global Fourier phase. There seems to be general agreement that reconstruction from phase is simpler than reconstruction from local magnitude.

In this paper we are concerned with the latter. It must be noted that the study of reconstruction is not the most important issue for an object recognition system, where most of the visual information must be *discarded*. Nevertheless, it is of importance to study as well as possible the course of information through the system in order to understand its mechanism. If images of different objects could be reconstructed from the data format used by a recognition algorithm, this algorithm would be rather limited in its value. Additionally, our results may be of importance to other applications of Gabor wavelets.

In the following we briefly outline the theory of wavelets and frames necessary to understand reconstruction from the linear transform, introduce the Gabor wavelet transform, and review some literature on phase retrieval from power spectra of images. We then outline the proof of a theorem stating that, given the right transform parameters and appropriate band-limitation, no image information is lost by applying Gabor magnitudes to an image except the DC-value of the image and a global sign. The proof uses techniques from [5] and applies to all images except a subset of measure zero.

Finally, we explore the quality of reconstruction by numerical experiments based on the ideas of the proof. The results are not perfect because much of the proof depends on transform values being exactly zero, which does not translate very well into numerical computation. However, for all images we tested, we were able to retrieve very good approximations of either the image itself or its negative.

2 Gabor wavelets and frames

For the analysis of signal properties at various scales the *wavelet transform* has been introduced in [4]. The signal is projected onto a family of wavelet functions $\psi_{x_0,a}(x)$ derived from a so-called *mother wavelet* $\psi(x)$ by applying translations by x_0 and dilations by a factor $a > 0$. For image processing, we follow the proposal by Murenzi [10] and choose the 2-dimensional Euclidean group $IG(2)$ with dilations for the construction of a wavelet family. This leads to daughter wavelets which are translated, rotated and scaled versions of the mother.

The transition to wavelets defined on $\mathbf{R}^2$ leads to a wavelet family parameterization by the translation vector $\boldsymbol{x}_0 \in \mathbf{R}^2$, scale factor $a > 0$ and the orientation angle $\vartheta \in [0, 2\pi[$. This extension of the affine linear group to two spatial dimensions preserves the idea of scaling the transformation kernel ψ. Analogous to the 1D case the 2D wavelet transform consists of projection of the image data $I(\boldsymbol{x})$ onto the wavelet family ($\mathrm{Q}(\vartheta)$ stands for the 2D rotation matrix by the angle ϑ):

$$\mathcal{I}(\boldsymbol{x}_0, a, \vartheta) = \langle I, \psi_{\boldsymbol{x}_0,a,\vartheta} \rangle, \tag{1}$$

$$\psi_{\boldsymbol{x}_0,a,\vartheta}(\boldsymbol{x}) = a^{-1}\psi[a^{-1}\mathrm{Q}(\vartheta)(\boldsymbol{x} - \boldsymbol{x}_0)] \tag{2}$$

The mother wavelet (and, consequently, all wavelets) must satisfy the *admissibility condition* [8]:

$$0 < C = 4\pi^2 \int_{\mathbf{R}^2} d^2\omega \, \frac{|\hat{\psi}(\boldsymbol{\omega})|^2}{\|\boldsymbol{\omega}\|^2} < \infty. \tag{3}$$

Consequently, the wavelets must have zero DC-value ($\hat{\psi}(\mathbf{0}) = 0$) and decay sufficiently quickly for increasing $\|\boldsymbol{\omega}\|$. This condition, together with implementing the translations of wavelets as convolutions means that wavelets are bandpass functions.

We now narrow our focus to the Gabor function as mother wavelet. As Gabor functions are not DC-free, an additional term must be introduced to ensure that wavelet property. Following Murenzi [10], we let:

$$\psi(\boldsymbol{x}) = \frac{1}{\sigma\tau} \exp\left(-\frac{1}{2}\|\mathrm{S}_{\sigma,\tau}\boldsymbol{x}\|^2\right)\left(\exp\left(j\boldsymbol{x}^T\boldsymbol{e}_1\right) - \exp\left(-\frac{\sigma^2}{2}\right)\right) \tag{4}$$

In these equations the diagonal matrix $\mathrm{S}_{\sigma,\tau} = Diag(1/\sigma, 1/\tau)$ controls the shape of the elliptical Gaussian relative to the wavelength. Different ways of removing the DC-value can be used, but the one above is the most elegant for analytical treatment.

Different Gabor functions are, in general, not orthogonal and the wavelet families are usually not linearly independent. That means that (1) yields an overcomplete representation of the image signal $I(\boldsymbol{x})$. To handle linear transforms of this nature the *frame* concept, which can be seen as a generalization of the basis in a linear space. We follow the description in [8].

For a Hilbert space $\mathcal{H}$ and a measure space $(\mathcal{M}, \mu)$ linear transforms H from $\mathcal{H}$ into $L^2(\mathcal{M}, \mu)$ are defined by the projection onto a family of functions $\mathcal{H}_\mathcal{M} = \{h_\xi \in \mathcal{H} : \xi \in \mathcal{M}\}$ via $\mathrm{H}f(\xi) = \langle h_\xi, f\rangle$, such that H is measurable for every $f \in \mathcal{H}$. This family is called a *frame* if there exist positive finite constants A and B such that for every $f \in \mathcal{H}$

$$A\|f\|^2_\mathcal{H} \leq \|\mathrm{H}f\|^2_{L^2(\mathcal{M},\mu)} \leq B\|f\|^2_\mathcal{H}\,. \tag{5}$$

Such constants are called *frame bounds*. If $A = B$ the frame is called *tight*. The freedom in the choice of μ can be put to different uses, e.g., the frame elements can be normalized or one of the frame bounds can be fixed at 1. Furthermore, it allows a coherent formulation of discrete and continuous wavelet transforms. In our concrete case, the measure space is $\mathcal{M} = \mathbf{R}^2 \times \mathbf{R}^+ \times \mathcal{U}$ for the two spatial dimensions, scale and orientation, and the accompanying measure is given by

$$d\mu = d^2x_0 \; a^{-3} da \; d\vartheta\,. \tag{6}$$

In the continuous case the so constructed inverse 2D wavelet transform becomes:

$$I(\boldsymbol{x}) = \frac{1}{C} \int_{\mathbf{R}^+} \frac{da}{a^3} \int_0^{2\pi} d\vartheta \int_{\mathbf{R}^2} d^2x_0 \, \mathcal{I}(\boldsymbol{x}_0, a, \vartheta)\psi_{\boldsymbol{x}_0,a,\vartheta}(\boldsymbol{x})\,, \tag{7}$$

with the C from (3).

For practical purposes, it is not desirable to expand the image representation from a function on $\mathbf{R}^2$ to one on $\mathbf{R}^4$, so *sampling* of translations, scales and orientations $\boldsymbol{x}_0 = \boldsymbol{n}_0\boldsymbol{\Delta}$, $a = a_{\min}a_0^m$, $\vartheta = 2\pi l/L$ with $\boldsymbol{n}_0 \in \mathbf{Z}^2$, $m \in \{0, 1, \ldots, M-1\}$, and $l \in \{0, 1, \ldots, L-1\}$, becomes inevitable.

We now switch from continuous functions to discretely sampled images of $N_1 \times N_2$ pixels. The underlying finite lattice will be called $\mathcal{S}_\boldsymbol{N}$. Now, the discrete Gabor wavelet transform can be computed in either domain by the inner product $\mathcal{I}(\boldsymbol{n}_0, m, l) = \langle I, \psi_{\boldsymbol{n}_0,m,l}\rangle$.

3 From Fourier to Gabor magnitudes

In order to state theorems about the reconstructability of an image from its Gabor magnitudes $|\mathcal{I}(\boldsymbol{n}_0, m, l)|$ we choose a collection of theorems on Fourier magnitudes as a starting point.

In general the Fourier transform $\hat{I}(\boldsymbol{\omega})$ is a complex-valued function which can be described in terms of a magnitude and a phase. The fact that the inverse DFT applied to a modified transform with all magnitudes set to 1 and original phases preserves essential image properties [11] is frequently interpreted as saying that the Fourier magnitudes contain "less" image information than the phases. However, analytical results and existing phase retrieval algorithms provide hints that the situation is not as simple.

These theorems are based on the fact that the Fundamental Theorem of Algebra does not hold for polynomials in more than one variable. More precisely, the set of polynomials in more than one variable which can be factored in a nontrivial way are of measure zero in the vector space of all polynomials of the same degree [6]. A nontrivial factorization is very undesirable because the number of ambiguities caused by phase removal increases exponentially with the number of factors.

Hayes's theorem identifies the 2D z-Transform,

$$\check{I}(\boldsymbol{z}) = \frac{1}{2\pi} \sum_{\boldsymbol{n} \in \mathcal{S}_{\boldsymbol{N}}} I(\boldsymbol{n}) z_1^{-n_1} z_2^{-n_2}, \tag{8}$$

and the 2D discrete space Fourier transform (DSFT) on a compact support, with polynomials in two variables.

Theorem 1 (Hayes, [5]). *Let I_1, I_2 be 2D real sequences with support $\mathcal{S}_{\boldsymbol{N}} = \{0, \ldots, N_1 - 1\} \times \{0, \ldots, N_2 - 1\}$ and let Ω a set of $|\Omega|$ distinct points in $\mathcal{U}^2$ arranged on a lattice $\mathcal{L}(\Omega)$ with $|\Omega| \geq (2N_1 - 1)(2N_2 - 1)$. If $\check{I}_1(\boldsymbol{z})$ has at most one irreducible nonsymmetric factor and*

$$\left|\check{I}_1(\boldsymbol{\nu})\right| = \left|\check{I}_2(\boldsymbol{\nu})\right| \;\; \forall \boldsymbol{\nu} \in \mathcal{L}(\Omega) \tag{9}$$

then

$$\begin{aligned} I_1(\boldsymbol{n}) \in \big\{ & I_2(\boldsymbol{n}),\; I_2(\boldsymbol{N} - \boldsymbol{n} - \boldsymbol{1}), \\ & -I_2(\boldsymbol{n}),\; -I_2(\boldsymbol{N} - \boldsymbol{n} - \boldsymbol{1}) \big\}. \end{aligned} \tag{10}$$

Theorem 1 states that DSFT magnitudes-only reconstruction yields either the original, or a negated, a point reflected, or a negated and point reflected version of the input signal. Together with the main statement from [6] that the set of all reducible polynomials $\check{I}(\boldsymbol{z})$ is of measure zero, the technicality about the irreducible nonsymmetric factors can be omitted, and we generalize Theorem 1 to complex-valued sequences as follows:

Theorem 2. *Let I_1, I_2 be complex sequences defined on the compact support $\mathcal{S}_{\boldsymbol{N}}$ and let $\check{I}_1(\boldsymbol{\nu})$ and $\check{I}_2(\boldsymbol{\nu})$ be only trivially reducible (i.e. have only factors of the form $z_1^{p_1} z_2^{p_2}$), and*

$$\left|\check{I}_1(\boldsymbol{\nu})\right| = \left|\check{I}_2(\boldsymbol{\nu})\right| \;\; \forall \boldsymbol{\nu} \in \mathcal{L}(\Omega) \tag{11}$$

with $\mathcal{L}(\Omega), |\Omega|$ as in Theorem 1 then

$$I_1(\boldsymbol{n}) \in \{ \exp(j\eta)\, I_2(\boldsymbol{n}),\, \exp(j\eta)\, I_2^*(\boldsymbol{N} - \boldsymbol{n} - \boldsymbol{1}) \,|\, \eta \in [0, 2\pi[\} \;. \tag{12}$$

Transferring the modified Hayes theorem to the spatial magnitudes of the Gabor wavelet transform yields ambiguities which are reduced by inter- and intrasubband structures. More concretely, the Gabor magnitudes relate to the autocorrelation of the spectra of the subband images. However, due to the known localization of the Gabor responses in frequency space, the lost information can be recovered. This line of reasoning has allowed us to prove the following:

Theorem 3 (Gabor Magnitude Theorem). *Let $\mathcal{B}(N_1, N_2)$ be the space of all functions on the grid $\mathcal{S}_{\boldsymbol{N}}$ such that* $\mathrm{DFT}I(\boldsymbol{\rho}) = 0$ *for* $|\rho_1| \geq \frac{N_1}{4}, |\rho_2| \geq \frac{N_2}{4}$*, and let the wavelet family $\psi_{\boldsymbol{n}_0,m,l}$ constitute a frame in $\mathcal{B}(N_1, N_2)$. For all $I_1, I_2 \in \mathcal{B}(N_1, N_2)$ such that $\langle I_1, \psi_{\boldsymbol{n}_0,m,l}\rangle$ and $\langle I_2, \psi_{\boldsymbol{n}_0,m,l}\rangle$ are only trivially reducible polynomials and $|\langle I_1, \psi_{\boldsymbol{n}_0,m,l}\rangle| = |\langle I_2, \psi_{\boldsymbol{n}_0,m,l}\rangle|$ $\forall \boldsymbol{n}_0, m, l$ it follows that $I_1(\boldsymbol{n}) = \pm I_2(\boldsymbol{n})$.*

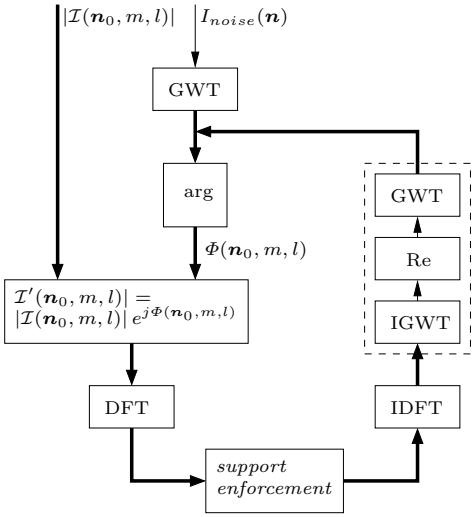

Fig. 1. Scheme for Gabor phase retrieval. Within one iteration loop each subband image is filtered according to its required signal energy concentration and boundary in frequency domain. In the next step the Gabor transform is computed which is nearest to the subspace of all Gabor-transformed real images. Last, the phases of the updated subband images are extracted and combined with the true magnitudes.

To complete the argument, Hayes' theorems can again be used to state that the Gabor transforms of almost all images are only trivially reducible. A detailed proof of the theorem can be found in [19]. Thus, we may conclude, that a low-pass filtered version of almost all images can be reconstructed from their Gabor transform magnitudes up to the sign.

The condition about the vanishing Fourier coefficients (band limitation) is not a restriction on the class of images to which the theorem applies, because each image of final resolution can be turned into a band-limited one by interpolation through zero-padding in the frequency domain. Put the other way around, from a Gabor wavelet transform of a certain spatial resolution, images of *half* that resolution can be reconstructed uniquely up to the sign.

4 Numerical Retrieval of Gabor Phases

In this section we construct a Gabor phase retrieval algorithm using one major idea from the proof of Theorem 3. In that proof, we have interchanged spatial and

Fig. 2. The first row shows some original images, the second their reconstructions from their magnitudes of their Gabor wavelet transform after 1300 iterations. The reconstruction of the Lena image actually yielded its negative. The original images from which the magnitudes are taken are 128×128 images interpolated to 256×256. For display the grey value range has been normalized to [0,255].

frequency domain for the application of Hayes's theorems, and the same can be done in the phase retrieval algorithm.

The given magnitudes for reconstruction are combined with the phases of an *arbitrary* Gabor-transformed image. Then, band limitation and subband localization are enforced by zeroing all frequencies above the boundary and outside the appropriate region of frequency concentration for a certain scale and orientation. That region is determined by applying a threshold of 0.1 to the Gabor kernel in frequency space. The result is transformed back into an image and transformed forward in order to project it onto the space of all Gabor wavelet transforms of real-valued images. Then the next cycle starts with the combination of the given magnitudes with the updated phases.The full course of the algorithm is shown in figure 1.

The main problem with the reconstruction from magnitudes is that the set of all transforms with given magnitudes but arbitrary phases is not convex, in contrast to the set of all transforms with given phases and variable magnitudes. Therefore, the iterative projection algorithm is not a POCS (projection onto convex sets) algorithm, and there is no straightforward convergence proof. This is in contrast to magnitude retrieval [15].

An alternative approach [16] uses a gradient descent algorithm to estimate an image minimizing an error functional. This minimization yields near to perfect results on bandpass filtered images.

5 Reconstruction experiments

We ran numerical reconstruction experiments on three different images with several hundred iterations. A real white noise "image" was chosen as initialization. The

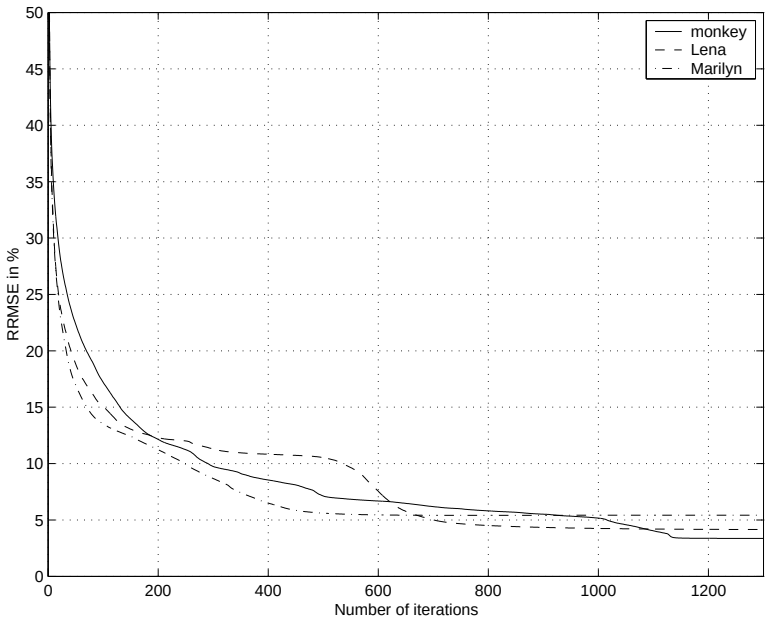

Fig. 3. The development of the error (RRMSE) for all three tested images

results of the reconstruction on some natural images are shown in figure 2. The transform parameters used to produce the images were $\sigma = 4$, $M = 8$, $L = 16$, $a_0 = \sqrt{2}$. With these parameters, reconstruction from the *linear* transform is perfect up to the DC-value. For display, all images are normalized to the gray value range [0,255], which dictates a DC-value.

In order to assess convergence we measured a relative RMSE defined as

$$\mathrm{RRMSE} = \sqrt{\frac{\sum_{m=0}^{M-1}\sum_{l=0}^{L-1}\sum_{\boldsymbol{n}_0\in\mathbf{S}_{\boldsymbol{N}}} a_0^{-2m}\left[|\mathcal{I}(\boldsymbol{n}_0,m,l)| - |\mathcal{I}^{\mathrm{rec}}(\boldsymbol{n}_0,m,l)|\right]^2}{\sum_{m=0}^{M-1}\sum_{l=0}^{L-1}\sum_{\boldsymbol{n}_0\in\mathbf{S}_{\boldsymbol{N}}} a_0^{-2m}|\mathcal{I}(\boldsymbol{n}_0,m,l)|^2}} \ . \quad (13)$$

As displayed in figure 3, the reconstruction does not converge to zero error. The remaining RRMSE corresponds to slight gray level deviations in uniform (low-frequency) image zones as can be seen comparing the reconstructions to the originals, (see figure 2). We interpret reconstruction errors as accumulated numerical errors from the SVD regularization of low frequencies in the IGWT, which is repeated in each iteration. However the reconstructed images retain local texture properties very well, which is crucial for image understanding based on representation by such features.

After a rapid improvement in some 10 iterations, which already yield a perfectly recognizable image, convergence becomes rather slow. There are local inversions of sign, which compete to dictate the global sign.

6 Discussion

We have shown that almost all images can be recovered from their Gabor magnitudes. As natural images, which are the only interesting ones for computer vision, constitute only a tiny subset of all functions with compact support, it is theoretically possible that many of them fall into the subset of images not represented uniquely

by their Gabor magnitudes, which we will call *ambiguous*. Although possible, this appears highly unlikely, because slight modifications of natural images still yield natural images. However, neither the set of natural images nor the precise form of the set of ambiguous images is known. The latter can not be uncovered with the simple dimensionality argument used in this paper and definitely requires further research. Furthermore, it is unclear how much the different reconstructions of ambiguous Gabor magnitudes will differ. If there should be two images with definitely different contents but nevertheless identical Gabor magnitudes, this would make the method problematic for image understanding. We have shown that this is very unlikely, but still have no absolute proof that it cannot happen.

For further evidence, we have implemented a numerical algorithm for Gabor phase retrieval, which is based on the ideas of the proof. In the cases we tested, we could always recover a good approximation of the image up to the sign and numerical errors in the low frequency contents. Our theorem suggests that twice the sampling rate is needed in each dimension for reconstruction from magnitudes only than for reconstruction from the full transform. As a simple rule of thumb, this looks very plausible in neuronal terms, if one considers a single complex number to be represented by *four* positive real numbers (because cell activities cannot be negative). Thus, four simple cells, which code for the linear wavelet coefficient, must be replaced by four complex cells at slightly different positions in order to convey the same information.

References

1. John G. Daugman. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Journal of the Optical Society of America A*, 2(7):1362–1373, 1985.
2. Benoît Duc, Stefan Fischer, and Josef Bigün. Face authentication with gabor information on deformable graphs. *IEEE Transactions on Image Processing*, 8(4):504 – 516, 1999.
3. I. Fogel and Dov Sagi. Gabor filters as texture discriminator. *Biological Cybernetics*, 61:103–113, 1989.
4. A. Grossmann and J. Morlet. Decomposition of Hardy functions into square integrable wavelets of constant shape. *SIAM Journal of Mathematical Analysis*, 15(4):723 – 736, July 1984.
5. Monson H. Hayes. The Reconstruction of a Multidimensional Sequence from the Phase or Magnitude of Its Fourier Transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 30(2):140 – 154, April 1982.
6. Monson H. Hayes and James H. McClellan. Reducible Polynomials in More Than One Variable. *Proceedings of the IEEE*, 70(2):197 – 198, February 1982.
7. J.P. Jones and L.A. Palmer. An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6):1233–1258, 1987.
8. Gerald Kaiser. *A Friendly Guide to Wavelets*. Birkhäuser, 1994.
9. Martin Lades, Jan C. Vorbrüggen, Joachim Buhmann, Jörg Lange, Christoph von der Malsburg, Rolf P. Würtz, and Wolfgang Konen. Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers*, 42(3):300–311, 1993.
10. R. Murenzi. Wavelet Transforms Associated to the n-Dimensional Euclidean Group with Dilations: Signal in More Than One Dimension. In J. M. Combes, A. Grossmann, and P. Tchamitchian, editors, *Wavelets – Time-Frequency Methods and Phase Space*, pages 239 – 246. Springer, 1989.
11. Alan V. Oppenheim and Jae S. Lim. The Importance of Phase in Signals. *Proceedings of the IEEE*, 96(5):529 – 541, May 1981.
12. Daniel A. Pollen and Steven F. Ronner. Visual cortical neurons as localized spatial frequency filters. *IEEE Transactions on Systems, Man, and Cybernetics*, 13(5):907–916, 1983.

13. Eero P. Simoncelli, William T. Freeman, Edward H. Adelson, and David J. Heeger. Shiftable Multiscale Transforms. *IEEE Transactions on Information Theory*, 38(2):587 – 607, March 1992.
14. Jochen Triesch and Christoph von der Malsburg. Robust classification of hand postures against complex backgrounds. In *Proceedings of the Second International Conference on Automatic Face and Gesture Recognition*, pages 170–175. IEEE Computer Society Press, 1996.
15. Sharon Urieli, Moshe Porat, and Nir Cohen. Optimal reconstruction of images from localized phase. *IEEE Trans. Image Processing*, 7(6):838–853, 1998.
16. Christoph von der Malsburg and Ladan Shams. Role of complex cells in object recognition. *Nature Neuroscience*, 2001. Submitted.
17. Laurenz Wiskott, Jean-Marc Fellous, Norbert Krüger, and Christoph von der Malsburg. Face recognition by elastic bunch graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):775–779, 1997.
18. Xing Wu and Bir Bhanu. Gabor Wavelet Representation for 3-D Object Recognition. *IEEE Transactions on Image Processing*, 6(1):47 – 64, January 1997.
19. Ingo J. Wundrich, Christoph von der Malsburg, and Rolf P. Würtz. Image representation by the magnitude of the discrete Gabor wavelet transform. *IEEE Transactions on Image Processing*, 1999. In revision.
20. Rolf P. Würtz. Object recognition robust under translations, deformations and changes in background. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):769–775, 1997.
21. Rolf P. Würtz and Tino Lourens. Corner detection in color images through a multiscale combination of end-stopped cortical cells. *Image and Vision Computing*, 18(6-7):531–541, 2000.