Structure Multivector for Local Analysis of Images*

Michael Felsberg and Gerald Sommer

Christian-Albrechts-University of Kiel
Institute of Computer Science and Applied Mathematics
Cognitive Systems
Preußerstraße 1-9, 24105 Kiel, Germany
Tel: +49 431 560433, Fax: +49 431 560481
mfe@ks.informatik.uni-kiel.de

Abstract. The structure multivector is a new approach for analyzing the local properties of a two-dimensional signal (e.g. image). It combines the classical concepts of the structure tensor and the analytic signal in a new way. This has been made possible using a representation in the algebra of quaternions. The resulting method is linear and of low complexity. The filter-response includes local phase, local amplitude and local orientation of intrinsically one-dimensional neighborhoods in the signal. As for the structure tensor, the structure multivector field can be used to apply special filters to it for detecting features in images.

1 Introduction

In image and image sequence processing, different paradigms of interpreting the signals exist. Regardless of they are following a constructive or an appearance based strategy, they all need a capable low-level preprocessing scheme. The analysis of the underlying structure of a signal is an often discussed topic. Several capable approaches can be found in the literature, among these the quadrature filters derived from the 2D analytic signal [7], the structure tensor [6, 8], and steerable filters [5].

Since the preprocessing is only the first link in a long chain of operations, it is useful to have a linear approach, because otherwise it would be nearly impossible to design the higher-level processing steps in a systematic way. On the other hand, we need a rich representation if we want to treat as much as possible in the preprocessing. Furthermore, the representation of the signal during the different operations should be complete, in order to prevent a loss of information. These constraints enforce us to use the framework of geometric algebra which is also advantageous if we combine image processing with neural computing and robotics (see [9]).

* This work has been supported by German National Merit Foundation and by DFG Graduiertenkolleg No. 357 (M. Felsberg) and by DFG Grant So-320-2-2 (G. Sommer).

In the one-dimensional case, quadrature filters are a frequently used approach for processing data. They are derived from the analytic signal by bandpass filtering. The classical extension to two dimensions is done by introducing a preference direction of the Hilbert transform [7] and therefore, the filter is not very satisfying because the orientation has to be sampled.

The alternative approach is to design a steerable quadrature filter pair [5], which needs an additional preprocessing step for estimating the orientation. As a matter of course, this kind of orientation adaptive filtering is not linear.

The structure tensor (see e.g. [8]) is a capable approach for detecting the existence and orientation of local, intrinsic one-dimensional neighborhoods. From the tensor field the orientation vector field can be extracted and by a normalized or differential convolution special symmetries can be detected [6]. The structure tensor can be computed with quadrature filters but the tensor itself does not possess the typical properties of a quadrature filter. Especially the linearity and the split of the identity is lost, because the phase is neglected.

In this paper, we introduce a new approach for the 2D analytic signal which enables us to substitute the structure tensor by an entity which is linear, preserves the split of the identity and has a geometrically meaningful representation: the structure multivector.

2 A New Approach for the 2D Analytic Signal

2.1 Fundamentals

Since we work on images, which can be treated as sampled intervals of $\mathbb{R}^2$, we use the geometric algebra $\mathbb{R}_{0,2}$ which is isomorphic to the algebra of quaternions $\mathbb{H}$. The whole complex signal theory naturally embeds in the algebra of quaternions, i.e. complex numbers are considered as a subspace of quaternions here. The basis of the quaternions reads $\{\mathbf{1}, \boldsymbol{i}, \boldsymbol{j}, \boldsymbol{k}\}$ while the basis of the complex numbers reads $\{\mathbf{1}, \boldsymbol{i}\}$. Normally, the basis vector $\mathbf{1}$ is omitted.
Throughout this paper, we use the following notations:

- vectors are bold face, e.g. $\boldsymbol{x} = x_1\boldsymbol{i} + x_2\boldsymbol{j}$
- the Fourier transform is denoted[1] $f(\boldsymbol{x}) \circ\!\!-\!\!\bullet\, F(\boldsymbol{u}) = \int f(\boldsymbol{x}) \exp(i2\pi\boldsymbol{u}\cdot\boldsymbol{x})\, d\boldsymbol{x}$
- the real part, the $\boldsymbol{i}$-part, the $\boldsymbol{j}$-part, and the $\boldsymbol{k}$-part of a quaternion q is obtained by $\mathcal{R}\{q\}$, $\mathcal{I}\{q\}$, $\mathcal{J}\{q\}$, and $\mathcal{K}\{q\}$, respectively

The 1D analytic signal is defined as follows. The signal which is obtained from $f(\boldsymbol{x})$ by a phase shift of $\pi/2$ is called the Hilbert transform $f_H(\boldsymbol{x})$ of $f(\boldsymbol{x})$. Since $f_H(\boldsymbol{x})$ is constrained to be real-valued, the spectrum must have an odd symmetry. Therefore, the transfer function has the form[2] $H(\boldsymbol{u}) = \boldsymbol{i}\,\mathrm{sign}(\boldsymbol{u})$. If

[1] Note that the dot product of two vectors is the negative scalar product ($\boldsymbol{x}\cdot\boldsymbol{u} = -\langle\boldsymbol{x},\boldsymbol{u}\rangle$).

[2] Since we use vector notation for 1D functions, we have to redefine some real-valued functions according to $\mathrm{sign}(\boldsymbol{u}) = \mathrm{sign}(u)$, where $\boldsymbol{u} = u\boldsymbol{i}$.

we combine a signal and its Hilbert transform corresponding to

$$f_A(\boldsymbol{x}) = f(\boldsymbol{x}) - f_H(\boldsymbol{x})\boldsymbol{i} \ , \tag{1}$$

we get a complex-valued signal, which is called the analytic signal of $f(\boldsymbol{x})$.

According to the transfer function of the Hilbert transform, the Fourier transform of the analytic signal $f_A(\boldsymbol{x})$ is located in the right half-space of the frequency domain, i.e. $f_A(\boldsymbol{x}) \circ\!\!-\!\!\bullet\, 2F(\boldsymbol{u})\delta_{-1}(\boldsymbol{u})$ (δ_{-1}: Heaviside function).

2.2 The 2D Phase Concept

We want to develop a new 2D analytic signal for intrinsically 1D signals (in contrast to the 2D analytic signal in [2] which is designed for intrinsically 2D signals), which shall contain three properties: local amplitude, local phase and local orientation. Compared to the 1D analytic signal we need one additional angle. We cannot choose this angle without constraints: if the signal is rotated by π, we obtain the same analytic signal, but conjugated. Therefore, we have the following relationship: negation of the local phase is identical to a rotation of π.

Note the difference between *direction* and *orientation* in this context; the direction is a value in $[0; 2\pi)$ and the orientation is a value in $[0; \pi)$.

Any value of the 2D analytic signal can be understood as a 3D vector. The amplitude fixes the sphere on which the value is located. The local phase corresponds to rotations on a great circle on this sphere. To be consistent, a rotation of the signal must then correspond to a rotation on a small circle (local orientation).

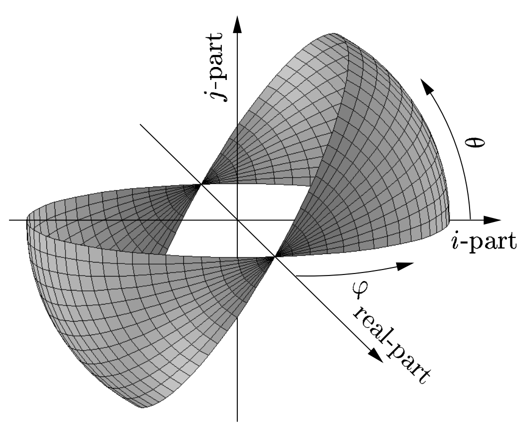

Fig. 1. Coordinate system of the 2D phase approach

The coordinate system defined in this way is displayed in figure 1. It is the same as in [6], but Granlund and Knutsson use the 2D phase in the context of orientation adaptive filtering.

The angles φ and θ are obtained by $\theta = \frac{1}{2}\arg((\mathcal{I}\{q\} + \mathcal{J}\{q\}\boldsymbol{i})^2)$ (*local orientation*) and $\varphi = \arg(\mathcal{R}\{q\} - (\mathcal{J}\{q\} - \mathcal{I}\{q\}\boldsymbol{i})e^{-\boldsymbol{i}\theta})$ (*local phase*), where $q \in \mathbb{H}$ with $\mathcal{K}\{q\} = 0$ and $\arg(z) \in [0; 2\pi)$.

Note that this definition of the quaternionic phase is different from that of the quaternionic Fourier transform (see e.g. [2]). The reason for this will be explained in section 2.3.

2.3 The Monogenic Signal

Now, having a phase concept which is rich enough to code all local properties of intrinsically 1D signals, we construct a generalized Hilbert transform and an analytic signal for the 2D case, which make use of the new embedding.

The following definition of the Riesz transform[3] is motivated by theorem 1, which establishes a correspondence between the Hilbert transform and the Riesz transform. The transfer function of the *Riesz transform* reads

$$H(\boldsymbol{u}) = \frac{\boldsymbol{u}}{|\boldsymbol{u}|} \ , \tag{2}$$

and $F_H(\boldsymbol{u}) = H(\boldsymbol{u})F(\boldsymbol{u}) \bullet\!\!-\!\!\circ f_H(\boldsymbol{x})$.

Example: the Riesz transform of $f(\boldsymbol{x}) = \cos(2\pi\boldsymbol{u}_0 \cdot \boldsymbol{x})$ is
$f_H(\boldsymbol{x}) = -\exp(\boldsymbol{k}\theta_0)\sin(2\pi\boldsymbol{u}_0 \cdot \boldsymbol{x})$ where $\theta_0 = \arg(\boldsymbol{u}_0)$.

Obviously, the Riesz transform yields a function which is identical to the 1D Hilbert transforms of the cosine function, except for an additional rotation in the $\boldsymbol{i} - \boldsymbol{j}$ plane (the exponential function).

Up to now, we have only considered a special example, but what about general signals? What kind of signals can be treated with this approach? The answer can be found easily: the orientation phase must be independent of the frequency coordinate. This sounds impossible, but in fact, the orientation phase is constant, if the spectrum is located on a line through the origin.

Signals which have a spectrum of this form are intrinsically 1D (i.e. they are constant in one direction). This is exactly the class of functions the structure tensor has been designed for and we have the following theorem:

Theorem 1 *Let $f(\boldsymbol{t})$ be a one-dimensional function with the Hilbert transform $f_H(\boldsymbol{t})$. Then, the Riesz transform of the two-dimensional function $f'(\boldsymbol{x}) = f((\boldsymbol{x}\cdot\boldsymbol{n})\boldsymbol{i})$ reads $f'_H(\boldsymbol{x}) = -\boldsymbol{n}\boldsymbol{i}\, f_H((\boldsymbol{x}\cdot\boldsymbol{n})\boldsymbol{i})$, where $\boldsymbol{n} = \cos(\theta)\boldsymbol{i} + \sin(\theta)\boldsymbol{j}$ is an arbitrary unit vector.*

[3] Originally, we used the term spherical Hilbert transform in [4]. We want to thank T. Bülow for alluding to the existence of the Riesz transform and for giving us the reference [10] which enabled us to identify the following definition with it.

Now we simply adapt (1) for the 2D case and obtain the *monogenic signal* of a 2D signal. Using this definition, we obtain for our example: $f_A = \exp(\frac{\boldsymbol{u}_0}{|\boldsymbol{u}_0|} 2\pi \boldsymbol{u}_0 \cdot \boldsymbol{x})$.

Hence, the monogenic signal uses the phase concept, which has been defined in section 2.2. According to theorem 1, the monogenic signal of an intrinsically one-dimensional signal $f'(\boldsymbol{x}) = f((\boldsymbol{x} \cdot \boldsymbol{n})\boldsymbol{i})$ reads

$$f'_A(\boldsymbol{x}) = f((\boldsymbol{x} \cdot \boldsymbol{n})\boldsymbol{i}) - \boldsymbol{n} f_H((\boldsymbol{x} \cdot \boldsymbol{n})\boldsymbol{i}) \ . \tag{3}$$

Of course, the monogenic signal can be computed for *all* functions which are Fourier transformable. However, for signals which do not have an intrinsic dimension of one[4], the correspondence to the 1D analytic signal is lost.

Independently of the intrinsic dimensionality of the signal, the analytic signal can also be calculated in a different way. The 1D analytic signal is obtained in the Fourier domain by the transfer function $1 + \mathrm{sign}(\boldsymbol{u})$. For the monogenic signal we have the same result if we modify the Fourier transform according to $\tilde{f}(\boldsymbol{x}) = \int_{\mathbb{R}^2} \exp(\boldsymbol{k}\theta/2) F(\boldsymbol{u}) \exp(\boldsymbol{i}2\pi \boldsymbol{u} \cdot \boldsymbol{x}) \exp(-\boldsymbol{k}\theta/2)\, d\boldsymbol{u}$, the *inverse spherical Fourier transform*. Then, we have $f_A(\boldsymbol{x}) = \tilde{f}(\boldsymbol{x})$ (see [4]).

Since the integrand of $\tilde{f}(\boldsymbol{x})$ is symmetric, we can also integrate over the half domain and multiply the integral by two. Therefore, we can use any transfer function of the form $1 + \mathrm{sign}(\boldsymbol{u} \cdot \boldsymbol{n})$ without changing the integral. By simply omitting half of the data, the redundancy in the representation is removed.

In order to calculate the energy of the monogenic signal, we need the transfer function, which changes $F(\boldsymbol{u})$ to $F_A(\boldsymbol{u})$: it is obtained from (2) and (1) and reads $1 - \frac{\boldsymbol{u}}{|\boldsymbol{u}|}\boldsymbol{i} = 1 + \cos(\theta) + \sin(\theta)\boldsymbol{k}$. The energy of the monogenic signal is

$$\int_{\mathbb{R}^2} |(1 + \cos(\theta) + \sin(\theta)\boldsymbol{k})F(\boldsymbol{u})|^2 \, d\boldsymbol{u} = 2 \int_{\mathbb{R}^2} |F(\boldsymbol{u})|^2 \, d\boldsymbol{u} \ , \tag{4}$$

i.e. it is two times the energy of the original signal[5].

From the group of similarity transformations (i.e. shifts, rotations and dilations) only the rotation really affects the monogenic signal; the orientation phase is changed according to the rotation. If we interpret the monogenic signal as a vector field in 3D (see also section 3.2), the group of 2D similarity transforms[6] even commutes with the operator that yields the monogenic signal.

The reader might ask, why do we use a quaternion-valued Fourier transform which differs from the QFT (see e.g. [2]). The reason is not obvious. The QFT covers more symmetry concepts than the complex Fourier transform. The classical transform maps a reflection through the origin onto the conjugation operator. The QFT maps a reflection in one of the axes onto one of the algebra

[4] The case of intrinsic dimension zero (i.e. a constant signal) is irrelevant, because the Hilbert transform is zero in both cases.

[5] This is only valid for DC free signals. The energy of the DC component is not doubled as in the case of the 1D analytic signal.

[6] Note that in the context of a 3D embedding, the group of 2D similarity transforms is the subgroup of the 3D transforms restricted to shifts in $\boldsymbol{i}$ or $\boldsymbol{j}$ direction, rotations around the real axis and a dilation of the $\boldsymbol{i}$ and $\boldsymbol{j}$ axes.

automorphisms. We can use $\mathbb{C}\otimes\mathbb{C}$ instead of $\mathbb{H}$ to calculate the QFT. Moreover, we have $\mathbb{C}\otimes\mathbb{C}\cong\mathbb{C}^2$ (see e.g. [3]) and consequently, we can calculate the QFT of a real signal by two complex transformations using the formula

$$F_q(\boldsymbol{u}) = F(\boldsymbol{u})\frac{\boldsymbol{1}-\boldsymbol{k}}{2} + F(u_1\boldsymbol{i} - u_2\boldsymbol{j})\frac{\boldsymbol{1}+\boldsymbol{k}}{2} \tag{5}$$

and the symmetry wrt. the axes is obvious.

In this paper, we want to present an isotropic approach which means that symmetry wrt. the axes is not sufficient. Therefore, we had to design the new transform. The design of isotropic discrete filters is a quite old topic, see e.g. [1].

3 Properties of the Monogenic Signal

3.1 The Spatial Representation

The definition of the Riesz transform in the frequency domain can be transformed into a spatial representation. The transfer function (2) can be split into two functions: $\frac{u_1}{|\boldsymbol{u}|}$ and $\frac{u_2}{|\boldsymbol{u}|}$. The only thing left is to calculate the inverse Fourier transform of these functions. In [10] the transform pairs can be found: $\frac{x_1}{2\pi|\boldsymbol{x}|^3} \circ\!\!-\!\!\bullet -\boldsymbol{i}\frac{u_1}{|\boldsymbol{u}|}$ and $\frac{x_2}{2\pi|\boldsymbol{x}|^3} \circ\!\!-\!\!\bullet - \boldsymbol{i}\frac{u_2}{|\boldsymbol{u}|}$.

The functions $\frac{x_1}{2\pi|\boldsymbol{x}|^3}$ and $\frac{x_2}{2\pi|\boldsymbol{x}|^3}$ are the kernels of the 2D Riesz transform in vector notation. From a mathematician's point of view, the Riesz transform is the multidimensional generalization of the Hilbert transform. Consequently, the monogenic signal is directly obtained by the convolution

$$f_A(\boldsymbol{x}) = \int_{\mathbb{R}^2}\left(\delta_0(\boldsymbol{t}) + \frac{\boldsymbol{t}}{2\pi|\boldsymbol{t}|^3}\right) f(\boldsymbol{x}-\boldsymbol{t})\,d\boldsymbol{t} \ .$$

The graph in figure 2 sums up all ways to calculate the monogenic signal from the preceding sections. The inverse spherical Fourier transform is denoted $\mathcal{F}_s^{-1}$.

3.2 The Structure Multivector

Normally, images are intrinsically two-dimensional, so the concepts described in section 2.3 cannot be applied globally. On the other hand, large areas of images are intrinsically one-dimensional, at least on a certain scale. Therefore, a local processing would take advance of the new approach.

The classical approach of the analytic signal has its local counterpart in the quadrature filters. A pair of quadrature filters (or a complex quadrature filter) is characterized by the fact that the impulse response is an analytic signal. On the other hand, both impulse responses are band-limited and of finite spatial extent, so that the problem of the unlimited impulse response of the Hilbert filter is circumvented.
An example of the output of a 1D quadrature filter can be found in figure 3.

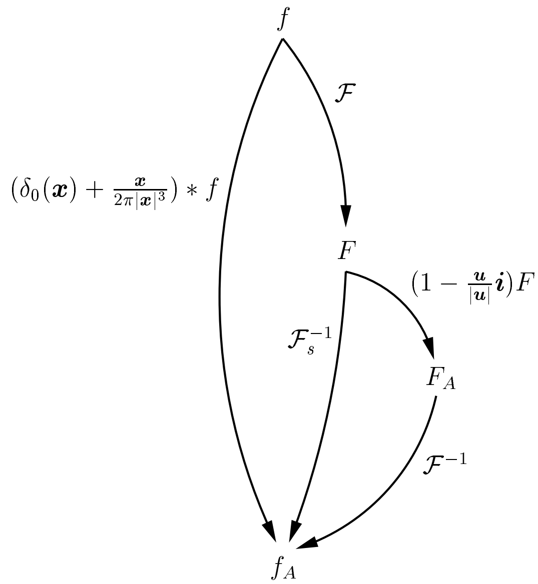

Fig. 2. Three ways to calculate the monogenic signal

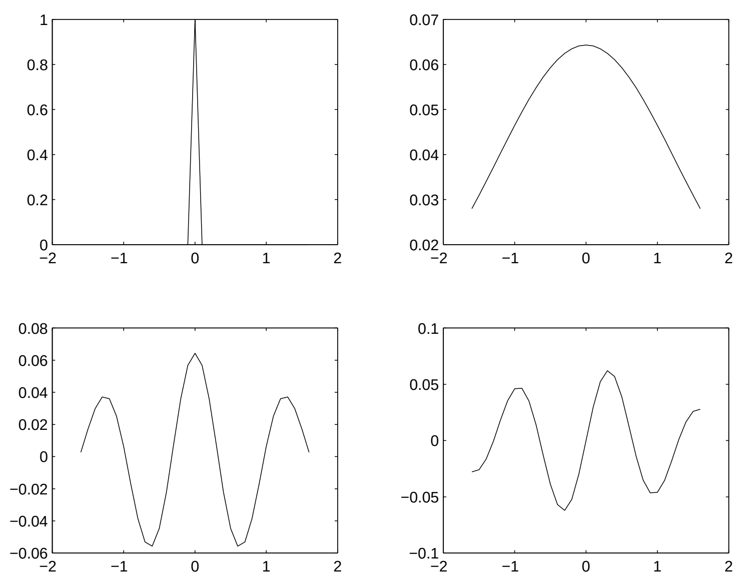

Fig. 3. Upper left: impulse, upper right: magnitude of filter output, lower left: real part of filter output, lower right: imaginary part of filter output

While the representation in figure 3 is very common, we will introduce now a different representation. One-dimensional signals can be interpreted as surfaces in 2D space. If we assign the real axis to the signal values and the imaginary axis to the abscissa, we obtain a representation in the complex plane. The analytic signal can be embedded in the same plane – it corresponds to a vector field which is only non-zero on the imaginary axis (see figure 4).

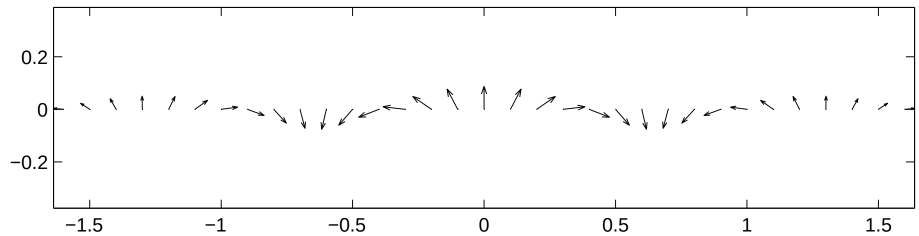

Fig. 4. Representation of the 1D analytic signal as a vector field

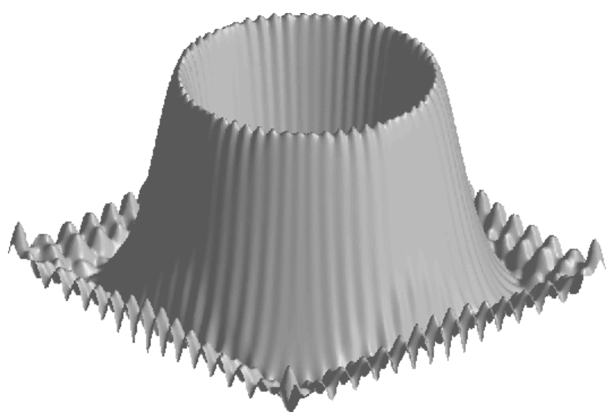

Fig. 5. Siemens-star convolved with a spherical quadrature filter (magnitude)

Based on the monogenic signal, we introduce the *spherical quadrature filters*. They are defined according to the 1D case as a hypercomplex filter whose impulse response is a monogenic signal.

It is remarkable that the spherical quadrature filters have isotropic energy and exactly choose the frequency bands they are designed for. In figure 5 it can be seen that the energy is isotropic and that it is maximal for the radius 77.8, which corresponds to a frequency of $\frac{16.3}{256}$. The used bandpass has a center frequency of $\frac{1}{16}$.

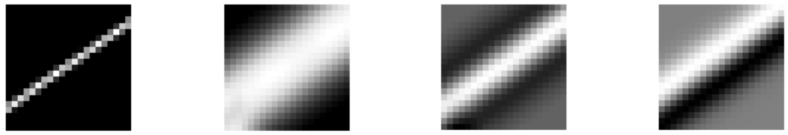

Fig. 6. From left to right: impulse-line; filter output: amplitude, real part, combined imaginary parts

In figure 6 the output of a spherical quadrature filter applied to an impulse-line is displayed. The lower right image shows the combined imaginary parts which means that instead of the imaginary unit $\boldsymbol{i}$ the unit vector $\boldsymbol{n}$ is used.

Same as for 1D signals, 2D signals can be embedded as a surface in 3D space. The signal values are assigned to the real axis and the spatial coordinates to the $\boldsymbol{i}$- and the $\boldsymbol{j}$-axis. The monogenic signal can be represented in the same embedding. It corresponds to a vector field, which is only non-zero in the plane spanned by $\boldsymbol{i}$ and $\boldsymbol{j}$ (see figure 7).

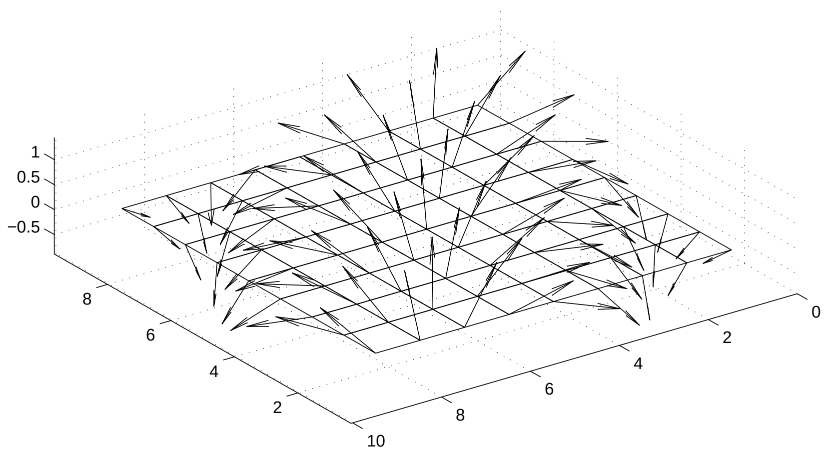

Fig. 7. Representation of the monogenic signal as a vector field

The result of filtering a signal with a spherical quadrature filter is a quaternion-valued field. Though the $\boldsymbol{k}$-component (*bivector*) of the field is always zero, we denote this field as a multivector field or the *structure multivector* of the signal. As already the name induces, the structure multivector is closely related to the structure tensor. The structure tensor as defined in [6] mainly includes the following information: the amplitude as a measurement for the existence of local structure and the orientation of the local structure.

Jähne [8] extracts an additional information: the coherence. The coherence is the relationship between the oriented gradients and all gradients, so it is a measurement for the degree of orientation in a structure and it is closely related to the variance of the orientation. The variance is a second order property. It includes a product of the arguments and therefore, it is not linear. Consequently, the coherence cannot be measured by a linear approach like the structure multivector. Two structures with different orientations simply yield the vector sum of both multivector fields.

The structure multivector consists of three independent components (local phase, local orientation and local amplitude) and it codes three properties. Consequently, there is no additional information possible. The structure tensor possesses three degrees of freedom (it is a symmetric tensor). Therefore, apart from the amplitude and the orientation one can extract a third information, the coherence.

4 Experiments and Discussion

4.1 Experiments

For the computation of the structure multivector we use a multi-scale approach, i.e. we couple the shift of the Gaussian bandpass with the variance as for the Gabor wavelets.

For the experiments, we chose some synthetic examples with letters as gray level or textured images.

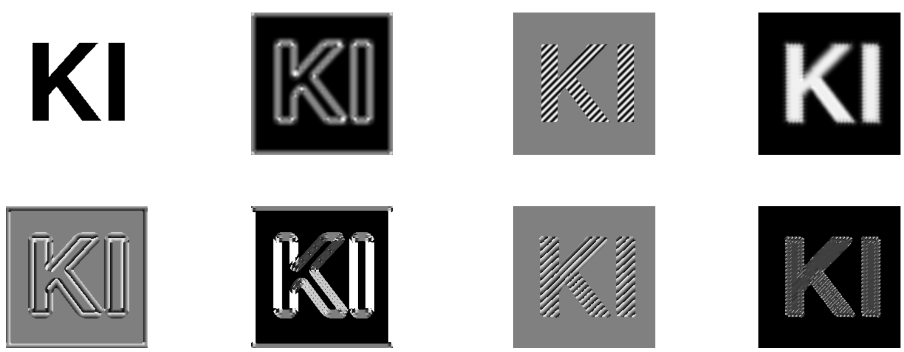

(a) (b)

Fig. 8. (a) Structure multivector of an image without texture. Upper left: original, upper right: amplitude, lower left: φ-phase, lower right: θ-phase.
(b) Structure multivector of an image with one texture. Upper left: original, upper right: amplitude, lower left: φ-phase, lower right: θ-phase

In figure 8(a) it can be seen that the structure multivector responds only at the edges. Therefore, the amplitude is a measure for the presence of structure.

The φ-phase is linear, which can only be guessed in this representation. But it can be seen that the φ-phase is monotonic modulo a maximal interval (which is in fact 2π). The θ-phase represents the orientation of the edge. Note that the highest and the lowest gray level (standing for π and zero, respectively) represent the same orientation.

In figure 8(b) it can be seen that the structure multivector responds also inside the object. The amplitude is nearly constant, which corresponds to a texture with constant energy. Of course, this property is lost if the wrong scale is considered. The φ-phase is linear, see notes above. The θ-phase represents the orientation of the texture. A constant gray level corresponds to a constant estimated orientation. The small spikes in the figure are produced by the extraction of the local orientation angle. The underlying quaternion-valued field does not show these artifacts.

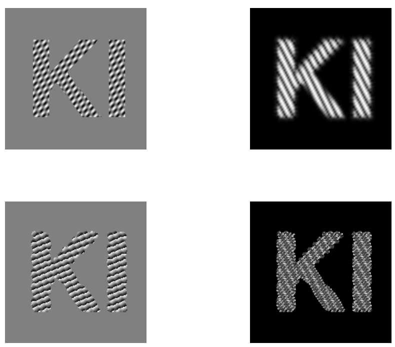

(a)

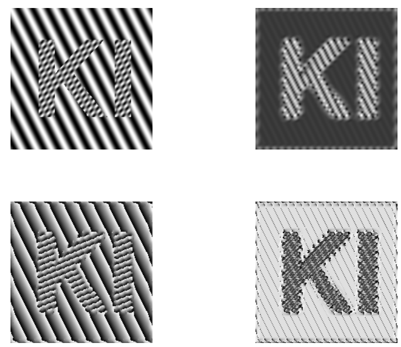

(b)

Fig. 9. (a) Structure multivector of an image with two superposed textures. Upper left: original, upper right: amplitude, lower left: φ-phase, lower right: θ-phase.
(b) Structure multivector of an image with two superposed textures and textured background. Upper left: original, upper right: amplitude, lower left: φ-phase, lower right: θ-phase

In figure 9(a) it can be seen that the structure multivector responds only with respect to the dominant texture (the one with higher frequency). The magnitude of the response is modulated with that component of the weaker texture that is normal to the dominant texture. This effect is even more obvious in figure 9(b). The φ-phase is always directed parallel to the dominant texture. The θ-phase represents the orientation of the dominant texture in each case.

4.2 Conclusion

We have presented a new approach to the 2D analytic signal: the monogenic signal. It has an isotropic energy distribution and deploys the same local phase approach as the 1D analytic signal. There is no impact of the orientation on the local phase which is one of the most important drawbacks of the classical 2D analytic signal.

Additionally, the monogenic signal includes information about the local orientation and therefore, it is related to the structure tensor. On the other hand, there are two differences compared to the latter approach: the structure multivector does not include coherence information and it is linear.

The local counterpart to the monogenic signal is the structure multivector. The latter is the response of the spherical quadrature filters which are a generalization of the 1D quadrature filters. From the structure multivector one obtains a stable orientation estimation (as stable as the orientation vector field of the structure tensor). In contrast to the classical quadrature filters, the spherical quadrature filters do not have a preference direction. Therefore, the orientation need not be sampled or steered.

We introduced an interpretation technique for the analytic and the monogenic signal in form of vector fields. Furthermore, we tried to explain the impact of simple structures and textures on the structure multivector. Applications can easily be designed, e.g. texture segmentation (see also [2]).

References

[1] Brady, J. M., and Horn, B. M. P. Rotationally symmetric operators for surface interpolation. *Computer Vision, Graphics, and Image Processing 22*, 1 (April 1983), 70–94.

[2] Bülow, T. *Hypercomplex Spectral Signal Representations for the Processing and Analysis of Images.* PhD thesis, Christian-Albrechts-University of Kiel, 1999.

[3] Felsberg, M., Bülow, T., and Sommer, G. Commutative Hypercomplex Fourier Transforms of Multidimensional Signals. In *Geometric Computing with Clifford Algebra*, G. Sommer, Ed. Springer, Heidelberg, 2001.

[4] Felsberg, M., and Sommer, G. Structure multivector for local analysis of images. Tech. Rep. 2001, Institute of Computer Science and Applied Mathematics, Christian-Albrechts-University of Kiel, Germany, February 2000. http://www.ks.informatik.uni-kiel.de/~mfe/Techn_Report.ps.gz.

[5] Freeman, W. T., and Adelson, E. H. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence 13*, 9 (September 1991), 891–906.

[6] Granlund, G. H., and Knutsson, H. *Signal Processing for Computer Vision.* Kluwer Academic Publishers, Dordrecht, 1995.

[7] Hahn, S. L. *Hilbert Transforms in Signal Processing.* Artech House, Boston, London, 1996.

[8] Jähne, B. *Digitale Bildverarbeitung.* Springer, Berlin, 1997.

[9] Sommer, G. The global algebraic frame of the perception-action cycle. In *Handbook of Computer Vision and Applications* (1999), B. Jähne, H. Haußecker, and P. Geissler, Eds., vol. 3, Academic Press, San Diego, pp. 221–264.

[10] Stein, E., and Weiss, G. *Introduction to Fourier Analysis on Euclidean Spaces.* Princeton University Press, New Jersey, 1971.