Lehrstuhl für Bildverarbeitung
Institute of Imaging & Computer Vision

Analysis of Superimposed Oriented Patterns

Til Aach, Cicero Mota, Ingo Stuke, Matthias Mühlich and Erhard Barth

in: IEEE Transactions on Image Processing. See also BIBTEX entry below.

BIBTEX:

```
@article{AAC06b,
author = {Til Aach and Cicero Mota and Ingo Stuke and Matthias M{\"u}hlich and Erhardt Barth},
title = {Analysis of Superimposed Oriented Patterns},
journal = {IEEE Transactions on Image Processing},
        publisher       = {IEEE},
volume = {15},
number = {12},
year = {2006},
pages = {3690--3700}}
```


document created on: November 21, 2006
created from file: orientip_cover.tex
cover page automatically created with CoverPage.sty
(available at your favourite CTAN mirror)

RWTH AACHEN
RHEINISCH-WESTFÄLISCHE TECHNISCHE HOCHSCHULE AACHEN

Analysis of Superimposed Oriented Patterns

Til Aach, *Senior Member, IEEE,* Cicero Mota, Ingo Stuke, Matthias Mühlich, and Erhard Barth

***Abstract*— Estimation of local orientation in images may be posed as the problem of finding the minimum gray level variance axis in a local neighborhood. In bivariate images, the solution is given by the eigenvector corresponding to the smaller eigenvalue of a 2×2 tensor. For an ideal single orientation, the tensor is rank-deficient, i.e., the smaller eigenvalue vanishes. A large minimal eigenvalue signals the presence of more than one local orientation, what may be caused by non-opaque additive or opaque occluding objects, crossings, bifurcations, or corners. We describe a framework for estimating such superimposed orientations. Our analysis is based on the eigensystem analysis of suitably extended tensors for both additive and occluding superpositions. Unlike in the single orientation case, the eigensystem analysis does not directly yield the orientations, rather, it provides so-called mixed orientation parameters. We therefore show how to decompose the mixed orientation parameters into the individual orientations. We also show how to use tensor invariants to increase efficiency, and derive a new feature for describing local neighborhoods which is invariant to rigid transformations. Applications are, e.g., in texture analysis, directional filtering and interpolation, feature extraction for corners and crossings, tracking, and signal separation.**

I. Introduction

ESTIMATION of local orientation is essential in a variety of image filtering and analysis tasks, such as directional filtering by spectral-domain [1], [2] or steerable [3], [4] filtering, directional interpolation [5], [6], texture and pattern analysis [7], [8], [9], [10], [11], or compression [12], [13]. In these applications, orientation estimation seeks to find local onedimensional (1D) structures such as lines in bivariate [8], [1], [2] or tri- and higher-variate image data [5], [6], [7]. Similarly, in the sense of finding locally oriented structures in spatiotemporal image sequences, orientation estimation is a key ingredient of estimating optical flow [14]. Approaches to local orientation estimation include quadrature filters [15], [11], [16], steerable filters [17], and inertia-tensor based methods [18], [19], [20], [21], [22], [23]. Regarding 1D local orientation as the direction in which the gray level profile varies least [19], [24], [25], the latter find the orientation by analysing the eigensystem of a symmetric matrix, or tensor, computed from the gradients of the image. In bivariate images, the orientation is given by the eigenvector corresponding to the tensor's lower eigenvalue, while the other (orthogonal) eigenvector points in the direction of maximum variation. The lower and larger eigenvalue quantify the gray level variations along and perpendicular to the orientation axis, respectively. This framework assumes that only a single orientation is present. In case of two or more superimposed orientations, the eigenvectors no longer represent orientation. Such neighborhoods — generated for instance, by corners [26] — are characterized by the lower eigenvalue being reasonably large, indicating that no axis along which the gray level variation is low can be found. Superimposed orientations appear in X-ray projection imaging [27], in multioriented textures and fabrics, or locally as corners and junctions. Superimposed orientations also occur ubiquituously in Radon space ("sinograms" in Computer Tomography [28]), where sine curves of different objects cross. Corners and junctions are a rich source of information: L- and Y-junctions represent object corners, T-junctions occur at occluding object boundaries, X-junctions at object crossings, while Ψ-junctions are caused by bending object surfaces [29], [30].

In the following, we describe a framework for identifying and estimating superimposed orientations. After reviewing single orientation estimation in section II, we develop the approaches for the estimation of overlaid orientations. Starting from the problem description for transparently superimposed optical flows in [31] and multi-orientation fields in [32], we show that both additively overlaid *and* occluding orientations can be treated in the same manner. The superimposed orientations can be estimated by the eigensystem analysis of suitably extended tensors, yielding so-called mixed orientation parameters. When normalized to unit length, the mixed orientation parameter (MOP) vector forms a unique descriptor for double-orientation neighborhoods in bivariate images, but does not provide the orientations explicitly. We therefore show how to decompose the MOP vector into the individual orientations for both bivariate images and higher-variate data. We furthermore derive an hierarchical algorithm which successively tests for zero-, single-, double- and higher orientation neighborhoods based on tensor invariants *without* an explicit eigensystem analysis. Moreover, once the appropriate orientation model for a neighborhood is identified, its MOP parameters can also be estimated *without* eigensystem analysis. We then expand the approach towards multispectral data and derive a new set of features for describing local neighborhoods, which are invariant to similarity transforms such as rotation. Subsequently, we show how both additively superimposed and occluding patterns can be separated. We provide results for both synthetic and real images, and conclude by a discussion.

T. Aach and M. Mühlich are with RWTH Aachen University, Germany. C. Mota is with the Federal University of Amazonas, Brazil. I. Stuke and E. Barth are with the University of Lübeck, Germany. This work was supported in part by DFG grant Ba 1176/9-3 (LOCOMOTOR). C. Mota was partially supported by grants from CNPq/FAPEAM.

Parts of this work were published at the 6th IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI-2004), Lake Tahoe, March 28-30, 2004, at ICASSP-2004, Montreal, May 2004, at ICIP-2004, Singapore, October 2004, and at the DAGM Symposium Tübingen, August 2004.

II. Estimation of Single Orientations

A. Bivariate gray level images

Let $f(\mathbf{x})$, $\mathbf{x}=(x,y)^T$, with $f:\mathbb{R}^2\rightarrow\mathbb{R}$ denote a bivariate gray level image. The image $f(\mathbf{x})$ is oriented in a local region

Ω if and only if

$$f(\mathbf{x}) = f(\mathbf{x}+k\mathbf{v}) \quad \forall k \in \mathbb{R} \text{ and } \forall\, \mathbf{x}, \mathbf{x}+k\mathbf{v} \in \Omega \quad (1)$$

where we exclude the case that $f(\mathbf{x})$ is constant over Ω. The unit vector $\mathbf{v} = (\cos\theta, \sin\theta)^T = \mathbf{v}(\theta)$ describes the orientation of $f(\mathbf{x})$ in terms of the angle θ. With the directional derivative operator along $\mathbf{v}(\phi)$ defined as $\alpha(\phi) = \frac{\partial}{\partial \mathbf{v}(\phi)} = \cos(\phi)\frac{\partial}{\partial x} + \sin(\phi)\frac{\partial}{\partial y}$, condition (1) is equivalent to [19], [20]

$$\frac{\partial f(\mathbf{x})}{\partial \mathbf{v}(\theta)} = \alpha(\theta) f(\mathbf{x}) = \mathbf{v}^T(\theta) \cdot \nabla f(\mathbf{x}) = 0 \quad \forall\, \mathbf{x} \in \Omega \ , \quad (2)$$

where $\nabla f = (f_x, f_y)^T$ is the gradient of $f(\mathbf{x})$. When met for θ, this condition holds also for $\theta \pm \pi$, we therefore restrict θ to lie within $(-\pi/2, \pi/2]$. Condition (2), in turn, is equivalent to

$$Q(\theta) = \int_\Omega \left(\frac{\partial f(\mathbf{x})}{\partial \mathbf{v}(\theta)} \right)^2 d\Omega = \int_\Omega \left(\mathbf{v}^T(\theta) \nabla f(\mathbf{x}) \right)^2 d\Omega = 0 \ . \quad (3)$$

$Q(\phi)$ is a measure of the gray level variation in the direction determined by ϕ. (In practice, the integrand in (3) is often weighted by a function $w(\mathbf{x})$ emphasizing the central pixels in Ω, and with a continuous roll-off towards its borders. Since this weighting does not influence our considerations, we drop it for ease of notation.) Generally, due to noise and the fact that the orientation may not be perfectly constant over Ω, it is impossible to find θ such that $Q(\theta) = 0$. We thus seek to minimize $Q(\theta)$ [19], [20], i.e.

$$\theta = \arg \min_{-\pi/2 < \phi \leq \pi/2} Q(\phi) \ . \quad (4)$$

This can be interpreted as seeking the axis of minimum intensity variation, or, equivalently, the direction which is "most orthogonal" to all gradients of $f(\mathbf{x})$ in Ω. $Q(\phi)$ can be rewritten to

$$Q(\phi) = \mathbf{v}^T(\phi) \mathbf{J} \mathbf{v}(\phi) \ , \quad \mathbf{v}^T\mathbf{v} = 1 \ . \quad (5)$$

The symmetric and positive semi-definite tensor $\mathbf{J}$ is calculated from $f(\mathbf{x})$ by summing over the pixelwise products of its gradient components f_x, f_y according to

$$\begin{aligned} \mathbf{J} &= \mathbf{J}(f) = \int_\Omega (\nabla f)(\nabla f)^T d\Omega \\ &= \int_\Omega \begin{bmatrix} f_x^2 & f_x f_y \\ f_x f_y & f_y^2 \end{bmatrix} d\Omega = \int_\Omega \nabla f \otimes \nabla f d\Omega \ , \quad (6) \end{aligned}$$

where $\otimes$ denotes the tensor product or outer product. On a discrete grid, computing $\mathbf{J}$ can be realized by linear shift invariant derivative filtering and pixelwise multiplication followed by lowpass filtering. Minimizing the composite criterion $L(\mathbf{v}) = \mathbf{v}^T\mathbf{J}\mathbf{v} + \lambda_2(1 - \mathbf{v}^T\mathbf{v})$ is equivalent to finding $\mathbf{v}$ such that

$$\mathbf{J}\mathbf{v} = \lambda_2 \mathbf{v} \ , \quad \mathbf{v}^T\mathbf{v} = 1 \quad (7)$$

where the constraint $\mathbf{v}^T\mathbf{v} = 1$ excludes the trivial solution $\mathbf{v} = \mathbf{0}$. The solution $\mathbf{v}$ is the normalized eigenvector of $\mathbf{J}$ corresponding to its lower eigenvalue λ_2. Note that $\mathbf{v}(\theta)$ is then uniquely determined up to its sign, which in turn is determined by the constraint $\theta \in (-\pi/2, \pi/2]$. The residual intensity variation over Ω is equal to the lower eigenvalue, since $Q(\theta) = \mathbf{v}^T(\theta)\mathbf{J}\mathbf{v}(\theta) = \mathbf{v}^T\lambda_2\mathbf{v} = \lambda_2$. For $f(\mathbf{x})$ ideally oriented in the sense of (1), the smaller eigenvalue thus vanishes, and $\text{rank}(\mathbf{J}) = 1$. Compliance with the single orientation assumption is hence indicated by a low value of λ_2, and a high value for the large eigenvalue λ_1. Violation of the single orientation hypothesis, i.e., the presence of more than one orientation, is indicated by a high λ_2, or $\mathbf{J}$ having full rank. For $f(\mathbf{x})$ being constant over Ω, both eigenvalues vanish, and $\text{rank}(\mathbf{J}) = 0$.

B. Multivariate and vector-valued image data

Let the mapping $\mathbf{f} : \mathbb{R}^p \to \mathbb{R}^q$ now denote a multivariate, multispectral image signal, with, e.g., $p = 3$ and $q = 3$ in color movies. The image $\mathbf{f}(\mathbf{x})$ is oriented in the region Ω if there is a subspace $E \subset \mathbb{R}^p$ such that

$$\mathbf{f}(\mathbf{x}+\mathbf{v}) = \mathbf{f}(\mathbf{x}) \quad \forall \mathbf{x}, \mathbf{v};\, \mathbf{x}, \mathbf{x}+\mathbf{v} \in \Omega,\, \mathbf{v} \in E. \quad (8)$$

This can equivalently be expressed as saying that the intrinsic dimension [33] of $\mathbf{f}(\mathbf{x})$ over Ω is $p - \dim(E)$. For oriented bivariate image neighborhoods, $\dim(E)$ is always equal to one, what leads to (1). The goal of orientation estimation now is to obtain E. Condition (8) is equivalent to [7] $\frac{\partial \mathbf{f}(\mathbf{x})}{\partial \mathbf{v}} = \mathbf{0} \quad \forall\ \mathbf{x} \in \Omega$ and $\mathbf{v} \in E$ with $\mathbf{f} = (f_1, \ldots, f_q)^T$ and $\partial \mathbf{f}(\mathbf{x}) / \partial \mathbf{v} = (\mathbf{v}^T \nabla f_1, \ldots, \mathbf{v}^T \nabla f_q)^T$. Equation (3) is thus replaced by

$$Q = \int_\Omega \left| \frac{\partial \mathbf{f}(\mathbf{x})}{\partial \mathbf{v}} \right|^2 d\Omega = 0 \ . \quad (9)$$

Let $\mathbf{v} = (v_1, \ldots, v_p)^T$ and $\mathbf{f}_{x_j}(\mathbf{x}) = \partial \mathbf{f}(\mathbf{x}) / \partial x_j$. Eq. (9) can then be rewritten as

$$\mathbf{v}^T \mathbf{J} \mathbf{v} = 0, \ \ \mathbf{v}^T\mathbf{v} = 1 \quad (10)$$

where the $p \times p$ tensor $\mathbf{J} = \mathbf{J}(\mathbf{f})$ is now given by

$$\mathbf{J} = \int_\Omega \begin{bmatrix} \mathbf{f}_{x_1}^T(\mathbf{x}) \mathbf{f}_{x_1}(\mathbf{x}) & \ldots & \mathbf{f}_{x_1}^T(\mathbf{x}) \mathbf{f}_{x_p}(\mathbf{x}) \\ \vdots & & \vdots \\ \mathbf{f}_{x_p}^T(\mathbf{x}) \mathbf{f}_{x_1}(\mathbf{x}) & \ldots & \mathbf{f}_{x_p}^T(\mathbf{x}) \mathbf{f}_{x_p}(\mathbf{x}) \end{bmatrix} d\Omega. \quad (11)$$

Consequently, the integrand can be rewritten as

$$\left| \frac{\partial \mathbf{f}(\mathbf{x})}{\partial \mathbf{v}} \right|^2 = \mathbf{v}^T \left(\sum_{k=1}^q \nabla f_k (\nabla f_k)^T \right) \mathbf{v} \quad (12)$$

which implies that the tensor $\mathbf{J}$ extends towards the sum over the tensors calculated for each component $f_k(\mathbf{x})$ of $\mathbf{f}(\mathbf{x})$ (cf. [18] for color images):

$$\mathbf{J}_q = \sum_{k=1}^q \mathbf{J}(f_k) \ . \quad (13)$$

Since $\mathbf{J}_q$ is symmetric and positive semi-definite, Eq. (10) is equivalent to $\mathbf{J}\mathbf{v} = \mathbf{0}$, which means that E is the null-eigenspace of $\mathbf{J}_q$. As in (5), E is estimated by minimizing $\mathbf{v}^T \mathbf{J}_q \mathbf{v}$ subject to $\mathbf{v}^T\mathbf{v} = 1$, leading to an eigensystem analysis as in (7). If $\mathbf{f}(\mathbf{x})$ is ideally oriented in Ω in the sense of (8), $\text{rank}(\mathbf{J}_q) = p - \dim(E)$. Compliance with the single orientation hypothesis can thus be derived from the eigenvalues of $\mathbf{J}_q$.

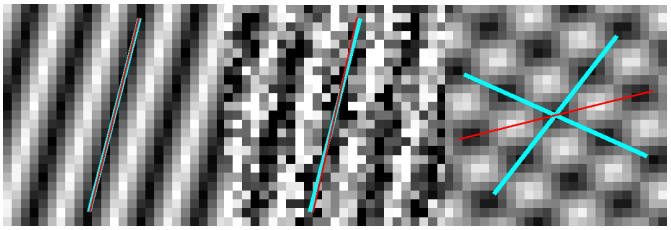

Fig. 1. Synthetically generated ideal orientation neighborhoods, 25×25 pixel, with true and estimated orientations shown in light and dark gray, respectively. Left: Single orientation, noisefree. Estimation error: $0.43°$, $\lambda_2/\lambda_1 = 0.00212$. Center: With additive Gaussian noise, SNR 0dB. Estimation error $1.72°$, $\lambda_2/\lambda_1 = 0.23$. Right: Double orientation, noisefree. The single orientation estimate captures neither of the true orientations. $\lambda_2/\lambda_1 = 0.6$.

Fig. 1 illustrates this approach: Left, the ideal single orientation is estimated with very low error. Compliance of the analyzed neighborhood with the single orientation hypothesis is indicated by a low λ_2 relative to the larger eigenvalue λ_1. Additive noise (center) leads to an increased but still relatively low λ_2. The additive double orientations on the right-hand side of Fig. 1 violate the single orientation model. Consequently, the estimate captures neither orientation, but signals the model violation by a relatively large λ_2. In these examples, the derivatives f_x and f_y were calculated with the finite-difference kernel $[-1 \;\; 0 \;\; 1]$ in horizontal resp. vertical direction combined with simultaneous orthogonal smoothing by the kernel $[1 \;\; 3 \;\; 1]$.

III. Estimation of Double Orientations

A. Bivariate gray level images

1) Modelling and eigensystem analysis: Let us now assume that within Ω, $f(\mathbf{x})$ is composed additively of two non-opaque oriented subimages according to

$$f(\mathbf{x}) = f_1(\mathbf{x}) + f_2(\mathbf{x}) \;, \quad f_1, f_2 : \mathbb{R}^2 \to \mathbb{R} \;\; . \tag{14}$$

The orientations of $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ are given by the unit vectors $\mathbf{u} = (\cos\theta, \sin\theta)^T = (u_x, u_y)^T$ and $\mathbf{v} = (\cos\gamma, \sin\gamma)^T = (v_x, v_y)^T$, respectively. $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ thus obey

$$\begin{aligned} \alpha(\theta) f_1(\mathbf{x}) &= \mathbf{u}^T \nabla f_1(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \Omega \\ \alpha(\gamma) f_2(\mathbf{x}) &= \mathbf{v}^T \nabla f_2(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \Omega \end{aligned} \tag{15}$$

The composite image $f(\mathbf{x})$ then satisfies [32]

$$\alpha(\theta)\alpha(\gamma) f(\mathbf{x}) = \frac{\partial^2 f(\mathbf{x})}{\partial \mathbf{u} \partial \mathbf{v}} = 0 \quad \forall \mathbf{x} \in \Omega \;\; . \tag{16}$$

This constraint can be rewritten as the inner product $\mathbf{a}^T \; \mathbf{d}f(\mathbf{x}) = 0 \;\; \forall \mathbf{x} \in \Omega$, where $\mathbf{a}$ is a threedimensional vector given by

$$\begin{aligned} \mathbf{a}^T &= (u_x v_x, u_x v_y + u_y v_x, u_y v_y) \\ &= (\cos\theta\cos\gamma, \sin(\theta+\gamma), \sin\theta\sin\gamma) = (a, b, c) \end{aligned} \tag{17}$$

and

$$\mathbf{d}f = (f_{xx}, f_{xy}, f_{yy})^T \;\; . \tag{18}$$

The components of $\mathbf{a}$ are the so-called mixed orientation parameters (MOP) resulting from the concatenation of two directional derivatives. The constraint is equivalent to

$$Q_2(\mathbf{a}) = \int_\Omega \left[\mathbf{a}^T \mathbf{d}f\right]^2 d\Omega = \mathbf{a}^T \mathbf{T} \mathbf{a} = 0, \;\; \mathbf{a}^T\mathbf{a} > 0 \;\; . \tag{19}$$

The 3×3 tensor $\mathbf{T}$ is the result of the outer product of the vector $\mathbf{d}f$ with itself followed by integration over Ω:

$$\begin{aligned} \mathbf{T} &= \int_\Omega (\mathbf{d}f)(\mathbf{d}f)^T d\Omega \\ &= \int_\Omega \begin{bmatrix} f_{xx}^2 & f_{xx}f_{xy} & f_{xx}f_{yy} \\ f_{xx}f_{xy} & f_{xy}^2 & f_{xy}f_{yy} \\ f_{xx}f_{yy} & f_{xy}f_{yy} & f_{yy}^2 \end{bmatrix} d\Omega. \end{aligned} \tag{20}$$

Alternatively, $f(\mathbf{x})$ might be composed from two oriented signals $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ within Ω which *occlude* each other. Then, in some part Ω_1 of Ω, we have $f(\mathbf{x}) = f_1(\mathbf{x})$, while in its complement Ω_2, $f(\mathbf{x}) = f_2(\mathbf{x})$, whereby $\Omega_1 \cup \Omega_2 = \Omega$ and $\Omega_1 \cap \Omega_2 = \emptyset$. Our model thus is

$$f(\mathbf{x}) = \begin{cases} f_1(\mathbf{x}) & \forall \mathbf{x} \in \Omega_1 \\ f_2(\mathbf{x}) & \forall \mathbf{x} \in \Omega_2 \end{cases} \tag{21}$$

with $\frac{\partial f_1(\mathbf{x})}{\partial \mathbf{u}} = 0 \quad \forall \mathbf{x} \in \Omega_1$ and $\frac{\partial f_2(\mathbf{x})}{\partial \mathbf{v}} = 0 \quad \forall \mathbf{x} \in \Omega_2$. Evidently, condition (16) also holds in this case. However, an alternative constraint can be found to

$$\frac{\partial f(\mathbf{x})}{\partial \mathbf{u}} \cdot \frac{\partial f(\mathbf{x})}{\partial \mathbf{v}} = \mathbf{a}^T \mathbf{d}_o f = 0 \tag{22}$$

where the vector $\mathbf{d}_o f$ is formed from first-order derivatives only according to

$$\mathbf{d}_0 f = (f_x^2, f_x f_y, f_y^2)^T \;\; . \tag{23}$$

Conditions (16) and (22) hold for all $\mathbf{x} \in \Omega$ except for the boundary $\partial\Omega$ between Ω_1 and Ω_2, where they may differ from zero. Neglecting $\partial\Omega$, (22) leads to the tensor

$$\begin{aligned} \mathbf{T}_o &= \int_\Omega (\mathbf{d}_o f)(\mathbf{d}_o f)^T d\Omega \\ &= \int_\Omega \begin{bmatrix} f_x^4 & f_x^3 f_y & f_x^2 f_y^2 \\ f_x^3 f_y & f_x^2 f_y^2 & f_x f_y^3 \\ f_x^2 f_y^2 & f_x f_y^3 & f_y^4 \end{bmatrix} d\Omega \end{aligned} \;\; , \tag{24}$$

which is structurally very similar to $\mathbf{T}$, and replaces $\mathbf{T}$ in (19). The following discussion therefore applies to both $\mathbf{T}$ and $\mathbf{T}_o$ in the same manner.

As for the estimation of single orientations, we cannot expect to find a MOP vector such that $Q_2(\mathbf{a}) = 0$. Minimizing $Q_2(\mathbf{a})$ subject to the constraint $\mathbf{a}^T\mathbf{a} > 0$ leads to

$$\mathbf{T}\mathbf{a} = \lambda_3 \mathbf{a}, \;\; \mathbf{a}^T\mathbf{a} = 1 \;\; . \tag{25}$$

The sought MOP vector is the eigenvector corresponding to the lowest eigenvalue λ_3 of the 3×3 tensor $\mathbf{T}$. The residual error of this solution is equal to λ_3. Confidence in the double orientation hypothesis is high if λ_3 is small and the other two eigenvalues are large. If the image $f(\mathbf{x})$ exhibits two ideal orientations in Ω as defined by (14) and (15), we have $\lambda_3 = 0$, and $\mathrm{rank}(\mathbf{T}) = 2$. Since therefore $\mathbf{T}\mathbf{a} = \lambda_3\mathbf{a} = \mathbf{0}$, $\mathbf{a}$ is a so-called homogeneous vector which can only be determined up

to scale and sign [34], [35]. Homogeneous vectors are elements of a projective space, where two vectors are considered equivalent when they differ only in sign and norm [36], [34]. We may therefore constrain $\mathbf{a}$ to unit length, hence reducing its number of degrees of freedom (DoF) from three to two — a number which is equal to the number of DoF of two orientations in a bivariate image neighborhood. Except for its sign, the normalized MOP vector $\mathbf{a}$ therefore is an unambiguous — albeit implicit — descriptor of two-orientation neighborhoods. We show later how from this descriptor features for, e.g., tracking or classification purposes, can be computed. Explicit orientation analysis, though, requires a decomposition of the MOP vector into the two orientations.

2) MOP vector decomposition: To obtain the orientation angles θ and γ from the MOP vector $\mathbf{a}$, one could observe that $a \pm c = R\cos(\theta \mp \gamma)$, and $b = R\sin(\theta + \gamma)$, where R is an unknown scaling factor. Even though θ and γ are restricted to lie within $(-\pi/2, \pi/2]$, solving these equations for $\theta \pm \gamma$ leads to ambiguities, since the latter lie within $-\pi < \theta \pm \gamma \leq \pi$. On this interval, sine and cosine cannot be unambiguously inverted. Therefore, to decompose $\mathbf{a}$ into the unit vectors $\mathbf{u}$ and $\mathbf{v}$, we define $z_1 = u_x v_y$ and $z_2 = u_y v_x$. Then, the product $z_1 z_2$ is equal to the product of the first and third component of the MOP vector $\mathbf{a}$ in (17), i.e. $z_1 z_2 = ac$. Similarly, we obtain $z_1 + z_2 = u_x v_y + u_y v_x = b$. It follows that z_1 and z_2 are the roots of the polynomial

$$P(z) = (z - z_1)(z - z_2) = z^2 - bz + ac \tag{26}$$

which can be formed from the components of the MOP vector $\mathbf{a}$. The orientation vectors are now obtained by normalizing $(a, z_1)^T$ and $(a, z_2)^T$ to unit length:

$$\begin{aligned} \mathbf{u} &= \frac{(a, z_1)^T}{\sqrt{a^2 + z_1^2}} = \frac{(z_2, c)^T}{\sqrt{z_2^2 + c^2}} \\ \mathbf{v} &= \frac{(a, z_2)^T}{\sqrt{a^2 + z_2^2}} = \frac{(z_1, c)^T}{\sqrt{z_1^2 + c^2}} \end{aligned} \tag{27}$$

Both alternatives to compute $\mathbf{u}$ and $\mathbf{v}$ can equally be applied except when one orientation is horizontal and the other vertical: for $\theta = \pi/2$ and $\gamma = 0$, we have $\mathbf{a} = (0, 1, 0)^T$, $a = c = z_1 = 0$ and $b = z_2 = 1$. Vice versa, for $\theta = 0$ and $\gamma = \pi/2$, $\mathbf{a} = (0, 1, 0)^T$, $z_2 = 0$ and $z_1 = 1$. These cases are easily detected, and can be handled by either selecting the appropriate alternative from (27), or by directly setting $\mathbf{u} = (1, 0)^T$ and $\mathbf{v} = (0, 1)^T$ whenever $\mathbf{a} = (0, 1, 0)^T$. Scaling $\mathbf{a}$ does not change the results of (27) except for the sign of $\mathbf{u}$ and $\mathbf{v}$, what confirms that we may indeed constrain $\mathbf{a}$ to unit length.

Fig. 2 illustrates this approach: Left, a single orientation is estimated based on the double orientation hypothesis. One of the two vectors captures the true orientation. The very small lower eigenvalues λ_2 and λ_3 of $\mathbf{T}$ indicate that the *double* orientation assumption is too complex for the *single* orientation neighborhood. The neighborhood in the center exhibits two ideal orientations, which are both estimated with low error. Compared to the largest eigenvalue λ_1, only λ_3 is very small, thus confirming the double orientation hypothesis. The right-hand side shows the same double orientation structure in noise.

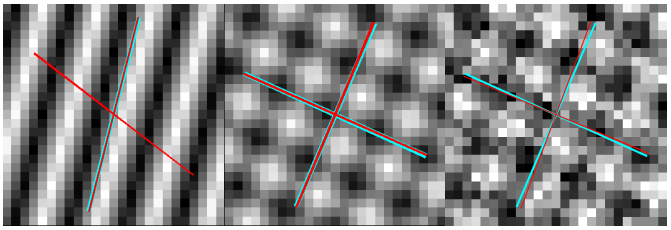

Fig. 2. Neighborhoods as in Fig. 1, with true and estimated orientations shown in light and dark gray, respectively. All orientations are estimated based on the additive double orientation model. Left: Single orientation, noisefree, $\lambda_2/\lambda_1 = 0.0018$, $\lambda_3/\lambda_1 = 0.0004$. Center: noisefree double orientation, angular estimation error $0.46°$ for each orientation, $\lambda_2/\lambda_1 = 0.74$, $\lambda_3/\lambda_1 = 0.00016$. Right: with additive Gaussian noise, SNR 3dB. Estimation errors $1.72°$ and $2.73°$, $\lambda_2/\lambda_1 = 0.64$, $\lambda_3/\lambda_1 = 0.078$.

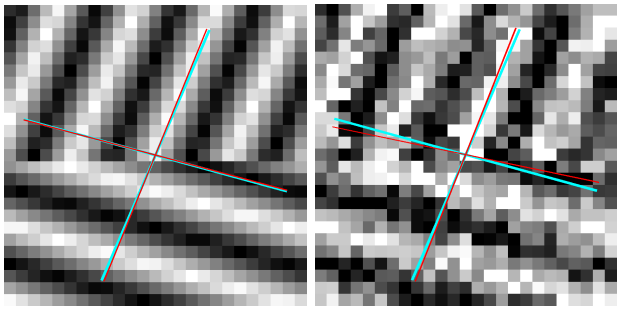

Fig. 3. Neighborhood with two occluding orientations, 25×25 pixel, with true and estimated orientations shown in light and dark gray, respectively. Left: noisefree, estimation errors $0.29°$ and $1.11°$, $\lambda_2/\lambda_1 = 0.78$, $\lambda_3/\lambda_1 = 0.017$. Right: with additive Gaussian noise, SNR 3dB. Estimation errors $1.58°$ and $3.56°$, $\lambda_2/\lambda_1 = 0.69$, $\lambda_3/\lambda_1 = 0.078$.

Both λ_3 and the average estimation error increase, the estimate is, though, still reasonable. Derivatives were calculated as in Fig. 1, with the first derivative operator applied twice to obtain the second order derivatives. (Replacing the simple derivative filters by a Derivative-of-Gaussian with $\sigma = 1.2$ reduced the estimation errors in the noisy neighborhood to $0.73°$ and $0.23°$, and the lowest eigenvalue to $\lambda_3/\lambda_1 = 0.0035$.) Figure 3 shows a similar illustration for two occluding orientations analyzed using constraint (22), where the estimation was based on Derivatives-of-Gaussian filters with $\sigma = 1.2$.

3) Hierarchical Orientation Estimation: Single and double orientation estimation is ideally only possible if the tensors $\mathbf{J}$ and $\mathbf{T}$ (or $\mathbf{T}_o$) exhibit exactly a single zero eigenvalue, respectively. Using the eigenvalues for testing whether or not to accept the single or double orientation hypothesis, however, requires an eigensystem analysis. In this section, we describe a hierarchical algorithm to test for single or double orientations, and show how the eigensystem analysis can be avoided using tensor invariants such as determinant and trace, what, for 2×2 and 3×3 tensors, leads to faster and simpler implementations. We start by developing a test for the presence of a single local orientation. The determinant of $\mathbf{J}$ is $K_1 = \det(\mathbf{J}) = \lambda_1 \lambda_2 = J_{11} J_{22} - J_{12}^2$, and its trace is $H_1 = \mathrm{trace}(\mathbf{J}) = \lambda_1 + \lambda_2 = J_{11} + J_{22}$, where J_{kl} is the element in row k and column l of $\mathbf{J}$. Clearly, $H_1 = 0$ if and only if $\lambda_1 = \lambda_2 = 0$, this indicates a homogeneous neighborhood Ω without oriented structure. Presence of a single ideal orientation in Ω results in $H_1 > 0$, $K_1 = 0$. Furthermore, since always $K_1 \leq \frac{1}{4} H_1^2$, with equality

only for $K_1 = H_1 = 0$, confidence for the single orientation assumption is high if

$$\frac{K_1}{H_1^2} = \frac{\lambda_1\lambda_2}{(\lambda_1+\lambda_2)^2} = \frac{J_{11}J_{22} - J_{12}^2}{(J_{11}+J_{22})^2} < \epsilon_1 \qquad (28)$$

where ϵ_1, $0 < \epsilon_1 \leq \frac{1}{4}$, is a confidence parameter.

In regions where the single orientation model is rejected, we test for double orientations. For the double-orientation tensors $\mathbf{T}$ or $\mathbf{T}_o$, we use the invariants $K_2 = \det(\mathbf{T}) = \lambda_1\lambda_2\lambda_3 = T_{11}M_{11} - T_{12}M_{12} + T_{13}M_{13}$ and $S_2 = \lambda_1\lambda_2 + \lambda_2\lambda_3 + \lambda_1\lambda_3 = M_{11} + M_{22} + M_{33}$, where M_{ij} is the minor [37] obtained from $\mathbf{T}$ by deleting row i and column j. Since we already tested for a single orientation based on $\mathbf{J}$, only one eigenvalue of $\mathbf{T}$ can be zero. Since $K_2^2 \leq S_2^3$, a suitable confidence criterion for the double-orientation assumption therefore is

$$\frac{K_2^2}{S_2^3} < \epsilon_2 \quad , \qquad (29)$$

where ϵ_2, $0 < \epsilon_2 < (\frac{1}{3})^3 = \frac{1}{27}$, is another confidence parameter. All quantities needed for these tests can be computed directly from $\mathbf{J}$ and $\mathbf{T}$.

If either one of the tensors $\mathbf{J}$ and $\mathbf{T}$ is thus found to possess exactly one (close to) zero eigenvalue, the corresponding eigenvector can also be found directly: an eigenvector $\mathbf{v}$ such that $\mathbf{J}\mathbf{v} = \mathbf{0}$ as in (7) is evidently orthogonal to both rows of $\mathbf{J}$. Furthermore, since $\mathrm{rank}(\mathbf{J}) = 1$, both rows are linear dependent. The sought eigenvector $\mathbf{v}$ is thus given by a unit vector orthogonal to the first row of $\mathbf{J}$, i.e. by

$$\mathbf{v} = \frac{(-J_{12}, J_{11})^T}{\sqrt{J_{12}^2 + J_{22}^2}} \qquad (30)$$

Similarly, for the double-orientation tensor, condition (25) means we seek the eigenvector orthogonal to all rows of $\mathbf{T}$, which are linear dependent. At least one pair of rows, however, must be linear independent. If the double orientation assumption holds, the MOP vector $\mathbf{a}$ can thus be calculated as the normalized cross product of two linear independent rows of $\mathbf{T}$ [1]. An even more robust estimate of the MOP vector could be obtained as the average of the eigenvectors computed from two (or more) different pairs of linear independent rows of $\mathbf{T}$.

B. Multivariate and vector valued image data

Extension of the described approach from scalar images towards multispectral images $\mathbf{f} : \mathbb{R}^2 \to \mathbb{R}^q$ is straightforward [40]: as in (13), it suffices to replace the double orientation tensor $\mathbf{T}$ by the sum $\mathbf{T}_q = \sum_{k=1}^{q} \mathbf{T}(f_k)$ of tensors computed for each spectral component f_k. Similarly, the tensor $\mathbf{T}_o$ in (24) extends to [29]

$$\mathbf{T}_o = \int_\Omega \begin{bmatrix} |\mathbf{f}_x|^4 & |\mathbf{f}_x|^2\mathbf{f}_x^T\mathbf{f}_y & |\mathbf{f}_x^T\mathbf{f}_y|^2 \\ |\mathbf{f}_x|^2\mathbf{f}_x^T\mathbf{f}_y & \frac{1}{2}(|\mathbf{f}_x|^2|\mathbf{f}_y|^2 + |\mathbf{f}_x^T\mathbf{f}_y|^2) & |\mathbf{f}_y|^2\mathbf{f}_x^T\mathbf{f}_y \\ |\mathbf{f}_x^T\mathbf{f}_y|^2 & |\mathbf{f}_y|^2\mathbf{f}_x^T\mathbf{f}_y & |\mathbf{f}_y|^4 \end{bmatrix} d\Omega \qquad (31)$$

[1] More generally, it can be shown that for a matrix with only one zero eigenvalue, the eigenvector corresponding to the zero eigenvalue can be found directly from the minors of that matrix [38], [39].

For tri- and higher-variate signals, however, ambiguities may arise which require modifications especially in the decomposition of the MOP vector into the individual orientations. We therefore consider now the p-variate mapping $f : \mathbb{R}^p \to \mathbb{R}$, which, as in (14), is assumed to be composed from two single-oriented subimages $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ by $f(\mathbf{x}) = f_1(\mathbf{x}) + f_2(\mathbf{x})$, $f_1(\mathbf{x}), f_2(\mathbf{x}) : \mathbb{R}^p \to \mathbb{R}$. In bivariate images, the subimages $f_i(\mathbf{x})$ can only be oriented along lines. For trivariate data, we may have $f(\mathbf{x}) = g_1(\mathbf{n}_1^T\mathbf{x}) + g(\mathbf{n}_2^T\mathbf{x}) = g(\mathbf{n}_1^T\mathbf{x}, \mathbf{n}_2^T\mathbf{x})$, what can be regarded as a single signal oriented along a line orthogonal to both (unit) vectors $\mathbf{n}_1$ and $\mathbf{n}_2$, or as the superposition of two signals oriented along planes with normals $\mathbf{n}_1$ and $\mathbf{n}_2$, respectively. We focus in the following on orientation along lines. In other words, if $f_1(\mathbf{x})$ is oriented along $\mathbf{u}$ and $f_2(\mathbf{x})$ along $\mathbf{v}$ according to (8), with $\mathbf{u} \in E_1 \subset \mathbb{R}^p$ and $\mathbf{v} \in E_2 \subset \mathbb{R}^p$, we have $\dim(E_1) = \dim(E_2) = 1$. The subimages then obey $\partial f_1(\mathbf{x})/\partial\mathbf{u} = \partial f_2(\mathbf{x})/\partial\mathbf{v} = 0$, and $f(\mathbf{x})$ satisfies $\partial^2 f(\mathbf{x})/(\partial\mathbf{u}\partial\mathbf{v}) = 0$. With the notation $\partial_i = \partial/\partial x_i$ for the partial derivative operators, and $f_{ij} = \partial_i\partial_j f$, the constraint expands to

$$\begin{aligned} \frac{\partial^2 f(\mathbf{x})}{\partial\mathbf{u}\partial\mathbf{v}} &= \left(\sum_{i=1}^{p} u_i\partial_i \cdot \sum_{j=1}^{p} v_j\partial_j\right) f(\mathbf{x}) \qquad (32) \\ &= \sum_{i=1}^{p}\sum_{j=1}^{i} a_{ij} f_{ij} = 0 \end{aligned}$$

since $f_{ij} = f_{ji}$. The right-hand side exhibits $k = p(p+1)/2$ components. The mixed-orientation parameters a_{ij} depend on the components u_i and v_j of $\mathbf{u}$ and $\mathbf{v}$ according to

$$a_{ij} = \begin{cases} u_j v_j & \text{for } i = j \\ u_i v_j + u_j v_i & \text{else} \end{cases} \qquad (33)$$

what is consistent with (17) for $p = 2$. Ordering the MOPs a_{ij} into the k-dimensional MOP vector

$$\mathbf{a} = \mathbf{a}(\mathbf{u}, \mathbf{v}) = [a_{ij}]_{j\leq i}^T \qquad (34)$$

and with the vector $\mathbf{d}f = [f_{ij}]_{j\leq i}^T$ of second derivatives, (32) becomes

$$\mathbf{a}^T\mathbf{d}f = 0 \;\; \forall\, \mathbf{x} \in \Omega \quad . \qquad (35)$$

As before, $\mathbf{a}$ is a homogeneous vector which we may constrain to $\mathbf{a}^T\mathbf{a} = 1$. Furthermore, since we focus on orientations along lines, i.e., $\dim(E_1) = \dim(E_2) = 1$, any two vectors $\mathbf{u}_0,\ \mathbf{u}_1 \in E_1$ are linear dependent. The same holds for any two vectors $\mathbf{v}_0,\ \mathbf{v}_1 \in E_2$. Thus, $\mathbf{u}_1 = c_u\mathbf{u}_0$ and $\mathbf{v}_1 = c_v\mathbf{v}_0$ hold, with $c_u,\ c_v \in \mathbb{R}$. This implies that any two MOP vectors $\mathbf{a}_0 = \mathbf{a}(\mathbf{u}_0, \mathbf{v}_0)$ and $\mathbf{a}_1 = \mathbf{a}(\mathbf{u}_1, \mathbf{v}_1)$ out of the set $\mathcal{A} = \{\mathbf{a}(\mathbf{u}, \mathbf{v}) : \mathbf{u} \in E_1, \mathbf{v} \in E_2\}$ are also linear dependent, since $\mathbf{a}_1 = \mathbf{a}(\mathbf{u}_1, \mathbf{v}_1) = \mathbf{a}(c_u\mathbf{u}_0, c_v\mathbf{v}_0) = c_u \cdot c_v \cdot \mathbf{a}_0$. In other words, the subspace spanning the set of admissible solutions for $\mathbf{a}(\mathbf{u}, \mathbf{v})$ according to (33) is also onedimensional.

With the symmetric $p \times p$ tensor

$$\mathbf{T} = \int_\Omega \mathbf{d}f \otimes \mathbf{d}f \; d\Omega \qquad (36)$$

the MOP vector $\mathbf{a}$ satisfies $\mathbf{a}^T\mathbf{T}\mathbf{a} = 0$, $\mathbf{a}^T\mathbf{a} = 1$. Similarly, the alternative constraint (22) for occluding orientations

expands to

$$\frac{\partial f(\mathbf{x})}{\partial \mathbf{u}} \cdot \frac{\partial f(\mathbf{x})}{\partial \mathbf{v}} = \sum_{i=1}^{p} \sum_{j=1}^{i} a_{ij} f_i f_j = 0 \qquad (37)$$

since $f_i f_j = f_j f_i$. This can be written as $\mathbf{a}^T \mathbf{d}_o f = 0$, where the vector $\mathbf{d}_o f$ is constructed from first-order derivatives by $\mathbf{d}_o f = [f_i f_j]_{j \le i}^T$, yielding the tensor $\mathbf{T}_o = \int_\Omega \mathbf{d}_o f \otimes \mathbf{d}_o f d\Omega$.

Let us now examine the relation between the MOP vector $\mathbf{a}(\mathbf{u}, \mathbf{v})$ and the sought orientations more closely. As a k-dimensional unit vector, the MOP vector has $k - 1 = p(p + 1)/2 - 1$ DoF, whereas two onedimensional orientations in p-variate signals exhibit $2(p - 1)$ DoF. Only for $p = 2$ (and $p = 1$, i.e. univariate signals where no orientation can be defined), these numbers are identical. For trivariate signals, $\mathbf{a}$ has 5 DoF, while we seek two threedimensional unit vectors with together only 4 DoF. This discrepancy increases with increasing p. For $p > 2$, solving

$$\hat{\mathbf{a}} = \arg\min_{\mathbf{s}} \left(\mathbf{s}^T \mathbf{T} \mathbf{s} + \lambda (1 - \mathbf{s}^T \mathbf{s}) \right) \qquad (38)$$

will thus not necessarily lead to a vector $\hat{\mathbf{a}}$ which complies with (33) and (34). However, from solutions $\mathbf{a}$ consistent with (33), the symmetrized tensor $\mathbf{A} = \frac{1}{2}(\mathbf{u}\mathbf{v}^T + \mathbf{v}\mathbf{u}^T)$ can be constructed by setting $[\mathbf{A}]_{jj} = a_{jj}$ and $[\mathbf{A}]_{ij} = \frac{1}{2} a_{ij}$, $i \neq j$, $i, j = 1, \ldots, p$. Thus, $\mathbf{A} \in \mathbb{R}^{p \times p}$ holds. The following theorem is proven in the appendix:

Theorem 1 *Let $\mathbf{u}, \mathbf{v}$ be unit vectors with angle β between them. Then $\mathbf{A} = \frac{1}{2}(\mathbf{u}\mathbf{v}^T + \mathbf{v}\mathbf{u}^T)$ is of rank two, and its two non-zero eigenvalues are*

$$\begin{aligned} \sigma_+ &= \frac{1}{2}(1 + \mathbf{u}^T \mathbf{v}) = \cos^2 \frac{\beta}{2} \geq 0 \qquad (39) \\ \sigma_- &= \frac{1}{2}(\mathbf{u}^T \mathbf{v} - 1) = -\sin^2 \frac{\beta}{2} \leq 0 \ . \end{aligned}$$

The corresponding eigenvectors are

$$\mathbf{e}_+ = \frac{\mathbf{u} + \mathbf{v}}{n_+} \ , \mathbf{e}_- = \frac{\mathbf{u} - \mathbf{v}}{n_-} \ , \qquad (40)$$

where n_+ and n_- are scaling factors.

Instead of computing the eigensystem of the matrix $\mathbf{A}$, we can alternatively analyze the coefficients c_i, $i = 0, \ldots, p$ of its characteristic polynomial

$$\begin{aligned} \pi(\sigma) &= \det(\sigma \mathbf{I} - \mathbf{A}) \\ &= \sum_{i=0}^{p} c_i \sigma^{p-i} = \prod_{i=1}^{p} (\sigma - \sigma_i) \ , \qquad (41) \end{aligned}$$

where $\mathbf{I}$ is the identity matrix and σ_i, $i = 1, \ldots, p$, denote the eigenvalues of $\mathbf{A}$. Expanding the lower row of (41) shows that the coefficients c_i are symmetric polynomials of the eigenvalues. If $\mathbf{A}$ is formed as in theorem 1, we have $\sigma_1 = \sigma_+$, $\sigma_p = \sigma_-$, and, for $p > 2$, $\sigma_2, \ldots, \sigma_{p-1} = 0$. For the coefficients c_i follows

$$\begin{aligned} c_1 &= -\sum_{i=1}^{p} \sigma_i = -(\sigma_+ + \sigma_-) = -\text{trace}(\mathbf{A}) = -\cos\beta \\ c_2 &= \sum_{i \neq j} \sigma_i \sigma_j = \sigma_+ \sigma_- = -\frac{\sin^2 \beta}{4} \\ c_i &= 0 \text{ for } i > 2 \qquad (42) \end{aligned}$$

The last equality holds since it consists of sums over products of three and more different eigenvalues, of which at least one is always equal to zero. Thus, we can conclude that $\mathbf{a}$ is consistent with (33) if and only if the coefficients of the characteristic polynomial of the matrix $\mathbf{A}$ formed as in theorem 1 satisfy $c_2 < 0$ and $c_i = 0$, $i > 2$ [40]. Since the c_i can be computed analytically from the entries of $\mathbf{A}$, these can be tested without an eigensystem analysis of $\mathbf{A}$.

1) MOP Vector Decomposition: As a rank-two matrix, $\mathbf{A}$ can thus alternatively be composed from its principal components by $\mathbf{A} = \sigma_+ \mathbf{e}_+ \mathbf{e}_+^T + \sigma_- \mathbf{e}_- \mathbf{e}_-^T$. With (40) and (39), this leads to $n_+ = 2\sqrt{\sigma_+}$ and $n_- = 2\sqrt{-\sigma_-}$ for the normalization factors. From (40), the sought orientation vectors can then be recovered by

$$\begin{aligned} \mathbf{u} &= \sqrt{\sigma_+} \mathbf{e}_+ + \sqrt{-\sigma_-} \mathbf{e}_- \\ \mathbf{v} &= \sqrt{\sigma_+} \mathbf{e}_+ - \sqrt{-\sigma_-} \mathbf{e}_- \qquad (43) \end{aligned}$$

and subsequent normalization. Double orientations in general p-variate image data can thus be estimated as follows:

- Calculate $\mathbf{T}$ from the second-order derivatives over Ω.
- Estimate the MOP vector $\hat{\mathbf{a}}$ as in (38).
- Form the tensor $\mathbf{A}$ from the entries of $\hat{\mathbf{a}}$.
- Perform an eigensystem analysis of $\mathbf{A}$. If $f(\mathbf{x})$ is consistent with the double-orientation hypothesis (up to small deviations and noise), $\mathbf{A}$ will exhibit $p - 2$ eigenvalues close to zero between σ_+ and σ_-. Ignoring these in the subsequent steps is equivalent to approximating $\mathbf{A}$ by the closest rank(2)-matrix in terms of the Frobenius-norm. Alternatively, test the coefficients c_i for $c_2 < 0$ and $c_i = 0$ for $i > 2$.
- Recover the sought orientations from the eigenvectors $\mathbf{e}_+$, $\mathbf{e}_-$ of $\mathbf{A}$ by transform (43).

Based on the above derivation, we state in the following which invariant features are encoded in the MOP vector $\mathbf{a}$.

C. Invariant Features

A rotation $\mathbf{P}$ applied to $f(\mathbf{x})$ locally within Ω leads to the transformed vectors $\tilde{\mathbf{u}} = \mathbf{P}\mathbf{u}$, $\tilde{\mathbf{v}} = \mathbf{P}\mathbf{v}$. The tensor $\mathbf{A}$ then transforms to $\tilde{\mathbf{A}} = \mathbf{P}\mathbf{A}\mathbf{P}^T$. Since, for a rotation, $\mathbf{P}^T = \mathbf{P}^{-1}$, $\tilde{\mathbf{A}}$ is obtained from $\mathbf{A}$ by a (rigid) similarity transformation, and therefore has the *same* eigenvalues as $\mathbf{A}$. All scalar invariants under local rigid transformations are thus generated by σ_+ and σ_- or, in other words, by the angle β between the orientations. As shown in the appendix, the cosine of β, for instance, is given by

$$\cos\beta = \frac{\sigma_+ + \sigma_-}{\sigma_+ - \sigma_-} \ . \qquad (44)$$

Invariant features for double-orientation neighborhoods, for instance for tracking purposes, can hence be found from the MOP vector *without* explicitly computing the orientation vectors. Even the eigensystem analysis of $\mathbf{A}$ can be avoided: For bivariate image data, we can compute $\cos\beta$ analytically from the MOP vector in (17) according to

$$\cos\beta = \frac{a + c}{\sqrt{(a - c)^2 + b^2}} \ . \qquad (45)$$

For tri- and highervariate data, we exploit the fact that, being formed from the eigenvalues according to (42), the coefficients c_i are also invariants. Since $c_1^2 = \cos^2\beta$ and $4c_2 = -\sin^2\beta$, we have as an alternative invariant feature

$$\tan^2\beta = -\frac{4c_2}{c_1^2} \tag{46}$$

where c_1^2 is the squared trace of $\mathbf{A}$, while c_2 is a linear combination of its 2×2 minors. Again, these quantities can be computed analytically from the entries of $\mathbf{A}$.

IV. Signal Separation

In both additive and occluding superpositions, the individual signals $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ can be separated from each other once the orientations $\mathbf{u}$ and $\mathbf{v}$ are known. To this end, we apply the directional derivatives along θ and γ, respectively, to the composite signal $f(\mathbf{x})$. With (15), this reduces one signal to zero. For the additive case, we thus obtain

$$\begin{aligned} \alpha(\theta)f(\mathbf{x}) &= \frac{\partial f(\mathbf{x})}{\partial \mathbf{u}} = \frac{\partial f_2(\mathbf{x})}{\partial \mathbf{u}} = \alpha(\theta)f_2(\mathbf{x}) \\ \alpha(\gamma)f(\mathbf{x}) &= \frac{\partial f(\mathbf{x})}{\partial \mathbf{v}} = \frac{\partial f_1(\mathbf{x})}{\partial \mathbf{v}} = \alpha(\gamma)f_1(\mathbf{x}) \ , \end{aligned} \tag{47}$$

i.e. filtered versions of $f_2(\mathbf{x})$ and $f_1(\mathbf{x})$ extending over all of Ω. For occlusions, we get

$$\alpha(\theta)f(\mathbf{x}) = \begin{cases} 0 & \text{for} \quad \mathbf{x} \in \Omega_1 \\ \alpha(\theta)f_2(\mathbf{x}) & \text{for} \quad \mathbf{x} \in \Omega_2 \end{cases} \tag{48}$$

and

$$\alpha(\gamma)f(\mathbf{x}) = \begin{cases} \alpha(\gamma)f_1(\mathbf{x}) & \text{for} \quad \mathbf{x} \in \Omega_1 \\ 0 & \text{for} \quad \mathbf{x} \in \Omega_2 \end{cases} \ . \tag{49}$$

Evaluating expressions (48) and (49) allows us to estimate the subregions Ω_1 and Ω_2. Calculating $|\partial f/\partial \mathbf{u}|^2$ and $|\partial f/\partial \mathbf{v}|^2$ $\forall \mathbf{x} \in \Omega$, where in practice we average over a 3×3 neighborhood around each pixel $\mathbf{x}$, $\mathbf{u}$ and $\mathbf{v}$ are assigned as follows:

$$\begin{aligned} &\text{if} \quad |\tfrac{\partial f}{\partial \mathbf{u}}|^2 < |\tfrac{\partial f}{\partial \mathbf{v}}|^2 \quad \text{decide } \mathbf{x} \in \Omega_1 \ , \ \text{orientation } \mathbf{u} \\ &\text{if} \quad |\tfrac{\partial f}{\partial \mathbf{v}}|^2 < |\tfrac{\partial f}{\partial \mathbf{u}}|^2 \quad \text{decide } \mathbf{x} \in \Omega_2 \ , \ \text{orientation } \mathbf{v} \end{aligned} \tag{50}$$

V. Results

Fig. 4 shows the additive superposition of a diagonal and a circular sine pattern in additive Gaussian (pseudo)noise with variance $\sigma^2 = 100$ (PSNR 28dB, image size 256×256 pixel). The upper right image depicts the result of double orientation estimation. First, the MOP vector $\mathbf{a}$ was estimated for each pixel as described in section III-A.1, which was then decomposed into the orientations as discussed in section III-A.2. The size of the sliding local neighborhood Ω was 13×13 pixel. The derivatives f_x and f_y were calculated by convolving with the finite-difference kernel $[1 \quad 0 \quad -1]$ in horizontal resp. vertical direction combined with simultaneous orthogonal smoothing by the kernel $[1 \ 3 \ 1]$. The orientation fields are subsampled and depicted as black arrows. Evidently, the estimated orientations correspond well to the underlying true ones. This observation is confirmed by the lower row of Fig. 4, which shows the separation between the two overlaid orientation patterns by directional derivative filtering according to (47). In both cases, one of the patterns vanishes almost completely as expected, while a linearly filtered version of the other one is retained. It should be noted, however, that the estimation algorithm only provides the two locally estimated orientations $\mathbf{u}$ and $\mathbf{v}$ for each pixel $\mathbf{x}$, without sorting these into the global "diagonal" and "circular" patterns. Before directional filtering along each one of the estimated orientations, we therefore sorted these globally by first computing the mean vectors for both $\mathbf{u}$ and $\mathbf{v}$ over the entire image plane, tacitly assuming that each pattern is represented by one of the mean vectors. For each pixel $\mathbf{x}$, one of the orientation vectors was then assigned to the pattern with mean vector closest to it, while the other orientation vector was assigned to the other pattern. This straightforward procedure worked for all examples given here. Alternatively, the global sorting of the orientation fields could be carried out by, e.g., k-means clustering with $k = 2$, or by optimizing a homogeneity criterion [41], [42] of the vector fields. Results for images of real textures are given in Figs. 5 and 6. Fig. 7 shows orientation estimation results for a part of an image of an X-rayed tire, which reveals its internal metal gratings. The number of gratings and hence the number of orientations varies over the image plane, and is a priori unknown for each pixel. The number of orientations was therefore detected according to (28) and (29), with $\epsilon_1 = 0.025$ and $\epsilon_2 = 0.01$, and 19×19 pixel for Ω. The detection results are shown gray-level coded in the left image of the lower row, distinguishing between areas with single (dark-gray), double (light-gray) and more than two orientations (white). For the single and double orientation regions, the estimated orientations are shown in the upper right image. The results of nulling out one of two the metal gratings in the double orientation regions by directional derivative filtering are given in the middle and right image of the lower row.

Fig. 8 shows an image of two differently oriented corrugated cardboards occluding each other. The occluding boundary was detected invoking (28) with $\epsilon_1 = 0.025$. Around the boundary, double orientations were then estimated using the tensor $\mathbf{T}_o$ in (24), resulting in two orientations for each pixel in the vicinity of the boundary, as shown in the middle of Fig. 8. The estimated orientations were then separated according to (50), retaining one orientation vector for each pixel, as also shown in Fig. 8.

Finally, Fig. 9 illustrates the invariance properties of the difference angle $\cos\beta$ as calculated directly from the MOP vector $\mathbf{a}$ by (44), i.e., *without* explicit orientation computation. The upper row shows rotated versions of two doubly oriented sine patterns in noise with $\beta = -45°$ and $\beta = -50°$, respectively. The values of $\cos\beta$ estimated for each pixel according to (44) are shown in the gray level plots of the middle row. Evidently, there is a slight sensitivity to noise. Still, even though the true values of $\cos\beta$ between left and right image pair are very close to each other, the difference between them is well captured. This observation is confirmed by the lineplot in Fig. 9, which moreover proves that the rotation has indeed almost no effect on the estimated values. The lineplot also reveals that the estimates exhibit a slight bias caused by effects of differentiation on the discrete image

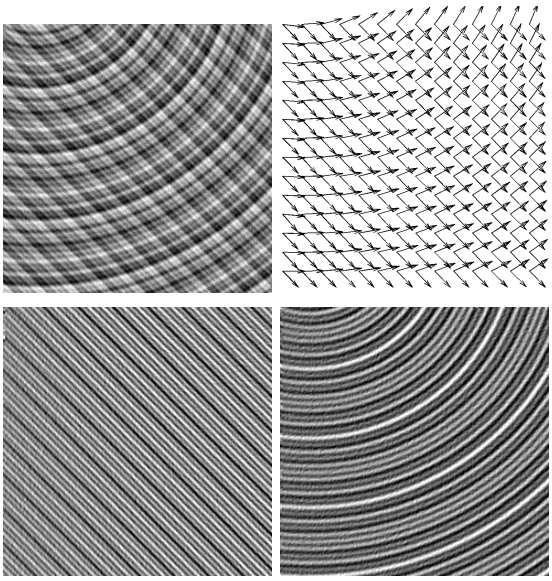

Fig. 4. Synthetic additive double orientation pattern in noise (PSNR 28dB). Upper row: Original image (left) and estimated orientation fields (right). Lower row: Result of local directional derivative filtering along each one of the estimated orientations.

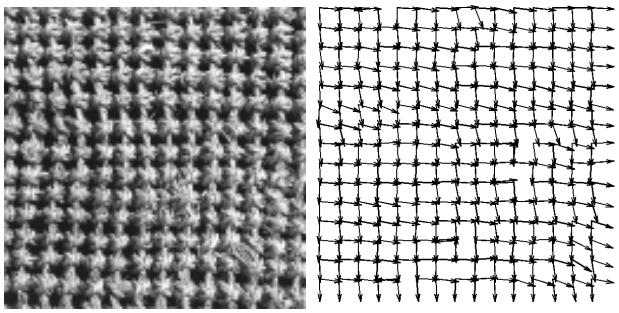

Fig. 5. Double orientation estimation applied to a texture image.

grids. In fact, the mean of the estimated values of $\cos\beta$ for the left image pair is 0.715 (true value: $\cos 45^\circ = 0.71$), and for the right image pair 0.66 (true value: $\cos 50^\circ = 0.64$). This corresponds to a systematic error for β of about 1°.

VI. Discussion

We have described a framework for the detection and estimation of superimposed double orientations. The first step is always the formation of a symmetric positive semi-definite tensor from image derivatives. Its eigensystem analysis provides the mixed orientation parameter vector, which encodes the sought orientations. For bivariate image data, the MOP vector can be regarded as an unambiguous albeit implicit descriptor of double-orientation neighborhoods, from which the orientations can be obtained by a closed-form approach. For higher-variate data, the MOP vector becomes more and more overparametrized. Its decomposition then requires another eigensystem analysis of an additional matrix formed from the MOP vector components. These considerations also led to a new set of rotation-invariant features for the description of double orientation neighborhoods in bi- and higher-variate image data, which can be calculated without explicit

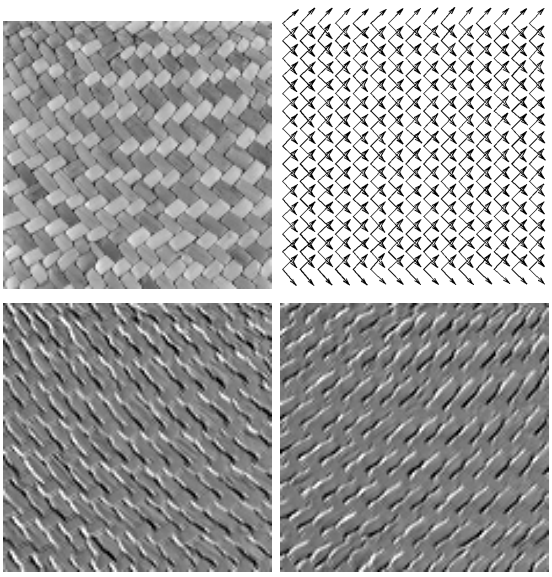

Fig. 6. Double orientation estimation applied to a texture image. Upper row: Original (left) and estimated orientations (right). Lower row: Result of local directional derivative filtering along each one of the estimated orientations.

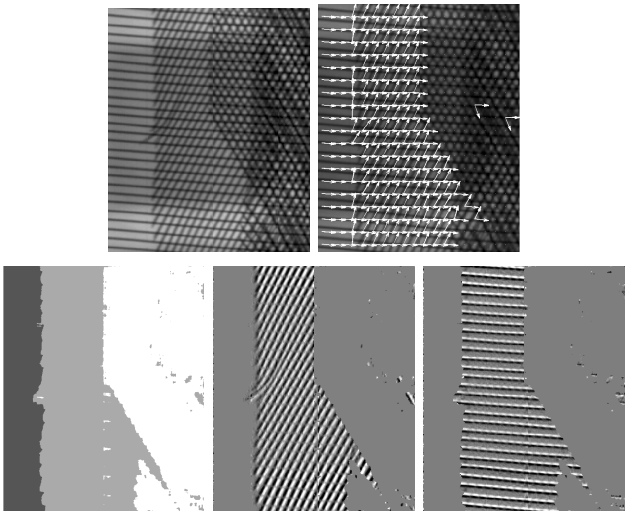

Fig. 7. Part of an X-ray image of a tire showing its internal metal grating structure. Upper row: original image (529×639 pixel, left) and estimated single and double orientations superimposed onto the original (right). Lower row: detected orientation structure (left), and results of nulling out one of the gratings in the double orientation regions (middle and right).

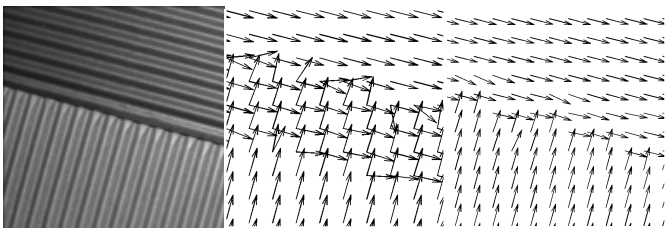

Fig. 8. Left: Two corrugated cardboards occluding each other (128×128 pixel). Middle: Result of orientation detection and estimation, yielding two orientations for each pixel in the vicinity of the occluding boundary. Right: Result after separating the double orientations as described in section IV, retaining one orientation for each pixel (these results are shown enlarged for better perceptibility).

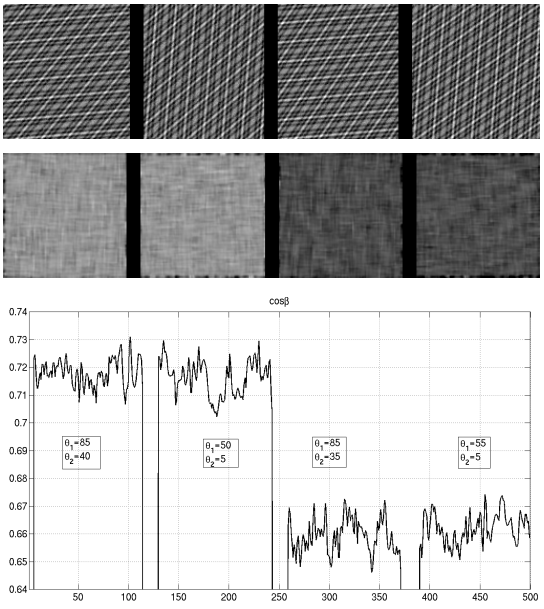

Fig. 9. Upper row: Additively superimposed sine patterns in white Gaussian (pseudo)noise, SNR 28dB, 128×128 pixel. Orientations: $\theta = 5°$, $\gamma = 50°$ (left), $\theta = 40°$, $\gamma = 85°$ (middle left), $\theta = 5°$, $\gamma = 55°$ (middle right), $\theta = 35°$, $\gamma = 85°$ (right). The cosine of the true difference angle $\beta = \theta - \gamma$ for the left image pair is $\cos\beta = \cos 45° = 0.71$, and for the right image pair $\cos\beta = \cos 50° = 0.64$. Middle row: Estimated values of $\cos\beta$ shown as gray level plot scaled to $[0.6, 0.8]$. The lineplot shows the values of $\cos\beta$ along row 64.

orientation analysis, and even without eigensystem analyses. Being implicitly based on orientations, however, such features may help to solve tasks such as tracking more robustly, since the relative orientations of a rigid junction will change less than the appearance and other features of the junction. Assessing the suitability of such features for, e.g., tracking purposes will be a future research topic.

Furthermore, we have developed tests for the detection of flat, single and double orientation neighborhoods. We have shown that the quantities needed in these tests can be computed directly from the tensors involved, without the need to actually perform the eigensystem analyses. Once the number of orientations is known, we have shown that the MOP vector can also be calculated without explicit eigensystem analysis as long as the corresponding tensor is (approximately) rank deficient.

For additive superpositions, the tensor consisted of second-order derivatives. We showed that the same tensor can also applied to the analysis of oriented patterns which occlude each other. Alternatively, in this case a structurally similar tensor formed from products of first-order derivatives of the image signal can be used. We have here computed this latter tensor at the original image resolution. As shown in [23], an improvement in the accuracy of the tensor entries can be achieved by first upsampling the image signal before computing the tensor, since the multiplications correspond to convolutions in the frequency domain, which extend the bandwidth beyond the original Nyquist frequency. Furthermore, our experiments on textured images confirm that the framework can successfully be applied to data which comply with neither the additive superposition hypothesis nor the occlusion hypothesis perfectly.

For both additively superimposed and occluding orientation patterns, a separation of the patterns can be accomplished by directional filtering along one of the estimated orientations, thus nulling out the corresponding subsignal while retaining a filtered version of the other subsignal. In the case of occluding patterns, this operation also permits to estimate the subregions Ω_1 and Ω_2 occupied by each subsignal $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$.

Our experiments show that even simple derivative filters perform reasonably well. Nonetheless, future work will, e.g., evaluate the accuracy and noise resistance of other, optimized filters, such as in [43], [44]. Another topic for future research is to improve the segmentation into regions with different number of orientations by adaptive thresholding, for instance based on a Markov random field model [45].

Finally, the extension towards estimation of more than two orientations will be of interest. This would, for instance, allow to deal with the problem of generally non-vanishing derivatives along the boundary between two oriented patterns occluding each other, by introducing a third directional derivative along the occluding boundary. Besides allowing the analysis of three or more superimposed patterns, we expect that this will also lead to additional invariant features. We believe, however, that this extension is more challenging than it appears at a first glance. First, it will require the computation of derivatives of order three and higher in the tensor $\mathbf{T}$, which may become impractical. An interesting question in this respect is whether the tensor $\mathbf{T}_o$ for the analysis of occluding patterns, which utilizes only first-order derivatives, will be more stable in this respect. Secondly, for three and more orientations in bivariate images, the quadratic polynomial (26) will be replaced by polynomials of degrees three and higher. Closed-form solutions for up to four orientations thus exist, as shown for multiple motions in [39], or, for multiple subspace estimation, in [46]. In the general case of three or more orientations in tri- and higher-variate data, the matrices $\mathbf{T}$ in (20) or $\mathbf{T}_o$ in (24) will extend to three- and higher-dimensional tensors, to which the discussed framework, which is heavily based on matrix algebra, may not be directly applicable anymore.

VII. Appendix

To prove Theorem 1, we multiply $\mathbf{A}$ by $\mathbf{e}_+$, obtaining

$$\begin{aligned} \mathbf{A}\mathbf{e}_+ &= \frac{1}{2}(\mathbf{u}\mathbf{v}^T + \mathbf{v}\mathbf{u}^T) \cdot \frac{\mathbf{u}+\mathbf{v}}{n_+} \\ &= \frac{1}{2}(1 + \mathbf{u}^T\mathbf{v})\frac{\mathbf{u}+\mathbf{v}}{n_+} = \sigma_+\mathbf{e}_+ \end{aligned} \tag{51}$$

and similarly for σ_- and $\mathbf{e}_-$. Setting $\mathbf{u}^T\mathbf{v} = \cos\beta$ concludes the proof. Any other eigenvalue must be zero, since its eigenvector $\mathbf{e}$ must be orthogonal to both $\mathbf{e}_+$ and $\mathbf{e}_-$. From $\mathbf{e}_+^T\mathbf{e} = (\mathbf{u}+\mathbf{v})^T\mathbf{e} = 0$ and $\mathbf{e}_-^T\mathbf{e} = (\mathbf{u}-\mathbf{v})^T\mathbf{e} = 0$ follows that $\mathbf{e}$ is also orthogonal to both $\mathbf{u}$ and $\mathbf{v}$, i.e., $\mathbf{u}^T\mathbf{e} = \mathbf{v}^T\mathbf{e} = 0$. Hence, $\mathbf{A}\mathbf{e} = \mathbf{0}$. Thus, $\mathbf{A}$ has only two non-zero eigenvalues, and its rank is two.

To calculate $\cos\beta$ directly from σ_+ and σ_-, we have to bear in mind that the matrix $\mathbf{A}$ is formed from the components of the MOP vector $\mathbf{a}$, which, being a homogeneous vector, is

only determined up to an unknown scale factor R. The same holds thus for $\mathbf{A}$ as well as for σ_+ and σ_-. We therefore have $\sigma_+ + \sigma_- = R \cdot \cos\beta$ and $\sigma_+ - \sigma_- = R$, from which (44) follows.

REFERENCES

[1] T. Aach and D. Kunz, "Anisotropic spectral magnitude estimation filters for noise reduction and image enhancement," in *Proceedings ICIP-96*. Lausanne, Switzerland: IEEE, September 16–19 1996, pp. 335–338.

[2] ——, "A lapped directional transform for spectral image analysis and its application to restoration and enhancement," *Signal Processing*, vol. 80, no. 11, pp. 2347–2364, 2000.

[3] W. Freeman and E. Adelson, "The design and use of steerable filters," *IEEE Trans. PAMI*, vol. 13, no. 9, pp. 891–906, 1991.

[4] M. Jacob and M. Unser, "Design of steerable filters for feature detection using Canny-like criteria," *IEEE Trans. PAMI*, vol. 26, no. 8, pp. 1007–1019, 2004.

[5] J. Hladuvka and E. Gröller, "Direction-driven shape-based interpolation of volume data," in *Proceedings Vision, Modeling and Visualization*. Stuttgart: November 21–23, 2001, pp. 113–120, 521.

[6] M. Bertram, G. Rose, D. Schäfer, J. Wiegert, and T. Aach, "Directional interpolation of sparsely sampled cone-beam CT sinogram data," in *Proceedings IEEE International Symposium on Biomedical Imaging (ISBI)*. Arlington, VA: IEEE, April 15–18 2004, pp. 928–931.

[7] J. Bigün, G. H. Granlund, and J.Wiklund, "Multidimensional orientation estimation with applications to texture analysis and optical flow," *IEEE Trans. PAMI*, vol. 13, no. 8, pp. 775–790, 1991.

[8] J. Bigün, "Frequency and orientation sensitive texture measures using linear symmetry," *Signal Processing*, vol. 29, pp. 1–16, 1992.

[9] J. D. Costa, F. L. Pouliquen, C. Germain, and P. Baylou, "New operators for optimized orientation estimation," in *Proc. International Conference on Image Processing*, vol. 3, Oct. 2001, pp. 744–747.

[10] F. L. Pouliquen, J. D. Costa, C. Germain, and P. Baylou, "A new adaptive framework for unbiased orientation estimation in textured images," *Pattern Recognition*, vol. 38, pp. 2032–2046, 2005.

[11] G. H. Granlund and H. Knutsson, *Signal Processing for Computer Vision*. Dordrecht: Kluwer, 1995.

[12] E. Barth, T. Caelli, and C. Zetzsche, "Image encoding, labeling, and reconstruction from differential geometry," *CVGIP: Graphical Model and Image Processing*, vol. 55, no. 6, pp. 428–46, November 1993.

[13] C. Mota and E. Barth, "On the uniqueness of curvature features," in *Dynamische Perzeption*, Proceedings in Artificial Intelligence, G. Baratoff and H. Neumann, Eds., vol. 9. Infix Verlag, 2000, pp. 175–178.

[14] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185–203, 1981.

[15] H. Knutsson and G. H. Granlund, "Texture analysis using two-dimensional quadrature filters," in *IEEE Workshop on Computer Architecture for Pattern Analysis and Image Data Base Management*. Pasadena, California: IEEE, October 1983.

[16] T. Aach, A. Kaup, and R. Mester, "On texture analysis: Local energy transforms versus quadrature filters," *Signal Processing*, vol. 45, no. 2, pp. 173–181, 1995.

[17] E. Simoncelli and H. Farid, "Steerable wedge filters for local orientation analysis," *IEEE Trans. Image Processing*, vol. 5, no. 9, pp. 1377–1382, 1996.

[18] S. Di Zenzo, "A note on the gradient of a multi-image," *Computer Vision, Graphics, and Image Processing*, vol. 33, pp. 116–125, 1986.

[19] J. Bigün and G. H. Granlund, "Optimal orientation detection of linear symmetry," in *Proceedings IEEE First International Conference on Computer Vision*, London, June 1987, pp. 433–438.

[20] M. Kass and A. Witkin, "Analyzing oriented patterns," *Computer Vision, Graphics, and Image Processing*, vol. 37, pp. 362–385, 1987.

[21] C. G. Harris and M. J. Stevens, "A combined corner and edge detector," in *4th Alvey Vision Conference*, 1988.

[22] W. Förstner, "A feature based corresponding algorithm for image matching," *International Archive of Photogrammetry and Remote Sensing*, vol. 26, pp. 150–166, 1986.

[23] U. Köthe, "Edge and junction detection with an improved structure tensor," in *DAGM03: 25th Pattern Recognition Symposium*, B. Michaelis and G. Krell, Eds., German Association for Pattern Recognition. Magdeburg, Germany: Springer Verlag: LNCS 2781, 2003, pp. 25–32.

[24] T. Aach, I. Stuke, C. Mota, and E. Barth, "Estimation of multiple local orientations in image signals," in *Proceedings ICASSP-2004*. Montreal: IEEE, May 17–21 2004, pp. III 553–556.

[25] B. Rieger and L. J. Vliet, "A systematic approach to nD orientation representation," *Image and Vision Computing*, vol. 22, pp. 453–459, 2004.

[26] E. Trucco and A. Verri, *Introductory Techniques for 3D Computer Vision*. Upper Saddle River: Prentice-Hall, 1998.

[27] T. Aach, U. Schiebel, and G. Spekowius, "Digital image acquisition and processing in medical x-ray imaging," *Journal of Electronic Imaging*, vol. 8, no. Special Section on Biomedical Image Representation, pp. 7–22, 1999.

[28] J. Hsieh, *Computed Tomography. Principles, Design, Artifacts, and Recent Advances*. Bellingham, Washington: SPIE Press, 2003.

[29] C. Mota, I. Stuke, T. Aach, and E. Barth, "Estimation of multiple orientations at corners and junctions," in *DAGM04: 26th Pattern Recognition Symposium*, German Association for Pattern Recognition. Tübingen: Springer Verlag: LNCS 3175, Aug. 30 – Sept. 1 2004, pp. 163–170.

[30] M. Felsberg and G. Granlund, "POI detection using channel clustering and the 2D energy tensor," in *DAGM04: 26th Pattern Recognition Symposium*, German Association for Pattern Recognition. Tübingen: Springer Verlag: LNCS 3175, Aug. 30 – Sept. 1 2004, pp. 103–110.

[31] M. Shizawa and K. Mase, "A unified computational theory for motion transparency and motion boundaries based on eigenenergy analysis," in *Proceedings IEEE International Conference on Computer Vision and Pattern Recognition*. Maui: IEEE, June 3–6 1991, pp. 289–295.

[32] M. Shizawa and T. Iso, "Direct representation and detection of multi-scale, multi-orientation fields using local differentiation filters," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'93)*. Piscataway: IEEE, June 1993, pp. 508–514.

[33] C. Zetzsche and E. Barth, "Fundamental limits of linear filters in the visual processing of two-dimensional signals," *Vision Research*, vol. 30, pp. 1111–1117, 1990.

[34] M. Mühlich, *Estimation in Projective Spaces and Applications in Computer Vision*. J.-W. Goethe University Frankfurt, Germany: Ph.D. Thesis, 2005.

[35] M. Mühlich and R. Mester, "Optimal estimation of homogeneous vectors," in *Image Analysis: Proceedings 14th Scandinavian Conference (SCIA)*, H. Kalviainen, J. Parkkinen, and A. Kaarna, Eds., vol. LNCS 3540. Joensuu: Springer, June 19–22 2005, pp. 322–332.

[36] D. A. Forsyth and J. Ponce, *Computer Vision. A Modern Approach*. Upper Saddle River: Prentice-Hall, 2003.

[37] E. Kreyszig, *Advanced Engineering Mathematics*. New York, Chichester: John Wiley, 8th ed., 1999.

[38] E. Barth, "The minors of the structure tensor," in *Mustererkennung 2000*, G. Sommer, Ed. Berlin: Springer, 2000, pp. 221–228.

[39] C. Mota, I. Stuke, and E. Barth, "Analytic solutions for multiple motions," in *Proceedings IEEE International Conference on Image Processing (ICIP), Vol. II*, Thessaloniki: IEEE, October 7–10 2001, pp. 917–920.

[40] C. Mota, T. Aach, I. Stuke, and E. Barth, "Estimation of multiple orientations in multi-dimensional signals," in *IEEE International Conference on Image Processing (ICIP)*. Singapore: IEEE, Oct. 24–27 2004, pp. 2665–2668, and on CD–ROM: ISBN 0–7803–8555–1.

[41] T. Aach and D. Kunz, "Bayesian motion estimation for temporally recursive noise reduction in x-ray fluoroscopy," *Philips Journal of Research*, vol. 51, no. 2, pp. 231–251, 1998.

[42] C. Stiller, "Motion estimation for coding of moving video at 8kb/s with Gibbs modeled vector field smoothing," in *Proceedings Visual Communications and Image Processing 90*, M. Kunt, Ed., vol. 1360. Lausanne, Switzerland: SPIE, October 1990, pp. 468–476.

[43] H. Scharr and B. Jähne, "Optimization of spatio-temporal filter families for fast and accurate motion estimation," in *Image Sequence Analysis to Investigate Dynamic Processes, Lecture Notes in Computer Science*. Springer, 2003.

[44] M. Elad, P. Teo, and Y. Hel-Or, "On the design of optimal filters for gradient-based motion estimation," *International Journal on Mathematical Imaging and Vision*, vol. 23, no. 3, pp. 245–265, 2005.

[45] T. Aach and A. Kaup, "Bayesian algorithms for change detection in image sequences using Markov random fields," *Signal Processing: Image Communication*, vol. 7, no. 2, pp. 147–160, 1995.

[46] R. Vidal, Y. Ma, and S. Sastry, "Generalized principal component analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 12, pp. 1–15, 2005.

Til Aach (M 1994 – SM 2002) received his diploma and Doctoral degrees, both in EE, from RWTH Aachen University in 1987 and 1993, respectively. While working towards his Doctoral Degree, he was a research scientist with the Institute for Communications Engineering, RWTH Aachen University, being in charge of several projects in image analysis, 3D-television and medical image processing. From 1993 to 1998, he was with Philips Research Labs, Aachen, Germany, where he was responsible for several projects in X-Ray imaging and processing. In 1996, he was also an independent lecturer with the University of Magdeburg, Germany. In 1998, he was appointed a Full Professor and Director of the Institute for Signal Processing, University of Luebeck. In 2004, he became Head of the Institute of Imaging and Computer Vision, RWTH Aachen University. His research interests are in medical and industrial image processing, signal processing, pattern recognition, and computer vision. He has authored or co-authored over 160 papers, and received several awards, among these the award of the German "Informationstechnische Gesellschaft" (ITG/VDE), for a paper published in these Transactions in 1998. Dr. Aach is named as a co-inventor in about 20 patents. He is an Associate Editor of the IEEE Transactions on Image Processing, and was Technical Program Co-Chair for the IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI) in 2000, 2002, 2004, and 2006. He is a member of the Bio-Imaging and Signal Processing Committee (BISP-TC) of the IEEE Signal Processing Society.

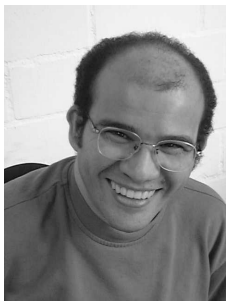

Cicero Mota studied Mathematics at the Federal University of Amazonas (B.Math. 1986) and the Institute for Pure and Applied Mathematics, IMPA (Dr.Sc. 1999) in Brazil. His main research interests are in computer vision and information theory. He has been with the Institute for Signal Processing and the Institute for Neuro- and Bio-informatics, University of Luebeck, and with the Institute for Applied Physics, University of Frankfurt, Germany. He is currently a Lecturer at the Federal University of Amazonas and serves as Head for the Mathematics Master program. He was awarded fellowships by the Brazilian agencies CNPq and CAPES and the German Academic Exchange Service (DAAD).

Ingo Stuke is currently a Scientist at the Institute for Signal Processing, University of Luebeck, Germany, where he is working towards his Doctoral Degree. His main research interests are in the field of motion and orientation estimation for image processing and computer vision applications. ¿From 1995 to 2001 he studied computer science at the University of Luebeck.

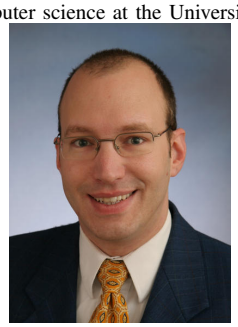

Matthias Mühlich received his diploma and Ph.D. degree from Johann Wolfgang Goethe University Frankfurt, Germany, in 1999 and 2005, respectively. He is currently working at RWTH Aachen University, Aachen, Germany. His current research interests are the estimation of orientations, motions and subspaces.

Erhardt Barth is currently a senior scientist and lecturer at the University of Luebeck in Germany. His main research interests are in the areas of human and computer vision, multi-dimensional signal processing, and pattern recognition. He obtained a Ph.D. in Electrical Engineering from the Technical University of Munich in 1994 and has been with the Universities of Melbourne and Munich, the Institute for Advanced Study in Berlin, and the NASA Vision Science and Technology Group in California.