Universität des Saarlandes

Fachrichtung 6.1 – Mathematik

Preprint Nr. 113

Nonlinear Structure Tensors

Thomas Brox, Joachim Weickert,
Bernhard Burgeth and Pavel Mrázek

Saarbrücken 2004

Fachrichtung 6.1 – Mathematik
Universität des Saarlandes

Preprint No. 113
submitted: October 12, 2004

Nonlinear Structure Tensors

Thomas Brox
Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science,
Saarland University, Building 27.1,
66041 Saarbrücken, Germany
brox@mia.uni-saarland.de

Joachim Weickert
Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science,
Saarland University, Building 27.1,
66041 Saarbrücken, Germany
weickert@mia.uni-saarland.de

Bernhard Burgeth
Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science,
Saarland University, Building 27.1,
66041 Saarbrücken, Germany
burgeth@mia.uni-saarland.de

Pavel Mrázek
UPEK Prague R&D Center,
Husinecka 7, 13000 Praha 3, Czech Republic
pavel.mrazek@upek.com

Edited by
FR 6.1 – Mathematik
Universität des Saarlandes
Postfach 15 11 50
66041 Saarbrücken
Germany

Fax: + 49 681 302 4443
e-Mail: preprint@math.uni-sb.de
WWW: http://www.math.uni-sb.de/

Abstract

In this article we introduce nonlinear versions of the popular structure tensor, also known as second moment matrix. These nonlinear structure tensors replace the Gaussian smoothing of the classical structure tensor by discontinuity-preserving nonlinear diffusions. While nonlinear diffusion is a well-established tool for scalar and vector-valued data, it has not often been used for tensor images so far. Two types of nonlinear diffusion processes for tensor data are studied: an isotropic one with a scalar-valued diffusivity, and its anisotropic counterpart with a diffusion tensor. We prove that these schemes preserve the positive semidefiniteness of a matrix field and are therefore appropriate for smoothing structure tensor fields. The use of diffusivity functions of total variation (TV) type allows us to construct nonlinear structure tensors without specifying additional parameters compared to the conventional structure tensor. The performance of nonlinear structure tensors is demonstrated in three fields where the classic structure tensor is frequently used: orientation estimation, optic flow computation, and corner detection. In all these cases the nonlinear structure tensors demonstrate their superiority over the classical linear one. Our experiments also show that for corner detection based on nonlinear structure tensors, anisotropic nonlinear tensors give the most precise localisation.

Key Words: Structure tensor; PDEs; diffusion; orientation estimation; optic flow; corner detection

Contents

1 Introduction

The matrix field of the *structure tensor*, introduced by Förstner and Gülch [17] as well as by Bigün and Granlund [6] in an equivalent formulation, plays a fundamental role in today's image processing and computer vision, as it allows both orientation estimation and image structure analysis. It has proven its usefulness in many application fields such as corner detection [17], texture analysis [36, 7, 26], diffusion filtering [47, 48], and optic flow estimation [7, 23]. It has even been successfully employed in numerical mathematics for grid optimisation when solving hyperbolic differential equations [43]. A detailed description on structure tensor concepts can be found in the textbook of Granlund and Knutsson [19].

The structure tensor offers three advantages. Firstly, the matrix representation of the image gradient allows the integration of information from a local neighbourhood without cancellation effects. Such effects would appear if gradients with opposite orientation were integrated directly. Secondly, smoothing the resulting matrix field yields robustness under noise by introducing an integration scale. This scale determines the local neighbourhood over which an orientation estimation at a certain pixel is performed. Thirdly, the integration of local orientation creates additional information, as it becomes possible to distinguish areas where structures are oriented uniformly, like in regions with edges, from areas where structures have different orientations, like in corner regions.

The classical structure tensor applies a linear technique such as Gaussian convolution for averaging information within a neighbourhood. Although Gaussian smoothing is a simple and robust method, it is known to have two important drawbacks: It blurs and dislocates structures. This is a consequence of the fact that the local neighbourhood for the integration is fixed in both its size and its shape. Consequently, it cannot adapt to the data, and the orientation estimation of a pixel located close to the boundary of two different regions is disturbed by ambiguous information.

It is well-known that Gaussian convolution is equivalent to linear diffusion. Therefore is is natural to address the limitations of Gaussian convolution by using nonlinear diffusion techniques which smooth the data while respecting discontinuities [35, 47]. For the structure tensor this means that the local neighbourhood, originally defined by the Gaussian kernel, is now adapted to the data and avoids smoothing across discontinuities. However, the structure tensor is a matrix field, and until recent time, techniques for nonlinear diffusion have only been available for scalar-valued and vector-valued data sets.

Tschumperlé and Deriche have introduced an isotropic[1] nonlinear diffusion scheme for matrix-valued data [44], while the anisotropic counterpart to their technique has been presented by Weickert and Brox [50]. These new methods allow to replace the Gaussian smoothing of the original linear structure tensor by a nonlinear diffusion method.

However, the application of one of these techniques is not perfectly straightforward. It has to be ensured that the nonlinear diffusion schemes do not violate the positive semidefiniteness property of the structure tensor. This will be proven in this paper. In contrast to an earlier conference publication [10], the nonlinear structure tensor, as it is proposed here, applies the original matrix-valued diffusion techniques from [44] and [50], thus using all available information for steering the diffusion. Moreover, it employs diffusivity functions based on *total variation (TV) flow* [2, 14], the diffusion filter corresponding to TV regularisation [40]. This flow offers a number of favourable properties, and it does not require additional contrast parameters such as most other diffusivity functions.

In principle, it makes sense to use nonlinear structure tensors in any application in which the classic structure tensor has already proven its usefulness and where discontinuities in the data play a role or delocalisation effects should be avoided. For this paper we focus on orientation analysis, optic flow estimation, and corner detection. Our experiments in these fields allow a direct comparison between the performance of the nonlinear structure tensors and the classic linear one.

Paper Organisation. The following section starts with a brief review of the conventional linear structure tensor, its properties and shortcomings. In Section 3 we then discuss isotropic and anisotropic nonlinear diffusion filters for matrix-valued data, and in Section 4 we prove that these nonlinear filters preserve the positive semidefiniteness if the original data field is positive semidefinite. Tensor-valued nonlinear diffusion filtering is used in Section 5 for constructing isotropic and anisotropic nonlinear structure tensors. The Sections 6–8 deal with applications of the nonlinear structure tensors to orientation analysis, optic flow estimation, and corner detection. The paper is concluded with a summary in Section 9.

Related Work. There are several proposals in the literature that intend to avoid the blurring effects of the conventional structure tensor across discontinuities. Nagel and Gehrke [33] introduced a structure tensor for optic flow estimation using local information in order to adapt the Gaussian kernel to

[1]In our notation, isotropic nonlinear diffusion means nonlinear diffusion driven by a scalar-valued diffusivity, in contrast to anisotropic nonlinear diffusion, which is driven by a matrix-valued diffusion tensor.

the data. This work has been further extended in [30, 31]. While nonlinear diffusion filtering and adaptive Gaussian smoothing are similar for small amounts of smoothing, significant differences arise when more substantial smoothing is performed. In this case, nonlinear diffusion based on the iterative appliciation of very small averaging kernels can realise highly complex adaptive kernel structures.

An orientation estimation method based on robust statistics has been proposed by van den Boomgaard and van de Weijer [46]. Another related method is proposed by Köthe [25]. In order to detect edges and corners, an adaptive, hour-glass shaped filter is used for smoothing the structure tensor. Analysing the differences and understanding the relations between such adaptive filters, robust estimation and nonlinear diffusion methods is a topic of current research; see e.g. [32] for the scalar case and [9] for the tensor case.

Our article comprises and extends earlier work presented at conferences [50, 10]. These extensions contain: (i) the diffusion of the structure tensor by means of diffusivities based on TV flow, (ii) a proof that the used schemes preserve the positive semidefiniteness of the original matrix field also in the continuous setting, (iii) an extensive comparison of linear, isotropic, and anisotropic diffusion of the structure tensor, and (iv) the application of the nonlinear structure tensor to corner detection.

2 Linear Structure Tensor

Let $\Omega \subset \mathbb{R}^m$ denote our m-dimensional image domain, and let us consider some greyscale image $h : \Omega \to \mathbb{R}$. Then the structure tensor is a field of symmetric $m \times m$ matrices that contains in each element information on orientation and intensity of the surrounding structure of h. The initial matrix field is computed from the gradient of h by applying the tensor product $J_0 = \nabla h \nabla h^\top$. Although this tensor product contains no more information than the gradient itself, it has the advantage that it can be smoothed without cancellation effects in areas where gradients have opposite signs, since $\nabla h \nabla h^\top = (-\nabla h)(-\nabla h^\top)$: Consider, for instance, a thin line. It has a positive gradient on one side, and a negative gradient on the other side. Any smoothing operation on the gradient directly would cause both gradients to mutually cancel out. Smoothing the matrix field, however, avoids this cancellation effect.

The smoothing is usually performed by convolution of the matrix components with a Gaussian kernel K_ρ with standard deviation ρ:

$$J_\rho = K_\rho * (\nabla h \nabla h^\top). \tag{1}$$

Since convolution is a linear operation, we refer to the classic structure tensor as *linear structure tensor*. It is a symmetric, positive semidefinite matrix, since it results from averaging of symmetric positive semidefinite matrices. Gaussian smoothing not only improves the orientation information with regard to noise, but also creates a scale-space with the *integration scale* ρ. This scale parameter determines the size of the neighbourhood considered for the structure analysis.

The structure tensor can also be defined for vector-valued data sets, colour images for instance [13]. Let $h : \Omega \to \mathbb{R}^n$ be a vector-valued data set and h_i its i-th component. Then the structure tensor is given by

$$J_\rho = K_\rho * \sum_{i=1}^{n} \nabla h_i \nabla h_i^\top . \tag{2}$$

Besides the information on orientation and magnitude of structures, which is already present in the gradient, the structure tensor contains some further information. This additional information has been obtained by the smoothing process and measures the homogeneity of orientations within the neighbourhood of a pixel.

This information can be extracted from the structure tensor by means of a principal axis transformation $J_\rho = S^\top \Lambda S$, where the eigenvectors of J_ρ are the rows of S, and the corresponding eigenvalues λ_i with $\lambda_1 \geq \cdots \geq \lambda_m$, are the elements of the diagonal matrix $\Lambda = \operatorname{diag}(\lambda_i)$. The eigenvector to the smallest eigenvalue then determines the dominant orientation of the local structure, while the trace $\mathrm{tr} J_\rho$ (sum of the diagonal elements) of the structure tensor J_ρ determines its magnitude. The coherence in 2-D image data is often expressed by the condition number of J_ρ (largest eigenvalue divided by smallest eigenvalue) or by the measure $(\lambda_1 - \lambda_2)^2$, yet also other measures based on the eigenvalues may be reasonable.

Magnitude and coherence of the structure tensor can be used for structure analysis. Homogeneous areas in an image cause the magnitude to be small. In areas around edges the structure tensor has a large magnitude as well as a large coherence, while corners result in a large magnitude but small coherence. For higher dimensional data, this structure analysis becomes more complicated, as more cases must be distinguished.

Fig. 1 shows the most important properties of the linear structure tensor. Fig. 1a depicts a synthetic test image distorted by Gaussian noise with $\sigma = 30$. In Fig. 1b the matrix product $\nabla h \nabla h^\top$ is shown as a coloured orientation plot. Its orientation information is expressed by the colour whereas the magnitude is encoded by the intensity of the plot. Fig. 1c finally shows the linear structure tensor for $\rho = 3$. The following properties can be observed:

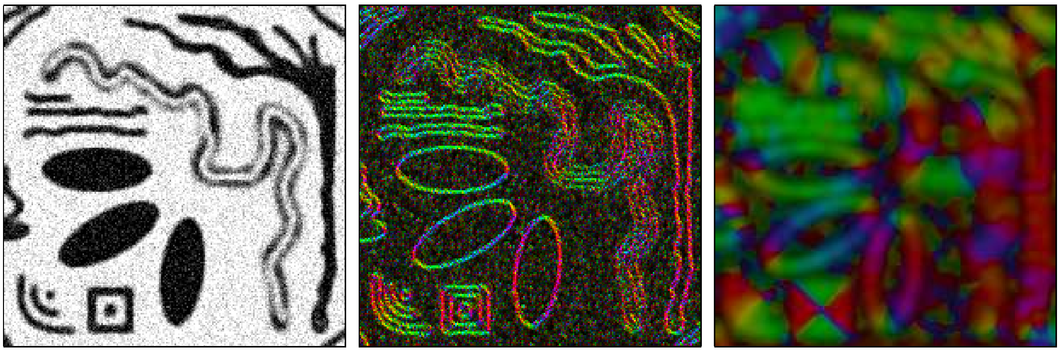

Figure 1: LEFT: (a) Synthetic image with Gaussian noise. CENTER: (b) $J_0 = \nabla h \nabla h^\top$. RIGHT: (c) Linear structure tensor J_ρ with $\rho = 3$.

- **Noise removal**
  Most of the noise present in the initial matrix field has been removed due to the smoothing.

- **Propagation of orientation information**
  In most applications of the structure tensor it is desirable that there is a filling-in effect of orientation information from structured areas into areas without structure as far as these areas are small in respect of a certain scale. By means of the structures in the lower left of Fig. 1 it can be seen that the linear structure tensor fulfills this requirement appropriately. This subsequent simplification results in a scale-space property with the scale parameter ρ.

- **Dislocation of discontinuities and blurring effects**
  Fig. 1c reveals a blurring effect that is typical for Gaussian smoothing. Edges disappear with increasing ρ and the remaining edges dislocate. A smoothing method based on nonlinear diffusion should be able to preserve these discontinuities. This shall be discussed next.

3 Nonlinear Diffusion Filtering of Tensor Data

In this section we study isotropic and anisotropic nonlinear diffusion filters that will allow us to construct nonlinear structure tensors later on.

3.1 Isotropic Nonlinear Diffusion

The goal of nonlinear diffusion filtering is to reduce smoothing in the presence of edges [35]. This can be achieved by a decreasing *diffusivity function* g which correlates the amount of smoothing with the image gradient magnitude (suitable functions will be discussed in Subsection 3.3). Nonlinear diffusion filtering creates a family of simplified images $\{u(x,t) \mid t \geq 0\}$ of some scalar initial image $f(x)$ by solving the partial differential equation (PDE)

$$\partial_t u = \operatorname{div}\left(g(|\nabla u|^2)\nabla u\right) \quad \text{on} \quad \Omega \times (0,\infty), \tag{3}$$

with f as initial condition,

$$u(x,0) = f(x) \quad \text{on} \quad \Omega, \tag{4}$$

and reflecting (homogeneous Neumann) boundary conditions:

$$\partial_\nu u = 0 \quad \text{on} \quad \partial\Omega \times (0,\infty), \tag{5}$$

where ν denotes the outer normal on the image boundary $\partial\Omega$. The diffusion time t determines the amount of simplification: For $t = 0$ the original image f is recovered, and larger values for t result in more pronounced smoothing.

An extension of nonlinear diffusion filtering to vector-valued data $f = (f_i) : \Omega \to \mathbb{R}^n$ has been proposed in [18]. It evolves f under the diffusion equations

$$\partial_t u_i = \operatorname{div}\left(g\Big(\sum_{k=1}^{n} |\nabla u_k|^2\Big)\nabla u_i\right) \qquad (i = 1, ..., n) \tag{6}$$

where u is a vector with n components. Note that all vector channels are coupled in this scheme: They are smoothed with a joint diffusivity taking into account the edges of all channels. This synchronisation avoids that edges evolve at different locations in different channels: A discontinuity in one channel inhibits also smoothing in the others.

The coupled vector-valued diffusion scheme is also a good basis for smoothing a matrix field $F = (f_{i,j}) : \Omega \to \mathbb{R}^{n\times n}$. When regarding the components of an $n \times n$ matrix as components of an n^2-dimensional vector, which is not unnatural since e.g. the Frobenius norm of a matrix equals the Euclidean norm of the resulting vector, it is possible to diffuse also a matrix field with Eq. 6. In fact, this leads to the following PDEs for matrix-valued diffusion [44]:

$$\partial_t u_{i,j} = \operatorname{div}\left(g\Big(\sum_{k,l=1}^{n} |\nabla u_{k,l}|^2\Big)\nabla u_{i,j}\right) \qquad (i,j=1,...,n). \tag{7}$$

In Section 4 we will see that the coupling of the tensor channels guarantees that the evolving matrix field $U(x,t) = (u_{i,j}(x,t))$ remains positive semidefinite if its initial value $F(x) = (f_{i,j}(x))$ is positive semidefinite.
It is easy to verify that the diffusion equations (7) can be regarded as a steepest descend method for minimising the energy functional

$$E(U) = \int_\Omega \Psi\Big(\sum_{k,l=1}^{m} |\nabla u_{k,l}|^2\Big)\, dx \tag{8}$$

with a penaliser $\Psi(s^2)$ whose derivative satisfies $\Psi'(s^2) = g(s^2)$.

3.2 Anisotropic Nonlinear Diffusion

Besides these isotropic diffusion schemes, there exist also anisotropic counterparts. In the anisotropic case not only the amount of diffusion is adapted locally to the data but also the direction of smoothing. It allows for example to smooth along image edges while inhibiting smoothing across edges. This can be achieved by replacing the scalar-valued diffusivity function by a matrix-valued diffusion tensor.
Vector-valued anisotropic diffusion evolves the original image $f(x) = (f_i(x))$ under the PDE [49]

$$\partial_t u_i = \operatorname{div}\left(g\Big(\sum_{k=1}^{n} \nabla u_k \nabla u_k^\top\Big)\nabla u_i\right) \qquad (i=1,...,n), \tag{9}$$

subject to the reflecting boundary conditions

$$\partial_\nu\left(g\Big(\sum_{k=1}^{n} \nabla u_k \nabla u_k^\top\Big)\, \nabla u_i\right) = 0 \qquad (i=1,...,n). \tag{10}$$

Here the scalar-valued function g has been generalised to a matrix-valued function in the following way: Let $A = S\operatorname{diag}(\lambda_i)S^\top$ denote the principal axis transformation of some symmetric matrix A, with the eigenvalues λ_i as the elements of the diagonal matrix $\operatorname{diag}(\lambda_i)$ and the normalised eigenvectors as the columns of the orthogonal matrix S. Then we set $g(A) := S \mathrm{diag}\,(g(\lambda_i))S^\top$. The diffusivity $g(s^2)$ is the same decreasing function as in the isotropic case. Simply speaking, in the anisotropic setting the diffusivity function is applied to the eigenvalues of the *matrix* obtained from the outer product of the gradient. This gives a *diffusion tensor* $g(\sum_k \nabla u_k \nabla u_k^\top)$. In the isotropic setting, the diffusivity function is applied to the *scalar-valued* squared gradient magnitude, or the scalar product of the gradient. This yields a *scalar-valued* diffusivity $g(\sum_k \nabla u_k^\top \nabla u_k)$. Note that the transition from the isotropic to the anisotropic setting simply consists of exchanging the order of ∇u and $\nabla u^\top$.

Anisotropic diffusion offers the advantage of smoothing in a direction-specific way: Along the i-th eigenvector of $\sum_k \nabla u_k \nabla u_k^\top$ with corresponding eigenvalue λ_i, the eigenvalue of the diffusion tensor is given by $g(\lambda_i)$. In eigendirections with large variation of local structure, λ_i is large and $g(\lambda_i)$ is small. This avoids smoothing across discontinuities. Along discontinuities, λ_i is small such that $g(\lambda_i)$ is large and full diffusion is performed. For more information about anisotropic diffusion in general, we refer to [47].

In [50] this vector-valued scheme has been generalised to matrix-valued data by considering the PDEs

$$\partial_t u_{i,j} = \mathrm{div}\left(g\Big(\sum_{k,l=1}^{n} \nabla u_{k,l} \nabla u_{k,l}^\top\Big) \nabla u_{i,j}\right) \qquad (i,j = 1, ..., n). \tag{11}$$

In a similar way as in [51], one can prove that this process can be regarded as a gradient descend method for mimimising the energy functional

$$E(U) = \int_\Omega \mathrm{tr}\,\Psi\Big(\sum_{k,l=1}^{n} \nabla u_{k,l} \nabla u_{k,l}^\top\Big)\, dx. \tag{12}$$

Note the structural similarity to the isotropic functional (8) which may be rewritten as

$$E(U) = \int_\Omega \Psi\Big(\mathrm{tr} \sum_{k,l=1}^{n} \nabla u_{k,l} \nabla u_{k,l}^\top\Big)\, dx. \tag{13}$$

Thus, for going from the isotropic to the anisotropic functional, all one has to do is to exchange the order of the penaliser Ψ and the trace operator.

3.3 Diffusivity Functions

The choice of the diffusivity function g has a rather large impact on the outcome of diffusion. The following family of diffusivity functions is very interesting:

$$g(|\nabla u|^2) = \frac{1}{|\nabla u|^p} \tag{14}$$

with $p \in \mathbb{R}$ and $p \geq 0$. These diffusivities offer the advantage that they do not require any image specific contrast parameters. Moreover, they lead to scale invariant filters [1], for which even some analytical results have been established [45].

For $p = 0$, linear homogeneous diffusion is obtained, which is equivalent to Gaussian smoothing with standard deviation $\sqrt{2t}$, and forms the basis of Gaussian scale-space theory [22, 42].

For $p = 1$ one obtains the *total variation (TV) flow* [2, 14], the diffusion filter that corresponds to TV minimisation [40] with a penaliser $\Psi(|\nabla u|^2) = 2|\nabla u|$. TV flow offers a number of interesting properties such as finite extinction time [3], shape-preserving qualities [5], and equivalence to TV regularisation in 1-D [11].

Finally, for $p > 1$ the diffusion not only preserves edges but even enhances them. A diffusivity with $p = 2$ has been considered in [24] for the so-called *balanced forward–backward* diffusion filtering. While a complete well-posedness theory exists for $p \leq 1$, some theoretical questions are a topic of ongoing research for the edge-enhancing case $p > 1$.

In the present paper we focus on TV flow ($p = 1$), since it is theoretically well-founded [3, 15], and it offers a good compromise between the smoothing properties for small values of p, and the edge preserving qualities for large p. We introduce a small regularisation with some fixed parameter $\epsilon > 0$ that avoids singularities and creates a differentiable diffusivity function:

$$g(|\nabla u|^2) = \frac{1}{\sqrt{\epsilon^2 + |\nabla u|^2}}. \tag{15}$$

4 Preservation of Positive Semidefiniteness

When applying a diffusion process to matrix-valued data it is by no means clear that the positive (semi-)definiteness of the original data is preserved. However, for the diffusion schemes we use here, we now prove a maximum-minimum principle for the field of eigenvalues associated with a matrix field.

Let $F = (f_{i,j}) : \Omega \longrightarrow \mathbb{R}^{n\times n}$ denote the initial field of $n \times n$-matrices. Accordingly $U(x,t) = (u_{i,j}(x,t))$ stands for the diffused matrix field, while $F(x)$ serves as initial value for the isotropic diffusion equation (7), or the anisotropic diffusion process (11) with the diffusivity function (15). Furthermore, let $\lambda_k^F(x)$ resp. $\lambda_k^U(x,t)$ be the k-th eigenvalue of the initial matrix field $F(x)$ and the diffused field $U(x,t)$ with $k = 1,\dots,n$. Denoting by $\lambda_{min}^F(x)$ and $\lambda_{max}^F(x)$ the smallest and the largest eigenvalue of the matrix $F(x)$, $x \in \Omega$, we have the following result.

Theorem 1: (Extremum Principle for the Eigenvalues.)
For $t > 0$, the eigenvalues of the diffused matrix field $U(.,t)$ are bounded by the eigenvalues of the initial matrix field F:

$$\inf_{y\in\Omega} \lambda_{min}^F(y) \;\le\; \lambda_k^U(x,t) \;\le\; \sup_{y\in\Omega} \lambda_{max}^F(y) \qquad (\forall x \in \Omega,\; k = 1,\dots,n). \tag{16}$$

Proof: We consider the anisotropic case, the arguments carry over to the isotropic case (7) essentially verbatim. For any *unit* column vector $v \in \mathbb{R}^n$ it follows from (11) by linearity properties of the matrix multiplications and differential operators involved that

$$\begin{aligned}
\partial_t(v^\top U v) &= \partial_t\big(\sum_{i,j} v_i \cdot u_{i,j} \cdot v_j\big) \\
&= \sum_{i,j} v_i \cdot \partial_t u_{i,j} \cdot v_j \\
&= \sum_{i,j} v_i \cdot \left(\operatorname{div} \left(g\Big(\sum_{k,l} \nabla u_{k,l} \nabla u_{k,l}^\top\Big) \nabla u_{i,j}\right)\right) \cdot v_j \\
&= \operatorname{div} \left(g\Big(\sum_{k,l} \nabla u_{k,l} \nabla u_{k,l}^\top\Big)\, \nabla\Big(\sum_{i,j} v_i \cdot u_{i,j} \cdot v_j\Big)\right) \\
&= \operatorname{div} \left(g\Big(\sum_{k,l} \nabla u_{k,l} \nabla u_{k,l}^\top\Big) \nabla(v^\top U v)\right).
\end{aligned}$$

Due to the properties of the regularised TV diffusivity function g, the associated matrix $g\big(\sum_{k,l} \nabla u_{k,l} \nabla u_{k,l}^\top\big)$ fits into the framework for a scalar-valued continuous nonlinear diffusion scale-space [47]. As a consequence, the scalar valued functions $(x,t) \mapsto v^\top U(x,t)v$ are smooth and an extremum principle holds. Therefore,

$$\begin{aligned}
\inf_{y\in\Omega} \lambda^F_{min}(y) &\leq \inf_{y\in\Omega} v^\top F(y)v \\
&\leq v^\top\, U(x,t)\, v \\
&\leq \sup_{y\in\Omega}\; v^\top F(y)\, v \\
&\leq \sup_{y\in\Omega} \lambda^F_{max}(y)
\end{aligned}$$

for $t > 0$ and all *unit* vectors $v \in \mathbb{R}^n$. Since the expression $v^\top U(x,t)v$ is the well-known Rayleigh quotient, one can choose for each eigenvalue $\lambda^U_k(x,t)$ a unit eigenvector $v^k = v^k(x,t)$ such that

$$\lambda^U_k(x,t) = v^{k\top}(x,t)\, U(x,t) v^k(x,t).$$

Hence the assertion follows from the sequence of inequalities above. ■

An immediate consequence is the following corollary.

Corollary 1: (Preservation of Positive (Semi-)Definiteness.)
Under the assumptions for Theorem 1, positive (semi-)definiteness of the initial matrix field F implies positive (semi-)definiteness of the diffused matrix field $U(.,t)$ for $t > 0$.

This corollary is crucial for many applications such as diffusion tensor MRI or structure tensor smoothing, since it guarantees that the positive semidefinitess of the initial field is not destroyed by diffusion filters of type (7) or (11). A discrete reasoning why such filters preserve positive semidefiniteness can be found in [50], where it is argued that convex combinations of positive semidefinite matrices are computed in each iteration. Both results show that already the channel coupling of the diffusion processes via a joint diffusivity or a joint diffusion tensor is sufficient to preserve positive semidefiniteness. Thus, from a viewpoint of preservation of positive semidefiniteness, it is not required to consider more sophisticated constrained flows [12] or functionals with Cholesky decomposition [29].

5 Nonlinear Structure Tensors

Now that we have understood how nonlinear diffusion filtering of tensor fields works we are in the position of using this knowledge for constructing nonlinear structure tensors.

General Idea. Given some image $h : \Omega \to \mathbb{R}$ with $\Omega \subset \mathbb{R}^m$ we consider the tensor product

$$F := (f_{ij}) := \nabla h \nabla h^\top. \tag{17}$$

Then the classic stucture tensor applies componentwise Gaussian convolution to the matrix field $F = (f_{ij}) : \Omega \to \mathbb{R}^{m\times m}$. This is equivalent to regarding F as initial value for the linear matrix-valued diffusion equation

$$\partial_t u_{ij} \;=\; \Delta u_{ij} \qquad (i,j = 1, ..., m) \tag{18}$$

where the diffusion time t is related to the standard deviation ρ of the Gaussian via $t = \rho^2/2$. Nonlinear structure tensors replace this diffusion equation either by the isotropic diffusion scheme

$$\partial_t u_{i,j} = \operatorname{div}\Big(g\Big(\sum_{k,l=1}^{m} |\nabla u_{k,l}|^2\Big)\,\nabla u_{i,j}\Big) \qquad (i,j = 1, ..., m) \tag{19}$$

or the anisotropic diffusion process

$$\partial_t u_{i,j} = \operatorname{div}\Big(g\Big(\sum_{k,l=1}^{m} \nabla u_{k,l} \nabla u_{k,l}^\top\Big)\,\nabla u_{i,j}\Big) \qquad (i,j = 1, ..., m), \tag{20}$$

both in combination with diffusivity functions such as (14). The result $U(x,t) = (u_{ij}(x,t))$ gives the desired isotropic or anisotropic nonlinear structure tensor field. Since F is a positive semidefinite matrix field, Corollary 1 guarantees that $U(x,t)$ is also positive semidefinite for all $t > 0$, provided that the underlying scalar diffusion process satisfies a maximum–minimum principle.

Role of the Parameters. We have two parameters for a nonlinear structure tensor. Firstly, there is the diffusion time t that determines the amount of smoothing, i.e. the size of the neighbourhood. It corresponds directly to the integration scale ρ of the linear structure tensor, since the Gaussian convolution in the linear structure tensor equals linear diffusion with diffusion time $t = \rho^2/2$. Thus, t is not a conceptually new parameter. Secondly, there is the parameter p that determines the amount of edge preservation. Note that this latter parameter is implicitly also present in the classic structure tensor: The classic linear structure tensor is a special case of the nonlinear structure tensor for $p = 0$ where the diffusivity g becomes equal to 1. Since we favour the TV diffusion case, we usually fix p to 1. In this case, p does not constitute an additional parameter. Consequently, going from linear to nonlinear structure tensors does to introduce new problems of parameter selection.

Implementation. Compared to scalar-valued nonlinear diffusion filters, their tensor-valued counterparts do not involve additional difficulties with respect to implementations. In our experiments we apply standard space discretisations by means of central finite differences (see e.g. [49]). With respect to the time discretisation, an efficient semi-implicit additive operator splitting (AOS) scheme is used [27, 52]. Since it is absolutely stable, it is possible to choose significantly larger time step sizes than for the widely used explicit (Euler-forward) discretisations.
Application Areas. All applications of the classic structure tensor are also potential applications for its nonlinear variants. It should be clear, however, that the nonlinear structure tensors have only advantages in the presence of discontinuities or when dislocalisation problems appear. If this is not the case, nonlinear structure tensors cannot be any better than the conventional one. In the presence of important discontinuities in the data, on the other hand, the accuracy of the results should improve with the usage of nonlinear structure tensors. We will now present experiments in three fields of application where the conventional structure tensor is very popular and where discontinuities or delocalisation effects can play a major role: orientation analysis, optic flow estimation, and corner detection. This list is not complete: For experiments on texture analysis by means of nonlinear structure tensors, we refer to [8].

6 Application to Orientation Estimation

In this section we analyse the use of nonlinear structure tensors for orientation estimation by applying them to the test image from Fig. 1.
Fig. 2 depicts the different versions of the structure tensor. For the linear structure tensor in Fig. 2a, Gaussian smoothing has been used ($p = 0$). Fig. 2b shows the nonlinear structure tensor smoothed with the isotropic scheme from (7) and TV flow ($p = 1$). Finally, Fig. 2c depicts the nonlinear structure tensor employing the anisotropic diffusion scheme from (11), again with $p = 1$. It can be observed that both nonlinear structure tensors succeed in avoiding the blurring effects that are the decisive drawback of the original linear structure tensor. The isotropic nonlinear structure tensor performs best at orientation discontinuities, while the anisotropic nonlinear structure tensor is slighly better at smoothing within one homogeneous region.
Fig. 3 depicts the results for different diffusivities from the family of (14). As the differences between the isotropic and anisotropic nonlinear structure tensor are small, only the isotropic version is shown. Fig. 3a depicts the result for $p = 0.8$, where the diffusion is closer to Gaussian smoothing. In

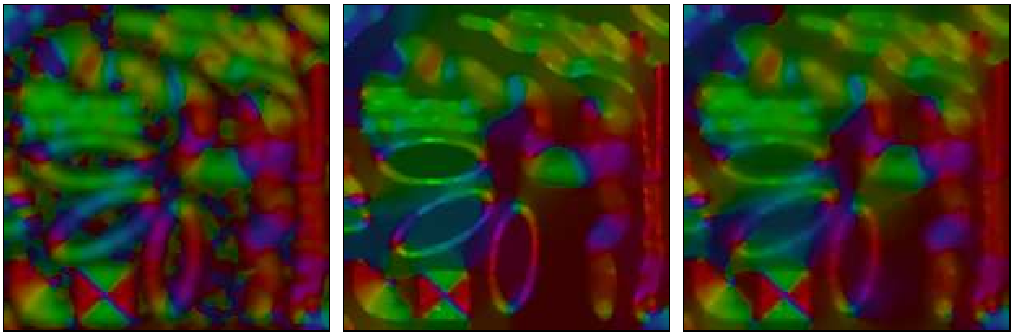

Figure 2: Left: (a) Linear structure tensor, $\rho = 3$, corresponding to $t = 4.5$. Center: (b) Isotropic nonlinear structure tensor (Eq. 7), $t = 3200$. Right: (c) Anisotropic nonlinear structure tensor (Eq. 11), $t = 3600$.

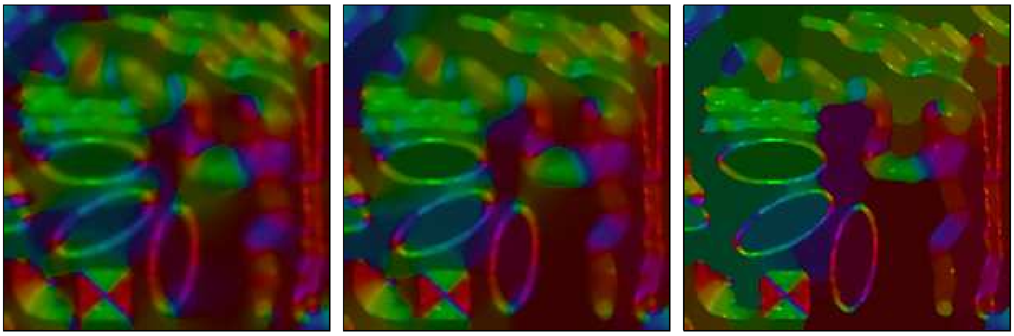

Figure 3: Isotropic nonlinear structure tensor for different p. Left: (a) $p = 0.8$ and $t = 800$. Center: (b) $p = 1$ and $t = 3200$. Right: (c) $p = 1.2$ and $t = 12000$.

contrast to that, the diffusion for $p = 1.2$ is edge enhancing. Hence, the result in Fig. 3c reveals sharper edges. Fig. 3b depicts the result achieved with TV flow, which is a good compromise between edge preservation and closing of structures.

Fig. 4 illustrates that the diffusion time t can be regarded as a scale parameter for nonlinear structure tensors: By increasing t the orientation field becomes simpler and larger regions of homogeneous orientation are formed. Thus t plays the same role for nonlinear structure tensors as standard deviation ρ of the Gaussian for the linear structure tensor.

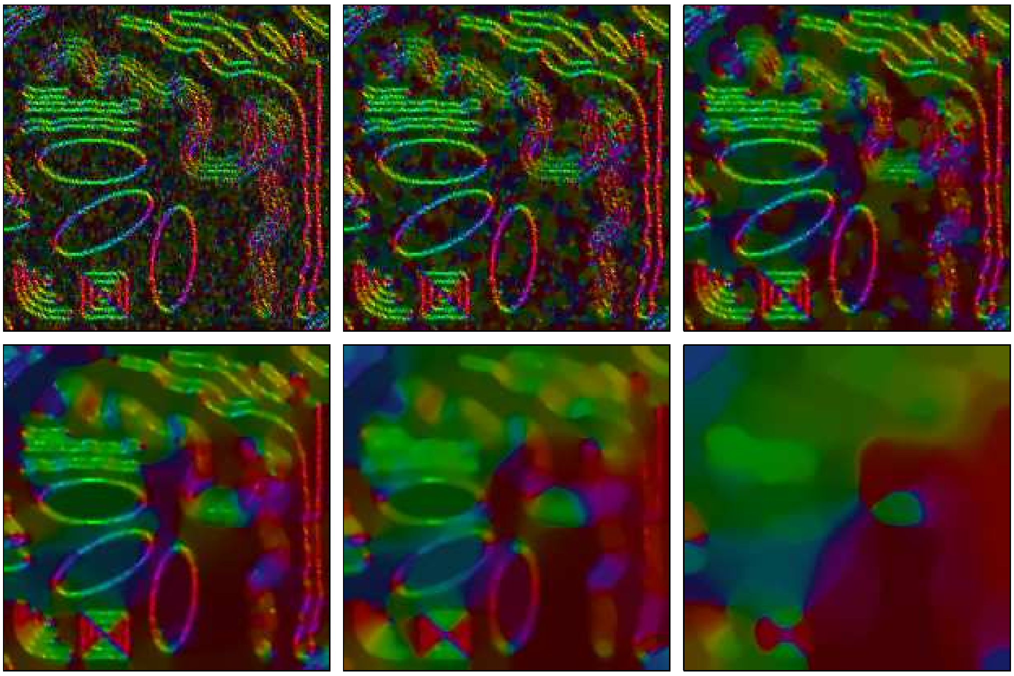

Figure 4: Temporal evolution of the isotropic nonlinear structure tensor ($p = 1$). From Left to Right, Top to Bottom: (a) $t = 250$. (b) $t = 500$. (c) $t = 1000$. (d) $t = 2000$. (e) $t = 4000$. (f) $t = 8000$.

7 Application to Optic Flow Estimation

Optic flow estimation by means of the structure tensor has first been investigated by Bigün *et al.* [7]. However, already the well-known method of Lucas and Kanade [28] implicitly used the structure tensor components. Both methods are very similar, and we will stick here to the method of Lucas and Kanade.

The goal in optic flow estimation is to find the displacement field (u, v) between two images of an image sequence $f(x, y, z)$ where (x, y) denotes location and z denotes time. Frequently it is assumed that image structures do not alter their grey values during their movement. This can be expressed by the optic flow constraint [21]

$$f_x u + f_y v + f_z = 0 \tag{21}$$

where subscripts denote partial derivatives. As this is only one equation for two unknown flow components, the optic flow is not uniquely determined by this constraint *(aperture problem)*. A second assumption has to be made. Lucas and Kanade proposed to assume the optic flow vector to be constant within some neighbourhood. Often one uses a Gaussian-weighted neighbourhood K_ρ where ρ is the standard deviation of the Gaussian. The optic flow in some point (x_0, y_0) can then be estimated by the minimiser of the local energy function

$$E(u,v) = \frac{1}{2} K_\rho * (f_x u + f_y v + f_z)^2. \tag{22}$$

A minimum (u,v) of E satisfies $\partial_u E = 0$ and $\partial_v E = 0$, leading to the linear system

$$\begin{pmatrix} K_\rho * (f_x^2) & K_\rho * (f_x f_y) \\ K_\rho * (f_x f_y) & K_\rho * (f_y^2) \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} -K_\rho * (f_x f_z) \\ -K_\rho * (f_y f_z) \end{pmatrix}. \tag{23}$$

Obviously, the entries of this linear system are five of the six different components of the spatio-temporal linear structure tensor

$$J_\rho = K_\rho * \left(\nabla f \nabla f^\top \right) = K_\rho * \begin{pmatrix} f_x^2 & f_x f_y & f_x f_z \\ f_x f_y & f_y^2 & f_y f_z \\ f_x f_z & f_y f_z & f_z^2 \end{pmatrix}. \tag{24}$$

With the nonlinear structure tensor available, we can introduce a *nonlinear* version of the Lucas–Kanade method by replacing the components of the linear structure tensor in (23) by those of the nonlinear one. This means that the fixed neighbourhood of the original method is replaced by an adaptive neighbourhood which respects discontinuities in the data.

Evaluation in optic flow estimation. In order to see the effect of the adaptive neighbourhood on the quality of the results, we tested all three versions of the structure tensor: the conventional linear structure tensor, the nonlinear structure tensor based on isotropic nonlinear diffusion, and the one based on anisotropic diffusion.
For optic flow estimation, a frequently used quantitative quality measure is the so-called *average angular error (AAE)* introduced in [4]. Given the estimated flow field (u_e, v_e) and ground truth (u_c, v_c), the AAE is defined as

$$AAE = \frac{1}{N} \sum_{i=1}^{N} \arccos \left(\frac{u_{ci} u_{ei} + v_{ci} v_{ei} + 1}{\sqrt{(u_{ci}^2 + v_{ci}^2 + 1)(u_{ei}^2 + v_{ei}^2 + 1)}} \right) \tag{25}$$

where N is the total number of pixels. In contrast to its indication, this quality measure does not only measure the angular error between the estimated flow vector and the correct vector, but also differences in the magnitude of both vectors, since it measures the angular error of the *spatio-temporal* vector $(u, v, 1)$.

For our experiments we used two different sequences from the literature with the correct flow field available: the famous *Yosemite* sequence[2] and the *Marble* sequence[3]. They both contain discontinuities, so they are well-suited to test the improvements that can be achieved with the nonlinear structure tensors.

Yosemite sequence.

Technique	t	AAE	Standard dev.
Linear structure tensor	21	8.78°	12.76°
Isotropic nonlinear structure tensor	400	7.67°	11.02°
Anisotropic nonlinear structure tensor	200	7.68°	11.84°

Marble sequence.

Technique	t	AAE	Standard dev.
Linear structure tensor	222	5.82°	4.48°
Isotropic nonlinear structure tensor	250	5.19°	2.98°
Anisotropic nonlinear structure tensor	163	5.10°	3.22°

Table 1. Comparison between the results. AAE = average angular error.

A quantitative comparison between the results obtained with the Lucas–Kanade method and the three different versions of the structure tensor is provided in Table 1. The second row indicates the only free parameter that was optimised for each sequence: the diffusion time. In the nonlinear cases, we always used $p = 1$, so TV flow was applied.
It can be observed that the nonlinear structure tensors can clearly outperform the conventional linear one. The difference between the isotropic and anisotropic scheme, however, is only marginal in this application.
Comparing the visual impression, the improvement achieved with the nonlinear structure tensor is even larger. This is because the nonlinear structure tensor is beneficial especially in the areas around discontinuities. Such areas are relatively small compared to the whole image, so most of the improvements are hidden by a global measure such as the AAE. Fig. 5 and 6 show the results obtained with the different versions of the structure tensor together with the correct flow field. Again colour expresses the orientation of the flow vectors while the intensity shows their magnitude. In our case, this kind of representation is preferable to the common representation method using vector plots, because no subsampling is necessary and so the quality of the results at discontinuities becomes better visible. Note that the black parts

[2]Created by Lynn Quam at SRI, available from `ftp://csd.uwo.ca/pub/vision`

[3]Created by Otte and Nagel [34], available from `http://i21www.ira.uka.de/image_sequences`

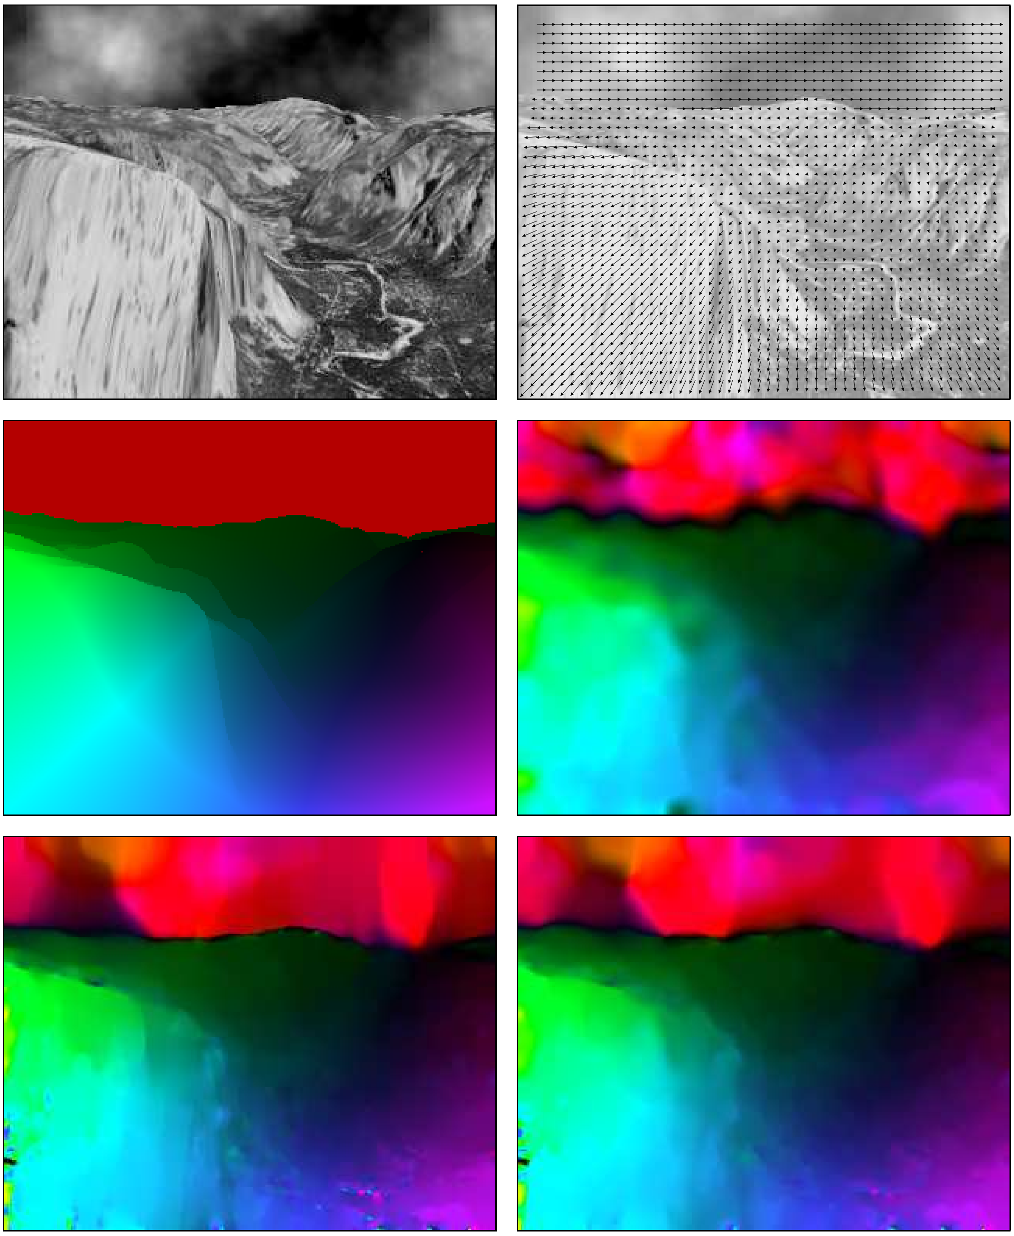

Figure 5: Yosemite sequence ($316 \times 252 \times 15$). FROM LEFT TO RIGHT, TOP TO BOTTOM: (a) Frame 8. (b) Ground truth optic flow field as vector plot. (c) Ground truth, where the orientations are represented by colours. (d) Lucas–Kanade with linear structure tensor. (e) Lucas–Kanade with isotropic nonlinear structure tensor. (f) Lucas–Kanade with anisotropic nonlinear structure tensor.

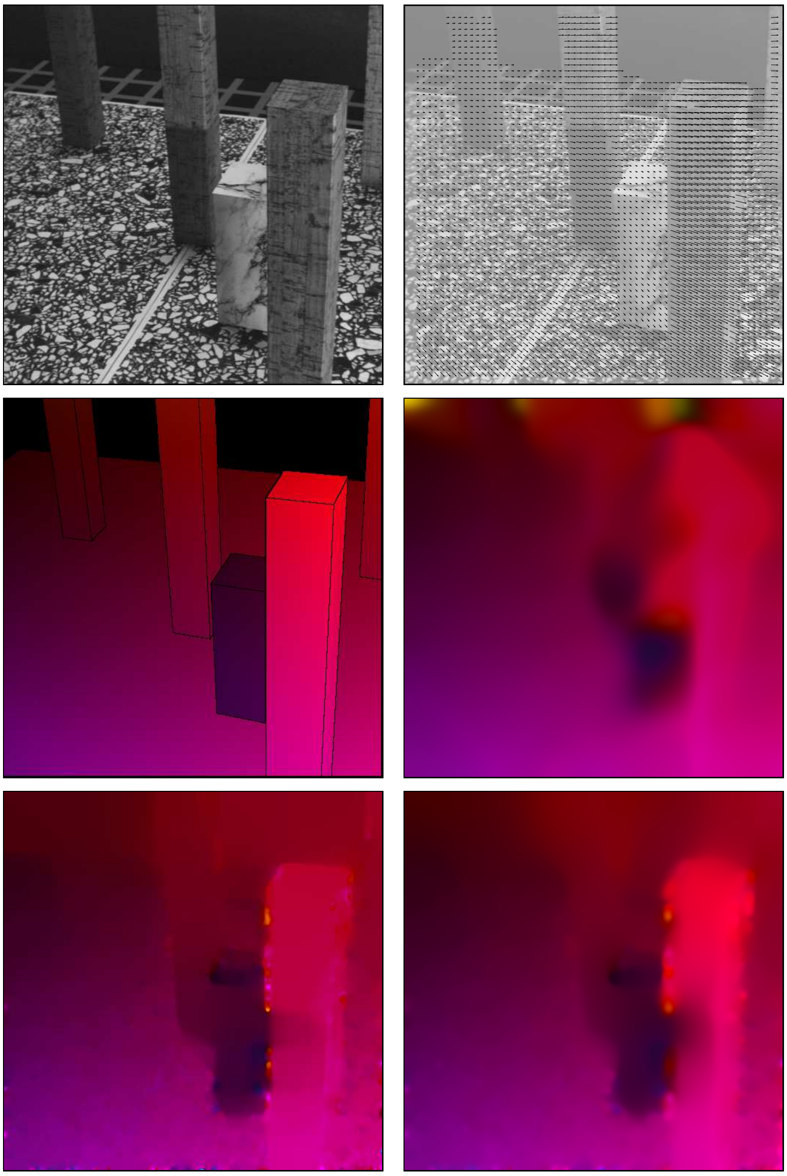

Figure 6: Marble sequence ($512 \times 512 \times 32$). From Left to Right, Top to Bottom: (a) Frame 16. (b) Ground truth (vector plot). (c) Ground truth (colour plot). (d) Lucas–Kanade with linear structure tensor. (e) Lucas–Kanade with isotropic nonlinear structure tensor. (f) Lucas–Kanade with anisotropic nonlinear structure tensor.

in Fig. 6c are excluded from the calculation of the AAE, because there is no ground truth available for these areas.

8 Application to Corner Detection

When looking for some important, distinguished locations of an image, one often considers points where two or more edges meet. Such locations have been named *corners*, *junctions* or *interest points*, and a range of possible approaches exists to detect them in an image; see e.g. the reviews in [38, 41]. The methods based on the structure tensor are well established in this field, so it is interesting to see how the nonlinear structure tensors will perform.

At zero integration scale, the structure tensor J_0 as introduced in (1) or (2) contains information on intrinsically 1-dimensional features of the image, i.e. edges. For grey-scale images, only one eigenvalue of the structure tensor J_0 may attain nonzero values (equal to the squared gradient magnitude), while its corresponding eigenvector represents the gradient direction.

Two-dimensional features of an image (corners) can be detected by integrating the local 1-D information of J_0 within some neighbourhood. The classical method is to smooth J_0 linearly using convolution with a Gaussian, which yields the linear structure tensor. Alternatively, one can consider a nonlinear structure tensor which is obtained by the integration within a data-adaptive neighbourhood by means of nonlinear diffusion. If two differently oriented edges appear in the neighbourhood, the smoothed structure tensor J will possess two nonzero eigenvalues $\lambda_1, \lambda_2 \gg 0$. Several possibilities have been proposed to convert the information from J into a measure of 'cornerness', e.g. by Förstner [16], Harris and Stephens [20], Rohr [37], or Köthe [25]. In our experiments we employ the last approach, and detect corners at local maxima of the smaller eigenvalue of the smoothed structure tensor.

Like in optic flow estimation, we will employ and compare three different smoothing strategies leading to three different versions of the structure tensor:

- Linear smoothing according to (1) with a scale parameter ρ leads to the linear structure tensor J_ρ.
- Isotropic nonlinear diffusion according to (7) with TV diffusivity g^{TV} (Eq. 15) gives the isotropic nonlinear structure tensor J_t^{TV} at time t.
- Anisotropic nonlinear diffusion with a diffusion tensor

  $$D = S\,\mathrm{diag}\,(g^{TV}(\lambda_1), 1)\,S^\top \tag{26}$$

  where λ_1 is the larger eigenvalue of the structure tensor J_ρ calculated from the evolving data U, and S contains the eigenvectors as columns. This process is a combination of a linear smoothing along edges (i.e.

in the direction where fast smoothing and exchange of information is desired), and TV diffusion along the gradient (i.e. smoothing is slower across discontinuities). The resulting structure tensor will be denoted $J^{\mathrm{A}}_{\rho,t}$.

Corner detection using the linear structure tensor J_ρ is the basic choice. It is robust under noise, but the localisation of the detected features is less precise. Because of the linear smoothing, the detected location of a corner tends to shift as the scale ρ increases: see Fig. 8 top.
With isotropic TV flow and J^{TV}_t, the local amount of smoothing is an inverse function of the gradient magnitude. Therefore, the feature blurring and displacement is slowed down when compared to the linear method, and corners remain well localised even for higher diffusion times when any possible noise or small-scale features would have been removed; see the example in the middle of Fig. 8.
As the anisotropic method producing $J^{\mathrm{A}}_{\rho,t}$ prefers smoothing along edges, the exchange of information at places where two edges meet is much faster. A small diffusion time suffices to produce significant corner features which are well localised; see Fig. 8 bottom.
The localisation precision of each method is evaluated on the test image of Fig. 9 left where the ideal locations of corners are known. The parameters of each method were tuned so that they detect the corners reliably and accurately: We evaluate the average distance between the strongest 16 detected points and the 16 ground truth corners. The best result with the linear structure tensor J_ρ gives an average error of 1.92 pixels. The isotropic nonlinear method J^{TV}_t produces a mean error of 1.51 pixels, while the anisotropic structure tensor $J^{\mathrm{A}}_{\rho,t}$ is the most precise corner detector: Its mean error is only 0.97 pixels.
The three methods (without any change of parameters) were then employed at a frequently used 'lab' test image; the results are presented in Fig. 10. We observe that all the three methods detect corners which seem to correspond well to real corners and interest points present in the image. Also in this case, the nonlinear methods outperform their linear counterpart in the precision of corner localisation, and the anisotropic nonlinear corner detector gives the best results. An example is shown at the bottom right of Fig. 10.

9 Conclusions

A number of image processing and computer vision tasks make use of a structure tensor based on Gaussian smoothing, or – equivalently – linear

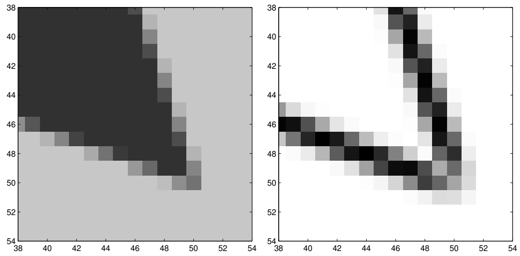

Figure 7: LEFT: detail of a test image with ideal corner position $(50, 50)$. RIGHT: larger eigenvalue of the structure tensor J_0.

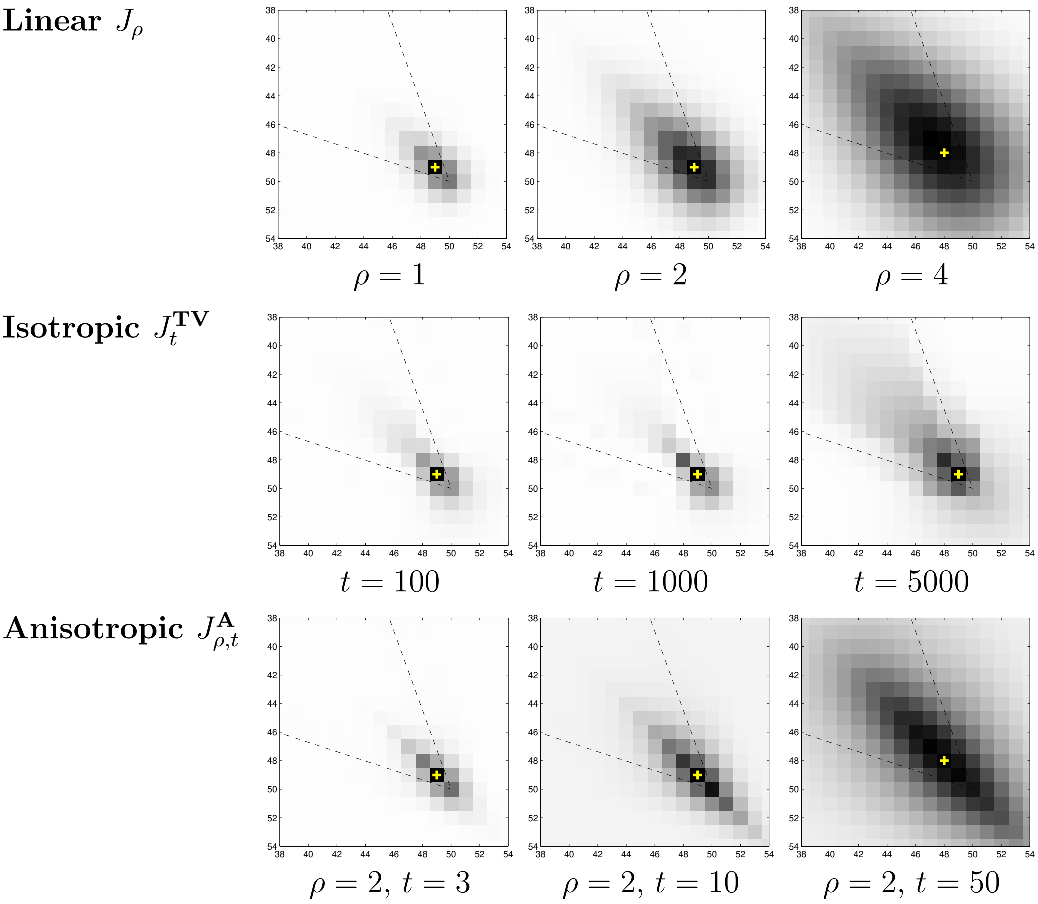

Figure 8: Cornerness measured by the smaller eigenvalue of a smoothed structure tensor J, and the detected corner. TOP: linear smoothing. MIDDLE: isotropic nonlinear diffusion with TV diffusivity. BOTTOM: anisotropic nonlinear diffusion.

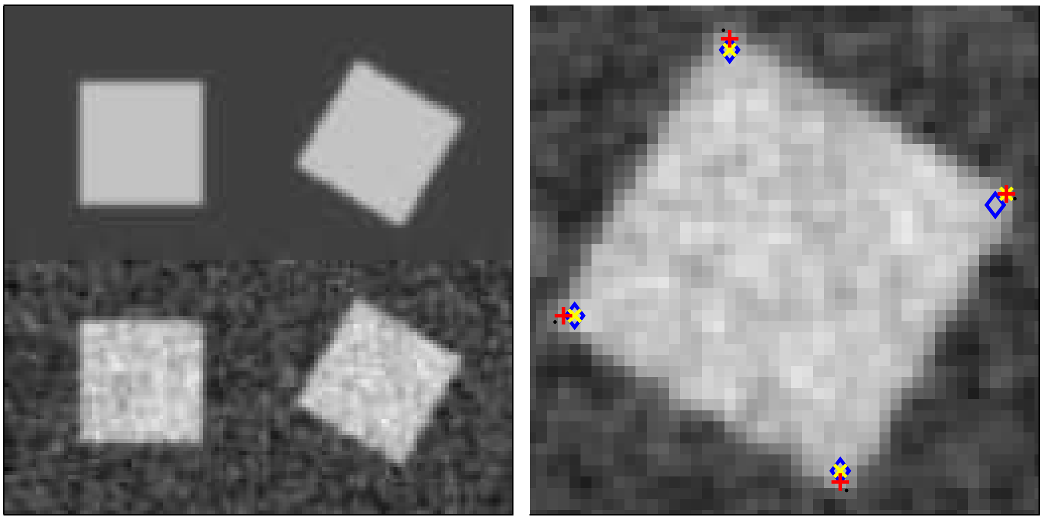

Figure 9: A test image (left) and results of three corner detectors (right, detail).
Blue diamond: linear smoothing, $\rho = 1.5$, mean error 1.92.
Yellow 'x': isotropic smoothing with TV flow, $T = 1400$, mean error 1.51.
Red '+': anisotropic smoothing, $T = 5$, $\rho = 2$, mean error 0.97.
The ideal corner locations are shown by black dots.

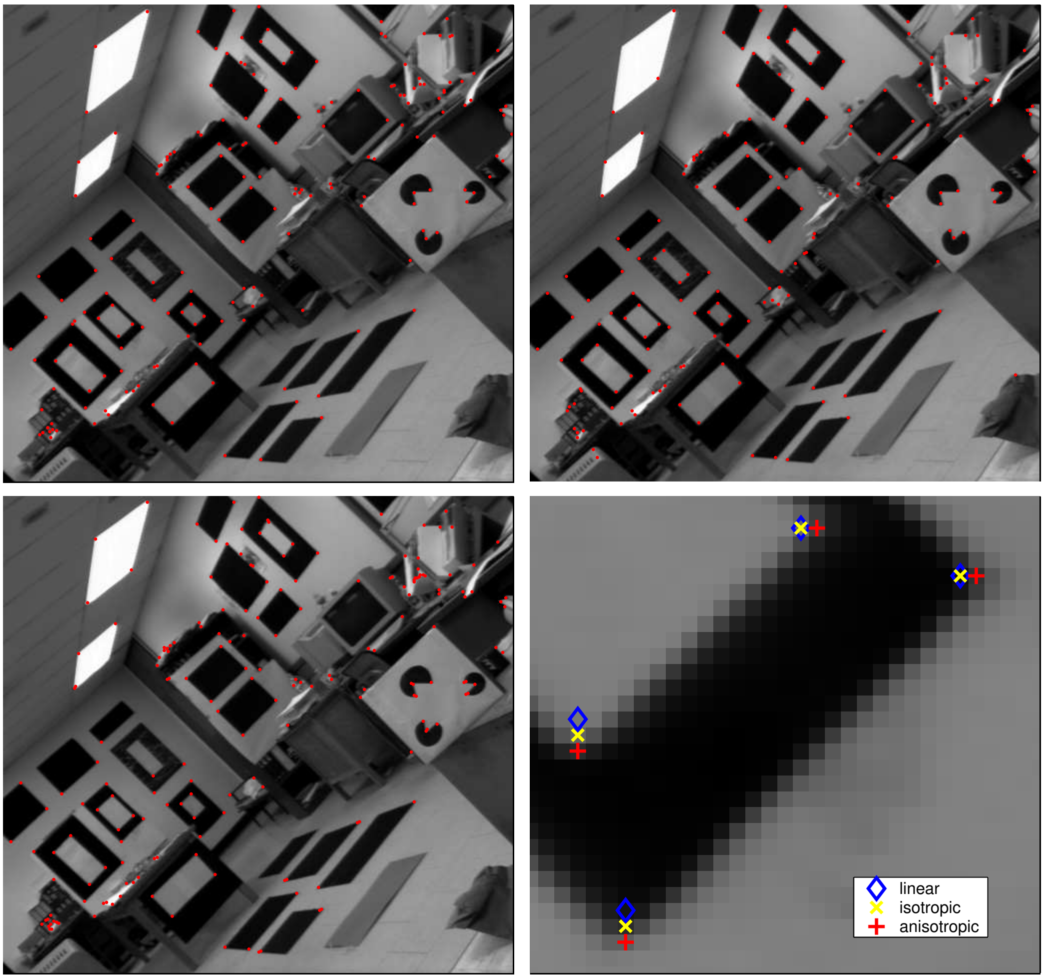

Figure 10: Corners detected in the 'lab' test image. Top left: linear smoothing of the structure tensor. Top right: isotropic TV flow. Bottom left: anisotropic smoothing. Bottom right: a detail for comparison. For all methods, the smoothing parameters are identical to those for the 'squares' test image in Fig. 9, and the 200 strongest corners are shown.

diffusion. Unfortunately, linear diffusion is well-known to destroy important structures such as discontinuities, while other structures may be dislocated. To address these problems, we introduced nonlinear structure tensors that are based on isotropic or anisotropic diffusion filters for matrix-valued data. These data-adaptive smoothing processes avoid averaging of ambiguous structures across discontinuities. Our nonlinear structure tensors contain the conventional linear structure tensor as a special case, and we proved that the matrix-valued nonlinear diffusion filters do not destroy positive semidefiniteness. By using nonlinear diffusion filters with TV diffusivities, nonlinear structure tensors do not involve more parameters than the linear structure tensor. Applying the structure tensor to orientation estimation, optic flow computation, and corner detection allowed a direct comparison between the performance of the linear structure tensor and its nonlinear extensions. The higher accuracy of the results confirmed the superiority of the nonlinear structure tensors. For corner detection, it turned out that specific structure tensors based on anisotropic nonlinear diffusion offer advantages over the ones using isotropic nonlinear diffusion.
We would like to emphasise that these three application fields serve as proof-of-concept only. We are convinced that nonlinear structure tensors are of more general usefulness in all kinds of problems where preservation of discontinuities or avoidance of dislocation effects are desirable, e.g. texture segmentation [39, 8]. In our future research, we intend to perform comparisons and analyse connections between diffusion-based nonlinear structure tensors and other data-adaptive variants of structure tensors. First results in this direction are reported in [9].

Acknowledgements

Our research has partly been funded by the projects WE 2602/1-1 and WE 2602/2-1 of the *Deutsche Forschungsgemeinschaft (DFG)*. This is gratefully acknowledged.

References

[1] L. Alvarez, F. Guichard, P.-L. Lions, and J.-M. Morel. Axioms and fundamental equations in image processing. *Archive for Rational Mechanics and Analysis*, 123:199–257, 1993.

[2] F. Andreu, C. Ballester, V. Caselles, and J. M. Mazón. Minimizing total variation flow. *Differential and Integral Equations*, 14(3):321–360, Mar. 2001.

[3] F. Andreu, V. Caselles, J. I. Diaz, and J. M. Mazón. Qualitative properties of the total variation flow. *Journal of Functional Analysis*, 188(2):516–547, Feb. 2002.

[4] J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, Feb. 1994.

[5] G. Bellettini, V. Caselles, and M. Novaga. The total variation flow in R^N. *Journal of Differential Equations*, 184(2):475–525, 2002.

[6] J. Bigün and G. H. Granlund. Optimal orientation detection of linear symmetry. In *Proc. First International Conference on Computer Vision*, pages 433–438, London, England, June 1987. IEEE Computer Society Press.

[7] J. Bigün, G. H. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790, Aug. 1991.

[8] T. Brox, M. Rousson, R. Deriche, and J. Weickert. Unsupervised segmentation incorporating colour, texture, and motion. In N. Petkov and M. A. Westenberg, editors, *Computer Analysis of Images and Patterns*, volume 2756 of *Lecture Notes in Computer Science*, pages 353–360. Springer, Berlin, 2003.

[9] T. Brox, R. van den Boomgaard, F. Lauze, J. van de Weijer, J. Weickert, P. Mrázek, and P. Kornprobst. Adaptive structure tensors and their applications. In J. Weickert and H. Hagen, editors, *Visualization and Image Processing of Tensor Fields*. Springer, Berlin, 2005. Submitted.

[10] T. Brox and J. Weickert. Nonlinear matrix diffusion for optic flow estimation. In L. Van Gool, editor, *Pattern Recognition*, volume 2449 of *Lecture Notes in Computer Science*, pages 446–453. Springer, Berlin, 2002.

[11] T. Brox, M. Welk, G. Steidl, and J. Weickert. Equivalence results for TV diffusion and TV regularisation. In L. D. Griffin and M. Lillholm, editors, *Scale-Space Methods in Computer Vision*, volume 2695 of *Lecture Notes in Computer Science*, pages 86–100, Berlin, 2003. Springer.

[12] C. Chefd'Hotel, D. Tschumperlé, R. Deriche, and O. Faugeras. Constrained flows of matrix-valued functions: Application to diffusion tensor regularization. In A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, editors, *Computer Vision – ECCV 2002*, volume 2350 of *Lecture Notes in Computer Science*, pages 251–265. Springer, Berlin, 2002.

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

[14] F. Dibos and G. Koepfler. Global total variation minimization. *SIAM Journal on Numerical Analysis*, 37(2):646–664, 2000.

[15] X. Feng and A. Prohl. Analysis of total variation flow and its finite element approximations. Technical Report 1864, Institute of Mathematics and its Applications, University of Minnesota, Minneapolis, MN, July 2002. Submitted to Communications on Pure and Applied Mathematics.

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

[17] W. Förstner and E. Gülch. A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *Proc. ISPRS Intercommission Conference on Fast Processing of Photogrammetric Data*, pages 281–305, Interlaken, Switzerland, June 1987.

[18] G. Gerig, O. Kübler, R. Kikinis, and F. A. Jolesz. Nonlinear anisotropic filtering of MRI data. *IEEE Transactions on Medical Imaging*, 11:221–232, 1992.

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

[20] C. G. Harris and M. Stephens. A combined corner and edge detector. In *Proc. Fouth Alvey Vision Conference*, pages 147–152, Manchester, England, Aug. 1988.

[21] B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

[22] T. Iijima. Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bulletin of the Electrotechnical Laboratory*, 26:368–388, 1962. In Japanese.

[23] B. Jähne. *Spatio-Temporal Image Processing*, volume 751 of *Lecture Notes in Computer Science*. Springer, Berlin, 1993.

[24] S. L. Keeling and R. Stollberger. Nonlinear anisotropic diffusion filters for wide range edge sharpening. *Inverse Problems*, 18:175–190, Jan. 2002.

[25] U. Köthe. Edge and junction detection with an improved structure tensor. In B. Michaelis and G. Krell, editors, *Pattern Recognition*, volume 2781 of *Lecture Notes in Computer Science*, pages 25–32, Berlin, 2003. Springer.

[26] T. Lindeberg. *Scale-Space Theory in Computer Vision*. Kluwer, Boston, 1994.

[27] T. Lu, P. Neittaanmäki, and X.-C. Tai. A parallel splitting up method and its application to Navier–Stokes equations. *Applied Mathematics Letters*, 4(2):25–29, 1991.

[28] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proc. Seventh International Joint Conference on Artificial Intelligence*, pages 674–679, Vancouver, Canada, Aug. 1981.

[29] T. McGraw, B. C. Vemuri, Y. Chen, M. Rao, and T. H. Mareci. DT-MRI denoising and neuronal fiber tracking. *Medical Image Analysis*, 8:95–111, 2004.

[30] M. Middendorf and H.-H. Nagel. Estimation and interpretation of discontinuities in optical flow fields. In *Proc. Eighth International Conference on Computer Vision*, volume 1, pages 178–183, Vancouver, Canada, July 2001. IEEE Computer Society Press.

[31] M. Middendorf and H.-H. Nagel. Empirically convergent adaptive estimation of grayvalue structure tensors. In L. Van Gool, editor, *Pattern Recognition*, volume 2449 of *Lecture Notes in Computer Science*, pages 66–74. Springer, Berlin, 2002.

[32] P. Mrázek, J. Weickert, and A. Bruhn. On robust estimation and smoothing with spatial and tonal kernels. Technical Report 51, Series SPP-1114, Department of Mathematics, University of Bremen, Germany, June 2004.

[33] H.-H. Nagel and A. Gehrke. Spatiotemporally adaptive estimation and segmentation of OF-fields. In H. Burkhardt and B. Neumann, editors, *Computer Vision – ECCV '98*, volume 1407 of *Lecture Notes in Computer Science*, pages 86–102. Springer, Berlin, 1998.

[34] M. Otte and H.-H. Nagel. Estimation of optical flow based on higher-order spatiotemporal derivatives in interlaced and non-interlaced image sequences. *Artificial Intelligence*, 78:5–43, 1995.

[35] P. Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:629–639, 1990.

[36] A. R. Rao and B. G. Schunck. Computing oriented texture fields. *CVGIP: Graphical Models and Image Processing*, 53:157–185, 1991.

[37] K. Rohr. Modelling and identification of characteristic intensity variations. *Image and Vision Computing*, 10(2):66–76, 1992.

[38] K. Rohr. Localization properties of direct corner detectors. *Journal of Mathematical Imaging and Vision*, 4:139–150, 1994.

[39] M. Rousson, T. Brox, and R. Deriche. Active unsupervised texture segmentation on a diffusion based feature space. In *Proc. 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 699–704, Madison, WI, June 2003. IEEE Computer Society Press.

[40] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

[41] S. M. Smith and J. M. Brady. SUSAN: A new approach to low-level image processing. *International Journal of Computer Vision*, 23(1):45–78, May 1997.

[42] J. Sporring, M. Nielsen, L. Florack, and P. Johansen, editors. *Gaussian Scale-Space Theory*, volume 8 of *Computational Imaging and Vision*. Kluwer, Dordrecht, 1997.

[43] I. Thomas. Anisotropic adaptation and structure detection. Technical Report F11, Institute for Applied Mathematics, University of Hamburg, Germany, Aug. 1999.

[44] D. Tschumperlé and R. Deriche. Orthonormal vector sets regularization with PDE's and applications. *International Journal of Computer Vision*, 50(3):237–252, Dec. 2002.

[45] V. I. Tsurkov. An analytical model of edge protection under noise suppression by anisotropic diffusion. *Journal of Computer and Systems Sciences International*, 39(3):437–440, 2000.

[46] R. van den Boomgaard and J. van de Weijer. Robust estimation of orientation for texture analysis. In *Proc. Second International Workshop on Texture Analysis and Synthesis*, Copenhagen, June 2002.

[47] J. Weickert. *Anisotropic Diffusion in Image Processing.* Teubner, Stuttgart, 1998.

[48] J. Weickert. Coherence-enhancing diffusion filtering. *International Journal of Computer Vision*, 31(2/3):111–127, Apr. 1999.

[49] J. Weickert. Nonlinear diffusion filtering. In B. Jähne, H. Haußecker, and P. Geißler, editors, *Handbook on Computer Vision and Applications, Vol. 2: Signal Processing and Pattern Recognition*, pages 423–450. Academic Press, San Diego, 1999.

[50] J. Weickert and T. Brox. Diffusion and regularization of vector- and matrix-valued images. In M. Z. Nashed and O. Scherzer, editors, *Inverse Problems, Image Analysis, and Medical Imaging*, volume 313 of *Contemporary Mathematics*, pages 251–268. AMS, Providence, 2002.

[51] J. Weickert and C. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264, Dec. 2001.

[52] J. Weickert, B. M. ter Haar Romeny, and M. A. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Transactions on Image Processing*, 7(3):398–410, Mar. 1998.