Universität des Saarlandes

Fachrichtung 6.1 – Mathematik

Preprint Nr. 143

PDEs for Tensor Image Processing

Joachim Weickert, Christian Feddern, Martin Welk,
Bernhard Burgeth and Thomas Brox

Saarbrücken 2005

Fachrichtung 6.1 – Mathematik
Universität des Saarlandes

Preprint No. 143
submitted: 28.06.2005

PDEs for Tensor Image Processing

Joachim Weickert
Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Saarland University, Building 27,
66041 Saarbrücken, Germany
weickert@mia.uni-saarland.de

Christian Feddern
Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Saarland University, Building 27,
66041 Saarbrücken, Germany
feddern@mia.uni-saarland.de

Martin Welk
Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Saarland University, Building 27,
66041 Saarbrücken, Germany
welk@mia.uni-saarland.de

Bernhard Burgeth
Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Saarland University, Building 27,
66041 Saarbrücken, Germany
burgeth@mia.uni-saarland.de

Thomas Brox
Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Saarland University, Building 27,
66041 Saarbrücken, Germany
brox@mia.uni-saarland.de

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

Methods based on partial differential equations (PDEs) belong to those image processing techniques that can be extended in a particularly elegant way to tensor fields. In this survey paper the most important PDEs for discontinuity-preserving denoising of tensor fields are reviewed such that the underlying design principles becomes evident. We consider isotropic and anisotropic diffusion filters and their corresponding variational methods, mean curvature motion, and self-snakes. These filters preserve positive semidefiniteness of any positive semidefinite initial tensor field. Finally we discuss geodesic active contours for segmenting tensor fields. Experiments are presented that illustrate the behaviour of all these methods.

Keywords: matrix-valued images, denoising, regularisation, segmentation, partial differential equations, nonlinear diffusion, mean curvature motion, self-snakes, active contours.

Contents

1 Introduction

In the last 15 years, partial differential equations (PDEs) have become increasingly popular in image processing. This has a number of reasons: PDE-based methods are mathematically well-understood techniques, they allow a reinterpretation of several classical methods under a unifying framework, they have led to novel methods with more invariances, and they are the natural framework for scale-space analysis. Moreover, the PDE formulation reflects the continuous structure of space. Thus, PDE approximations aim to be independent of the underlying grid and may reveal good rotational invariance. In a number of image processing and computer vision areas, PDE-based methods and related variational approaches and level-set techniques belong to the best performing methods; see e.g. the books [2, 5, 19, 21, 26, 31] and the references therein.

Interestingly, PDEs are also among the first image processing techniques that have been extended from scalar- and vector-valued images to matrix-valued data. One of the reasons for this is the fact that these extensions are not too difficult, once the scalar-valued processes are mastered.

In this paper we give a survey on some of the most important PDE methods for discontinuity-preserving denoising of tensor images, namely nonlinear diffusion filters and their corresponding regularisation methods, mean curvature motion, and self-snakes. Moreover, we describe an extension of geodesic active contours for tensor images. In order to keep things as simple as possible, we focus on 2-D methods. It should be noted, however, that these concepts can be extended in a natural way to higher dimensions [14]. Parts of our description follow the original papers [13, 14, 32]. We would like to emphasise that we focus on methods for genuine tensor processing. Thus we do not consider scalar- or vector-valued PDE methods working on the eigensystem or filtering channels that are measured prior to computing tensors [11, 28, 29].

The paper is organised as follows. In Section 2 we introduce a generalised structure tensor for matrix fields. It is used for steering all nonlinear PDE methods that are discussed in the course of this paper. Section 3 describes nonlinear diffusion filters for tensor data, both in the isotropic case with a scalar diffusivity as well as in the anisotropic case with a diffusion tensor. Closely related regularisation methods are presented in Section 4. In Section 5 we design a mean curvature type evolution for tensor-valued data. Modifying tensor-valued mean curvature motion by a suitable edge stopping function leads us to tensor-valued self-snakes. They are discussed in Section 6. In Section 7 we use the self-snake model in order to derive geodesic active contours for tensor fields. The paper is concluded with a summary in Section 8.

2 Structure Analysis of Tensor-Valued Data

In this section we generalise the concept of an image gradient to the tensor-valued setting [13]. This may be regarded as a tensor extension of Di Zenzo's method for vector-valued data [12].

Let us consider some rectangular image domain $\Omega \in \mathbb{R}^2$ and some tensor image $F = (f_{i,j}) : \Omega \to \mathbb{R}^{2\times 2}$, where the indices (i,j) specify the tensor channel. We intend to define an "edge direction" for such a matrix-valued function. In the case of some scalar-valued image f, we would look for the direction v which is orthogonal to the gradient of a Gaussian-smoothed version of f:

$$0 = v^\top \nabla f_\sigma \tag{1}$$

where $f_\sigma := K_\sigma * f$ and K_σ denotes a Gaussian with standard deviation σ. Gaussian convolution makes the structure detection more robust against noise. The parameter σ is called *noise scale*.

In the general tensor-valued case, we cannot expect that all tensor channels yield the same edge direction. Therefore we proceed as follows. Let $F_\sigma = (f_{\sigma,i,j})$ be a Gaussian-smoothed version of $F = (f_{i,j})$, where the smoothing is performed componentwise. Then we define the edge direction as the unit vector v that minimises

$$E(v) \;:=\; \sum_{i=1}^{2}\sum_{j=1}^{2}(v^\top \nabla f_{\sigma,i,j})^2 \;=\; v^\top\Big(\sum_{i=1}^{2}\sum_{j=1}^{2}\nabla f_{\sigma,i,j}\nabla f_{\sigma,i,j}^\top\Big)v.$$

This quadratic form is minimised when v is an eigenvector to the smallest eigenvalue of the *structure tensor*

$$J(F_\sigma) := \sum_{i=1}^{2}\sum_{j=1}^{2}\nabla f_{\sigma,i,j}\nabla f_{\sigma,i,j}^\top. \tag{2}$$

The eigenvalues of this positive semidefinite matrix measure the local contrast in the directions of the eigenvectors. Its trace

$$\operatorname{tr} J(F_\sigma) = \sum_{i=1}^{2}\sum_{j=1}^{2}|\nabla f_{\sigma,i,j}|^2 \tag{3}$$

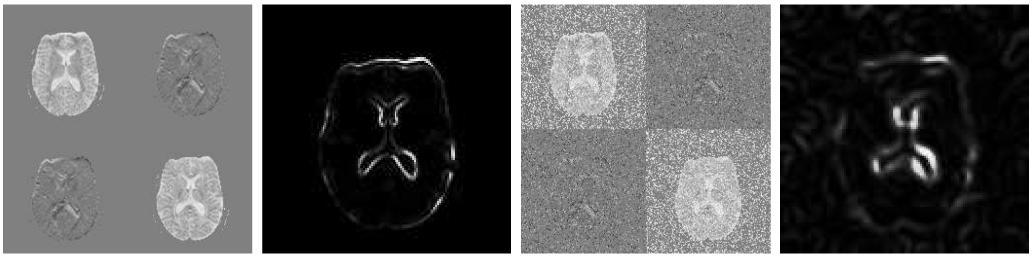

Figure 1: Edge detection with a structure tensor for matrix-valued data. **From left to right: (a)** Original 2-D tensor field extracted from a 3-D DT-MRI data set by using the channels $(1,1)$, $(1,2)$, $(2,1)$ and $(2,2)$. Each channel is of size 128×128. The channels $(1,2)$ and $(2,1)$ are identical for symmetry reasons. **(b)** Trace of the structure tensor of (a) with $\sigma = 1$. **(c)** Noisy version of (a) with 30 % noise. **(d)** Trace of the structure tensor from (c) with $\sigma = 3$. From [14].

sums up all eigenvalues. It can be regarded as a tensor-valued generalisation of the squared gradient magnitude. The matrix $J(F_\sigma)$ will allow us to generalise a number of PDE methods to the tensor-valued setting.

Figure 1 illustrates the concept of edge detection with the structure tensor. The test image we use for our experiments is obtained from a DT-MRI data set of a human brain. We have extracted a 2-D section from the 3-D data. The 2-D image consists of four quadrants which show the four tensor channels of a 2×2 matrix. The top right channel and bottom left channel are identical since the matrix is symmetric. To test the robustness under noise we have replaced 30 % of all data by random matrices: The angles of their eigensystem obey a uniform distribution on $[0, \pi]$, while their eigenvalues are random numbers uniformly distributed in $[0, 127]$. Figure 1 shows the outcome of using $\operatorname{tr} J(F_\sigma)$ for detecting edges in tensor-valued images. We observe that this method gives good results for the original data set. When increasing the noise scale σ, it is also possible to handle situations where substantial noise is present.

3 Diffusion Filtering

3.1 Linear Diffusion

Linear diffusion filtering is the oldest PDE method for image denoising [16]. It creates a family of simplified images $\{u(x,t) \mid t \geq 0\}$ from some scalar initial image $f(x)$ by solving the PDE

$$\partial_t u = \Delta u \quad \text{on} \quad \Omega \times (0, \infty), \tag{4}$$

with f as initial condition,

$$u(x, 0) = f(x) \quad \text{on} \quad \Omega, \tag{5}$$

and reflecting (homogeneous Neumann) boundary conditions:

$$\partial_\nu u = 0 \quad \text{on} \quad \partial\Omega \times (0, \infty). \tag{6}$$

Here ∂_ν denotes differentiation in the direction of the outer normal of the image boundary $\partial\Omega$. The diffusion time t determines the degree of simplification: For $t = 0$ the original image f is recovered, and larger values for t result in more pronounced smoothing. On an infinitely extended image domain, linear diffusion filtering with stopping time T is equivalent to Gaussian convolution with standard deviation $\sigma = \sqrt{2T}$.

It is straightforward to extend linear diffusion filtering to tensor images: All one has to do is to apply this process channelwise.

Figure 2 zooms into the corpus callosum region of Fig. 1(a),(c), and displays the evolution of this region under linear diffusion. The tensors are visualised by ellipses with colour-coded orientation. We observe that linear diffusion is well-suited for removing noise, but suffers from blurring important features such as discontinuities in the tensor field.

3.2 Isotropic Nonlinear Diffusion

The goal of nonlinear diffusion filtering is to smooth an image while respecting its discontinuities [23, 7]. Nonlinear diffusion filtering replaces the linear diffusion equation (4) by

$$\partial_t u = \operatorname{div}\left(g(|\nabla u_\sigma|^2)\, \nabla u\right) \quad \text{on} \quad \Omega \times (0, \infty). \tag{7}$$

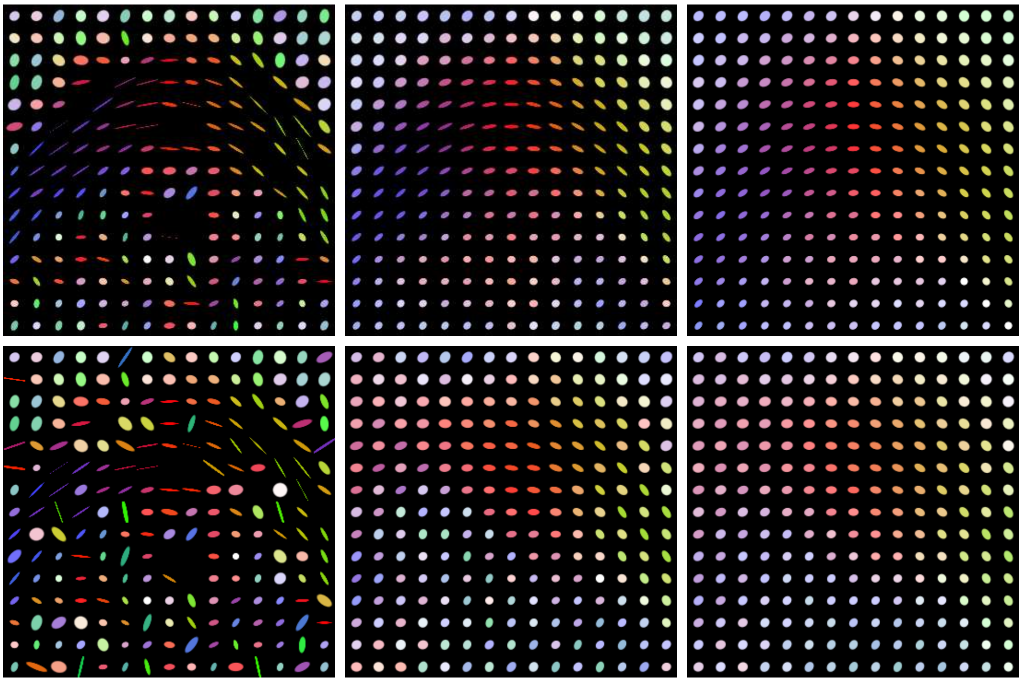

Figure 2: Tensor-valued linear diffusion. **Top row, from left to right:** Detail from a DT-MR image (size 15×15), at time $t = 0.96$, at time $t = 2.4$. **Bottom row, from left to right:** Same experiment with 30 % noise.

The *diffusivity function* g is a decreasing nonnegative function of the squared gradient magnitude of u_σ, a Gaussian smoothed version of u. One may choose e.g. [23]

$$g(|\nabla u_\sigma|^2) \;=\; \frac{1}{1 + |\nabla u_\sigma|^2/\lambda^2} \tag{8}$$

with some contrast parameter $\lambda > 0$. We observe that $|\nabla u_\sigma|^2$ serves as an edge detector: Locations where $|\nabla u_\sigma| \gg \lambda$ are regarded as edges where diffusion is inhibited, while locations with $|\nabla u_\sigma| \ll \lambda$ are considered to belong to the interior of a segment, where full diffusion is performed.

This scalar-valued diffusion scheme can also be generalised for smoothing a matrix field $F = (f_{i,j}) : \Omega \to \mathbb{R}^{2\times 2}$. Tschumperlé and Deriche [28] have proposed a PDE system for matrix-valued diffusion where a joint diffusivity function is used that depends on the trace of the structure tensor (in their case with $\sigma = 0$):

$$\partial_t u_{i,j} = \operatorname{div}\left(g(\operatorname{tr} J(U_\sigma))\, \nabla u_{i,j}\right) \qquad (i,j \in \{1,2\}). \tag{9}$$

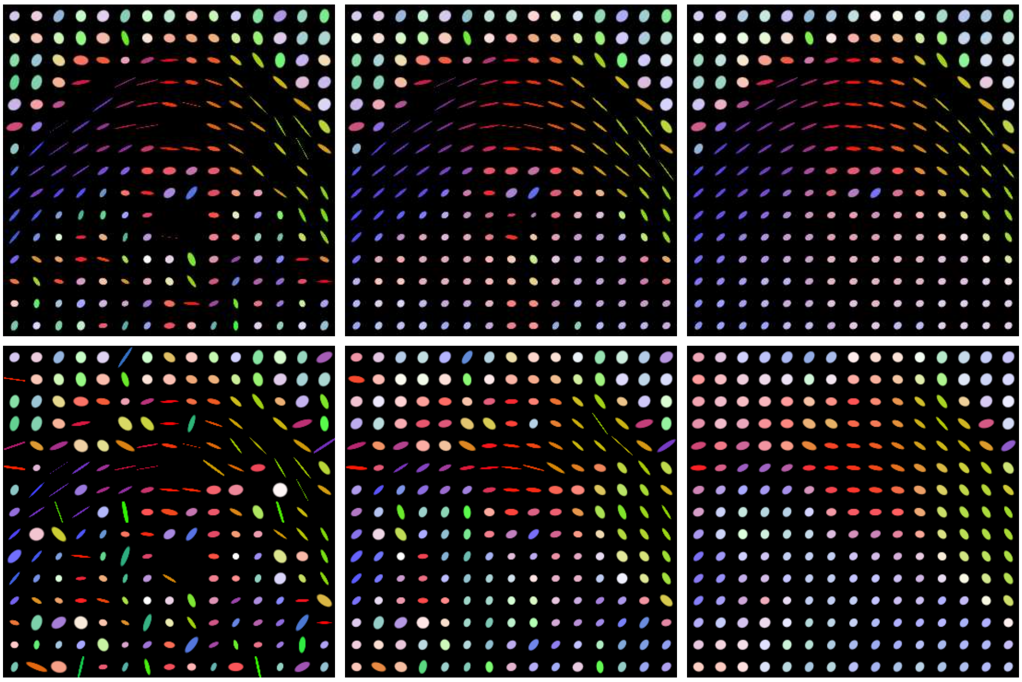

Figure 3: Tensor-valued isotropic nonlinear diffusion. **Top row, from left to right:** Detail from a DT-MR image (size 15×15), at time $t = 4.8$, and at $t = 12.0$ ($\lambda = 0.5$, $\sigma = 0.5$). **Bottom row, from left to right:** Same experiment with 30 % noise ($\lambda = 0.5$, $\sigma = 1$).

The synchronised channel evolution with a joint diffusivity avoids that edges are formed at different locations for the different tensor channels. This synchronisation of channel smoothing is also a frequently used strategy in vector-valued diffusion filtering [15].

Figure 3 illustrates the evolution under isotropic nonlinear diffusion. Discontinuities are well-preserved, but noise at discontinuities is removed rather slowly.

3.3 Anisotropic Nonlinear Diffusion

Besides isotropic diffusion schemes with a scalar-valued diffusivity, 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 to encourage smoothing along discontinuities rather than across them.

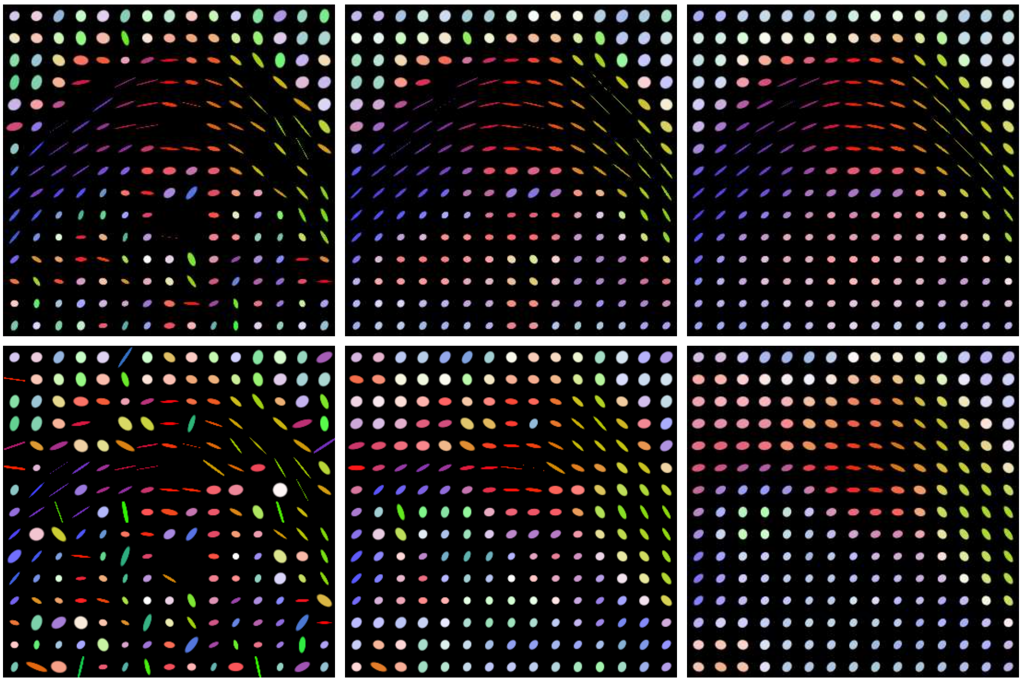

Figure 4: Tensor-valued anisotropic nonlinear diffusion. **Top row, from left to right:** Detail from a DT-MR image (size 15×15), at time $t = 1.92$, and at $t = 4.8$ ($\lambda = 0.5$, $\sigma = 0.5$). **Bottom row, from left to right:** Same experiment with 30 % noise, at time $t = 1.92$ and $t = 4.8$ ($\lambda = 0.5$, $\sigma = 1$).

This can be achieved by replacing the scalar-valued diffusivity function by a matrix-valued diffusion tensor.[1]

Tensor-valued anisotropic diffusion regards the original image $F(x) = (f_{i,j}(x))$ as initial value for the coupled PDE system [32]

$$\partial_t u_{i,j} = \operatorname{div}\left(g(J(U_\sigma))\,\nabla u_{i,j}\right) \qquad (i, j \in \{1, 2\}) \tag{10}$$

subject to the reflecting boundary conditions

$$\partial_\nu \left(g(J(U_\sigma))\,\nabla u_{i,j}\right) = 0 \qquad (i, j \in \{1, 2\}). \tag{11}$$

Here the scalar-valued function g is generalised to a matrix-valued function in the following way: Let $J(U_\sigma) = Q\ \operatorname{diag}(\lambda_i) Q^\top$ denote the principal axis decomposition of $J(U_\sigma)$, with the eigenvalues λ_i as the elements of the diagonal

[1]In the diffusion filtering literature, the word anisotropic is often already used for space-variant diffusion processes with a scalar diffusivity function.

matrix $\mathrm{diag}(\lambda_i)$, and the normalised eigenvectors as the columns of the orthogonal matrix Q. Then it is common to set $g(J(U_\sigma)) := Q\ \mathrm{diag}(g(\lambda_i))Q^\top$, where the scalar diffusivity g is the same decreasing function as in the isotropic case.

Anisotropic diffusion offers the advantage of smoothing in a direction-specific way: Along the i-th eigenvector of $J(U_\sigma)$ 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. Hence, $g(\lambda_i)$ is large and full diffusion is performed. For more information about anisotropic diffusion in general, we refer to [31].

Interestingly, the bare coupling of the tensor channels via a joint diffusion tensor guarantees – without additional projection steps – 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. In the discrete case, this follows from the fact that convex combinations of positive semidefinite matrices are computed [32]. In the continuous case, going from matrices to their quadratic forms allows to prove preservation of positive semidefiniteness by means of a scalar-valued maximum-minimum principle [4]. This reasoning also holds for linear and isotropic nonlinear diffusion. Hence one does not have to consider more sophisticated constrained flows [9] if one is only interested in preserving positive semidefiniteness.

The effects of anisotropic nonlinear diffusion are illustrated in Figure 4. It seems to combine the advantages of linear and isotropic nonlinear diffusion: Noise is removed efficiently while discontinuities are preserved for a long time.

Tensor-valued nonlinear diffusion has also led to nonlinear structure tensors, a refinement of the structure tensor concept itself [4, 3]. They offer advantages for optic flow estimation, texture analysis, and corner detection.

4 Regularisation Methods

Regularisation methods belong to the class of variational approaches for image restoration. Typically one calculates a restoration of some degraded scalar-valued image f as the minimiser of an energy functional

$$E(u) := \int_\Omega \left(|u-f|^2 + \alpha\,\Psi(|\nabla u|^2)\right)\,dx \tag{12}$$

where the penaliser $\Psi : [0, \infty) \to \mathbb{R}$ is an increasing function [20]. The first summand encourages similarity between the restored image and the original one, while the second one rewards smoothness. The smoothness weight $\alpha > 0$ is called *regularisation parameter*. From variational calculus it follows that a minimiser of $E(u)$ satisfies the Euler–Lagrange equation

$$\frac{u - f}{\alpha} \;=\; \operatorname{div}\,\big(\Psi'(|\nabla u|^2)\,\nabla u\big) \tag{13}$$

with homogeneous Neumann boundary conditions. This elliptic PDE can be regarded as a fully implicit time discretisation of the diffusion filter

$$\partial_t u \;=\; \operatorname{div}\,\big(\Psi'(|\nabla u|^2)\,\nabla u\big) \tag{14}$$

with initial image f and stopping time α; see [27] for more details.

In the tensor case, Deriche and Tschumperlé [28] consider the energy

$$E(U) = \int_\Omega \big(\|U - F\|^2 + \alpha\,\Psi(\operatorname{tr} J(U))\big)\,dx \tag{15}$$

where $\|\cdot\|$ is the Frobenius norm for matrices. Then the corresponding Euler–Lagrange equations are given by

$$\frac{u_{i,j} - f_{i,j}}{\alpha} \;=\; \operatorname{div}\,(\Psi'(\operatorname{tr} J(U))\,\nabla u_{i,j}) \qquad (i, j \in \{1, 2\}). \tag{16}$$

They can be regarded as an approximation to the isotropic nonlinear diffusion filter (9) if one chooses $\Psi' := g$ and $\sigma := 0$.

Weickert and Brox [32], on the other hand, consider

$$E(U) = \int_\Omega \big(\|U - F\|^2 + \alpha\,\operatorname{tr}\Psi(J(U))\big)\,dx. \tag{17}$$

It leads to the Euler–Lagrange equation

$$\frac{u_{i,j} - f_{i,j}}{\alpha} \;=\; \operatorname{div}\,(\Psi'(J(U))\,\nabla u_{i,j}) \qquad (i, j \in \{1, 2\}). \tag{18}$$

which is an approximation to the anisotropic nonlinear diffusion filter (10).

These considerations show that regularisation methods are closely related to nonlinear diffusion filtering. In practise they can lead to results that are hardly distinguishable from diffusion results [27]. For this reason we refrain from presenting specific experiments for this filter class.

5 Mean Curvature Motion

In this section we describe tensor-valued mean curvature motion [13]. To this end, we first have to sketch some basic ideas behind scalar-valued mean curvature motion.

We start with the observation that the Laplacian Δu of an isotropic linear diffusion model may be decomposed into two orthogonal directions $\xi \perp \nabla u$ and $\eta \parallel \nabla u$:

$$\partial_t u \; = \; \Delta u \; = \; \partial_{\xi\xi} u + \partial_{\eta\eta} u \tag{19}$$

where $\partial_{\xi\xi} u$ describes smoothing parallel to edges and $\partial_{\eta\eta}$ smoothes perpendicular to edges. *Mean curvature motion (MCM)* uses an anisotropic variant of this smoothing process by permitting only smoothing along the level lines:

$$\partial_t u \; = \; \partial_{\xi\xi} u. \tag{20}$$

This can be rewritten as

$$\partial_t u \; = \; |\nabla u| \; \mathrm{div} \left(\frac{\nabla u}{|\nabla u|} \right). \tag{21}$$

Alvarez et al. have used this evolution equation for denoising highly degraded images [1]. It is well-known from the mathematical literature that under MCM convex level lines remain convex, nonconvex ones become convex, and in finite time they vanish by approximating circular shapes while converging to points.

If we want to use an MCM-like process for processing tensor-valued data $F = (f_{i,j})$, it is natural to replace the second directional derivative $\partial_{\xi\xi} u$ in (20) by $\partial_{vv} u$, where v is the eigenvector to the smallest eigenvalue of the structure tensor $J(U)$. This leads us to the evolution

$$\begin{aligned} \partial_t u_{i,j} = \partial_{vv} u_{i,j} \qquad & \text{on} \qquad \Omega \times (0, \infty) & (22)\\ u_{i,j}(x, 0) = f_{i,j}(x) \qquad & \text{on} \qquad \Omega, & (23)\\ \partial_\nu u_{i,j} = 0 \qquad & \text{on} \qquad \partial\Omega \times (0, \infty) & (24) \end{aligned}$$

for all tensor channels (i, j). Note that this process synchronises the smoothing direction in all channels. It may be regarded as a tensor-valued generalisation of the vector-valued mean curvature motion by Chambolle [8].

Our notion of edge directions as eigenvectors of the structure tensor is equivalent to the generalised level lines in [10]. Hence, tensor-valued MCM can be regarded as smoothing along these lines. Reinterpreting the concept of curve

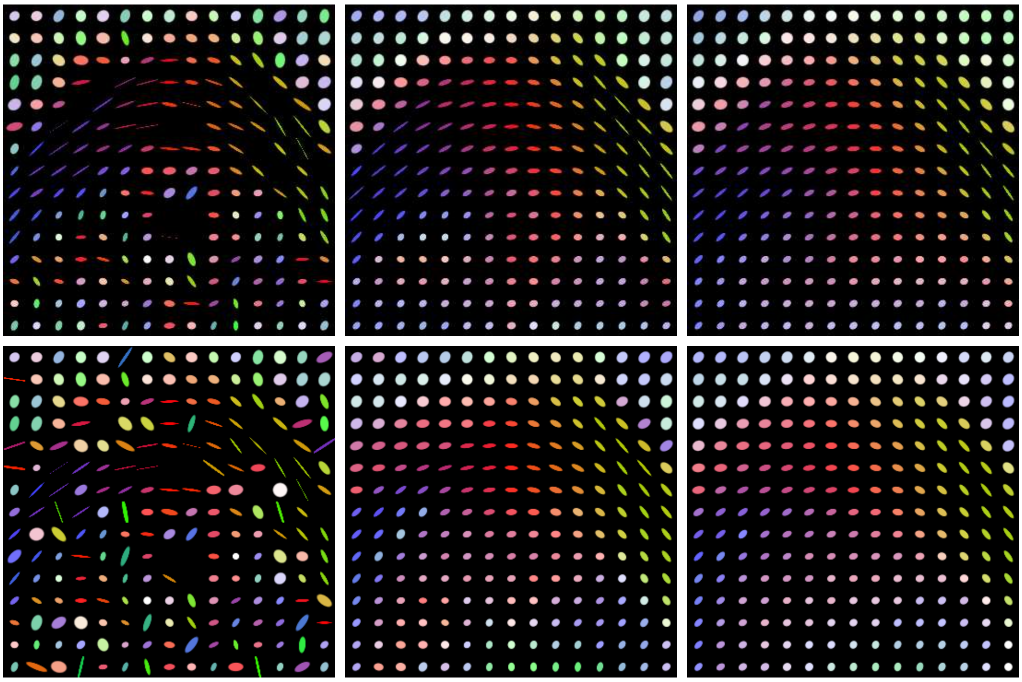

Figure 5: Tensor-valued mean curvature motion. **Top row, from left to right:** Detail of size 15×15 from a DT-MR image; at time $t = 2.4$, at time $t = 6$. **Bottom row, from left to right:** Same experiment with 30 % noise.

evolution in the same spirit allows even to gain a theoretical foundation of MCM as shortening flow for its level lines [14]. It is also possible to show that tensor MCM preserves positive semidefiniteness [14].

Figure 5 illustrates the tensor-valued mean curvature model. As can be seen in the first row, this process regularises the tensor field while it is capable of respecting anisotropies in a better way than linear diffusion. The second row of Figure 5 shows the same algorithm applied to the noisy image. It displays a fairly high robustness to noise: For increasing evolution times the results for the original and the noisy images approach each other.

6 Self-Snakes

In [25], Sapiro has proposed a specific variant of MCM that is well-suited for image enhancement. This process – which he names *self-snakes* – introduces

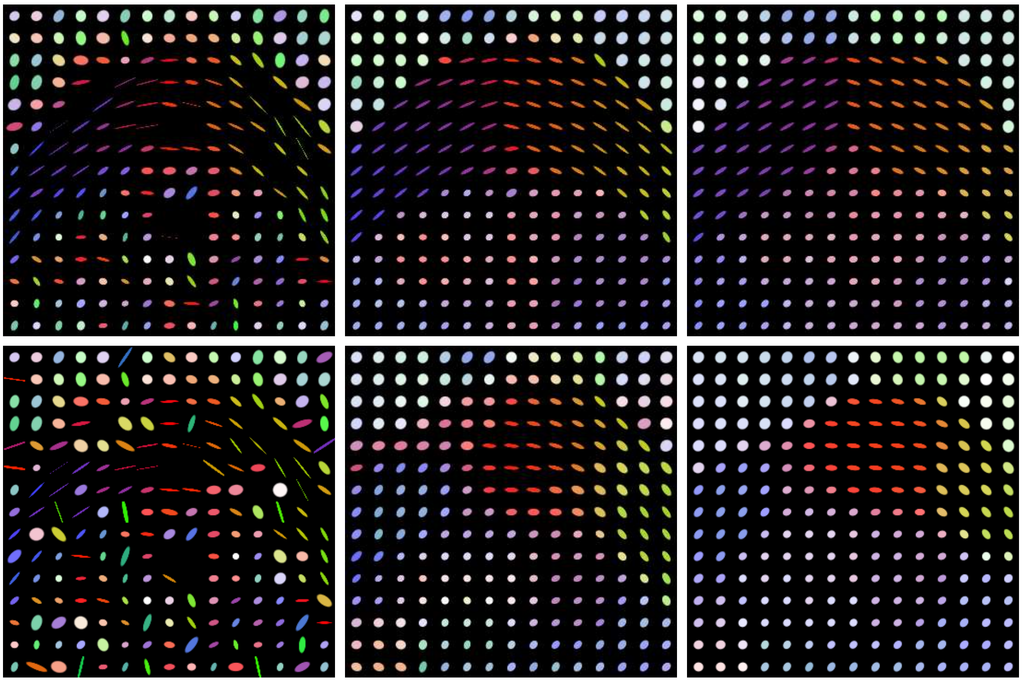

Figure 6: Tensor-valued self-snakes. **Top row, from left to right:** Detail of size 15×15 from a DT-MR image; at time $t = 2.4$, at time $t = 6$ ($\sigma = 0.5$, $\lambda = 2$). **Bottom row, from left to right:** Same experiment with 30 % noise ($\sigma = 1$, $\lambda = 2$).

an edge-stopping function into mean curvature motion in order to prevent further shrinkage of the level lines once they have reached important image edges. In the scalar-valued setting, a self-snake $u(x,t)$ of some image $f(x)$ is generated by the evolution process

$$\partial_t u = |\nabla u| \operatorname{div}\left(g(|\nabla u_\sigma|^2)\,\frac{\nabla u}{|\nabla u|}\right) \quad \text{on} \quad \Omega \times (0,\infty), \tag{25}$$

$$u(x,0) = f(x) \quad \text{on} \quad \Omega, \tag{26}$$

$$\partial_\nu u = 0 \quad \text{on} \quad \partial\Omega \times (0,\infty), \tag{27}$$

where g is a decreasing function such as the diffusivity (8). Self-snakes have been advocated as alternatives to nonlinear diffusion filters [33], they can be used for vector-valued images [25], and related processes have also been proposed for filtering 3-D images [24].

Using the product rule of differentiation, we may rewrite (25) as

$$\partial_t u \;=\; g(|\nabla u_\sigma|^2)\,\partial_{\xi\xi} u + \nabla^\top(g(|\nabla u_\sigma|^2))\,\nabla u \tag{28}$$

with $\nabla^\top := (\partial_x, \partial_y)$. This formulation suggests a straightforward generalisation to the tensor-valued setting. All we have to do is to replace $|\nabla u_\sigma|^2$ by $\operatorname{tr} J(U_\sigma)$, and $\partial_{\xi\xi}$ by ∂_{vv}, where v is the eigenvector to the smallest eigenvalue of $J(U)$. This leads us to the following tensor-valued PDE:

$$\partial_t u_{i,j} \;=\; g(\operatorname{tr} J(U_\sigma))\, \partial_{vv} u_{i,j} + \nabla^\top (g(\operatorname{tr} J(U_\sigma)))\, \nabla u_{i,j}. \tag{29}$$

We observe that the main difference to tensor-valued MCM consists in the additional term $\nabla^\top (g(\operatorname{tr} J(U_\sigma)))\, \nabla u_{i,j}$. It can be regarded as a shock term [22] that is responsible for the edge-enhancing properties of self-snakes.

With only minor modifications, it is possible to extend the semidefiniteness preservation proof for tensor-valued MCM also to the case of tensor-valued self-snakes.

Experimental results for the tensor-valued self-snake technique are shown in Figure 6. Compared to tensor-valued MCM, self-snakes offer increased sharpness at discontinuities due to the additional shock term. The filtered tensor fields look segmentation-like.

7 Geodesic Active Contour Models

Active contours go back to Kass et al. [17]. They play an important role in interactive image segmentation, in particular for medical applications. The underlying idea is that the user specifies an initial guess of an interesting contour (organ, tumour, ...). Then this contour is moved by image-driven forces to the edges of the object in question.

So-called geodesic active contour models [6, 18] achieve this by applying a specific kind of level set ideas. In its simplest form, a geodesic active contour model consists of the following steps. One embeds the user-specified initial curve $C_0(s)$ as a zero level curve into a function $f(x)$, for instance by using the distance transformation. Then f is evolved under a PDE that includes knowledge about the original image h:

$$\partial_t u = |\nabla u| \operatorname{div}\left(g(|\nabla h_\sigma|^2)\, \frac{\nabla u}{|\nabla u|} \right) \qquad \text{on} \qquad \Omega \times (0, \infty), \tag{30}$$

$$u(x, 0) = f(x), \qquad \text{on} \qquad \Omega, \tag{31}$$

$$\partial_\nu u = 0 \qquad \text{on} \qquad \partial\Omega \times (0, \infty), \tag{32}$$

where g inhibits evolution at edges of f. One may choose decreasing functions such as the diffusivity (8). Experiments indicate that, in general, (30) will

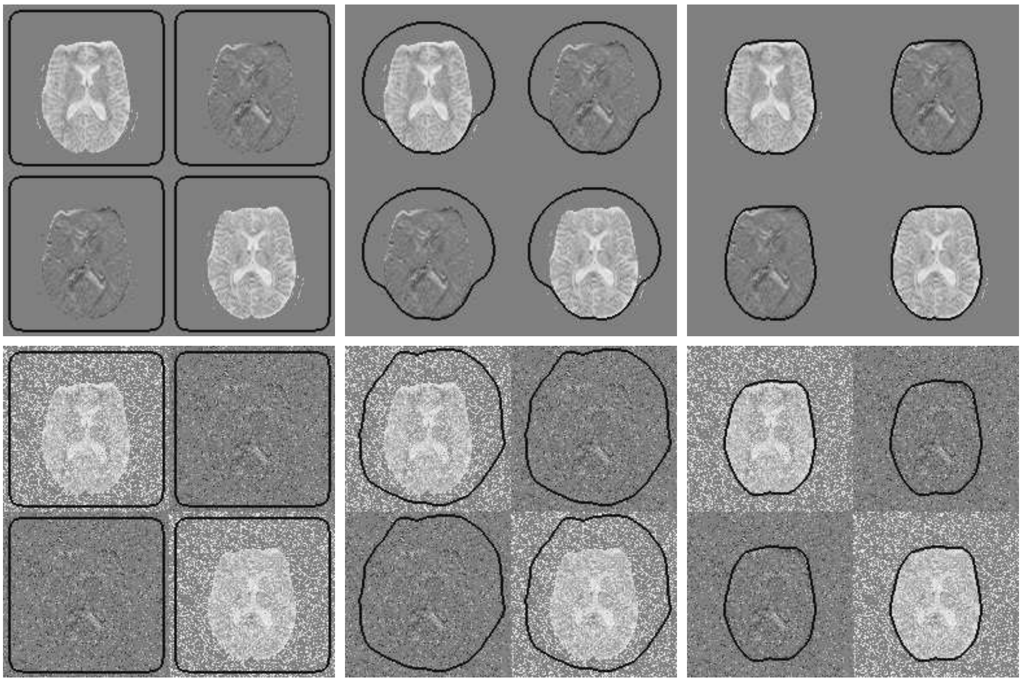

Figure 7: Tensor-valued geodesic active contours ($\sigma = 3$, $\lambda = 1$). **Top row, from left to right:** Tensor image of size 128×128 including contour at time $t = 0$, $t = 960$ and $t = 9600$. **Bottom row, from left to right:** Same experiment with 30 % noise. From [14].

have nontrivial steady states. The evolution is stopped at some time T, when the process does hardly alter anymore, and the final contour C is extracted as the zero level curve of $u(x, T)$. This contour turns out to be a shortest path with respect to an image-induced metric which motivates the notion of *geodesic* active contours.

In [13] geodesic active contours have been extended to tensor valued data $H = (h_{i,j})$ by using $\mathrm{tr}\,(J(H_\sigma))$ as argument of the stopping function g:

$$\partial_t u \;=\; |\nabla u| \, \mathrm{div} \left(g(\mathrm{tr}\, J(H_\sigma)) \, \frac{\nabla u}{|\nabla u|} \right). \qquad (33)$$

Note that, in contrast to the processes in the previous section, this equation is still scalar-valued, since the goal is to find a contour that segments all channels simultaneously. In [14] it is shown that the final contour is a geodesic in a metric that now depends on $J(H_\sigma)$. The PDE (33) may also be rewritten as

$$\partial_t u \;=\; g(\mathrm{tr}\, J(H_\sigma)) \, \partial_{\xi\xi} u \;+\; \nabla^\top (g(\mathrm{tr}\, J(H_\sigma))) \, \nabla u. \qquad (34)$$

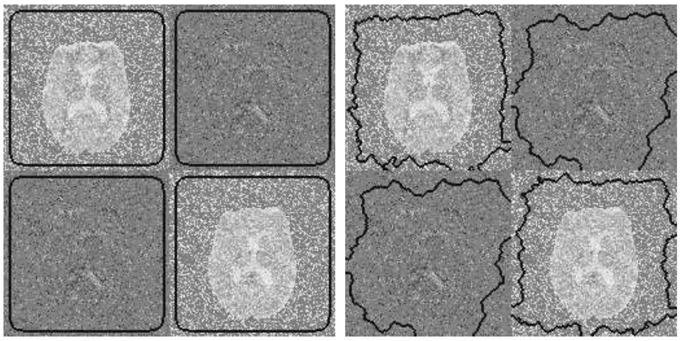

Figure 8: **Left:** Noisy tensor image from Fig. 7 with initial contour. **Right:** Result when no channel coupling is used. ($\sigma = 3$, $\lambda = 1$, $t = 9600$). From [14].

Since a tensor-valued image involves more channels than a scalar-valued one, we can expect that this additional information stabilises the process when noise is present.

Figure 7 shows the temporal evolution of the active contour model for tensor fields. The goal was to extract the contour of the human brain shown in the original image. First one notices that the evolution is slower in the noisy case. The reason is that noise creates large values in the trace of the structure tensor. As a consequence, the evolution is slowed down. For larger times, however, both results become very similar. This shows the high noise robustness of the active contour model for tensor-valued data sets. A comparison with an uncoupled active contour model in Figure 8 demonstrates the crucial role of channel coupling.

Alternative active contour models based on the Mumford–Shah functional have been considered recently in [30].

8 Summary and Conclusions

We have surveyed a number of discontinuity-preserving PDEs for denoising or segmenting tensor fields. They include isotropic and anisotropic nonlinear diffusion and their corresponding regularisation methods, mean curvature motion, self-snakes, and geodesic active contours. We have seen that they arise as natural extensions of their scalar-valued predecessors provided that uniform design principles are obeyed: The evolution of the different channels has to be coupled by a joint diffusivity, diffusion tensor or smoothing direction. Instead of adapting the nonlinear PDEs to the evolving image structure via the squared gradient magnitude, the trace of a structure tensor is used. In those cases where the edge direction is required, it is replaced by the eigenvector direction for the smallest eigenvalue of the structure tensor. Apart from these natural design principles, PDEs for tensor fields offer additional qualities: Their continuous nature supports rotationally invariant models. Last but not least, by virtue of the channel coupling a maximum–minimum principle for the scalar-valued PDEs translates into preservation of positive semidefiniteness in the tensor setting.

Acknowledgements

We are grateful to Anna Vilanova i Bartrolí (Eindhoven Institute of Technology) and Carola van Pul (Maxima Medical Center, Eindhoven) for providing us with the DT-MRI dataset and for discussing questions concerning data conversion. Our research is partly funded by the *DFG* under the projects We2602/1-1 and We2602/1-2.

References

[1] L. Alvarez, P.-L. Lions, and J.-M. Morel. Image selective smoothing and edge detection by nonlinear diffusion. II. *SIAM Journal on Numerical Analysis*, 29:845–866, 1992.

[2] G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*, volume 147 of *Applied Mathematical Sciences*. Springer, New York, 2002.

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

[4] T. Brox, J. Weickert, B. Burgeth, and P. Mrázek. Nonlinear structure tensors. Technical Report 113, Dept. of Mathematics, Saarland University, Saarbrücken, Germany, October 2004.

[5] F. Cao. *Geometric Curve Evolution and Image Processing*, volume 1805 of *Lecture Notes in Mathematics*. Springer, Berlin, 2003.

[6] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22:61–79, 1997.

[7] F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal on Numerical Analysis*, 32:1895–1909, 1992.

[8] A. Chambolle. Partial differential equations and image processing. In *Proc. 1994 IEEE International Conference on Image Processing*, volume 1, pages 16–20, Austin, TX, November 1994. IEEE Computer Society Press.

[9] 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.

[10] D. H. Chung and G. Sapiro. On the level lines and geometry of vector-valued images. *IEEE Signal Processing Letters*, 7(9):241–243, 2000.

[11] O. Coulon, D. C. Alexander, and S. A. Arridge. Diffusion tensor magnetic resonance image regularisation. *Medical Image Analysis*, 8(1):47–67, 2004.

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

[13] C. Feddern, J. Weickert, and B. Burgeth. Level-set methods for tensor-valued images. In O. Faugeras and N. Paragios, editors, *Proc. Second IEEE Workshop on Geometric and Level Set Methods in Computer Vision*, pages 65–72, Nice, France, October 2003. INRIA.

[14] C. Feddern, J. Weickert, B. Burgeth, and M. Welk. Curvature-driven PDE methods for matrix-valued images. *International Journal of Computer Vision*, 2005. In press.

[15] 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.

[16] 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.

[17] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1988.

[18] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Conformal curvature flows: from phase transitions to active vision. *Archive for Rational Mechanics and Analysis*, 134:275–301, 1996.

[19] R. Kimmel. *Numerical Geometry of Images: Theory, Algorithms, and Applications*. Springer, New York, 2003.

[20] N. Nordström. Biased anisotropic diffusion – a unified regularization and diffusion approach to edge detection. *Image and Vision Computing*, 8:318–327, 1990.

[21] S. Osher and N. Paragios, editors. *Geometric Level Set Methods in Imaging, Vision and Graphics*. Springer, New York, 2003.

[22] S. Osher and L. I. Rudin. Feature-oriented image enhancement using shock filters. *SIAM Journal on Numerical Analysis*, 27:919–940, 1990.

[23] 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.

[24] T. Preußer and M. Rumpf. A level set method for anisotropic geometric diffusion in 3D image processing. *SIAM Journal on Applied Mathematics*, 62(5):1772–1793, 2002.

[25] G. Sapiro. Vector (self) snakes: a geometric framework for color, texture and multiscale image segmentation. In *Proc. 1996 IEEE International Conference on Image Processing*, volume 1, pages 817–820, Lausanne, Switzerland, September 1996.

[26] G. Sapiro. *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press, Cambridge, UK, 2001.

[27] O. Scherzer and J. Weickert. Relations between regularization and diffusion filtering. *Journal of Mathematical Imaging and Vision*, 12(1):43–63, February 2000.

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

[29] B. Vemuri, Y. Chen, M. Rao, T. McGraw, Z. Wang, and T. Mareci. Fiber tract mapping from diffusion tensor MRI. In *Proc. First IEEE Workshop on Variational and Level Set Methods in Computer Vision*, pages 73–80, Vancouver, Canada, July 2001. IEEE Computer Society Press.

[30] Z. Wang and B. C. Vemuri. An affine invariant tensor dissimilarity measure and its applications to tensor-valued image segmentation. In *Proc. 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 228–233, Washington, DC, June 2004. IEEE Computer Society Press.

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

[32] 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.

[33] R. T. Whitaker and X. Xue. Variable-conductance, level-set curvature for image denoising. In *Proc. 2001 IEEE International Conference on Image Processing*, pages 142–145, Thessaloniki, Greece, October 2001.