

From Tensor-Driven Diffusion to Anisotropic Wavelet Shrinkage

Martin Welk¹, Joachim Weickert¹, and Gabriele Steidl²

¹ Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science,
Saarland University, 66041 Saarbrücken, Germany
{welk, weickert}@mia.uni-saarland.de
<http://www.mia.uni-saarland.de>

² Faculty of Mathematics and Computer Science,
A5 University of Mannheim, 68131 Mannheim, Germany
steidl@math.uni-mannheim.de
<http://kiwi.math.uni-mannheim.de>

Abstract. Diffusion processes driven by anisotropic diffusion tensors are known to be well-suited for structure-preserving denoising. However, numerical implementations based on finite differences introduce unwanted blurring artifacts that deteriorate these favourable filtering properties. In this paper we introduce a novel discretisation of a fairly general class of anisotropic diffusion processes on a 2-D grid. It leads to a locally semi-analytic scheme (LSAS) that is absolutely stable, simple to implement and offers an outstanding sharpness of filtered images. By showing that this scheme can be translated into a 2-D Haar wavelet shrinkage procedure, we establish a connection between tensor-driven diffusion and anisotropic wavelet shrinkage for the first time. This result leads to coupled shrinkage rules that allow to perform highly anisotropic filtering even with the simplest wavelets.

1 Introduction

Anisotropy originates from physics where it describes a direction-dependent behaviour of material properties. In image analysis, anisotropic filters that act direction-adaptive are an adequate framework to process oriented structures such as edges¹. Since oriented features play a central role in many computer vision applications, it is not surprising that much research on anisotropic filtering has been carried out in the last decade.

One class of methods where anisotropy is used are anisotropic diffusion filters with a matrix-valued diffusion tensor instead of a scalar-valued diffusivity; see e.g. [25]. They include *edge-enhancing diffusion (EED)* that denoises images isotropically within regions and smoothes anisotropically along image edges, and *coherence-enhancing diffusion (CED)* that processes flow-like structures by smoothing along the flow direction.

¹ Sometimes the notion *anisotropic* is already used for space-variant filtering; see e.g. [18]. In our nomenclature, such a filter would be called *isotropic*.

Also in the wavelet community many efforts have been made to incorporate anisotropy in order to represent and process oriented structures in a better way, e.g. by *contourlets* [9], *ridgelets* [12] and *curvelets* [4]. They take the form of frame elements that exhibit very high directional sensitivity and are highly anisotropic.

Initially, anisotropic concepts have been derived in a *continuous* setting where they can be described most elegantly. However, in order to apply anisotropic filters to digital images, one has to find adequate *discrete* representations for them. In practice this may create substantial problems with respect to rotation invariance, since the digital geometry allows to represent only a very restricted set of directions in a precise manner. If a filter is supposed to perform e.g. smoothing along an arbitrary one-dimensional structure, even slight directional errors can introduce blurring artifacts that severely deteriorate its performance. Therefore, research became necessary to find adequate discrete realisations of anisotropic filters, both in the diffusion setting [24, 26] and in the wavelet framework [3, 10].

The goal of the present paper is to address these problems by deriving novel discrete anisotropic filters that unify the concepts of tensor-driven diffusion and anisotropic wavelet shrinkage. We start by presenting a new scheme for anisotropic diffusion filtering that uses solutions of diffusion processes on 2×2 pixel images with a fixed diffusion tensor as building blocks. We show that this so-called *locally semi-analytic scheme (LSAS)* is absolutely stable, that it is simple to implement, that it gives an excellent approximation of rotation invariance, and that it hardly suffers from numerical blurring artifacts. Afterwards we interpret this scheme as a new strategy for anisotropic shift-invariant Haar wavelet shrinkage on a single scale. This leads to novel, anisotropic shrinkage rules with coupling of the coefficients.

Our paper is organised as follows. In Section 2 we derive our method as a novel scheme for anisotropic diffusion filtering, while Section 3 is devoted to its interpretation in the wavelet context. Experimental results are presented in Section 4, and the paper is concluded with a summary in Section 5.

Related work. Early schemes for anisotropic, tensor-driven diffusion such as [14, 19, 25] did not pay specific attention to the problem of rotation invariance and avoidance of blurring artifacts. Weickert and Schar [26] addressed these problems by a scheme for coherence-enhancing diffusion filtering that uses optimised, Sobel-like approximations of all first order spatial derivatives. However, no stability theory was presented, and experiments showed only conditional stability. The same holds for the modified scheme of Wang [24] who used Simoncelli's derivative approximations [21] instead. Moreover, both schemes require stencil sizes of at least 5×5 pixels, while the scheme in the present paper is absolutely stable and comes down to a more local 3×3 stencil.

We notice that constructing numerical methods for diffusion filters from analytic solutions of simpler systems is also a feature of the method of short-time kernels, see e.g. [22], where a locally linearised diffusion equation is solved by Gaussian convolution.

Much research on relations between PDE-based filters and wavelets has been carried out in the *continuous* setting; see e.g. [1, 2, 5, 6, 15, 20]. Work on the relations between wavelet shrinkage and PDE-based denoising in the *discrete* framework include a pa-

per by Coifman and Sowa [8] where they proposed total variation (TV) diminishing flows that act along the direction of Haar wavelets. Weickert et al. describes connections between (semi-)discrete diffusion filtering and Haar wavelet shrinkage, including a locally analytic four-pixel scheme, but focussed on the 1-D or the isotropic 2-D case with a scalar-valued diffusivity; see [27] and the references therein. To the best of our knowledge, however, nobody has found connections between nonlinear diffusion and wavelet shrinkage in the practically relevant anisotropic case so far. With respect to its four-pixel building blocks, our scheme can be regarded as an anisotropic, 2-D extension of the 1-D two-pixel scheme of Steidl et al. [23], and the 2-D isotropic four-pixel scheme of Welk et al. [28]. It is also an anisotropic extension of the equivalence results between discrete diffusion filtering and single scale Haar wavelet shrinkage that have been established by Mrázek and Weickert in the 1-D case [17] and in the isotropic setting [16].

2 A Local Discretisation of Anisotropic Diffusion with Low Numerical Blurring

We consider a nonlinear anisotropic diffusion equation [25]

$$\partial_t u = \operatorname{div} (D(J) \cdot \nabla u) \tag{1}$$

where $D(J)$ is an anisotropic diffusion tensor which depends on the image via the so-called *structure tensor* [13]

$$J = J_\rho(\nabla u_\sigma) := K_\rho * (\nabla(K_\sigma * u) \nabla(K_\sigma * u)^T). \tag{2}$$

Here, K_ρ and K_σ are Gaussian convolution kernels. This equation can model a wide variety of anisotropic diffusion processes, including EED and CED, by adjusting the parameters ρ, σ , and the dependence of D on the structure tensor J .

In order to discretise (1) and (2) in a way that introduces as little numerical blurring artifacts as possible, we will base our discretisation of J on *four-pixel cells* consisting of 2×2 pixels. Furthermore, our discretisation of the divergence expression will allow for a decomposition into approximations on these cells.

2.1 Discretisation of the Diffusion Tensor

Discretising the diffusion tensor D means to discretise the structure tensor J . The main step herein is the discretisation of the gradients ∇v of the given pre-smoothed image $v := K_\sigma * u$. As for nonlinear isotropic diffusion [28] a good location to discretise these quantities most locally is in the centre of a four-pixel cell. We aim therefore at discretising $\nabla v = (\partial_x v, \partial_y v)^T$, and thus $\nabla v \nabla v^T$, at the centre $(\frac{3}{2}, \frac{3}{2})$ of a four-pixel cell $\{v_{ij}\}_{i,j=1,2}$ from a sampling of the spatial function v .

First, $\partial_x v$ and $\partial_y v$ can be approximated from the given pixels $v_{11}, v_{12}, v_{21}, v_{22}$ by central differences at midpoints between neighbouring pixel positions. By considering a quadratic grid with grid size 1 and taking arithmetic means of these expressions, we obtain approximations for the derivatives at $(\frac{3}{2}, \frac{3}{2})$:

$$\begin{aligned}
 (\partial_x v)_{\frac{3}{2}, \frac{3}{2}} &\approx \frac{1}{2}(v_{2,2} + v_{2,1} - v_{1,2} - v_{1,1}), \\
 (\partial_y v)_{\frac{3}{2}, \frac{3}{2}} &\approx \frac{1}{2}(v_{2,2} - v_{2,1} + v_{1,2} - v_{1,1}).
 \end{aligned}
 \tag{3}$$

Having discretised the gradient ∇v , one computes the outer product $\nabla v \nabla v^T$. The structure tensor field results from smoothing this componentwise by the Gaussian of standard deviation ϱ .

2.2 Discretisation of Anisotropic Diffusion with Given Diffusion Tensor Field

We turn now to consider the anisotropic diffusion equation

$$\partial_t u = \operatorname{div}(D \cdot \nabla u)
 \tag{4}$$

with an arbitrary diffusion tensor field represented by positive semidefinite symmetric matrices $D = \begin{pmatrix} a & c \\ c & b \end{pmatrix}$. We assume that u is sampled at the integer pixel positions (i, j) while D is sampled at inter-pixel positions $(i + \frac{1}{2}, j + \frac{1}{2})$.

In discretising the right-hand side of (4) at some pixel position (i, j) , we will use the values of u at positions $(i + \varepsilon_1, j + \varepsilon_2)$ where $\varepsilon_1, \varepsilon_2 \in \{-1, 0, +1\}$, and the diffusion tensors at $(i \pm \frac{1}{2}, j \pm \frac{1}{2})$. For abbreviation we set

$$\begin{aligned}
 D_{i-\frac{1}{2}, j-\frac{1}{2}} &:= \begin{pmatrix} a_{--} & c_{--} \\ c_{--} & b_{--} \end{pmatrix}, & D_{i-\frac{1}{2}, j+\frac{1}{2}} &:= \begin{pmatrix} a_{-+} & c_{-+} \\ c_{-+} & b_{-+} \end{pmatrix}, \\
 D_{i+\frac{1}{2}, j-\frac{1}{2}} &:= \begin{pmatrix} a_{+-} & c_{+-} \\ c_{+-} & b_{+-} \end{pmatrix}, & D_{i+\frac{1}{2}, j+\frac{1}{2}} &:= \begin{pmatrix} a_{++} & c_{++} \\ c_{++} & b_{++} \end{pmatrix}.
 \end{aligned}
 \tag{5}$$

To obtain a discretisation which is “as local as possible”, we decompose the differential operators div and ∇ herein according to the 45°-rotated coordinates ξ, η where

$$\begin{pmatrix} \xi \\ \eta \end{pmatrix} := H \begin{pmatrix} x \\ y \end{pmatrix}, \quad H := \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.
 \tag{6}$$

Note that $H = H^T = H^{-1}$. To express the diffusion tensor in the ξ - η coordinates, D must be transformed by

$$HDH^T = \frac{1}{2} \begin{pmatrix} a + b + 2c & a - b \\ a - b & a + b - 2c \end{pmatrix}.
 \tag{7}$$

Then we have

$$\begin{aligned}
 (\operatorname{div}(D \nabla u))_{i,j} &= ((\partial_\xi, \partial_\eta)(HDH(\partial_\xi u, \partial_\eta u)^T))_{i,j} \\
 &= \frac{1}{2} \left(\partial_\xi((a + b + 2c)\partial_\xi u + (a - b)\partial_\eta u) \right. \\
 &\quad \left. + \partial_\eta((a - b)\partial_\xi u + (a + b - 2c)\partial_\eta u) \right)_{i,j}
 \end{aligned}$$




Fig. 1. The four-pixel cells contributing to $\dot{u}_{i,j}$

$$\begin{aligned}
 &\approx \frac{1}{2\sqrt{2}} \left(((a+b+2c)\partial_\xi u + (a-b)\partial_\eta u)_{i+\frac{1}{2},j+\frac{1}{2}} \right. \\
 &\quad - ((a+b+2c)\partial_\xi u + (a-b)\partial_\eta u)_{i-\frac{1}{2},j-\frac{1}{2}} \\
 &\quad + ((a-b)\partial_\xi u + (a+b-2c)\partial_\eta u)_{i+\frac{1}{2},j-\frac{1}{2}} \\
 &\quad \left. - ((a-b)\partial_\xi u + (a+b-2c)\partial_\eta u)_{i-\frac{1}{2},j+\frac{1}{2}} \right). \tag{8}
 \end{aligned}$$

Expanding $\partial_\xi u$ and $\partial_\eta u$ finally yields the dynamical system

$$\begin{aligned}
 \dot{u}_{i,j} = &\frac{1}{4} \left((a_{+++} + b_{+++} + 2c_{+++})(u_{i+1,j+1} - u_{i,j}) \right. \\
 &+ (a_{+++} - b_{+++})(u_{i+1,j} - u_{i,j+1}) \\
 &- (a_{---} + b_{---} + 2c_{---})(u_{i,j} - u_{i-1,j-1}) \\
 &- (a_{---} - b_{---})(u_{i,j-1} - u_{i-1,j}) \\
 &+ (a_{+-} - b_{+-})(u_{i+1,j} - u_{i,j-1}) \\
 &+ (a_{+-} + b_{+-} - 2c_{+-})(u_{i+1,j-1} - u_{i,j}) \\
 &- (a_{-+} - b_{-+})(u_{i,j+1} - u_{i-1,j}) \\
 &\left. - (a_{-+} + b_{-+} - 2c_{-+})(u_{i,j} - u_{i-1,j+1}) \right), \tag{9}
 \end{aligned}$$

where the dot denotes differentiation with respect to the time t . One observes that each summand on the right-hand side contains only quantities from one of the four-pixel cells

$$\begin{aligned}
 (--) : & \{i-1, i\} \times \{j-1, j\}, & (+-) : & \{i, i+1\} \times \{j-1, j\}, \\
 (-+) : & \{i-1, i\} \times \{j, j+1\}, & (++) : & \{i, i+1\} \times \{j, j+1\}
 \end{aligned} \tag{10}$$

which allows to split up (9) into the average of four dynamical systems each of which only contains interactions within one four-pixel cell. For illustration see Figure 1. With

$D = \begin{pmatrix} a & c \\ c & b \end{pmatrix}$ denoting the diffusion tensor discretised in $(\frac{3}{2}, \frac{3}{2})$, one such four-pixel dynamical system for the cell $\{1, 2\} \times \{1, 2\}$ reads as follows:

$$\begin{aligned}
 \dot{u}_{1,1} &= (a+b+2c)(u_{2,2} - u_{1,1}) + (a-b)(u_{2,1} - u_{1,2}), \\
 \dot{u}_{2,1} &= (a+b-2c)(u_{1,2} - u_{2,1}) + (a-b)(u_{1,1} - u_{2,2}), \\
 \dot{u}_{1,2} &= (a+b-2c)(u_{2,1} - u_{1,2}) + (a-b)(u_{2,2} - u_{1,1}), \\
 \dot{u}_{2,2} &= (a+b+2c)(u_{1,1} - u_{2,2}) + (a-b)(u_{1,2} - u_{2,1}).
 \end{aligned} \tag{11}$$

2.3 Semi-analytical Solution of the Four-Pixel System

Next we want to solve the system (11) analytically where we assume that the diffusion tensors D are kept fixed during the image evolution². To this end it is useful to introduce new variables $w_{i,j}$ by

$$W := HUH, \tag{12}$$

where $U := \begin{pmatrix} u_{1,1} & u_{2,1} \\ u_{1,2} & u_{2,2} \end{pmatrix}$, $W := \begin{pmatrix} w_{1,1} & w_{2,1} \\ w_{1,2} & w_{2,2} \end{pmatrix}$, and H happens to be the same matrix as introduced by (6). Then we can rewrite (11) in terms of the new variables as

$$\begin{aligned} \dot{w}_{1,1} &= 0, \\ \dot{w}_{2,1} &= -4aw_{2,1} - 4cw_{1,2}, \\ \dot{w}_{1,2} &= -4cw_{2,1} - 4bw_{1,2}, \\ \dot{w}_{2,2} &= 0. \end{aligned} \tag{13}$$

While $w_{1,1}$ and $w_{2,2}$ are constant, the dynamical system for $\mathbf{w} := (w_{2,1}, w_{1,2})^T$ can be rewritten as

$$\dot{\mathbf{w}} = -4D\mathbf{w}. \tag{14}$$

Let the eigendecomposition of D be given by $D = \lambda_1 \mathbf{e}_1 \mathbf{e}_1^T + \lambda_2 \mathbf{e}_2 \mathbf{e}_2^T$ with eigenvalues $\lambda_{1,2} = \frac{1}{2}(a + b \pm \sqrt{(a - b)^2 + 4c^2})$ and orthonormal eigenvectors $\mathbf{e}_1, \mathbf{e}_2$. Then, remembering that D is kept constant, the solution of (14) is

$$\mathbf{w}(t) = e^{-4\lambda_1 t} (\mathbf{e}_1^T \mathbf{w}(0)) \mathbf{e}_1 + e^{-4\lambda_2 t} (\mathbf{e}_2^T \mathbf{w}(0)) \mathbf{e}_2. \tag{15}$$

By the inverse transform of (12),

$$U(t) = H W(t) H, \tag{16}$$

this analytical solution can be expressed with respect to the original variables.

2.4 Numerical Scheme for Anisotropic Diffusion


We use now the explicit solution (15) of our four-pixel system as a building block for a numerical scheme for anisotropic diffusion. Because of the aforementioned splitting of (9) into the contributions from the four cells (10), the solution of (9) can be approximated by averaging the solutions of systems of the type (11). These solutions have been studied in Subsection 2.3. A time step will then be executed by computing the analytical solution (15) (resp. its back-transformed analog) for the desired evolution time, i.e. the time step size τ .

For the anisotropic diffusion processes that we are interested in, the diffusion tensor D depends on the structure tensor J_ϱ which arises from smoothing the outer product matrices $\nabla v \nabla v^T$ with a suitable convolution kernel, as seen in (2). To evaluate (15) requires therefore to compute D from the current data u , and to determine the eigendecomposition of D .

² In analogy to semi-implicit schemes that keep the nonlinear diffusion fixed at the previous time level while discretising the remainder in an implicit fashion, we call a method semi-analytic if it freezes the diffusion tensor and searches for an analytic solution.

LSAS Algorithm for Anisotropic Diffusion

- Compute the pre-smoothed image $v := K_\sigma * u^k$ by convolution.
- For each four-pixel cell $\{i, i + 1\} \times \{j, j + 1\}$, compute the approximation of the gradient ∇v according to (3), and the tensor product $\nabla v \nabla v^T$.
- Compute the structure tensor field $J = K_\rho * (\nabla v \nabla v^T)$ by convolution.
- For each four-pixel cell, compute the diffusion tensor $D = D(J)$.
- For each four-pixel cell, compute one time step of anisotropic diffusion via the analytical solution (12), (15), (16).
- For each pixel (*) with coordinates (i, j) , consider the four cells



which lead to four approximations

$$u_{i,j,-}^{k+1}, u_{i,j,+}^{k+1}, u_{i,j,-+}^{k+1}, u_{i,j,++}^{k+1} .$$

Average:

$$u_{i,j}^{k+1} = \frac{1}{4}(u_{i,j,-}^{k+1} + u_{i,j,+}^{k+1} + u_{i,j,-+}^{k+1} + u_{i,j,++}^{k+1}) .$$

Fig. 2. One time step of the locally semi-analytic scheme for anisotropic diffusion, where u^k, u^{k+1} refer to the old and new time step, respectively

We have therefore arrived at a *locally semi-analytic scheme (LSAS)* for anisotropic diffusion, one time step of which is summarised in Figure 2. The use of our analytical solution ensures that the resulting scheme for our four-pixel cell is stable in the Euclidean norm for any time step size (note that $\lambda_1, \lambda_2 \geq 0$). Since solutions from four-pixel cells are combined by simple averaging, this *absolute stability* transfers to the discretised anisotropic diffusion on the entire grid³.

3 Anisotropic Wavelet Shrinkage

In [23], it was shown that one-dimensional nonlinear diffusion on two-pixel signals coincides with Haar wavelet shrinkage if the shrinkage function is chosen in accordance with the diffusivity and the threshold parameter is equal to the diffusion time. A generalisation of this result to isotropic two-dimensional nonlinear diffusion was proposed

³ As for all known explicit diffusion schemes with unconditional stability, this favourable stability property is always in conjunction with conditional consistency: For fixed spatial grid size and a time step size tending to infinity, our scheme approaches a local averaging on a checkerboard decomposition of our grid.

in [28] where the shrinkage step was based on a diffusion inspired shrinkage function introduced in [16]. This shrinkage function couples the individual wavelet coefficients which leads to an improved rotation invariance of the procedure. Here we want to extend these promising results to the anisotropic setting.

The key for the connection between our four-pixel scheme and Haar wavelet shrinkage is the fact that the two-dimensional Haar wavelet transform acts naturally on subsequent 2×2 -pixel tiles of an image. Let us choose one such tile, say $F := \begin{pmatrix} f_{1,1} & f_{2,1} \\ f_{1,2} & f_{2,2} \end{pmatrix}$, and explain how it changes under two-dimensional Haar wavelet shrinkage. One cycle of Haar wavelet shrinkage consists of three steps: the analysis step, the shrinkage step and the synthesis step.

In the *analysis step*, the four-pixel image F image is transformed into the wavelet domain. To this end, the low and high pass Haar filters are applied to the rows and columns of F . More precisely, F is multiplied from the left and the right by the matrix H from (6) which results in an image

$$C := H F H . \tag{17}$$

Obviously, with $U = F$ and $W = C$, this coincides with our variable transform (12).

The *shrinkage step* modifies the high-pass coefficients $c_{2,1}$, $c_{1,2}$ and $c_{2,2}$ of C by applying a shrinkage function S_θ depending on a threshold parameter θ . Let us consider two examples first before introducing a shrinkage rule inspired by anisotropic diffusion.

In ordinary wavelet shrinkage the thresholding depends on the individual coefficients. For example, *soft shrinkage* [11] shrinks the coefficients towards 0 by an amount that is given by a threshold parameter θ :

$$S_\theta(c_{i,j}) := \begin{cases} c_{i,j} - \theta \operatorname{sgn}(c_{i,j}) & \text{if } |c_{i,j}| \geq \theta , \\ 0 & \text{otherwise .} \end{cases} \tag{18}$$

In [16] a shrinkage function inspired by *isotropic* nonlinear diffusion filtering was introduced that leads to a *coupled* shrinking of the coefficients. More precisely, the thresholding applies with respect to $\gamma(C) := (c_{2,1}^2 + c_{1,2}^2 + c_{2,2}^2)^{\frac{1}{2}}$. For a soft shrinkage and $(i, j) \in \{(2, 1), (1, 2), (2, 2)\}$ this comes down to

$$S_\theta(c_{i,j}) := \begin{cases} c_{i,j} - \frac{\theta}{\gamma(C)} \operatorname{sgn}(c_{i,j}) & \text{if } \gamma(C) \geq \theta , \\ 0 & \text{otherwise .} \end{cases} \tag{19}$$

Now we want to introduce an *anisotropic shrinkage* procedure with respect to a diffusion tensor D . In accordance with (13) and (15), we set $S_\theta(c_{2,2}) := c_{2,2}$ and define a coupled shrinkage of the antidiagonal coefficients $c_{1,2}$ and $c_{2,1}$ by

$$S_\theta \left(\begin{pmatrix} c_{2,1} \\ c_{1,2} \end{pmatrix} \right) := Q \begin{pmatrix} e^{-4\lambda_1\theta} & 0 \\ 0 & e^{-4\lambda_2\theta} \end{pmatrix} Q^T \begin{pmatrix} c_{2,1} \\ c_{1,2} \end{pmatrix} , \tag{20}$$

where $Q := (\mathbf{e}_1, \mathbf{e}_2)$ denotes the eigenvector matrix of D , and the threshold parameter θ was identified with the diffusion time t . This shows that besides the low-pass

coefficient $c_{1,1}$ also the high-pass coefficient $c_{2,2}$ remains unaffected, while the anti-diagonal coefficients $c_{1,2}$ and $c_{2,1}$ are shrunk in a coupled way. Let us abbreviate this anisotropic shrinkage procedure by $S_\theta(C)$.

Finally, the *synthesis step* leads us from the wavelet domain back to the original image domain. To this end we perform the inverse transform of step 1,

$$F^{(1)} = HS_\theta(C)H \tag{21}$$

on the shrunken coefficients. This is just the analog of (16).

In summary, one cycle of the above anisotropic Haar wavelet shrinkage coincides with the solution of (11) with initial condition $U(0) = F$, where the threshold parameter plays the role of the diffusion time.

Expressing an image in terms of Haar wavelets leads to a natural decomposition into tiles of 2×2 pixels (decimated wavelet transform). Shrinking these tiles separately according to the preceding procedure is not translationally invariant. Fortunately this is cured by the averaging procedure (9). It can be interpreted as a so-called *cyclic spinning* [7] that is related to a shift-invariant undecimated wavelet transform.

Apart from shift invariance, the LSAS algorithm can also be seen as a simple approach to create *rotationally invariant* anisotropic Haar wavelet shrinkage: Since our novel anisotropic shrinkage rules are a numerical scheme for a rotationally invariant continuous diffusion filter, rotation invariance is approximated at no additional expense.

4 Experiments

In our first experiment (Fig. 3) we use our scheme to perform edge-enhancing diffusion [25]. In this case, there is no integration over the outer products, so $\varrho = 0$. The diffusion tensor D has the same eigenvectors as the outer product $J = \nabla v \nabla v^T$, namely ∇v and its orthogonal ∇v^\perp . The eigenvalue in direction ∇v is given by $g(|\nabla v|^2)$ where $g(s^2) = 1 - \exp(-3.31488\lambda^8/s^8)$ with a given threshold parameter $\lambda > 0$, which



Fig. 3. Left to right: (a) Test image with noise. (b) Denoised by edge-enhancing diffusion with standard explicit scheme, $\lambda = 5$, $\sigma = 1.8$, $\varrho = 0$, $\tau = 0.166$, $N = 200$ iterations. (c) Denoised by edge-enhancing diffusion with LSAS, $\lambda = 5$, $\sigma = 1.8$, $\varrho = 0$, $\tau = 1$, $N = 200$ iterations.




Fig. 4. Left to right: (a) One quadrant of a rotationally invariant test image, 64×64 pixels. (b) Exact solution for coherence-enhancing diffusion with $\alpha = 0.001$, $C = 1$, $\sigma = 0.5$, $\varrho = 4$ and $t = 250$. (c) Filtered with the nonnegativity scheme [25] with $\tau = 1/6$, and $N = 1500$ iterations. Average absolute error: 17.99. (d) Processed with our LSAS algorithm, same parameters. Average absolute error: 3.81.



Fig. 5. Left to right: (a) Fingerprint image, 100×100 pixels. (b) Filtered with the nonnegativity scheme [25] for CED with $C = 1$, $\sigma = 0.5$, $\varrho = 4$, $\tau = 1/6$, and $N = 60$ iterations. (c) Processed with our LSAS algorithm for CED, same parameters. (d) LSAS algorithm with $\tau = 1$ and $N = 10$ iterations.

means that g is applied to the first eigenvalue of J . The eigenvalue of D in direction ∇v^T is fixed to 1. – The noisy image (Fig. 3a) is denoised with a standard explicit scheme with central spatial differences, and with the locally semi-analytic scheme. It is observed that the denoising result with our new scheme is slightly sharper. Moreover, a look at the parameters shows that the effective evolution time used by the new scheme is six times larger than with the explicit scheme which demonstrates how much the latter is indeed dominated by numerical blurring artifacts.

In our second experiment we consider coherence-enhancing diffusion (CED) [25]. It uses an integration scale ϱ that is considerably larger than σ , thereby introducing into J a smoothing over eigenvector systems. If the structure tensor has the eigendecomposition $J = \mu_1 e_1 e_1^T + \mu_2 e_2 e_2^T$ with $\mu_1 \geq \mu_2$, then the diffusion tensor $D(J)$ has the decomposition $D(J) := \lambda_1 e_1 e_1^T + \lambda_2 e_2 e_2^T$ with eigenvalues

$$\lambda_1 := \alpha, \quad \lambda_2 := \begin{cases} \alpha & \text{if } \mu_1 = \mu_2, \\ \alpha + (1 - \alpha) \exp\left(\frac{-C}{(\mu_1 - \mu_2)^2}\right) & \text{else,} \end{cases} \quad (22)$$

some small regularisation parameter $\alpha > 0$ and a contrast parameter $C > 0$. This process smoothes along flow-like structures. For a rotationally invariant test image such as the one in Figure 4, only radial linear diffusion with diffusivity α takes place. Hence, the exact solution at time t is given by a convolution with a Gaussian of standard deviation $\sqrt{2\alpha t}$. By comparing the solutions of the so-called nonnegativity discretisation from [25] with our LSAS algorithm and the exact solution, we see that the LSAS does not suffer from visible blurring artifacts. It preserves rotation invariance very well and creates significantly lower errors than the nonnegativity scheme.

These quantitative findings are also confirmed in the fingerprint example in Figure 5. We observe that the LSAS gives much sharper results, and that it yields still realistic results for time step sizes far beyond the stability limit $1/6$ of the nonnegativity scheme.

5 Conclusions

The contributions in our paper are twofold: Firstly we have presented a novel scheme for anisotropic diffusion with a high degree of rotation invariance and practically invisible blurring artifacts. It is absolutely stable in the Euclidean norm and simple to implement due to its explicit nature. Therefore it can serve as the method of choice whenever a well-founded, highly accurate scheme for anisotropic, tensor-driven diffusion is required. Secondly, we have clarified the diffusion-wavelet connection in the anisotropic case for the first time in the literature. This has led to novel, anisotropic shrinkage rules with coupling of the coefficients. More importantly, it also demonstrates that sophisticated concepts such as ridgelets and curvelets are not the only way to perform advanced anisotropic wavelet-based shrinkage: Even the most elementary class of wavelets, namely Haar wavelets, are sufficient for implementing highly anisotropic filters in a rotationally invariant fashion. We hope that this novel connection can help to fertilise further research on simple, structure-adaptive anisotropic wavelet concepts and to gain new insights in the design of coupled shrinkage rules.

References

1. Y. Bao and H. Krim. Towards bridging scale-space and multiscale frame analyses. In A. A. Petrosian and F. G. Meyer, editors, *Wavelets in Signal and Image Analysis*, volume 19 of *Computational Imaging and Vision*, chapter 6. Kluwer, Dordrecht, 2001.
2. K. Bredies, D. A. Lorenz, P. Maaß, and G. Teschke. A partial differential equation for continuous non-linear shrinkage filtering and its application for analyzing MMG data. In F. Truchetet, editor, *Wavelet Applications in Industrial Processing*, volume 5266 of *Proceedings of SPIE*, pages 84–93. SPIE Press, Bellingham, 2004.
3. E. Candés, L. Demanet, D. Donoho, and L. Ying. Fast discrete curvelet transform. Technical report, Applied and Computational Mathematics, Caltech, Pasadena, CA, 2005.
4. E. J. Candés and D. L. Donoho. Curvelets: A surprisingly effective nonadaptive representation of objects with edges. In A. Cohen, C. Rabut, and L. L. Schumaker, editors, *Curve and Surface Fitting*, Saint-Malo, 2000. Vanderbilt University Press.
5. A. Chambolle and B. L. Lucier. Interpreting translationally-invariant wavelet shrinkage as a new image smoothing scale space. *IEEE Transactions on Image Processing*, 10(7):993–1000, 2001.

6. A. Cohen, W. Dahmen, I. Daubechies, and R. DeVore. Harmonic analysis in the space BV. *Revista Matemática Iberoamericana*, 19:235–262, 2003.
7. R. R. Coifman and D. Donoho. Translation invariant denoising. In A. Antoine and G. Oppenheim, editors, *Wavelets in Statistics*, pages 125–150. Springer, New York, 1995.
8. R. R. Coifman and A. Sowa. New methods of controlled total variation reduction for digital functions. *SIAM Journal on Numerical Analysis*, 39(2):480–498, 2001.
9. M. N. Do and V. Vetterli. Contourlets. In J. Stöckler and G. Welland, editors, *Beyond Wavelets*, pages 1–27. Academic Press, New York, 2003.
10. D. Donoho and A. Flesia. Digital ridgelet transform based on true ridge functions. In J. Stöckler and G. Welland, editors, *Beyond Wavelets*, pages 1–33. Academic Press, New York, 2003.
11. D. L. Donoho. De-noising by soft thresholding. *IEEE Transactions on Information Theory*, 41:613–627, 1995.
12. D. L. Donoho. Orthonormal ridglets and linear singularities. *SIAM Journal on Mathematical Analysis*, 31(5):1062–1099, 2000.
13. 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.
14. B. Jawerth, P. Lin, and E. Sinzinger. Lattice Boltzmann models for anisotropic diffusion of images. *Journal of Mathematical Imaging and Vision*, 11:231–237, 1999.
15. Y. Meyer. *Oscillating Patterns in Image Processing and Nonlinear Evolution Equations*, volume 22 of *University Lecture Series*. AMS, Providence, 2001.
16. P. Mrázek and J. Weickert. Rotationally invariant wavelet shrinkage. In B. Michaelis and G. Krell, editors, *Pattern Recognition*, volume 2781 of *Lecture Notes in Computer Science*, pages 156–163, Berlin, 2003. Springer.
17. P. Mrázek, J. Weickert, and G. Steidl. Diffusion-inspired shrinkage functions and stability results for wavelet denoising. *International Journal of Computer Vision*, 64(2/3):171–186, Sept. 2005.
18. 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.
19. T. Preußner and M. Rumpf. An adaptive finite element method for large scale image processing. *Journal of Visual Communication and Image Representation*, 11(2):183–195, June 2000.
20. J. Shen. A note on wavelets and diffusion. *Journal of Computational Analysis and Applications*, 5(1):147–159, 2003.
21. E. P. Simoncelli. Design of multidimensional derivative filters. In *Proc. 1994 IEEE International Conference on Image Processing*, volume 1, pages 790–793, Austin, TX, Nov. 1994.
22. N. Sochen, R. Kimmel, and F. Bruckstein. Diffusions and confusions in signal and image processing. *Journal of Mathematical Imaging and Vision*, 14(3):195–210, May 2001.
23. G. Steidl, J. Weickert, T. Brox, P. Mrázek, and M. Welk. On the equivalence of soft wavelet shrinkage, total variation diffusion, total variation regularization, and SIDes. *SIAM Journal on Numerical Analysis*, 42(2):686–713, 2004.
24. W. Wang. On the design of optimal derivative filters for coherence-enhancing diffusion filtering. In *Proc. 2004 International Conference on Computer Graphics, Imaging and Visualization*, pages 35–40, Penang, Malaysia, July 2004. IEEE Computer Society Press.
25. J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart, 1998.

26. J. Weickert and H. Scharr. A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance. *Journal of Visual Communication and Image Representation*, 13(1/2):103–118, 2002.
27. J. Weickert, G. Steidl, P. Mrázek, M. Welk, and T. Brox. Diffusion filters and wavelets: What can they learn from each other? In N. Paragios, Y. Chen, and O. Faugeras, editors, *The Handbook of Mathematical Models in Computer Vision*. Springer, New York, 2005.
28. M. Welk, J. Weickert, and G. Steidl. A four-pixel scheme for singular differential equations. In R. Kimmel, N. Sochen, and J. Weickert, editors, *Scale-Space and PDE Methods in Computer Vision*, volume 3459 of *Lecture Notes in Computer Science*, pages 585–597, Berlin, 2005. Springer.