

In A. Kuijper, K. Bredies, T. Pock, H. Bischof (Eds.): Scale-Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science, Vol. 7893, pp. 380-391, Springer, Berlin, 2013. The final publication is available at link.springer.com.

L^2 -Stable Nonstandard Finite Differences for Anisotropic Diffusion

Joachim Weickert¹, Martin Welk², and Marco Wickert¹

¹ Mathematical Image Analysis Group, Campus E1.7
Saarland University, 66041 Saarbrücken, Germany
{weickert, wickert}@mia.uni-saarland.de

² University for Health Sciences, Medical Informatics and Technology
Eduard-Wallnöfer-Zentrum 1, 6060 Hall/Tyrol, Austria
martin.welk@umit.at

Abstract. Anisotropic diffusion filters with a diffusion tensor are successfully used in many image processing and computer vision applications, ranging from image denoising over compression to optic flow computation. However, finding adequate numerical schemes is difficult: Implementations may suffer from dissipative artifacts, poor approximation of rotation invariance, and they may lack provable stability guarantees. In our paper we propose a general framework for finite difference discretisations of anisotropic diffusion filters on a 3×3 stencil. It is based on a gradient descent of a discrete quadratic energy where the occurring derivatives are replaced by classical as well as the widely unknown nonstandard finite differences in the sense of Mickens. This allows a large class of space discretisations with two free parameters. Combining it with an explicit or semi-implicit time discretisation, we establish a general and easily applicable stability theory in terms of a decreasing Euclidean norm. Our framework comprises as many as seven existing space discretisations from the literature. However, we show that also novel schemes are possible that offer a better performance than existing ones. Our experimental evaluation confirms that the space discretisation can have a very substantial and often underestimated impact on the quality of anisotropic diffusion filters.

Keywords: diffusion filtering, finite difference methods, stability, rotation invariance, dissipativity

1 Introduction

Anisotropic diffusion filters with a diffusion tensor instead of a scalar-valued diffusivity are flexible tools that permit to steer the diffusion process in a desired

direction [1]: This can be very useful for image processing tasks ranging from image denoising and enhancement (see e.g. [1, 2]) to lossy image compression [3]. Anisotropic diffusion terms also appear in computer vision applications, e.g. in the Euler-Lagrange equations of variational methods for optic flow computations [4, 5], for stereo reconstruction [6], and for range image integration [7]. An anisotropic diffusion filter in the sense of [1] computes a filtered version $u(\mathbf{x}, t)$ of some initial image $f(\mathbf{x})$ by solving the diffusion equation

$$u_t = \operatorname{div}(\mathbf{D} \nabla u) \quad (1)$$

with f as initial condition,

$$u(\mathbf{x}, 0) = f(\mathbf{x}), \quad (2)$$

and homogeneous Neumann boundary conditions. Here the lower index t denotes a time derivative, and the divergence and the nabla operators involve spatial derivatives only. The diffusion tensor \mathbf{D} is a positive definite (and thus also symmetric) 2×2 matrix that is space-variant and may even depend on derivatives of the evolving image $u(\mathbf{x}, t)$. For our discussions below, one can use positive semidefinite diffusion tensors as well.

A large number of numerical schemes has been proposed for anisotropic diffusion processes, including finite elements [8], finite volume methods [9], and lattice Boltzmann techniques [10]. However, mostly finite difference methods are used [1, 2, 11–14], sometimes realised as wavelet shrinkage [15, 16]. Apart from [13], all finite difference schemes approximate the divergence term on a 3×3 stencil.

Unfortunately, finding good finite difference schemes for anisotropic diffusion filters is much more challenging than for their isotropic counterparts with a scalar-valued diffusivity. While the time discretisation mainly influences the efficiency of the method and does not create specific difficulties, the major problem comes from the space discretisation: If the diffusion process is strongly anisotropic, the corresponding direction has to be approximated with very high accuracy in order to avoid undesired dissipative blurring effects. The approximation quality of the rotationally invariant model can also vary a lot even among schemes with identical order of consistency. Last but not least, it is difficult to establish a stability theory for anisotropic diffusion filters: Weickert [1] presents a discrete theory that analyses stability in terms of a maximum-minimum principle. However, he shows that on a 3×3 stencil, this can only be guaranteed if the spectral condition number of \mathbf{D} does not exceed 5.82 (see also [17]). In practice, one is usually interested in using more pronounced anisotropies. In this case, there is no L^∞ -stability guarantee for the nonnegativity scheme from [1] and its generalisations by Mrázek and Navara [12]. Thus, it would be desirable to have at least an L^2 -stability theory such as for the wavelet-inspired schemes from [15, 16]. However, none of the finite difference discretisations in [1, 2, 11–13] gives L^2 -stability results.

Our Contributions. The goal of the present paper is to provide a general framework for L^2 -stable discretisations of anisotropic diffusion filters on a 3×3 stencil. It is derived as gradient descent of a discretised energy functional with a positive definite quadratic form. By considering also the widely unknown non-standard discretisations in the sense of Mickens [18], we end up with a two-parameter family of space discretisations on a 3×3 stencil. Interestingly this family covers seven existing finite difference discretisations that have been proposed for such a stencil [1, 2, 11, 12, 15, 16]. Moreover, we establish stability results in the Euclidean norm for explicit and semi-implicit time discretisations, providing a theoretical foundation for many of these schemes that has not been available so far. Last but not least we present an experiment that illustrates the large impact that the free parameters can have. With a suitable parameter choice, one can design novel schemes with low dissipativity and an excellent approximation of rotation invariance.

Organisation of the Paper. In Section 2 we derive our general finite difference stencil from a discrete energy. Its theoretical properties are analysed in Section 3. In the fourth section, we evaluate various discretisations that arise as special cases, and we conclude our paper with a summary in Section 5.

2 General Discretisation

Quadratic Energy Model. To discretise the anisotropic diffusion process (1) in time, we will use a sequence $t_0 < t_1 < t_2 < \dots$ of discrete time nodes. Freezing the space-variant diffusion tensor \mathbf{D} within each time interval $[t_k, t_{k+1})$, $k \in \mathbb{N}_0$ then creates a sequence of linearised processes. In each interval, the evolution equation is a gradient descent of the quadratic energy

$$E(u) = \frac{1}{2} \int_{\Omega} \nabla^\top u \mathbf{D} \nabla u \, dx \, dy \tag{3}$$

with a space-variant but time-invariant positive definite diffusion tensor

$$\mathbf{D} = \begin{pmatrix} a(x, y) & b(x, y) \\ b(x, y) & c(x, y) \end{pmatrix}. \tag{4}$$

For discretisation in space, we adopt for u a regular grid $\{1, \dots, N\} \times \{1, \dots, M\}$ with mesh size h in both x and y direction, i.e. the index (i, j) refers to the location (x_i, y_j) with $x_i = x_0 + i h$, $y_j = y_0 + j h$. Following the proceeding in [15, 16], we assume that approximations for a , b , and c are available in the locations $(i + \frac{1}{2}, j + \frac{1}{2})$. Provided that also u_x^2 , $u_x u_y$, and u_y^2 are approximated in $(i + \frac{1}{2}, j + \frac{1}{2})$, a discrete version of the energy (3) is then given by

$$E(\mathbf{u}) = \frac{1}{2} \sum_{i=0}^N \sum_{j=0}^M (a u_x^2 + 2b u_x u_y + c u_y^2)_{i+\frac{1}{2}, j+\frac{1}{2}}. \tag{5}$$

Suitable approximations should be *local*, i.e. involve only the four pixels in the cell $\{i, i+1\} \times \{j, j+1\}$. In terms of *accuracy*, we require their consistency to be of second order. At boundary locations (rows $j \in \{0, M\}$, columns $i \in \{0, N\}$), values of a, b, c , and u_x, u_y must satisfy appropriate constraints to be compatible with Neumann boundary conditions.

Derivative Approximations. To derive approximations with these properties, we start by discretising u_x and u_y . To this end, we consider combinations of the forward differences

$$\boxed{_} := D_x u_{i,j} := \frac{u_{i+1,j} - u_{i,j}}{h}, \quad (6)$$

$$\boxed{_} := D_x u_{i,j+1} := \frac{u_{i+1,j+1} - u_{i,j+1}}{h}, \quad (7)$$

$$\boxed{|} := D_y u_{i,j} := \frac{u_{i,j+1} - u_{i,j}}{h}, \quad (8)$$

$$\boxed{|} := D_y u_{i+1,j} := \frac{u_{i+1,j+1} - u_{i+1,j}}{h}. \quad (9)$$

Nonstandard Finite Difference Approximations. Since our quadratic energy involves expressions in u_x^2, u_y^2 , and $u_x u_y$, let us study approximations of these terms with second order of consistency using the discretisations (6)–(9).

We approximate u_x^2 by affine combinations of the arithmetic mean and the geometric mean of the finite differences in x -direction:

$$\begin{aligned} u_x^2 \Big|_{i+\frac{1}{2}, j+\frac{1}{2}} &\approx (1 - \alpha_{i+\frac{1}{2}, j+\frac{1}{2}}) \cdot \frac{1}{2} (\boxed{_} \cdot \boxed{_} + \boxed{_} \cdot \boxed{_}) \\ &+ \alpha_{i+\frac{1}{2}, j+\frac{1}{2}} \cdot \boxed{_} \cdot \boxed{_}, \end{aligned} \quad (10)$$

where $\alpha_{i+\frac{1}{2}, j+\frac{1}{2}}$ is an arbitrary weight that may be space-variant.

Analogously, u_y^2 is approximated by affine combinations of the arithmetic mean and the geometric mean of the finite differences in y -direction:

$$\begin{aligned} u_y^2 \Big|_{i+\frac{1}{2}, j+\frac{1}{2}} &\approx (1 - \alpha_{i+\frac{1}{2}, j+\frac{1}{2}}) \cdot \frac{1}{2} (\boxed{|} \cdot \boxed{|} + \boxed{|} \cdot \boxed{|}) \\ &+ \alpha_{i+\frac{1}{2}, j+\frac{1}{2}} \cdot \boxed{|} \cdot \boxed{|}. \end{aligned} \quad (11)$$

To treat u_x^2 and u_y^2 equally, we have chosen the same weight $\alpha_{i+\frac{1}{2}, j+\frac{1}{2}}$.

Eventually, $u_x u_y$ involves all four combinations of the two finite differences in x -direction and the two finite differences in y -direction:

$$\begin{aligned} u_x u_y \Big|_{i+\frac{1}{2}, j+\frac{1}{2}} &\approx \frac{1 - \beta_{i+\frac{1}{2}, j+\frac{1}{2}}}{2} \cdot \frac{1}{2} (\boxed{_} \cdot \boxed{|} + \boxed{_} \cdot \boxed{|}) \\ &+ \frac{1 + \beta_{i+\frac{1}{2}, j+\frac{1}{2}}}{2} \cdot \frac{1}{2} (\boxed{_} \cdot \boxed{|} + \boxed{_} \cdot \boxed{|}) \end{aligned} \quad (12)$$

with a space-variant weight $\beta_{i+\frac{1}{2},j+\frac{1}{2}}$.

The approximations (10)–(12) deserve some further discussion. Note that for $\alpha_{i+\frac{1}{2},j+\frac{1}{2}} \neq 0$, the second summand in (10) approximates u_x^2 at the location $(i+\frac{1}{2},j+\frac{1}{2})$ by multiplying two different approximations for u_x , namely $D_x u_{i,j}$ and $D_x u_{i,j+1}$. This is in accordance with one of Mickens’ principles for so-called *nonstandard finite difference* schemes [18]: “Nonlinear terms must, in general, be modelled nonlocally on the computational grid or lattice”. Here, the term nonlocal means that both approximations refer to different grid points: $D_x u_{i,j}$ is a central difference approximation in $(i+\frac{1}{2},j)$, while $D_x u_{i,j+1}$ is centred in $(i+\frac{1}{2},j+1)$.

In a similar way, one sees that also (11) uses nonstandard finite differences for $\alpha_{i+\frac{1}{2},j+\frac{1}{2}} \neq 0$, and so does (12) for $\beta_{i+\frac{1}{2},j+\frac{1}{2}} \neq 0$. Note that for $\beta_{i+\frac{1}{2},j+\frac{1}{2}} = 0$, approximation (12) is equivalent to

$$u_x u_y \Big|_{i+\frac{1}{2},j+\frac{1}{2}} \approx \frac{1}{2} (\boxed{-} + \boxed{-}) \cdot \frac{1}{2} (\boxed{1} + \boxed{1}) \quad (13)$$

which is a standard approximation, since both factors are centred in $(i+\frac{1}{2},j+\frac{1}{2})$.

Mickens advocates his principle of nonlocal approximation of nonlinear terms as an ingredient for obtaining qualitatively correct discrete models of continuous equations. The evaluation in Section 4 will show the benefit of this idea.

Gradient Descent. Our space-discrete approximation of the anisotropic diffusion process in every pixel (i,j) is finally given by the gradient descent

$$\frac{du_{i,j}}{dt} = - \frac{\partial E(\mathbf{u})}{\partial u_{i,j}} \quad (14)$$

for the discrete energy (5) with the approximations (10)–(12). The right hand side gives the desired discretisation of $\text{div}(\mathbf{D} \nabla u)$. It can be represented by the weights in a (3×3) -stencil. For inner pixels $1 < i < N$, $1 < j < M$ one obtains after some tedious but straightforward calculations the stencil

$$\frac{1}{2h^2} \cdot \begin{array}{|c|c|c|} \hline \left[(\beta-1)b + \alpha(a+c) \right]_{i-\frac{1}{2},j+\frac{1}{2}} & \begin{array}{c} \left[(1-\alpha)c - \alpha a - \beta b \right]_{i+\frac{1}{2},j+\frac{1}{2}} \\ + \left[(1-\alpha)c - \alpha a - \beta b \right]_{i-\frac{1}{2},j+\frac{1}{2}} \end{array} & \left[(\beta+1)b + \alpha(a+c) \right]_{i+\frac{1}{2},j+\frac{1}{2}} \\ \hline \begin{array}{c} \left[(1-\alpha)a - \alpha c - \beta b \right]_{i-\frac{1}{2},j+\frac{1}{2}} \\ + \left[(1-\alpha)a - \alpha c - \beta b \right]_{i-\frac{1}{2},j-\frac{1}{2}} \end{array} & \begin{array}{c} - \left[(1-\alpha)(a+c) - (\beta-1)b \right]_{i+\frac{1}{2},j+\frac{1}{2}} \\ - \left[(1-\alpha)(a+c) - (\beta+1)b \right]_{i+\frac{1}{2},j-\frac{1}{2}} \\ - \left[(1-\alpha)(a+c) - (\beta+1)b \right]_{i-\frac{1}{2},j+\frac{1}{2}} \\ - \left[(1-\alpha)(a+c) - (\beta-1)b \right]_{i-\frac{1}{2},j-\frac{1}{2}} \end{array} & \begin{array}{c} \left[(1-\alpha)a - \alpha c - \beta b \right]_{i+\frac{1}{2},j+\frac{1}{2}} \\ + \left[(1-\alpha)a - \alpha c - \beta b \right]_{i+\frac{1}{2},j-\frac{1}{2}} \end{array} \\ \hline \left[(\beta+1)b + \alpha(a+c) \right]_{i-\frac{1}{2},j-\frac{1}{2}} & \begin{array}{c} \left[(1-\alpha)c - \alpha a - \beta b \right]_{i+\frac{1}{2},j-\frac{1}{2}} \\ + \left[(1-\alpha)c - \alpha a - \beta b \right]_{i-\frac{1}{2},j-\frac{1}{2}} \end{array} & \left[(\beta-1)b + \alpha(a+c) \right]_{i+\frac{1}{2},j-\frac{1}{2}} \end{array} \quad (15)$$

Table 1. Seven existing space discretisations as special cases of our general stencil.

discretisation	α	β
standard discretisation [11]	0	0
Cottet and El-Ayyadi [2]	0	-1
nonnegativity discretisation [1]	0	$\text{sign}(b)$
Mrázek and Navara II [12]	$\frac{\min(a,c)}{a+c}$	0
Mrázek and Navara III [12]	$\frac{\min(a,c)}{2(a+c)}$	$\frac{1}{2}\text{sign}(b)$
wavelet-inspired scheme I [15]	$\frac{1}{2}$	0
wavelet-inspired scheme II [16]	$[0, \frac{1}{2}]$	0

where the y -axis is oriented upwards. This stencil approximates $\text{div}(\mathbf{D}\nabla u)$ with consistency order 2. In boundary pixels, homogeneous Neumann boundary conditions can be taken into account just by mirroring the first and last rows and columns of u .

A General Framework for Existing Schemes. Interestingly our space discretisation subsumes a number of anisotropic diffusion stencils from the literature. Table 1 lists seven representatives with the corresponding weight parameters α , β of our general stencil. In three of the listed schemes the weights are chosen space-variant. All but the last two schemes have originally been stated with the diffusion tensor discretised either at locations (i, j) or $(i + \frac{1}{2}, j)$, $(i, j + \frac{1}{2})$. In these cases, full correspondence with our scheme is achieved by a suitable grid resampling with linear interpolation. In Section 4 we will see that our general stencil also contains new parameter settings with favourable performance.

3 Theoretical Properties

In the anisotropic diffusion process (1), the diffusion tensor field \mathbf{D} is required to consist of positive definite tensors. As a consequence, the quadratic form within the continuous energy (3) is nonnegative. It is therefore natural to ask whether also the discrete energy (5) retains this property. This will help to determine stability properties of the gradient descent.

3.1 Positive Semidefiniteness of the Discrete Energy

Introducing the notations

$$\mathbf{w}_{i+\frac{1}{2},j+\frac{1}{2}} := \left(\begin{array}{c} \square \\ \square \\ \square \\ \square \end{array} \right)^\top, \quad (16)$$

$$\mathbf{H}_{i+\frac{1}{2},j+\frac{1}{2}} := \left(\begin{array}{cc|cc} \frac{1-\alpha}{2} a_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{\alpha}{2} a_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1-\beta}{4} b_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1+\beta}{4} b_{i+\frac{1}{2},j+\frac{1}{2}} \\ \frac{\alpha}{2} a_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1-\alpha}{2} a_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1+\beta}{4} b_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1-\beta}{4} b_{i+\frac{1}{2},j+\frac{1}{2}} \\ \hline \frac{1-\beta}{4} b_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1+\beta}{4} b_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1-\alpha}{2} c_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{\alpha}{2} c_{i+\frac{1}{2},j+\frac{1}{2}} \\ \frac{1+\beta}{4} b_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1-\beta}{4} b_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{\alpha}{2} c_{i+\frac{1}{2},j+\frac{1}{2}} & \frac{1-\alpha}{2} c_{i+\frac{1}{2},j+\frac{1}{2}} \end{array} \right), \quad (17)$$

we can rewrite our discrete energy (5) as

$$E(\mathbf{u}) = \frac{1}{2} \sum_{i=0}^N \sum_{j=0}^M \mathbf{w}_{i+\frac{1}{2},j+\frac{1}{2}}^\top \mathbf{H}_{i+\frac{1}{2},j+\frac{1}{2}} \mathbf{w}_{i+\frac{1}{2},j+\frac{1}{2}}. \quad (18)$$

Now we state our main result on the discrete energy functional.

Proposition 1 (Positive Semidefiniteness of $\mathbf{H}_{i+\frac{1}{2},j+\frac{1}{2}}$). *The matrix $\mathbf{H}_{i+\frac{1}{2},j+\frac{1}{2}}$ is positive semidefinite for any positive definite diffusion tensor $\mathbf{D}_{i+\frac{1}{2},j+\frac{1}{2}}$ if and only if $|\beta| \leq 1 - 2\alpha$.*

Sketch of the proof. We decompose \mathbb{R}^4 into the subspaces

$$V := \text{span}\{(1, 1, 0, 0)^\top, (0, 0, 1, 1)^\top\} \quad \text{and} \quad (19)$$

$$V^\perp = \text{span}\{(1, -1, 0, 0)^\top, (0, 0, 1, -1)^\top\}. \quad (20)$$

On V , the matrix $\mathbf{H}_{i+\frac{1}{2},j+\frac{1}{2}}$ acts like $\frac{1}{2}\mathbf{D}_{i+\frac{1}{2},j+\frac{1}{2}}$: To see this, let $(x, y)^\top$ be an eigenvector of $\mathbf{D}_{i+\frac{1}{2},j+\frac{1}{2}}$ with eigenvalue λ . Then $(x, x, y, y)^\top$ is an eigenvector of $\mathbf{H}_{i+\frac{1}{2},j+\frac{1}{2}}$ with eigenvalue $\frac{\lambda}{2}$. This guarantees positive definiteness on V .

On V^\perp , it is easy to check that the action of $\mathbf{H}_{i+\frac{1}{2},j+\frac{1}{2}}$ is given by the matrix

$$\mathbf{T} := \frac{1}{2} \begin{pmatrix} (1 - 2\alpha)a & \beta b \\ \beta b & (1 - 2\alpha)c \end{pmatrix} \quad (21)$$

with respect to the basis vectors stated above. To ensure nonnegativity of the eigenvalues of \mathbf{T} for any positive definite \mathbf{D} , the inequality $|\beta| \leq 1 - 2\alpha$ is necessary and sufficient. \square

Remark. As a consequence, the largest value of α for which positive semidefiniteness can be established is $\alpha = 0.5$. However, we do not recommend using $\alpha = 0.5$, since the stencil can decouple into two checkerboard-like subgrids then. For $\alpha < 0.5$ one has strict positive definiteness and no decoupling problems.

3.2 Stability Results for Fully Discrete Diffusion Schemes

The positive semidefinite energy (18) can be rewritten as

$$E(\mathbf{u}) = -\frac{1}{2} \mathbf{u}^\top \mathbf{A} \mathbf{u} \quad (22)$$

with a negative semidefinite matrix $\mathbf{A} \in \mathbb{R}^{NM \times NM}$ where each row contains the nine stencil entries of the corresponding spatial node. Its gradient descent

$$\frac{d\mathbf{u}}{dt} = \mathbf{A} \mathbf{u} \quad (23)$$

is a space-discrete and time-continuous anisotropic diffusion process. Let us now consider two common time discretisations of this dynamical system.

Explicit Time Discretisation. An explicit scheme with step size τ is given by

$$\frac{\mathbf{u}^{k+1} - \mathbf{u}^k}{\tau} = \mathbf{A}^k \mathbf{u}^k, \quad (24)$$

where the upper index denotes the time level. It can be written as

$$\mathbf{u}^{k+1} = (\mathbf{I} + \tau \mathbf{A}^k) \mathbf{u}^k. \quad (25)$$

Stability in the Euclidean norm requires $\|\mathbf{u}^{k+1}\|_2 \leq \|\mathbf{u}^k\|_2$. This is guaranteed for $\rho(\mathbf{I} + \tau \mathbf{A}^k) \leq 1$, where ρ denotes the spectral norm. For a negative semidefinite \mathbf{A}^k , this comes down to

$$\tau \leq \frac{2}{\rho(\mathbf{A}^k)}. \quad (26)$$

An estimate for $\rho(\mathbf{A}^k)$ can be derived via Gershgorin's Theorem. The stability bound (26) also allows to design extremely efficient variants of (24), so-called *fast explicit diffusion (FED)* schemes [19]. They use cycles of varying time steps, preserve the L^2 stability of the underlying scheme, and are well-suited for GPUs.

Semi-implicit Time Discretisation. The semi-implicit scheme

$$\frac{\mathbf{u}^{k+1} - \mathbf{u}^k}{\tau} = \mathbf{A}^k \mathbf{u}^{k+1} \quad (27)$$

requires to solve a linear system of equations:

$$(\mathbf{I} - \tau \mathbf{A}^k) \mathbf{u}^{k+1} = \mathbf{u}^k. \quad (28)$$

For negative semidefinite \mathbf{A}^k , the matrix $\mathbf{I} - \tau \mathbf{A}^k$ has only eigenvalues ≥ 1 and is thus invertible. Since $\rho((\mathbf{I} - \tau \mathbf{A}^k)^{-1}) \leq 1$, the semi-implicit scheme

$$\mathbf{u}^{k+1} = (\mathbf{I} - \tau \mathbf{A}^k)^{-1} \mathbf{u}^k \quad (29)$$

is absolutely stable in the Euclidean norm.

Table 2. Performance of different space discretisations.

discretisation	α	β	PSNR [dB]
standard discretisation [11]	0	0	24.60
Cottet and El-Ayyadi [2]	0	-1	25.38
nonnegativity discretisation [1]	0	$\text{sign}(b)$	29.86
Mrázek and Navara II [12]	$\frac{\min(a,c)}{a+c}$	0	24.17
Mrázek and Navara III [12]	$\frac{\min(a,c)}{2(a+c)}$	$0.5 \text{ sign}(b)$	27.42
wavelet-inspired scheme I [15]	0.5	0	29.57
wavelet-inspired scheme II [16]	0.49	0	32.88
our nonstandard stencil	0.44	$0.118 \text{ sign}(b)$	33.99

4 Evaluation of Specific Discretisations

Now that we have derived a general class of L^2 -stable discretisations for anisotropic diffusion processes, let us study the performance of different parameter settings. Since more recent applications of anisotropic diffusion focus on its interpolation quality (see e.g. [3, 5]), we consider an idealised demosaicking scenario, where we know the ground truth solution and where subpixel accuracy w.r.t. the interpolation direction plays an important role. For other applications such as denoising and image enhancement we have found similar performance rankings.

Demosaicking addresses a problem of many camera sensors: They use a colour filter array which allows them to measure only one out of three colour channels at each pixel location: either red (R), green (G), or blue (B). Thus, exactly one third of the colour image information is available, and two thirds must be interpolated. Often the colour information is arranged in the order of the so-called *Bayer array* that consists of a periodic repetition of the pattern

$$\begin{array}{|c|c|} \hline \text{R} & \text{G} \\ \hline \text{G} & \text{B} \\ \hline \end{array} \quad (30)$$

For our evaluation we consider the synthetic test image in Figure 1(a). It consists of concentric circular structures of varying frequencies. Thus, it allows us to assess the directional and the frequency behaviour of anisotropic diffusion interpolation. By removing two thirds of the colour information by means of the Bayer mask, we obtain the image in Figure 1(b). Now we interpolate the missing information at the unspecified channels in each pixel by evolving an explicit anisotropic diffusion scheme to its steady state. The data at the specified pixels serve as Dirichlet boundary conditions. In this synthetic example we can design a positive semidefinite diffusion tensor in such a way that only diffusion in tangential direction is allowed. Thus, the directional error of \mathbf{D} is zero, and all reconstruction errors are caused by limitations of our space discretisation due to dissipative artifacts or deviations from rotation invariance.

Table 2 compares the peak signal-to-noise ratio (PSNR) between the interpolated image and the original image. In spite of the fact that we interpolate over

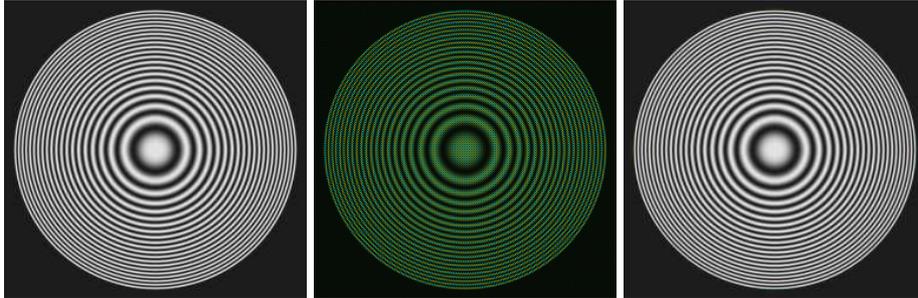


Fig. 1. (a) **Left:** Test image, 256×256 pixels. (b) **Middle:** After applying the Bayer colour filter array. (c) **Right:** Demosaicking result with our nonstandard scheme with $\alpha = 0.44$ and $\beta = 0.118 \operatorname{sign}(b)$.

a distance of at most two pixels and that we prescribe the correct interpolation direction, we observe that the seven stencils from the literature differ strongly in their performance: While the standard discretisation and the Mrázek–Navara scheme II perform fairly bad, the nonnegativity discretisation and the wavelet-inspired scheme II with $\alpha = 0.49$ give rather good results. It should be noted that the wavelet-inspired scheme I always uses $\alpha = 0.5$. Hence, it suffers from the before mentioned checkerboard decoupling which is particularly undesirable for demosaicking. Therefore, we recommend to use only stencils with $\alpha \leq 0.49$.

We see that all seven schemes from the literature can be outperformed by our nonstandard scheme with suitable parameters. Its demosaicking result is depicted in Figure 1(c). Since all stencils from Table 2 have the same consistency order (namely 2), it is remarkable that their actual performance is so different: The PSNR difference between the best and the worst stencil is 9.82 dB! This confirms the fundamental importance of a good discretisation when one wants to use anisotropic diffusion processes with a diffusion tensor. Similar findings have also been made in [12, 13, 15, 16].

The best 3×3 finite difference stencil in the anisotropic diffusion literature is given by the wavelet-inspired filter class II [16]. It contains α as a free parameter. Our nonstandard stencil class identifies β as a second degree of freedom. According to Proposition 1, β has to fulfil $|\beta| \leq 1 - 2\alpha$. Often it is advisable to use a β -value that has the same sign as b , since this can help to reduce the well-known over- and undershoots due to a lack of nonnegativity in the stencil. Thus, it is convenient to replace the parameter β by a parameter γ that is linked to $\operatorname{sign}(b)$ and can vary in the interval $[-1, 1]$ for all $\alpha < \frac{1}{2}$:

$$\beta = \gamma \cdot (1 - 2\alpha) \cdot \operatorname{sign}(b). \quad (31)$$

For example, $\beta = 0.118 \operatorname{sign}(b)$ in Table 2 can be expressed by $\gamma = 0.98$.

Table 3 illustrates the advantages of our nonstandard stencil over the wavelet-inspired stencil II. For small values of α , one can easily improve the PSNR in

Table 3. Comparison between the wavelet-inspired filter class II from [16] and our nonstandard filter class. The table depicts the PSNR for the demosaicking test scenario, and γ_{opt} refers to the γ -value where the nonstandard stencil yields the highest PSNR.

α	wavelet-insp. II	nonstandard	γ_{opt}
0	24.60	29.86	1
0.1	25.27	30.81	1
0.2	26.14	31.90	1
0.3	27.32	33.04	1
0.4	29.08	33.83	1
0.42	29.57	33.87	1
0.44	30.16	33.99	0.98
0.46	30.90	33.96	0.96
0.48	31.98	33.87	0.90
0.49	32.88	33.78	0.80
0.5	29.57	29.57	—

the demosaicking test case by more than 5 dB: All one has to do is to choose $\gamma = 1$ instead of $\gamma = 0$. The latter corresponds to the wavelet-inspired stencil II. As long as α is not too close to the critical value $\frac{1}{2}$ (which should be avoided anyway due to checkerboard artifacts) it turns out that $\gamma = 1$ gives the highest PSNR. However, even for $\alpha \in [0.43, 0.49]$, where $\gamma = 1$ is suboptimal, it still outperforms $\gamma = 0$. Thus, choosing $\gamma = 1$ works well in practice and reduces the parameter space to a single degree of freedom.

We see that within our stencil class, schemes with fairly large values for α and γ perform particularly well. This confirms Mickens’ principle of nonlocal approximation of nonlinear terms: The more α and γ differ from 0, the larger is the contribution of the nonstandard finite difference terms within (10)–(12).

5 Conclusions

We have shown that seven finite difference discretisations for anisotropic diffusion filtering with a diffusion tensor are special cases of a novel, unifying framework. It is derived systematically from a discrete energy formulation, and it exploits the widely unknown nonstandard finite differences of Mickens [18]. We have established general L^2 -stability results. Our framework does not only provide a theoretical foundation of existing schemes as L^2 -stable discrete energy minimisers, but also comprises novel stencils that outperform existing ones.

Our evaluation has shown that different discretisations of the same continuous model can give PSNR differences of almost 10 dB, even though they have identical consistency order. This confirms the widely underestimated fact that appropriate numerical algorithms are at least as important as good models.

We expect that the ideas in our paper can also be generalised to other anisotropic equations that create similar numerical challenges.

References

1. Weickert, J.: *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart (1998)
2. Cottet, G.H., El Ayyadi, M.: A Volterra type model for image processing. *IEEE Transactions on Image Processing* **7**(3) (March 1998) 292–303
3. Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., Seidel, H.P.: Image compression with anisotropic diffusion. *Journal of Mathematical Imaging and Vision* **31**(2–3) (July 2008) 255–269
4. Nagel, H.H., Enkelmann, W.: An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8** (1986) 565–593
5. Zimmer, H., Bruhn, A., Weickert, J.: Optic flow in harmony. *International Journal of Computer Vision* **93**(3) (July 2011) 368–388
6. Zimmer, H., Valgaerts, L., Bruhn, A., Breuß, M., Weickert, J., Rosenhahn, B., Seidel, H.P.: PDE-based anisotropic disparity-driven stereo vision. In Deussen, O., Keim, D., Saupe, D., eds.: *Vision, Modelling, and Visualization 2008*. AKA, Heidelberg (2008) 263–272
7. Schroers, C., Zimmer, H., Valgaerts, L., Bruhn, A., Demetz, O., Weickert, J.: Anisotropic range image integration. In Pinz, A., Pock, T., Bischof, H., Leberl, F., eds.: *Pattern Recognition*. Volume 7476 of *Lecture Notes in Computer Science*. Springer, Berlin (2012) 73–82
8. Preußner, T., Rumpf, M.: An adaptive finite element method for large scale image processing. *Journal of Visual Communication and Image Representation* **11**(2) (June 2000) 183–195
9. Drblíková, O., Mikula, K.: Convergence analysis of finite volume scheme for nonlinear tensor anisotropic diffusion in image processing. *SIAM Journal on Numerical Analysis* **46**(1) (2007) 37–60
10. Jawerth, B., Lin, P., Sinzinger, E.: Lattice Boltzmann models for anisotropic diffusion of images. *Journal of Mathematical Imaging and Vision* **11** (1999) 231–237
11. Weickert, J.: Nonlinear diffusion filtering. In Jähne, B., Haußecker, H., Geißler, P., eds.: *Handbook on Computer Vision and Applications, Vol. 2: Signal Processing and Pattern Recognition*. Academic Press, San Diego (1999) 423–450
12. Mrázek, P., Navara, M.: Consistent positive directional splitting of anisotropic diffusion. In Likar, B., ed.: *Proc. Sixth Computer Vision Winter Workshop, Bled, Slovenia* (February 2001) 37–48
13. Weickert, J., Scharr, H.: A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance. *Journal of Visual Communication and Image Representation* **13**(1/2) (2002) 103–118
14. Felsberg, M.: Autocorrelation-driven diffusion filtering. *IEEE Transactions on Image Processing* **20**(7) (July 2011) 1797–1806
15. Welk, M., Weickert, J., Steidl, G.: From tensor-driven diffusion to anisotropic wavelet shrinkage. In Bischof, H., Leonardis, A., Pinz, A., eds.: *Computer Vision – ECCV 2006, Part I*. Volume 3951 of *Lecture Notes in Computer Science*. Springer, Berlin (2006) 391–403
16. Welk, M., Steidl, G., Weickert, J.: Locally analytic schemes: A link between diffusion filtering and wavelet shrinkage. *Applied and Computational Harmonic Analysis* **24** (2008) 195–224
17. Dascal, L., Ditkowski, A., Sochen, N.: A study of the discrete maximum principle for the Beltrami color flow. *Journal of Mathematical Imaging and Vision* **29**(1) (2007) 63–77

18. Mickens, R.E.: Nonstandard Finite Difference Models of Differential Equations. World Scientific, Singapore (1994)
19. Grewenig, S., Weickert, J., Bruhn, A.: From box filtering to fast explicit diffusion. In Goesele, M., Roth, S., Kuijper, A., Schiele, B., Schindler, K., eds.: Pattern Recognition. Volume 6376 of Lecture Notes in Computer Science. Springer, Berlin (2010) 533–542