Computer Vision, Graphics, and Pattern Recognition Group
Department of Mathematics and Computer Science
University of Mannheim
D-68131 Mannheim, Germany

Reihe Informatik
4/2000

A Scheme for Coherence-Enhancing Diffusion Filtering with Optimized Rotation Invariance

Joachim Weickert and Hanno Scharr

Technical Report 4/2000
Computer Science Series
February 2000

The technical reports of the CVGPR Group are listed under
`http://www.ti.uni-mannheim.de/~bmg/Publications-e.html`

A Scheme for Coherence-Enhancing Diffusion Filtering with Optimized Rotation Invariance

Joachim Weickert[1] **and Hanno Scharr**[2]

[1]Computer Vision, Graphics, and Pattern Recognition Group,
Department of Mathematics and Computer Science,
University of Mannheim, 68131 Mannheim, Germany
E-Mail: Joachim.Weickert@ti.uni-mannheim.de
Web Page: http://www.ti.uni-mannheim/~bmg/weickert

[2]Interdisciplinary Center for Scientific Computing
Ruprecht Karls University
Im Neuenheimer Feld 368
69120 Heidelberg, Germany
E-Mail: Hanno.Scharr@iwr.uni-heidelberg.de
Web Page: http://www.klimt.iwr.uni-heidelberg.de/~hscharr

Abstract

For strongly directed anisotropic processes such as coherence-enhancing diffusion filtering it is crucial to use numerical schemes with highly accurate directional behavior. We show that this is not possible in a satisfactory way when discretizations are limited to 3×3 stencils. As a consequence, we investigate a novel algorithm based on 5×5 stencils. It utilizes recently discovered differentiation filters with optimized rotation invariance. By juxtaposing it with several common algorithms we demonstrate its superior behavior with respect to the following properties: rotation invariance, avoidance of blurring artifacts (dissipativity), and accuracy. The latter one is evaluated by deriving an analytical solution for coherence-enhancing diffusion filtering of images with circular symmetry. Furthermore, we show that the new scheme is 3 to 4 times more efficient than explicit schemes on 3×3 stencils. It does not require to solve linear systems of equations, and it can be easily implemented in any dimension.

Keywords: image processing, anisotropic diffusion, scale-spaces, rotation invariance, filter design, finite differences, analytical solutions, performance evaluation

1 Introduction

In this paper we present and evaluate a novel algorithm for coherence-enhancing anisotropic diffusion filtering. This scale-space and image restoration technique has been introduced in [40] for the enhancement of flow-like textures with line-like structures. It combines ideas of nonlinear diffusion filtering [31, 11, 39] with orientation analysis by means of the so-called structure tensor (second-moment matrix, scatter matrix) [9, 16, 23, 26, 34].

The basic idea is to smooth a degraded original image by applying a nonlinear diffusion process whose diffusion tensor allows anisotropic smoothing by acting mainly along the preferred structure direction. This so-called coherence orientation is determined by the eigenvector of the structure tensor with the smallest eigenvalue.

Such structure tensor driven nonlinear diffusion methods have demonstrated their usefulness for enhancing fingerprint images [40], for restoring old copper plates [14], for computer-aided quality control [40] and medical imaging [6, 45]. They have been applied to three-dimensional microscopy images [18], to the enhancement of corners [35, 39], ridges and valleys [38], and they have even been used for designing novel shock-capturing algorithms for hyperbolic conservation laws [19]. They create image-simplifying scale-space transformations in any dimension [44], and they can be extended to colour images [43, 28].

Since coherence-enhancing anisotropic diffusion filtering is essentially a one-dimensional smoothing strategy in a multidimensional image, it is of outmost importance to have a precise realization of the desired smoothing direction. When the goal consists e.g. of closing gaps in an interrupted line-like structure, it is clear that slight deviations from the correct smoothing direction will destroy any desired filter effect and result in a deterioration of the line by introducing blurring artifacts.

This direction sensitivity constitutes an additional problem for the design of appropriate algorithms for diffusion filtering that has not been addressed in the computer vision literature so far. In this paper we introduce a novel scheme that is especially designed to handle this difficulty. Its main ingredient is the consequent use of specific first-order derivative filters that have been optimized with respect to rotation invariance [37, 25]. We use these filters in an explicit (Euler forward) finite difference scheme. Since they come down to applying 3×3 convolution masks twice, they lead to a scheme which performs local averaging over a 5×5 mask. We shall see that such an algorithm does not only reveal visibly better performance with respect to rotation invariance, but also creates less blurring artifacts, is more accurate, and has a three to four times higher efficiency than other explicit schemes for coherence-enhancing diffusion. Another important feature of the algorithm is its simplicity: it can be easily extended to higher dimensional data sets and it does not require to solve linear systems of equations as is common for semi-implicit techniques.

The paper is organized as follows. In Section 2 we sketch the concept of coherence-enhancing diffusion filtering, we review two finite difference schemes that are often used for this purpose, and we discuss the optimality of one of them with respect to nonnegativity and rotation invariance, if the support is restricted to 3×3 pixels. In Section 3 we present our novel algorithm and give a short description of our optimized derivative filters. In Section 4 we compare its performance with other algorithms. We evaluate rotation invariance, dissipative effects resulting in blurring artifacts, quantitative errors compared to an analytical solution, and efficiency. Finally we conclude with a summary in Section 5.

Related work. Although there is a rich literature on partial differential methods for image processing and computer vision (see e.g. [8, 10, 20, 30, 41] for recent overviews) the design of efficient algorithms for nonlinear diffusion filtering has been addressed to a much smaller extend. For isotropic diffusion filtering with a scalar-valued diffusivity, one can find several alternatives to the conventional two-level explicit finite-difference scheme, for instance three-level methods [17], semi-implicit approaches [11] and their additive operator splitting variants [46], multigrid methods [1], adaptive finite element techniques [7], numerical schemes with wavelets as trial functions [17], spectral methods [17] and stochastic simulations [33]. Numerical techniques for anisotropic diffusion filters with a diffusion tensor include finite elements with grid adaptation [32] or multigrid acceleration [14], as well as lattice Boltzmann techniques [22]. Explicit finite difference schemes [12, 13, 42] have been applied for simplicity reasons, and semi-implicit stabilizations have been introduced to permit larger step sizes [44, 45]. For efficiency reasons they have been realized as multiplicative [45] or additive operator splittings [44]. A spatial discretization that allows to preserve nonnegativity for small condition numbers of the diffusion tensor has been derived in [41]. In this case absolutely stable semi-implicit schemes are possible. To the best of our knowledge there is, however, not a single publication that addresses the important problem of designing algorithms for anisotropic diffusion filtering with an optimized directional behavior. Although there are some publications on optimized filter design for anisotropic equations of mean curvature type [4, 5, 15], their results, however, cannot be transferred directly to coherence-enhancing anisotropic diffusion filtering.

2 Coherence-Enhancing Anisotropic Diffusion

2.1 General filter structure

Anisotropic diffusion filtering with a diffusion tensor evolves the initial image under an evolution equation of type

$$\frac{\partial u}{\partial t} = \nabla \cdot (D \nabla u) \tag{1}$$

where $u(x,t)$ is the evolving image, t denotes the diffusion time, and $D = \begin{pmatrix} a & b \\ b & c \end{pmatrix}$ is the diffusion tensor, a positive definite symmetric matrix that may be adapted to the local image structure.

This local image structure is measured by the so-called structure tensor (scatter matrix, second-moment matrix, Förstner interest operator) [9, 16, 23, 26, 34] which is given by

$$J_\rho(\nabla u_\sigma) = G_\rho * (\nabla u_\sigma \nabla u_\sigma^T).$$

The function G_ρ denotes a Gaussian with standard deviation ρ, and $u_\sigma := G_\rho * u$ is a regularized version of u that is obtained by convolution with a Gaussian G_σ. The eigenvectors of J_ρ give the preferred local orientations, and the corresponding eigenvalues denote the local contrast along these directions. The structure tensor is highly robust under isotropic additive Gaussian noise [24], and it can be implemented efficiently [21].

The eigenvalues of $J_\rho = \begin{pmatrix} J_{11} & J_{12} \\ J_{12} & J_{22} \end{pmatrix}$ are

$$\begin{aligned} \mu_1 &= \frac{1}{2}\left(J_{11} + J_{22} + \sqrt{(J_{11}-J_{22})^2 + 4J_{12}^2}\right), \\ \mu_2 &= \frac{1}{2}\left(J_{11} + J_{22} - \sqrt{(J_{11}-J_{22})^2 + 4J_{12}^2}\right) \end{aligned}$$

and the first eigenvector $(\cos\alpha, \sin\alpha)^T$ satisfies

$$\begin{pmatrix} \cos\alpha \\ \sin\alpha \end{pmatrix} \,||\, \begin{pmatrix} 2J_{12} \\ J_{22} - J_{11} + \sqrt{(J_{11}-J_{22})^2 + 4J_{12}^2} \end{pmatrix}$$

The diffusion tensor D of coherence-enhancing anisotropic diffusion uses the same eigenvectors as the structure tensor, and its eigenvalues are assembled via

$$\lambda_1 := c_1 \tag{2}$$

$$\lambda_2 := \begin{cases} c_1 & \text{if } \mu_1 = \mu_2, \\ c_1 + (1-c_1)\exp\left(\frac{-c_2}{(\mu_1-\mu_2)^2}\right) & \text{else,} \end{cases} \tag{3}$$

where $c_1 \in (0,1)$, $c_2 > 0$. The condition number of D is thus bounded by $1/c_1$, and the entries of D are

$$\begin{aligned} a &= \lambda_1 \cos^2\alpha + \lambda_2 \sin^2\alpha, \\ b &= (\lambda_1 - \lambda_2)\sin\alpha\cos\alpha, \\ c &= \lambda_1 \sin^2\alpha + \lambda_2 \cos^2\alpha\,. \end{aligned} \tag{4}$$

For more details on coherence-enhancing anisotropic diffusion we refer to [44].

2.2 Existing schemes

Equation (1) can be solved numerically using finite differences. Spatial derivatives are usually replaced by central differences, while the easiest way to discretize $\frac{\partial u}{\partial t}$ consists of using a forward difference approximation. The resulting so-called explicit scheme allows to calculate all values at a new time level directly from the ones in the previous level without solving linear or nonlinear systems of equations. An explicit scheme has the basic structure

$$\frac{u_{i,j}^{k+1} - u_{i,j}^{k}}{\tau} = A_{i,j}^{k} * u_{i,j}^{k} \tag{5}$$

where τ is the time step size and $u_{i,j}^k$ denotes the approximation of $u(x,t)$ in the pixel (i,j) at time $k\tau$. The expression $A_{i,j}^k * u_{i,j}^k$ is a discretization of $\nabla \cdot (D\nabla u)$. It comes down to the convolution of the image with a spatially and temporally varying mask $A_{i,j}^k$. Hence, we may calculate u at level $k+1$ directly from u at level k via

$$u_{i,j}^{k+1} = (I + \tau A_{i,j}^k) * u_{i,j}^k.$$

The stencil notation of two common discretizations for $A_{i,j}^k$ are shown in Figure 1. We assume that the pixels have length 1 in both directions. The so-called standard discretization [42] from Fig. 1(a) is the simplest way to discretize (1). The more complicated nonnegativity discretization from Fig. 1(b) offers one advantage [41]: all non-central weights remain nonnegative as long as the condition number of the diffusion tensor does not exceed $3 + 2\sqrt{2} \approx 5.828$. In this case the corresponding explicit scheme satisfies a very strict stability property which prevents blow-ups of the numerical solution: for a sufficiently small step size τ, the global minimum cannot decrease and the global maximum cannot increase.

Using models with restricted condition numbers of the diffusion tensor, however, means also to limit the anisotropy of the filter and is therefore not always desirable. In those cases where one is interested in strongly anisotropic behavior, the condition number may exceed $3+2\sqrt{2}$. Then the nonnegativity discretization has not necessarily advantages over the standard discretization.

A typical shortcoming of explicit schemes is that they usually require fairly small time steps in order to be stable. A common way to get better stability properties is to use some or all values from the unknown level $k+1$ in the approximation of the spatial derivatives as well. The resulting semi-implicit or fully implicit techniques require the solution of linear or nonlinear equations, which can become costly or complicated. An efficient semi-implicit strategy that leads to simple linear systems of equations stabilizes the explicit standard discretization by means of an additive operator splitting (AOS); see [44] for more details. The mixed derivative terms in this scheme, however, are still discretized explicitly. Such an AOS-stabilized standard discretization has a larger stability range than the explicit standard discretization, but it does not enjoy the absolute stability of pure AOS schemes like the ones in [46].

A pure AOS scheme based on the nonnegativity discretization is proposed in [41]. It is absolutely stable, if the condition number of the diffusion tensor remains below $3 + 2\sqrt{2}$, but cannot be applied beyond this limit. It is therefore not considered in our evaluations in Section 4, where condition numbers up to 1000 may appear.

2.3 A remark on optimization of 3×3-kernels

The existing schemes have been introduced for simplicity or stability reasons, but no attention has been paid to optimizing their directional behavior. Indeed, in Section 4 we shall see that their rotation invariance is not very good in certain situations.

In order to address this problem, we have used the optimization framework from [37] and [25] for optimizing 3×3-kernels for strongly directed anisotropic diffusion with respect to rotation invariance. For this task the good behavior of the kernels along

$\frac{-b_{i-1,j}-b_{i,j+1}}{4}$	$\frac{c_{i,j+1}+c_{i,j}}{2}$	$\frac{b_{i+1,j}+b_{i,j+1}}{4}$
$\frac{a_{i-1,j}+a_{i,j}}{2}$	$-\frac{a_{i-1,j}+2a_{i,j}+a_{i+1,j}}{2}-\frac{c_{i,j-1}+2c_{i,j}+c_{i,j+1}}{2}$	$\frac{a_{i+1,j}+a_{i,j}}{2}$
$\frac{b_{i-1,j}+b_{i,j-1}}{4}$	$\frac{c_{i,j-1}+c_{i,j}}{2}$	$\frac{-b_{i+1,j}-b_{i,j-1}}{4}$

$\frac{\lvert b_{i-1,j+1}\rvert-b_{i-1,j+1}}{4}+\frac{\lvert b_{i,j}\rvert-b_{i,j}}{4}$	$\frac{c_{i,j+1}+c_{i,j}}{2}-\frac{\lvert b_{i,j+1}\rvert+\lvert b_{i,j}\rvert}{2}$	$\frac{\lvert b_{i+1,j+1}\rvert+b_{i+1,j+1}}{4}+\frac{\lvert b_{i,j}\rvert+b_{i,j}}{4}$
$\frac{a_{i-1,j}+a_{i,j}}{2}-\frac{\lvert b_{i-1,j}\rvert+\lvert b_{i,j}\rvert}{2}$	$-\frac{a_{i-1,j}+2a_{i,j}+a_{i+1,j}}{2}-\frac{\lvert b_{i-1,j+1}\rvert-b_{i-1,j+1}+\lvert b_{i+1,j+1}\rvert+b_{i+1,j+1}}{4}-\frac{\lvert b_{i-1,j-1}\rvert+b_{i-1,j-1}+\lvert b_{i+1,j-1}\rvert-b_{i+1,j-1}}{4}+\frac{\lvert b_{i-1,j}\rvert+\lvert b_{i+1,j}\rvert+\lvert b_{i,j-1}\rvert+\lvert b_{i,j+1}\rvert+2\lvert b_{i,j}\rvert}{2}-\frac{c_{i,j-1}+2c_{i,j}+c_{i,j+1}}{2}$	$\frac{a_{i+1,j}+a_{i,j}}{2}-\frac{\lvert b_{i+1,j}\rvert+\lvert b_{i,j}\rvert}{2}$
$\frac{\lvert b_{i-1,j-1}\rvert+b_{i-1,j-1}}{4}+\frac{\lvert b_{i,j}\rvert+b_{i,j}}{4}$	$\frac{c_{i,j-1}+c_{i,j}}{2}-\frac{\lvert b_{i,j-1}\rvert+\lvert b_{i,j}\rvert}{2}$	$\frac{\lvert b_{i+1,j-1}\rvert-b_{i+1,j-1}}{4}+\frac{\lvert b_{i,j}\rvert-b_{i,j}}{4}$

Figure 1: **(a) Top:** standard discretization [42]. **(b) Bottom:** nonnegativity discretization [41]

the axes is required for all orientations. A point symmetric, zero mean 3×3-kernel has 4 degrees of freedom. As one can show, if consistency is required, 3 degrees of freedom are fixed. Surprisingly, optimizing the remaining parameter resulted in the nonnegativity discretization. For this reason, we refrain from presenting more details about the optimization. This result also shows that, if one is interested in better rotation invariance, one has to use larger stencil sizes than 3×3. This shall be done next.

3 A Novel Algorithm with Optimized Rotation Invariance

In order to explain our novel scheme which works on a 5×5 stencil, we rewrite the divergence operator in (1) as

$$\nabla \cdot (D\nabla u) = \partial_x(a\partial_x u + b\partial_y u) + \partial_y(b\partial_x u + c\partial_y u). \tag{6}$$

This expression is now evaluated in an explicit way, i.e. using only known values form the old time level k.

The key point is that we use first derivative operators with the stencil notations

$$F_x = \frac{1}{32}\begin{bmatrix} -3 & 0 & 3 \\ -10 & 0 & 10 \\ -3 & 0 & 3 \end{bmatrix} \quad \text{and} \quad F_y = \frac{1}{32}\begin{bmatrix} 3 & 10 & 3 \\ 0 & 0 & 0 \\ -3 & -10 & -3 \end{bmatrix}. \tag{7}$$

They have been derived recently in [37, 25], where it has been shown that they approximate rotation invariance significantly better than related popular stencils like the Sobel filters.

In order to motivate these derivative filters, it is useful to recall that the transfer function of an ideal first-order 1D derivative filter is given by $i\pi\tilde{k}$, where $\tilde{k}$ denotes the normalized wave number. The transfer function of the commonly used central derivative operator $[-\frac{1}{2}, 0, \frac{1}{2}]$, however, is $i\pi\tilde{k}\,\mathrm{sinc}(\pi\tilde{k})$. Multiplication with the sinc-function in the Fourier domain is equivalent to smoothing with a box filter. Indeed, if $\tilde{k} \to \pm 1$, the transfer function approaches 0. This means that high frequencies that are close to the Nyquist frequency are unaffected by this central difference operator, if they are in the direction of differentiation. In order to maintain rotation invariance, one should apply a similar smoothing operator perpendicular to the direction of the derivative. A weighted L_2-optimization in the Fourier domain showed that this effect is well approximated by the smoothing stencil

$$\frac{1}{16}\begin{bmatrix} 3 \\ 10 \\ 3 \end{bmatrix}.$$

The convolution between this smoothing stencil and the standard central derivative operator in x direction immediately gives the preceding convolution operator F_x. This filter design is illustrated in Figure 2. Since F_x and F_y contain only two different numbers, convolutions of this type are fairly inexpensive (2 multiplications, 4 additions, 1 subtraction per pixel). For more details about the optimization of the first order derivative filter, we refer to [37, 25].

Now we can turn our attention to the diffusion approximation again. We proceed in five steps:

1. Calculate the structure tensor using the optimized derivative filters from (7).

2. Assemble the diffusion tensor as a function of the structure tensor.

3. Calculate the flux components $j_1 := a\partial_x u + b\partial_y u$ and $j_2 := b\partial_x u + c\partial_y u$ with the optimized filters.

4. Calculate $\nabla \cdot (D\nabla u) = \partial_x j_1 + \partial_y j_2$ by means of the optimized derivative filters.

5. Update in an explicit way.

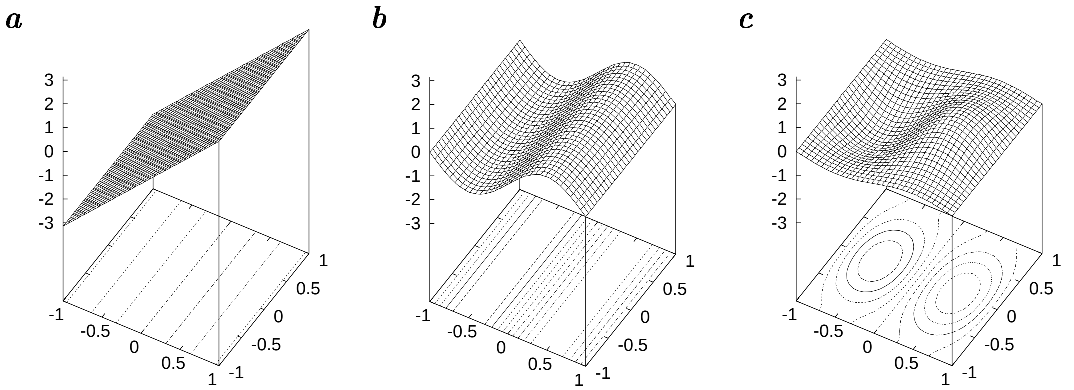

Figure 2: Transfer functions of derivatives. **a** Ideal operator. **b** Usual $[-\frac{1}{2},\ 0\ ,\frac{1}{2}]$ filter without smoothing in cross direction. Note the inherent smoothing in the direction of derivation. **c** Sobel-like 3×3-filter mask (7) with optimized smoothing in cross direction.

Since the resulting scheme makes consequent use of the derivative filters with optimized rotation invariance, we may expect good directional behavior. It should be noted that the total stencil of this scheme has size 5×5, since we are approaching the second order derivatives by consecutively applying first order derivatives of size 3×3. However, there is no need to write down a complicated 5×5 stencil, since it is nowhere required in the entire algorithm.

4 Performance Evaluation

In this section we shall juxtapose our method with existing ones in order to evaluate its performance. We focus on investigating rotation invariance, dissipative effects creating blurring artifacts, quantitative errors, and efficiency.

4.1 Rotation invariance and dissipativity

Test computations were performed on a ring image with varying frequencies (see Figure 3(a)). The maximum wave number is 0.5 . We apply three different explicit schemes for coherence-enhancing diffusion: the standard scheme, the nonnegativity scheme, and our novel one. As parameters we use $c_1 := 0.001$, $c_2 := 1$, $\sigma = 0.5$, $\rho = 10$, $\tau = 0.2$, and we apply 100 iterations.

Applying coherence enhancing anisotropic diffusion to this ring image should not deteriorate the rotation invariance, and it should hardly alter this image (only by some small isotropic diffusion caused by the parameter c_1; see Section 4.2).

Figure 4 shows the results for the upper right quadrant. The other quadrants look the same. One observes that the standard discretization introduces severe blurring artifacts for all directions except for the directions of the coordinate axes. This is caused by incorrect directional behavior and undesired damping of high frequencies, as can be seen if the transfer function for specific choices of λ_1, λ_2 and α are examined.

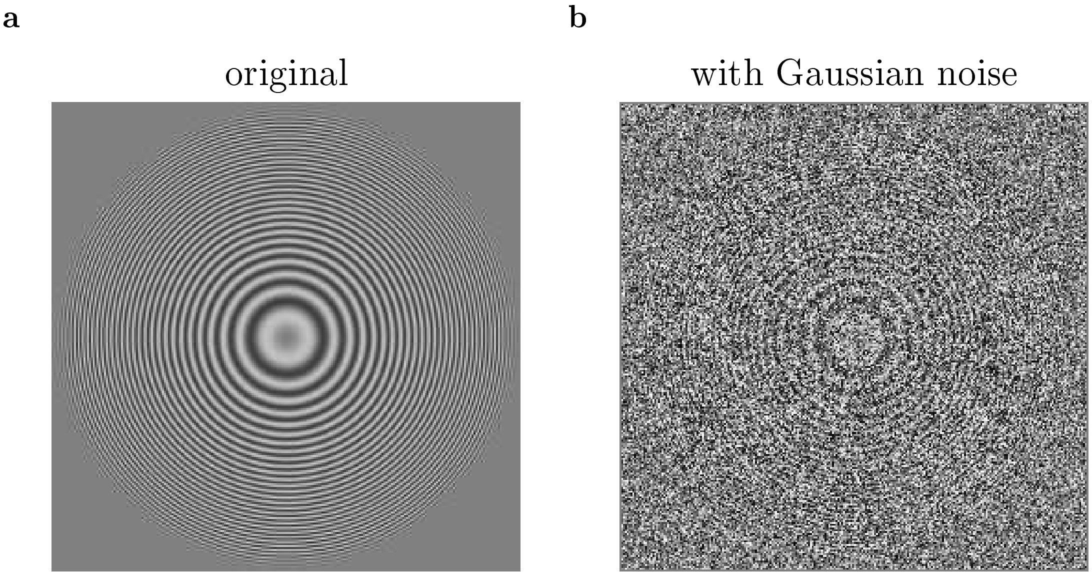

Figure 3: Test image: **a** Original, **b** with Gaussian noise

If the diffusion does not exactly follow the circles, it immediately blurs them. Similar dissipative effects can be observed for the nonnegativity discretization. However, it also performs well along the grid diagonals. As explained in Section 2.3 this is the optimal behavior for a 3×3-stencil. For the new scheme dissipative effects or deviations from rotational invariance cannot be observed, the result is visually indistinguishable from the original image. A quantitative analysis of these deviations is presented in the next section.

To demonstrate the importance of rotation invariance, we use our three implementations for reducing Gaussian noise that has been added to the test image (cf. Figure 3(b)). The Gaussian noise has zero mean, and the standard deviation has the same magnitude as the signal amplitude. Figure 5 shows the upper right quadrant using the same filter parameters as before. Only the new discretization reconstructs the signal satisfactory for all orientations and frequencies. For the other methods the signal information is so weak that the directional errors of the discretization dominate for high frequencies.

Smoothing effects result in loss of fine structures in the depicted ring images (Figure 4). For orientations parallel to the coordinate axes, all discretizations preserve the signal. Blurring artifacts therefore might be a directional failure only, with less effect for real data. In order to demonstrate that this error should not be neglected we applied the three algorithms to van Goghs "Road with Cypress and Star" (Auvers-sur-Oise, 1890; Otterlo, Rijksmuseum Kröller-Müller). This test image has been used in [44] for evaluating coherence-enhancing diffusion filtering. The general impression from Figure 6 is that the new scheme produces the sharpest and most detailed results. For all images processed with the new discretization one observes slight checkerboard-like structures. An explanation for this in terms of transfer functions has been given in Section 3: the derivative operators become less sensitive the more the frequencies approach the Nyquist

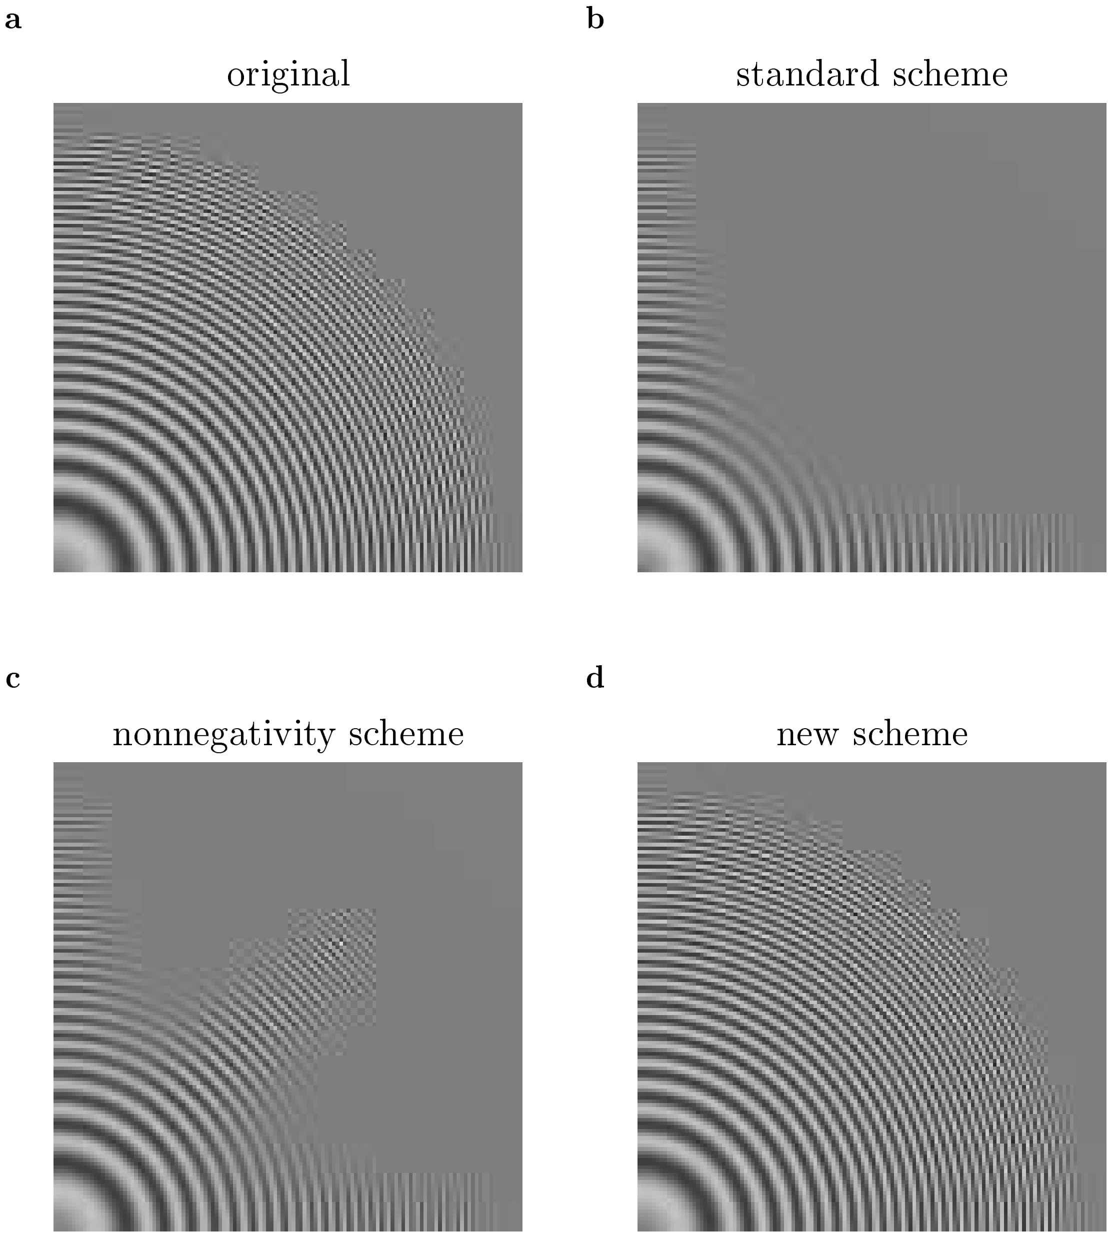

Figure 4: Rotation invariance and dissipativity test: Upper right quadrant of original ring image after applying three explicit schemes for coherence-enhancing diffusion filtering.

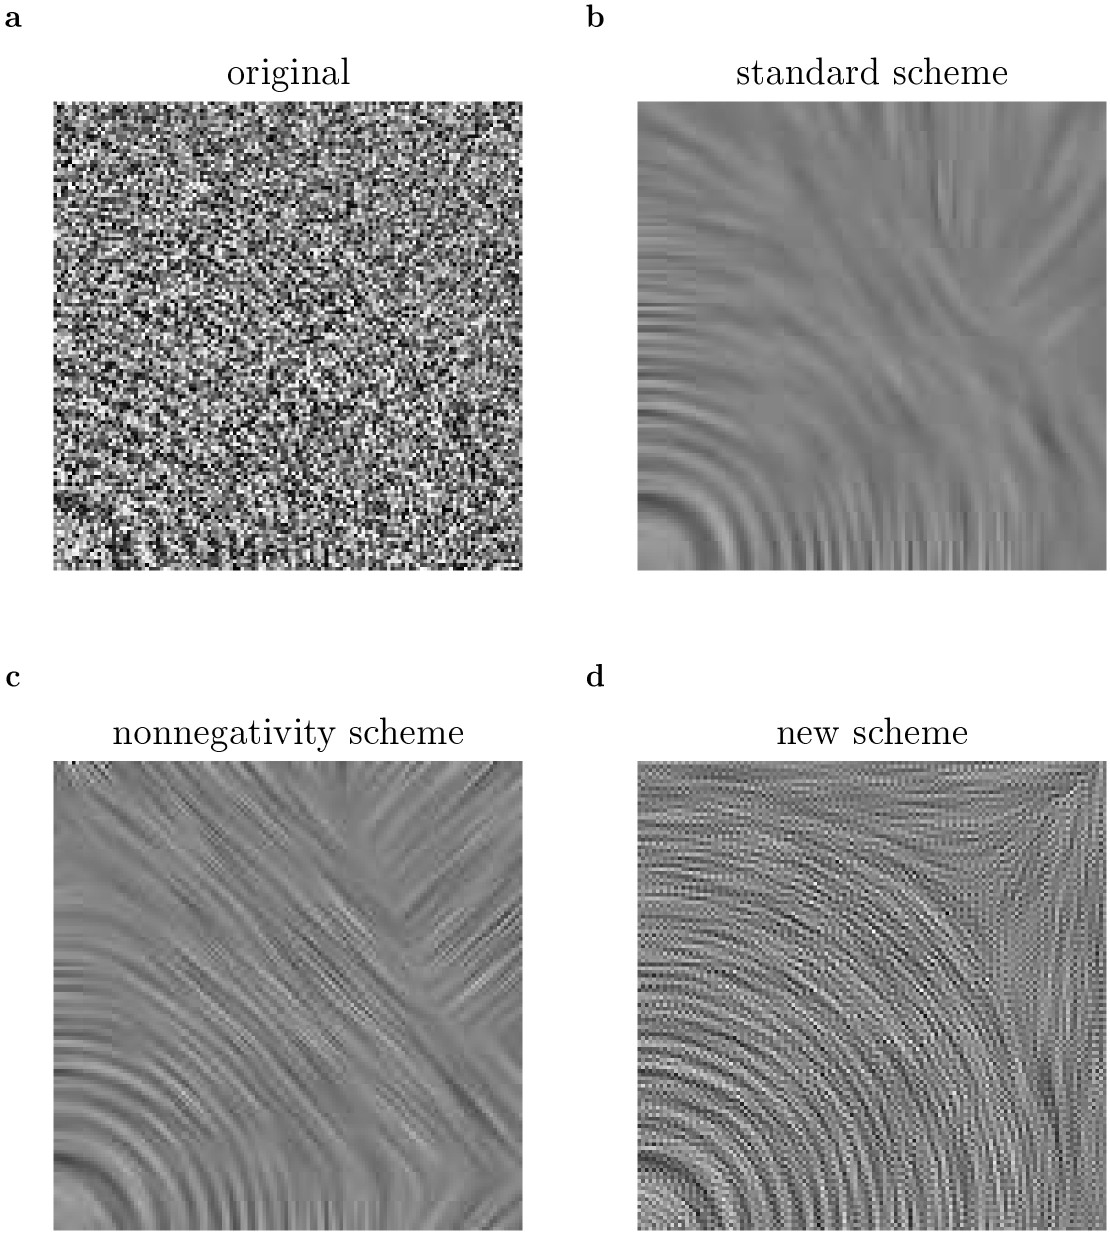

Figure 5: Reconstruction properties for the upper right quadrant of the noisy ring image from Figure 3(b).

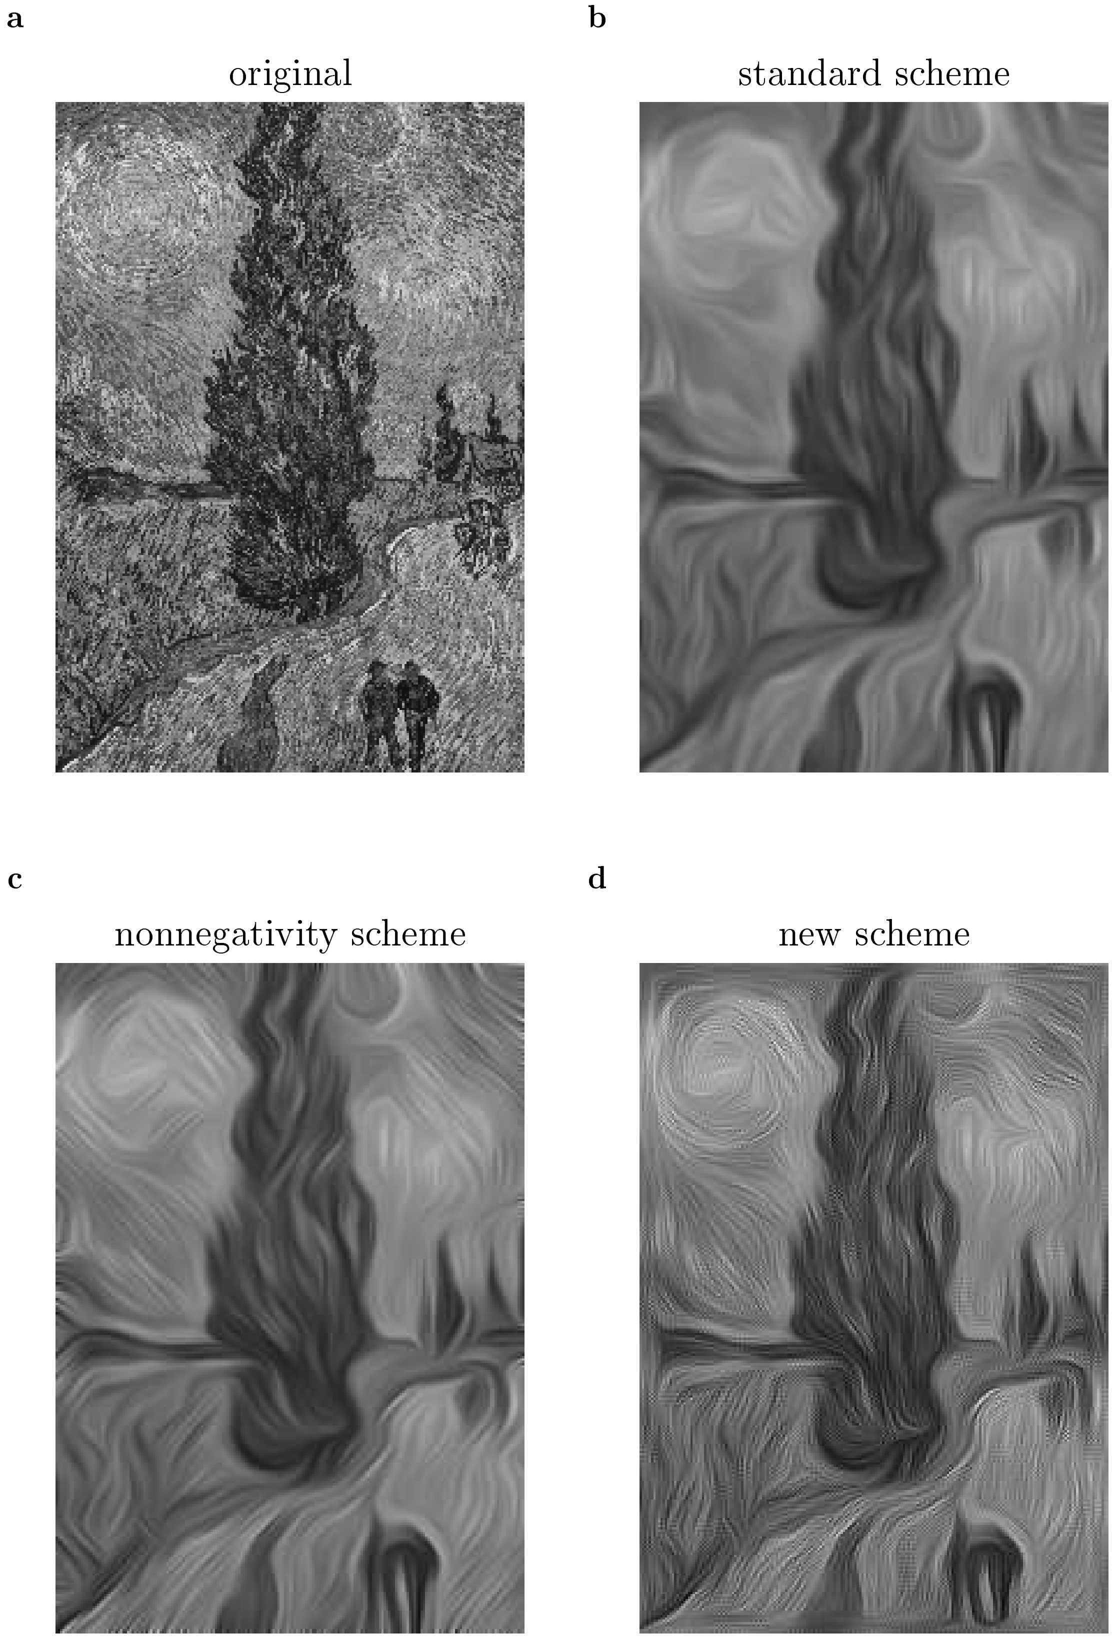

Figure 6: Dissipativity illustrated by means of van Goghs "Road with Cypress and Star". The filter parameters are $c_1 = 0.001$, $c_2 = 1$, $\sigma = 0.7$, $\rho = 4$, $\tau = 0.2$. In order to make the differences better visible, 100 iterations have been applied.

rate. Thus higher frequencies are smoothed to a smaller amount. A further analysis of this phenomenon is presented in Section 4.3. More detailed experiments showed that these structures remain bounded and decrease over time.

4.2 Quantitative analysis

Interestingly, coherence enhancing anisotropic diffusion can be simplified in specific situations such that even an analytical solution can be given. In fact, if we have images with linear symmetry (planar waves) or circular symmetry, then the eigenvectors of the structure tensor do not depend on the smoothing scales σ and ρ, and they are parallel or perpendicular to ∇u. Since the diffusion tensor has the same eigenvalues as the structure tensor, it follows that, for these images, coherence-enhancing diffusion degenerates to the isotropic linear diffusion equation

$$\frac{\partial u}{\partial t} = c_1 \, \triangle u. \tag{8}$$

This equation has analytical solutions. For circular symmetric initial images, they can be expressed in terms of the Bessel functions J_0 and the Weber functions Y_0 of order zero:

$$u(\mathbf{r}, t) = \int_{-\infty}^{\infty} (A(k)\, J_0(k|\mathbf{r}|) \; + \; B(k)\, Y_0(k|\mathbf{r}|))\; e^{-k^2 \lambda_1 t} \, \mathrm{d}k \, . \tag{9}$$

We may use these analytical solutions for evaluating the accuracy of our algorithms. We start with equation (9) with $A(k) = \delta(k - k_0)$ and $B(k) \equiv 0$ as initial image (t=0) and perform one iteration of our algorithms. The deviations from the analytical solution are plotted in Figure 7 for varying wave numbers k_0 and different explicit schemes.

As expected one observes that the errors of all methods increase with increasing frequencies. In most cases the nonnegativity discretization performs better than the standard discretization, while the novel scheme clearly outperforms both for all frequencies. Over a wide range of frequencies, the error of the new approach is smaller by 1.5 to 2.5 orders of magnitude than errors of the other discretizations. This shows that blurring artifacts resulting from violations of rotational invariance are severe sources for numerical errors. The fact that the novel scheme is also superior for high frequencies shows that its checkerboard-like structures are far less problematic than the blurring artifacts of other schemes.

4.3 Efficiency

Let us now evaluate the efficiency of our method in comparison with the explicit scheme with standard discretization, the explicit nonnegativity discretization, and an AOS-stabilized semi-implicit standard discretization [42]. We did not display results for the latter scheme in the previous experiments, because they were visually identical with those from the underlying explicit standard discretization. There are, however, clear efficiency differences, as we shall see now.

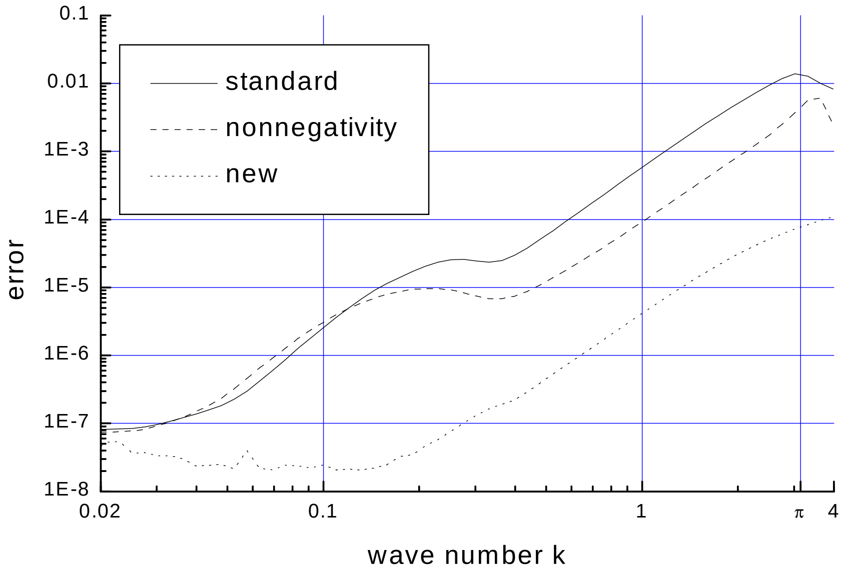

Figure 7: Quantitative error analysis: Mean square root of deviations per pixel between the numerical solution and the analytical solution (9). The parameters are $\sigma = 0$, $c_1 = 0.001$, $\rho = 1$, and $\tau = 0.24$. The error is normalized such that an error of 1 describes the deviation between the the analytical solution and a flat image with the same average grey value.

The total efficiency of a scheme is the product of the computational cost for one iteration and the number of iterations that are required for reaching a fixed diffusion time T. The latter depends on the largest time step size under which the scheme is stable.

Unfortunately, is is not possible to derive theoretical stability bounds for these methods. The reason is that the von Neumann stability using the Fourier transform is not applicable to our nonlinear equations, and stability reasonings based on maximum–minimum principles cannot be applied either, since the discretization of the mixed derivatives do in general create some negative weights. The only exception is the nonnegativity discretization provided that the condition number of the diffusion tensor remains below $3 + 2\sqrt{2} \approx 5.828$. In our case with $c_1 = 0.001$, however, condition numbers up to 1000 are possible, and in this case negative weights may also appear for the nonnegativity discretization.

Owing to the lack of strict theoretical results, we have to perform experimental stability measurements. As a stability criterion we use the temporal evolution of the variance of the filtered image. From theory one knows that the variance has to decrease monotonously for diffusion filters [41]. Thus, if a numerical scheme does increase the variance from one step to the next, it is a clear sign of instabilities. As an upper estimate for an experimentally stable behavior we have therefore searched for the largest time step size for which the variance decreases monotonously for the first 100 iterations.

Table 1: Efficiency of the different methods on a PC (Pentium II MMX, 440 MHz). All algorithms were implemented in a comparable way using ANSI C, and we used the same image and filter parameters as in Figure 6.

discretization	CPU time per iteration	stability bound
explicit, standard	0.2860 s	0.5
explicit, nonnegativity	0.2951 s	0.5
AOS-stabilized, standard	0.3068 s	2.1
our scheme	0.3415 s	2.1

The results of our efficiency analysis are depicted in Table 1. It should be observed that our stability criterion tends to overestimate the stability bounds: instabilities may also arise before the monotony of the variance is violated. In order to be on the safe side we recommend to use the time step size 0.25 for the explicit standard or nonnegativity discretizations, and to use the step size 1 for the AOS-stabilized standard discretization and our novel scheme. In this case the total CPU times for diffusing the van Gogh image until $T = 10$ on a PC (Pentium II MMX, 440 MHz) are as follows: 11.44 s for the explicit standard discretization, 11.80 s for the explicit nonnegativity discretization, 3.07 s for the AOS-stabilized standard discretization, and 3.42 s for the scheme proposed in this paper. This shows that our novel scheme is also very efficient. In spite of being explicit, it is almost as fast as the semi-implicit AOS-stabilized scheme. It is 3-4 times more efficient than the other explicit techniques.

This is caused by the fact that our scheme allows four times larger time step sizes than the other explicit techniques. An intuitive explanation for this can be given by considering an extremely simplified example: when discretizing the linear diffusion equation $\partial_t u = \partial_{xx} u$ in an explicit way we end up with the scheme

$$u_i^{k+1} = u_i^k + \frac{\tau}{h^2}(u_{i+1}^k - 2u_i^k + u_{i-1}^k),$$

where the upper indices denote the time discretization, the lower ones are the space discretization, τ is the time step size, and h is the spatial step size. From standard literature on numerical methods for PDEs [29] it is well-known that such a scheme is stable for $\tau \leq h^2/2$. A corresponding 1D version of our scheme would be

$$u_i^{k+1} = u_i^k + \frac{\tau}{(2h)^2}(u_{i+2}^k - 2u_i^k + u_{i-2}^k).$$

This method performs averaging over pixels that have the distance $2h$ instead of h as in the previous case. It is thus stable for step sizes $\tau \leq (2h)^2/2$, which is just four times the stability bound of the explicit scheme. These considerations suggest that it is the averaging over a 5×5 mask that is responsible to the increased stability bound of our method compared to classical explicit schemes which average over a 3×3 mask.

As already mentioned, this 1D example is extremely simplified. In particular it should be observed that the novel 2D scheme is much more than a traditional scheme on a subsampled grid: while such a simple subsampling would increase dissipative effects, the experiments in this section have shown that the opposite is the case for our scheme. Moreover, in the 2D case with optimized spatial derivative operators one obtains 5×5 masks that do in general contain no zero entries. Exceptions are those cases where the structure orientation coincides with the principal grid directions. But even here the stencil does not decouple into two subgrids. This explains the experimental fact that checkerboard-like structures decrease for larger times.

If checkerboard structures appear unpleasant, there are several ways to modify the scheme in order to shorten their lifetime. For instance, one can split the isotropic fraction $c_1 I$ of the diffusion tensor from the remainder, and discretize $c_1 \Delta u$ by means of a classic 3×3 stencil. It is also possible to alternate the novel scheme with one of the others that require a 3×3 stencil, since all schemes are consistent discretizations of coherence-enhancing diffusion filtering. It should, however, be noted that these modifications complicate the scheme, and that they may introduce additional dissipative effects and deviations from rotation invariance. In order to show the extremes, we focused in this paper on a pure 5×5 scheme without any blendings with existing 3×3 schemes.

Last but not least, it should be mentioned that it is possible to consider also implicit schemes that are based on the 5×5 stencil of the scheme in this paper. One can expect that they have an increased stability range. There is of coarse no free lunch: the price that one has to pay for this is the solution of a linear system of equations with up to 25 nonvanishing entries in each row.

5 Summary and Conclusions

In the present paper we have introduced an explicit discretization for coherence-enhancing anisotropic diffusion filtering which uses first-order derivative approximations that have been optimized for rotational invariance. In a detailed evaluation with existing schemes we have shown that the novel scheme has superior directional performance. This point is very important for anisotropic diffusion techniques, since directional errors introduce visible smoothing artifacts (dissipative effects) and large numerical errors. We have also shown that our scheme, which comes down to averaging over 5×5 masks, allows to use three to four times larger time step sizes than conventional explicit schemes that perform 3×3 averaging. With this efficiency gain it is almost as efficient as an AOS-stabilized semi-implicit technique. These performance characteristics render it the first choice in all situations where a simple algorithm for coherence-enhancing diffusion is needed that combines good quality with high efficiency.

It is our hope that the basic idea of using optimized first-order derivative operators in order to create good directional behavior of anisotropic filters can also be successfully applied to other important anisotropic PDEs such as mean curvature motion [3, 27] or the affine invariant morphological scale-space [2, 36].

References

[1] S.T. Acton, *Multigrid anisotropic diffusion*, IEEE Trans. Image Proc., Vol. 7, 280–291, 1998.

[2] L. Alvarez, F. Guichard, P.-L. Lions, J.-M. Morel, *Axioms and fundamental equations in image processing*, Arch. Rational Mech. Anal., Vol. 123, 199–257, 1993.

[3] L. Alvarez, P.-L. Lions, J.-M. Morel, *Image selective smoothing and edge detection by nonlinear diffusion. II*, SIAM J. Numer. Anal., Vol. 29, 845–866, 1992.

[4] L. Alvarez, F. Morales, *Affine morphological multiscale analysis of corners and multiple junctions*, Int. J. Comput. Vision., Vol. 25, 95–107, 1997.

[5] L. Alvarez, J.-M. Morel, *Formalization and computational aspects of image analysis*, Acta Numerica, 1–59, 1994.

[6] I. Bajla, I. Holländer, *Nonlinear filtering of magnetic resonance tomograms by geometry-driven diffusion*, Machine Vision and Applications, Vol. 10, 243–255, 1998.

[7] E. Bänsch, K. Mikula, *A coarsening finite element strategy in image selective smoothing*, Computation and Visualization in Science, Vol. 1, 53–61, 1997.

[8] M.-O. Berger, R. Deriche, I. Herlin, J. Jaffré, J.-M. Morel (Eds.), *ICAOS '96: Images, wavelets and PDEs*, Lecture Notes in Control and Information Sciences, Vol. 219, Springer, London, 1996.

[9] J. Bigün, G.H. Granlund, *Optimal orientation detection of linear symmetry*, Proc. First Int. Conf. on Computer Vision (ICCV '87, London, June 8–11, 1987), IEEE Computer Society Press, Washington, 433–438, 1987.

[10] V. Caselles, J.M. Morel, G. Sapiro, A. Tannenbaum (Eds.), *Special issue on partial differential equations and geometry-driven diffusion in image processing and analysis*, IEEE Trans. Image Proc, Vol. 7, No. 3, March 1998.

[11] F. Catté, P.-L. Lions, J.-M. Morel, T. Coll, *Image selective smoothing and edge detection by nonlinear diffusion*, SIAM J. Numer. Anal., Vol. 29, 182–193, 1992.

[12] G.-H. Cottet, M. El Ayyadi, *A Volterra type model for image processing*, IEEE Trans. Image Proc., Vol. 7, 292–303, 1998.

[13] G.-H. Cottet, L. Germain, *Image processing through reaction combined with nonlinear diffusion*, Math. Comp., Vol. 61, 659–673, 1993.

[14] A. Dressel, *Die nichtlineare Diffusion in der Bildverarbeitung*, M.Sc. thesis, Faculty of Mathematics, University of Heidelberg, Germany, 1999.

[15] L. Lucido, R. Deriche, L. Alvarez, V. Rigaud, *Sur quelques schémas numériques de résolution d'équations aux dérivées partielles pour le traitement d'images*, Report No. 3192, ROBOTVIS, INRIA, 2004 route des Lucioles, BP 93, 06902 Sophia–Antipolis Cedex, France, 1997.

[16] W. Förstner, E. Gülch, *A fast operator for detection and precise location of distinct points, corners and centres of circular features*, Proc. ISPRS Intercommission Conf. on Fast Processing of Photogrammetric Data (Interlaken, June 2–4, 1987), 281–305, 1987.

[17] J. Fröhlich, J. Weickert, *Image processing using a wavelet algorithm for nonlinear diffusion*, Report No. 104, Laboratory of Technomathematics, University of Kaiserslautern, P.O. Box 3049, 67653 Kaiserslautern, Germany, 1994.

[18] A.S. Frangakis, R. Hegerl, *Nonlinear anisotropic diffusion in three-dimensional electron microscopy*, M. Nielsen, P. Johansen, O.F. Olsen, J. Weickert (Eds.), Scale-space theories in computer vision, Lecture Notes in Computer Science, Vol. 1682, Springer, Berlin, 386–397, 1999.

[19] T. Grahs, A. Meister, T. Sonar, *Image processing for numerical approximations of conservation laws: nonlinear anisotropic artificial dissipation.* Technical Report F8, Institute for Applied Mathematics, University of Hamburg, Bundesstr. 55, 20146 Hamburg, Germany, 1998.

[20] B. ter Haar Romeny, L. Florack, J. Koenderink, M. Viergever (Eds.), *Scale-space theory in computer vision*, Lecture Notes in Computer Science, Vol. 1252, Springer, Berlin, 1997.

[21] H. Haußecker, B. Jähne, *A tensor approach for precise computation of dense displacement vector fields*, E. Paulus, F.M. Wahl (Eds.), Mustererkennung 1997, 199–208, Braunschweig, Springer, September, 1997.

[22] B. Jawerth, P. Lin, E. Sinzinger, *Lattice Boltzmann models for anisotropic diffusion of images*, M. Nielsen, P. Johansen, O.F. Olsen, J. Weickert (Eds.), Scale-space theories in computer vision, Lecture Notes in Computer Science, Vol. 1682, Springer, Berlin, 283–293, 1999.

[23] B. Jähne, *Spatio-temporal image processing*, Lecture Notes in Computer Science, Vol. 751, Springer, Berlin, 1993.

[24] B. Jähne, *Performance characteristics of low-level motion estimators in spatiotemporal images*, W. Förstner (Ed.), DAGM-Workshop Performance Characteristics and Quality of Computer Vision Algorithms, Braunschweig, September 18, 1997.

[25] B. Jähne, H. Scharr, S. Körkel, *Principles of filter design*, B. Jähne, H. Haußecker, P. Geißler (Eds.), Handbook on Computer Vision and Applications, Vol. 2: Signal Processing and Pattern Recognition, Academic Press, San Diego, 125–152, 1999.

[26] M. Kass, A. Witkin, *Analyzing oriented patterns*, Computer Vision, Graphics, and Image Processing, Vol. 37, 362–385, 1987.

[27] B.B. Kimia, A. Tannenbaum, S.W. Zucker, *On the evolution of curves via a function of curvature. I. The classical case*, J. Math. Anal. Appl., Vol. 163, 438–458, 1992.

[28] R. Kimmel, R. Malladi, N. Sochen, *Images as embedded maps and minimal surfaces: movies, color, texture, and volumetric medical images*, Int. J. Comput. Vision, to appear.

[29] K.W. Morton, D.F. Mayers, *Numerical solution of partial differential equations*, Cambridge University Press, Cambridge, 1994.

[30] M. Nielsen, P. Johansen, O.F. Olsen, J. Weickert (Eds.), *Scale-space theories in computer vision*, Lecture Notes in Computer Science, Springer, Berlin, Vol. 1682, 1999.

[31] P. Perona, J. Malik, *Scale space and edge detection using anisotropic diffusion*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 12, 629–639, 1990.

[32] T. Preußer, M. Rumpf, *An adaptive finite element method for large scale image processing*, M. Nielsen, P. Johansen, O.F. Olsen, J. Weickert (Eds.), Scale-space theories in computer vision, Lecture Notes in Computer Science, Vol. 1682, Springer, Berlin, 223–234, 1999.

[33] U.S. Ranjan, K.R. Ramakrishnan, *A stochastic scale space for multiscale image representation*, M. Nielsen, P. Johansen, O.F. Olsen, J. Weickert (Eds.), Scale-space theories in computer vision, Lecture Notes in Computer Science, Vol. 1682, Springer, Berlin, 441–446, 1999.

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

[35] K.O. Riedel, *Corner-preserving anisotropic diffusion and junction detection using the structure tensor*, W. Förstner, J.M. Buhmann, A. Faber, P. Faber (Eds.), Mustererkennung 1999, Springer, Berlin, 164–171, 1999.

[36] G. Sapiro, A. Tannenbaum, *Affine invariant scale-space*, Int. J. Comput. Vision, Vol. 11, 25–44, 1993.

[37] H. Scharr, S. Körkel, B. Jähne, *Numerische Isotropieoptimierung von FIR-Filtern mittels Querglättung*, E. Paulus, F.M. Wahl (Eds.), Mustererkennung 97, 367–374, Braunschweig, Springer, September, 1997.

[38] A.F. Solé, A. López, C. Cañero, P. Radeva, J. Saludes, *Crease enhancement diffusion*, M.I. Torres, A. Sanfeliu (Eds.), Pattern Recognition and Image Analysis (VIII NSPRIA, Bilbao, 1999), Vol. 1, 279–286, 1997.

[39] J. Weickert, *Scale-space properties of nonlinear diffusion filtering with a diffusion tensor*, Report No. 110, Laboratory of Technomathematics, University of Kaiserslautern, P.O. Box 3049, 67653 Kaiserslautern, Germany, 1994.

[40] J. Weickert, *Multiscale texture enhancement*, V. Hlaváč, R. Šára (Eds.), Computer analysis of images and patterns, Lecture Notes in Computer Science, Vol. 970, Springer, Berlin, 230–237, 1995.

[41] J. Weickert, *Anisotropic diffusion in image processing*, Teubner, Stuttgart, 1998.

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

[43] J. Weickert, *Coherence-enhancing diffusion of colour images*, A. Sanfeliu, J.J. Villanueva, J. Vitrià (Eds.), Pattern Recognition and Image Analysis (VII NSPRIA, Barcelona, April 21–25, 1997), Vol. 1, 239–244, 1997. Extended version in Image and Vision Computing, Vol. 17, 199–210, 1999.

[44] J. Weickert, *Coherence-enhancing diffusion filtering*, Int. J. Comput. Vision, Vol. 31, 111–127, 1999.

[45] J. Weickert, B.M. ter Haar Romeny, A. Lopez, W.J. van Enk, *Orientation analysis by coherence-enhancing diffusion*, Proc. Symp. Real World Computing (RWC '97, Tokyo, Jan. 29–31, 1997), 96–103, 1997.

[46] J. Weickert, B.M. ter Haar Romeny, M.A. Viergever, *Efficient and reliable schemes for nonlinear diffusion filtering*, IEEE Trans. Image Proc., Vol. 7, 398–410, 1998.