

Consistent Positive Directional Splitting of Anisotropic Diffusion

Pavel Mrázek and Mirko Navara*

Center for Machine Perception
Faculty of Electrical Engineering, Czech Technical University
Technická 2, 166 27 Praha 6, Czech Republic
www: <http://cmp.felk.cvut.cz>
e-mail: mrizekp@cmp.felk.cvut.cz, navara@vision.felk.cvut.cz

Abstract. Anisotropic nonlinear diffusion filters describe a large class of powerful image filtering operations. In order to be applicable to sampled data, the continuous diffusion equations have to be discretized, e.g. using additive operator splitting scheme which approximates the continuous equation by one-dimensional processes along several chosen directions. In this paper we concentrate on one step of the separation of the continuous diffusion into one-dimensional processes: we analyse, design and test consistent positive directional splitting on a 3×3 window. We derive a set of formulas for directional diffusivities with one free parameter, and demonstrate that the best results are obtained if the parameter is chosen from the interior of the admissible interval.

1 Introduction

Nonlinear diffusion has deservedly attracted much attention in the field of image processing for its ability to reduce noise while preserving (or even enhancing) important features of the image, such as edges or discontinuities; this can be opposed to linear diffusion (alias Gaussian filtering or linear scale-space representation [2]) which not only removes noise but also blurs and dislocates edges. A good introduction to NL diffusion can be found e.g. in [4,8] or [9].

The isotropic nonlinear diffusion was first introduced by Perona and Malik in [3], and put on solid mathematical grounds by Catté *et al.* [1]. Their filter has the form

$$\frac{\partial \mathbf{u}}{\partial t} = \operatorname{div}(g(|\nabla \mathbf{u}_\sigma|) \cdot \nabla \mathbf{u}) \quad (1)$$

where \mathbf{u} is the function (e.g. image data) which develops with time t . The diffusion process is controlled by diffusivity g , a function of the magnitude of the

* Pavel Mrázek has been supported by the Czech Ministry of Education under Project LN00B096; Mirko Navara acknowledges the support of the Czech Ministry of Education under Research Programme J04/98:212300013.

estimated image gradient, $\nabla \mathbf{u}_\sigma = \nabla(G_\sigma * \mathbf{u})$; G_σ can be any smoothing kernel (the Gaussian represents a classical example), and the symbol $*$ denotes convolution.

Isotropic nonlinear diffusion with a scalar diffusivity g is stopped near the object boundary; it preserves the important edges/discontinuities in the data, but also leaves the noise near such positions unfiltered. To mitigate this undesirable effect, Weickert [7,9,5] proposes to make the amount of diffusion dependent not only on the position in the image, but vary it also between various directions at a single location. The process can be designed so that the smoothing perpendicular to the image gradient, i.e. along coherent structures (such as edges or lines) is preferred to smoothing across edges. To obtain this behaviour, the flux cannot be parallel to the image gradient (as is the case of the ‘classical’, isotropic NL diffusion equation), the diffusivity controlling the process is a matrix, \mathbf{D} . If \mathbf{D} depends on the gradient of the evolving image itself, we obtain the equation of *anisotropic nonlinear diffusion*

$$\frac{\partial \mathbf{u}}{\partial t} = \operatorname{div}(\mathbf{D}(\nabla \mathbf{u}_\sigma) \cdot \nabla \mathbf{u}) \quad (2)$$

with the initial and boundary conditions, respectively,

$$\mathbf{u}(x, 0) = \mathbf{f}(x), \quad \langle \mathbf{D}(\nabla \mathbf{u}_\sigma) \cdot \nabla \mathbf{u}, n \rangle = 0 \quad \text{on } \partial\Omega \quad (3)$$

where n denotes the normal to the image boundary $\partial\Omega$. The diffusion starts from the input image \mathbf{f} at $t = 0$, and the boundary condition expresses the fact that no flux should pass through the image boundary.

The *diffusion tensor* $\mathbf{D}(\nabla \mathbf{u}_\sigma)$ and the way it is constructed have crucial influence on the properties of the resulting method, leading e.g. to edge-enhancing or coherence-enhancing procedures [9]. Whereas the latter can be used to close interrupted line-like structures, the former one, the edge-enhancing diffusion and similar methods derived from it are of importance for the task of additive noise filtering.

In this paper we concentrate on a later stage of the design of a diffusion filter: how to transform the continuous equation (2) to make it applicable to discrete image data. More precisely, we want to separate the continuous diffusion tensor into several diffusivities acting along one-dimensional directions, and we want the splitting to be consistent with the continuous formulation, positive, and rotationally symmetric. Before going into details, we review in section 2 some basics of additive operator splitting, an efficient discretization scheme solving equation (2), and show why the separation of \mathbf{D} into one-dimensional diffusivities is needed. Section 3 presents the splitting results of Joachim Weickert; then we analyse the possibilities of consistent positive splitting on a 3×3 window in section 4, derive the formulas and suggest that the anisotropic diffusion algorithm exhibits the best properties if the splitting parameter is chosen from the interior of the admissible interval.

2 AOS scheme for nonlinear diffusion

The simplest way to a numerical solution of equation (1) is to discretize it using finite differences, and summarize the equations for individual pixels using matrix notation into

$$\mathbf{u}^{k+1} = (\mathbf{I} + \tau \mathbf{A}(\mathbf{u}^k)) \mathbf{u}^k. \quad (4)$$

(see [6]). Here τ is the iteration time step, \mathbf{u}^k is the vector of image pixels at time instant $k \cdot \tau$, and the matrix $\mathbf{A}(\mathbf{u}^k)$ contains the diffusivity information for the connections between neighbouring pixels. The matrix \mathbf{A} is composed of matrices \mathbf{A}_l which store the diffusivities for one direction of pixel connections, $\mathbf{A} = \sum_l \mathbf{A}_l$; each of the matrices \mathbf{A}_l can be transformed into a tridiagonal form by some permutation of image pixels.

The advantage of the explicit scheme is that only very simple operations between neighbouring pixels are performed in every step. Its drawback lies in the fact that the time step τ has to be small in order for the equation (4) to be stable, more iterations are needed to reach a fixed stopping time T and the computational demands might prevent the method from being used in practical situations. In [6], Weickert *et al.* suggest to replace the explicit method by the *additive operator splitting (AOS)* scheme which separates and discretizes the diffusion equation (2) by

$$\mathbf{u}^{k+1} = \frac{1}{m} \sum_{l=1}^m (\mathbf{I} - m\tau \mathbf{A}_l(\mathbf{u}^k))^{-1} \mathbf{u}^k \quad (5)$$

with l the direction index, $l = 1, \dots, m$ (for isotropic diffusion, m is set equal to the dimensionality of the data).

Each of the summands in (5) represents a one-dimensional diffusion process along the direction l . The compound diffusion iteration is obtained as an average of these one-dimensional processes, regardless of the dimensionality of the input data.

Both the AOS and the explicit schemes are of the same approximation order (first order in τ , second order in the grid size Δx) to the continuous isotropic diffusion equation (1); this can be easily checked if you compare the equations (4) and (5), and use the equality $(\mathbf{I} - \alpha \mathbf{A})^{-1} = \mathbf{I} + \alpha \mathbf{A} + (\alpha \mathbf{A})^2 + \dots$. In this sense, the two discretizations are equivalent.

The matrices \mathbf{A}_l which store the diffusivity information for the diffusion direction l are formed in the same way as with the explicit algorithm. The matrices $\mathbf{I} - m\tau \mathbf{A}_l(\mathbf{u}^k)$ can be made tridiagonal and diagonally dominant by a simple rearrangement of the pixels, and then inverted efficiently by the Thomas algorithm [6]. This way, one iteration of the AOS scheme requires only about twice the computational effort needed for one iteration of the explicit scheme. What we gain for this price is absolute stability: the AOS scheme is stable and creates a discrete scale space for any choice of the discretization step τ . The freedom to select a larger τ means that fewer iterations are needed to reach a fixed stopping time T , and the algorithm becomes faster. Although a large τ

also weakens the filtering effect and the solution may become less precise an approximation to the ideal continuous solution, Weickert *et al.* [6] report that for typical precision requirements of 2%¹, the AOS scheme is at least 11 times faster than any stable explicit scheme.

The AOS scheme can be extended to any number of dimensions. Also, as each direction and each line in that direction can be processed independently from other lines/directions, a parallel implementation is straightforward [11].

Moving to the anisotropic diffusion of equation (2) for which the diffusivity need not be equal in all directions, we again want to approximate the continuous process by a discrete algorithm. Again, the AOS scheme will separate the 2D diffusion into several one-dimensional diffusion processes along chosen directions. However, the anisotropic diffusion will generally need more directions than the isotropic filter did. There exists a direct relation between the number of one-dimensional processes and the achievable anisotropy of the compound diffusion filter. For sake of simplicity and computational efficiency, we restrict the approximation to four directions defined by the boundary pixels of a 3×3 window. In this situation, the AOS discretization is computed according to

$$\mathbf{u}_l^{k+1} = \frac{1}{4} \sum_{l=-1}^2 (\mathbf{I} - 4\tau \mathbf{A}_l(\mathbf{u}^k))^{-1} \mathbf{u}^k. \quad (6)$$

The main difficulty is to split the 2D diffusion tensor \mathbf{D} correctly into one-dimensional diffusivities to fill the matrices \mathbf{A}_l . In order to obtain good properties of the resulting filter, such as maximum–minimum principle, stability and well-posedness, smoothing properties etc. (see e.g. [9]), we require that each directional diffusivity is nonnegative. The design of a nonnegative directional splitting of the diffusion equation forms the main topic of this paper.

3 Weickert’s positive directional splitting

In [9], Weickert gives a constructive proof of the following theorem: given \mathbf{D} , a symmetric positive definite matrix with a spectral condition number κ , there exists some $n = \nu(\kappa) \in \mathbb{N}$ such that $\text{div}(\mathbf{D} \cdot \nabla \mathbf{u})$ reveals a second-order nonnegative forward difference discretization on a $(2n + 1) \times (2n + 1)$ window.

The boundary pixels of a $(2n + 1) \times (2n + 1)$ window define $4n$ principal orientations $\beta_i \in (-\frac{\pi}{2}, \frac{\pi}{2}]$, $i = -2n + 1, \dots, 2n$. The theorem says that it is possible to separate the continuous process into $4n$ one-dimensional processes along these orientations. Moreover, Weickert showed that only three of these orientations are actually needed to guarantee the positive discretization at any single location, so that we end up with the approximation

$$\text{div}(\mathbf{D} \cdot \nabla \mathbf{u}) = \partial_{e_{\beta_0}}(\alpha_0 \partial_{e_{\beta_0}} \mathbf{u}) + \partial_{e_{\beta_k}}(\alpha_k \partial_{e_{\beta_k}} \mathbf{u}) + \partial_{e_{\beta_{2n}}}(\alpha_{2n} \partial_{e_{\beta_{2n}}} \mathbf{u}) \quad (7)$$

¹ The precision was computed with respect to an explicit scheme with a very small time step $\tau = 0.1$ which was proclaimed the ground truth.

where $e_{\beta_i} = (\cos \beta_i, \sin \beta_i)^T$, and $\alpha_0, \alpha_k, \alpha_{2n}$ are the nonnegative directional diffusivities along the orientations $\beta_0, \beta_k, \beta_{2n}$.

The latter result says that, given a spectral condition number κ , it suffices to use three one-dimensional processes. The first acts in the horizontal direction, the second in the vertical direction. The orientation of the third is chosen from the set of $4n = 4 \cdot \nu(\kappa)$ angles according to the direction of the eigenvectors of the diffusion tensor. A literal application of this principle would require the use of neighbourhoods of different sizes dependent on the spectral condition number of the diffusion tensor. This would cause problems in implementation, e.g. at the borders of the image where some data are not available. Also the discretization of the diffusion in (almost) arbitrary orientations would be difficult.

We usually proceed differently. We fix the size n of neighbourhoods that are used in the algorithm. Thus we restrict the selection of orientations, too, but the discretized diffusion is implemented easily. As a consequence, we can perform diffusion only with a limited spectral condition number of the diffusion tensor; its bound κ_{\max} is equal to the maximal κ such that $\nu(\kappa) \leq n$. If the spectral condition number of the diffusion tensor exceeds this value, our diffusion discretization cannot be both consistent and positive. Nevertheless, this case can be avoided by the choice of the diffusion tensor. Notice that in isotropic diffusion we had the spectral condition number equal to 1, so imposing a limit on anisotropy need not represent a severe restriction.

For the particular case of 2D data and a 3×3 window (implying four principal orientations² $\beta_{-1} = -\frac{\pi}{4}$, $\beta_0 = 0$, $\beta_1 = \frac{\pi}{4}$, $\beta_2 = \frac{\pi}{2}$) with the diffusion tensor $\mathbf{D} = \begin{pmatrix} a & b \\ b & c \end{pmatrix}$, its condition number bounded by $\kappa_{\max} = 3 + 2\sqrt{2}$, Weickert proposes the following directional diffusivities:

$$\alpha_{-1} = \frac{|b| - b}{2} \qquad \alpha_0 = a - |b| \qquad (8)$$

$$\alpha_1 = \frac{|b| + b}{2} \qquad \alpha_2 = c - |b|. \qquad (9)$$

The directional diffusivities are assembled into matrices \mathbf{A}_l , and the AOS discretization of the 2D continuous anisotropic diffusion process is computed using the equation (6). This directional splitting leads to an algorithm which reveals (discrete versions of) all theoretical properties of the continuous diffusion filter (see [9]). As for practical properties, some problems arise with the rotational symmetry of the discrete filter; we return to this point in the next section.

In a private communication, Joachim Weickert suggested that any diffusivity splitting method can only possess two of these three desirable properties:

1. positivity (max-min principle);
2. anisotropy (strong directionality without limits on the condition number of \mathbf{D});
3. consistency with the continuous equation (the splitting should converge to the continuous equation as the discretization steps approach zero, $\tau \rightarrow 0$, $\Delta x \rightarrow 0$).

² We assume equal grid size $\Delta x = 1$ in both dimensions.

Weickert and Scharr renounced the positivity to obtain rotational symmetry and good directionality (with less blurring in other than coherence directions) in [10]. The disadvantage of that choice is that the maximum-minimum principle is lost, some oscillations may appear, and the method is not suitable for noise filtering.

4 Consistent positive splitting on a 3×3 window

In this section we study in more detail the consistent directional splitting on a 3×3 window. Let us start from the diffusion tensor $\mathbf{D} = \begin{pmatrix} a & b \\ b & c \end{pmatrix}$. We want to approximate the (continuous) 2D diffusion by a diffusion composed of 1D processes acting along four directions $\beta_{-1} = -\frac{\pi}{4}$, $\beta_0 = 0$, $\beta_1 = \frac{\pi}{4}$, $\beta_2 = \frac{\pi}{2}$; we will call the approximation *consistent* if

$$\sum_{k=-1}^2 \partial_{e_{\beta_k}} (\alpha_k \partial_{e_{\beta_k}} \mathbf{u}) = \operatorname{div}(\mathbf{D} \cdot \nabla \mathbf{u}) \quad (10)$$

where α_k is the diffusivity along the direction β_k .

Let us expand the left hand side of equation (10). With $e_{\beta_k} = (\cos \beta_k, \sin \beta_k)^T$ and using the notation $u_x = \frac{\partial \mathbf{u}}{\partial x}$ (similarly for u_y , u_{xx} , etc.), we have

$$\partial_{e_{\beta_k}} = \langle e_{\beta_k}, \nabla \mathbf{u} \rangle = u_x \cos \beta_k + u_y \sin \beta_k \quad (11)$$

from which

$$\partial_{e_{\beta_k}} (\alpha_k \cdot \partial_{e_{\beta_k}} \mathbf{u}) = \alpha_k \cdot (u_{xx} \cos^2 \beta_k + 2u_{xy} \sin \beta_k \cos \beta_k + u_{yy} \sin^2 \beta_k). \quad (12)$$

The right hand side of (10) yields

$$\operatorname{div}(\mathbf{D} \cdot \nabla \mathbf{u}) = \left\langle \begin{pmatrix} \partial_x \\ \partial_y \end{pmatrix}, \begin{pmatrix} a u_x + b u_y \\ b u_x + c u_y \end{pmatrix} \right\rangle = a u_{xx} + 2b u_{xy} + c u_{yy}. \quad (13)$$

If we evaluate the trigonometric functions for the angles β_k in (12), sum up the contributions for all k , and assemble the elements corresponding to a given partial derivative of \mathbf{u} , we obtain the following set of linear equations:

$$\text{for } u_{xx} : \quad \frac{1}{2}\alpha_{-1} + \alpha_0 + \frac{1}{2}\alpha_1 = a \quad (14)$$

$$\text{for } u_{xy} : \quad -\alpha_{-1} + \alpha_1 = 2b \quad (15)$$

$$\text{for } u_{yy} : \quad \frac{1}{2}\alpha_{-1} + \frac{1}{2}\alpha_1 + \alpha_2 = c. \quad (16)$$

We want to solve this set of equations for the unknowns α_k with the restrictions $\alpha_k \geq 0$.

Let us take the second equation of the system, (15), and add another equation to it, formed for a parameter $p = \frac{\alpha_{-1} + \alpha_1}{2}$:

$$-\alpha_{-1} + \alpha_1 = 2b \quad (17)$$

$$\alpha_{-1} + \alpha_1 = 2p. \quad (18)$$

Summing and subtracting of the two equations lead to

$$\alpha_1 = p + b \geq 0 \quad (19)$$

$$\alpha_{-1} = p - b \geq 0 \quad (20)$$

from which we obtain the first solvability condition, $p \geq |b|$.

Using the parameter p , the solutions of (14)–(16) may be expressed as

$$\alpha_0 = a - p \quad \alpha_2 = c - p \quad (21)$$

$$\alpha_{-1} = p - b \quad \alpha_1 = p + b. \quad (22)$$

As we require $\alpha_k \geq 0$, the equation (21) provides directly the upper bound on our splitting parameter: $p \leq \min(a, c)$. To summarize, the consistent splitting (21)–(22) remains positive if and only if we select

$$p \in [|b|, \min(a, c)]. \quad (23)$$

The conditions ensuring that this interval is nonempty can be expressed in terms of the condition number of the diffusion tensor \mathbf{D} . This has been done by Weickert [9, pp. 88–95] and we have seen it in the previous section: the consistent positive splitting on a 3×3 window is possible if and only if the condition number of \mathbf{D} is less than or equal to $\kappa_{\max} = 3 + 2\sqrt{2}$.

Let us now return to an important point: which value from the admissible interval should be chosen for the parameter p ? In the following we offer three possibilities:

Splitting 1: $p = |b|$

Splitting 2: $p = \min(a, c)$

Splitting 3: $p = \frac{|b| + \min(a, c)}{2}$.

The first two suggestions take on the value of either of the limit cases; splitting 1 is clearly equivalent to Weickert’s positive splitting (8)–(9). The third splitting represents a compromise, an average of the two limits of the admissible interval (23).

The three splitting alternatives are tested for rotational symmetry in Fig. 2, showing a diffused Gaussian hill. You can observe that although all the three methods share the same property of consistency with the continuous formulation, splitting 1 and 2 does not transfer the continuous rotational symmetry well into the discrete situation. The artefacts become severe as the iteration time step τ increases. On the other hand, the compromising splitting 3 restricts the artefacts considerably, both for a large $\tau = 10$ and for a small $\tau = 1$ (the latter situation is stressed in figure 3 depicting the difference between the diffusion result and its copy rotated by 45 degrees). Figure 3d illustrates that if the diffusivity parameter λ is (too) small, the interval (23) becomes narrow on a large part of the image, and the three splitting alternatives perform equivalently in preferring some diffusion directions to others, which may create some irregular patterns in the data.

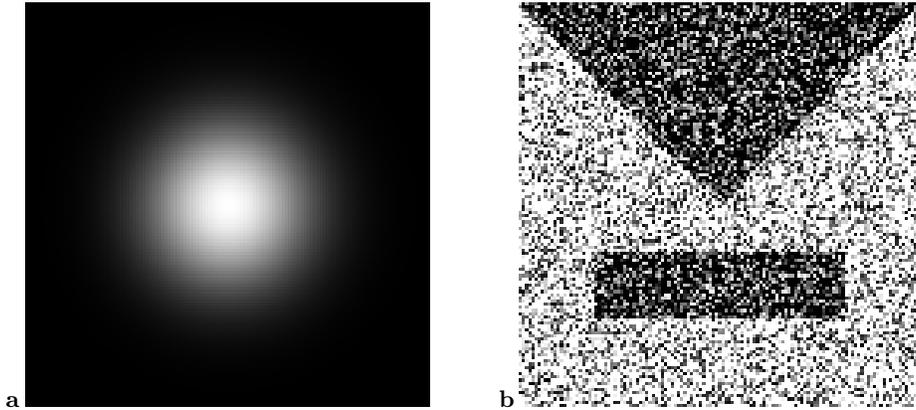


Fig. 1. Input images. Left: Gaussian hill, the input data to test the rotational symmetry of the methods. Right: noisy input image ($[0, 127]^2 \rightarrow [0, 255]$) for the ‘Triangle and rectangle’ experiment. Noise with uniform distribution in the range $[-255, 255]$ was added to two-valued synthetic data; the signal-to-noise ratio is 1.4.

The noise-filtering capabilities of splitting 1 and splitting 3, tested on the data from Fig. 1b, are compared in Fig. 4. You can see that some cloth-like artefacts remain in the filtered image with fewer iterations of splitting 1, and the output data for $\tau = 1$ and $\tau = 5$ differ significantly. In contrast with that, all the results of splitting 3 resemble, although – as usual for the AOS scheme – the filtering effect weakens as the discretization time step increases. These results prove that the differences between the splitting methods might be negligible if a small time step τ is employed. However, if you ask for a more efficient algorithm and wish to spend fewer iterations of the filtering procedure, thus needing a larger τ , splitting 3 becomes clearly superior.

5 Conclusion

We have analysed the possibilities for consistent positive directional splitting of anisotropic diffusion on a 3×3 window. We have shown that such a splitting exists if the interval $[b, \min(a, c)]$, formed from the elements of the diffusion tensor $\mathbf{D} = \begin{pmatrix} a & b \\ b & c \end{pmatrix}$, is nonempty. Moreover, we have derived the formulas for the directional diffusivities depending on a single diffusivity parameter, and demonstrated experimentally that the directional splitting reveals better properties (regarding e.g. rotational symmetry and sensitivity to the time step size) if the splitting parameter is chosen from the interior of the admissible interval.

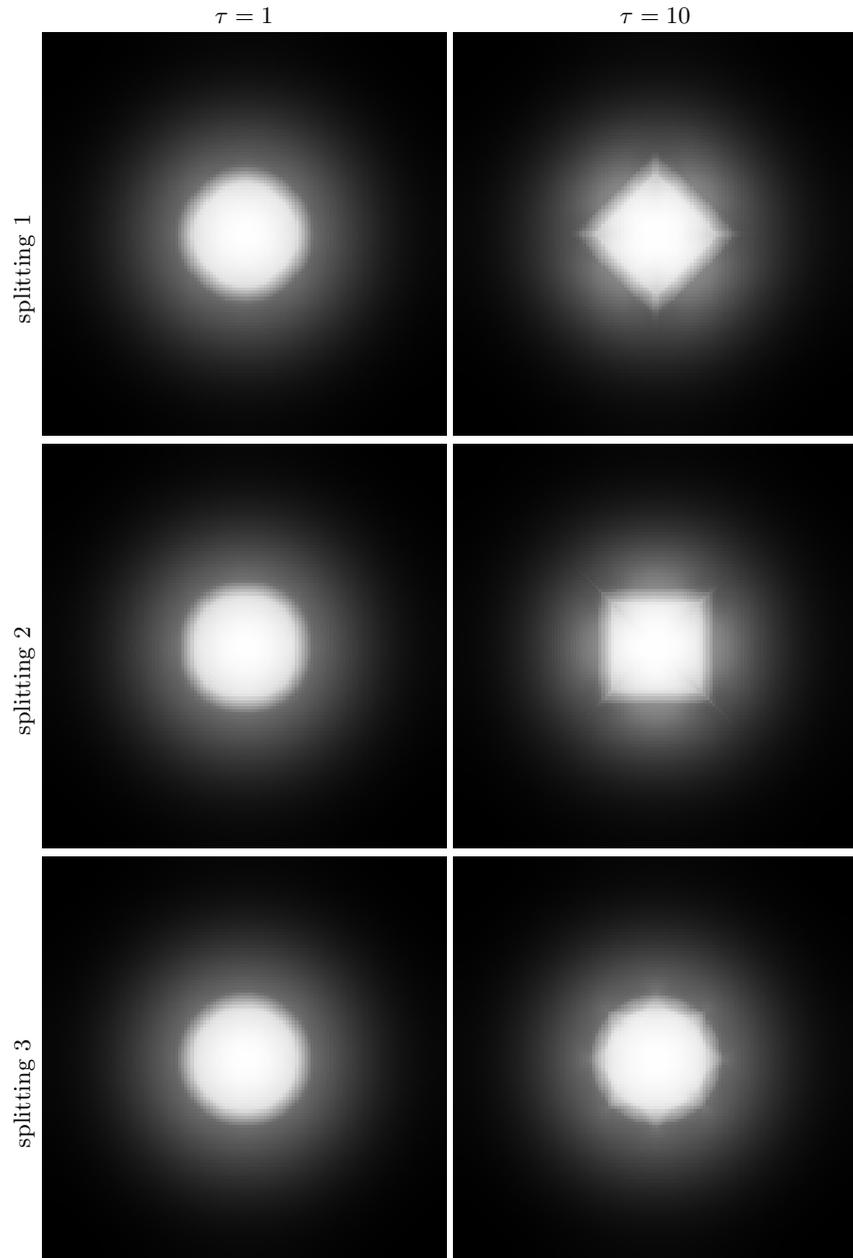


Fig. 2. A rotationally symmetric Gaussian hill from Fig. 1a filtered by the three directional splitting methods for the anisotropic NL diffusion AOS scheme. Left column contains the results for the discretization time step $\tau = 1$, the right for $\tau = 10$. The remaining diffusion parameters were given by $\sigma = 1$, $T = 200$, $\varphi_2 = 1$; the parameter λ was computed in each diffusion step as a 0.95-percentile of the image gradients.

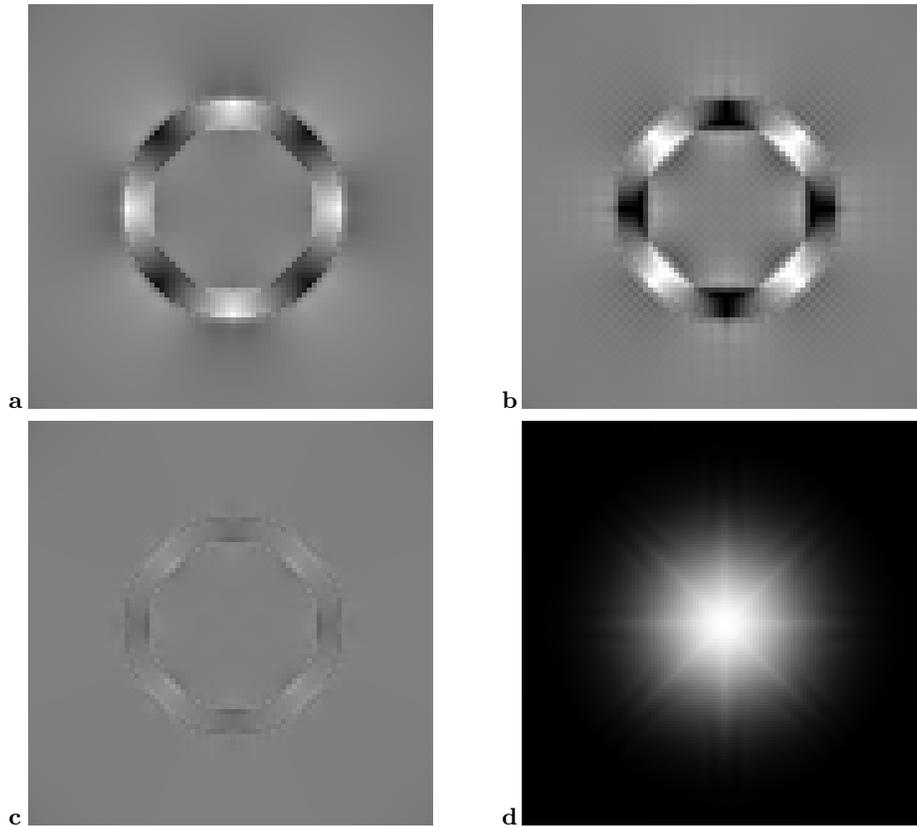


Fig. 3. The artefacts on rotational symmetry.

(a–c) the (amplified) difference between a diffusion result from Fig. 2, $\tau = 1$, and its copy rotated by 45 degrees: (a) splitting 1, (b) splitting 2, (c) splitting 3.

(d) Star-like patterns appear on a Gaussian diffused with a small diffusivity parameter λ (for any splitting method 1–3). Here $T = 200$, $\tau = 10$, the Perona-Malik λ was computed as a 0.1-percentile of the image gradients.

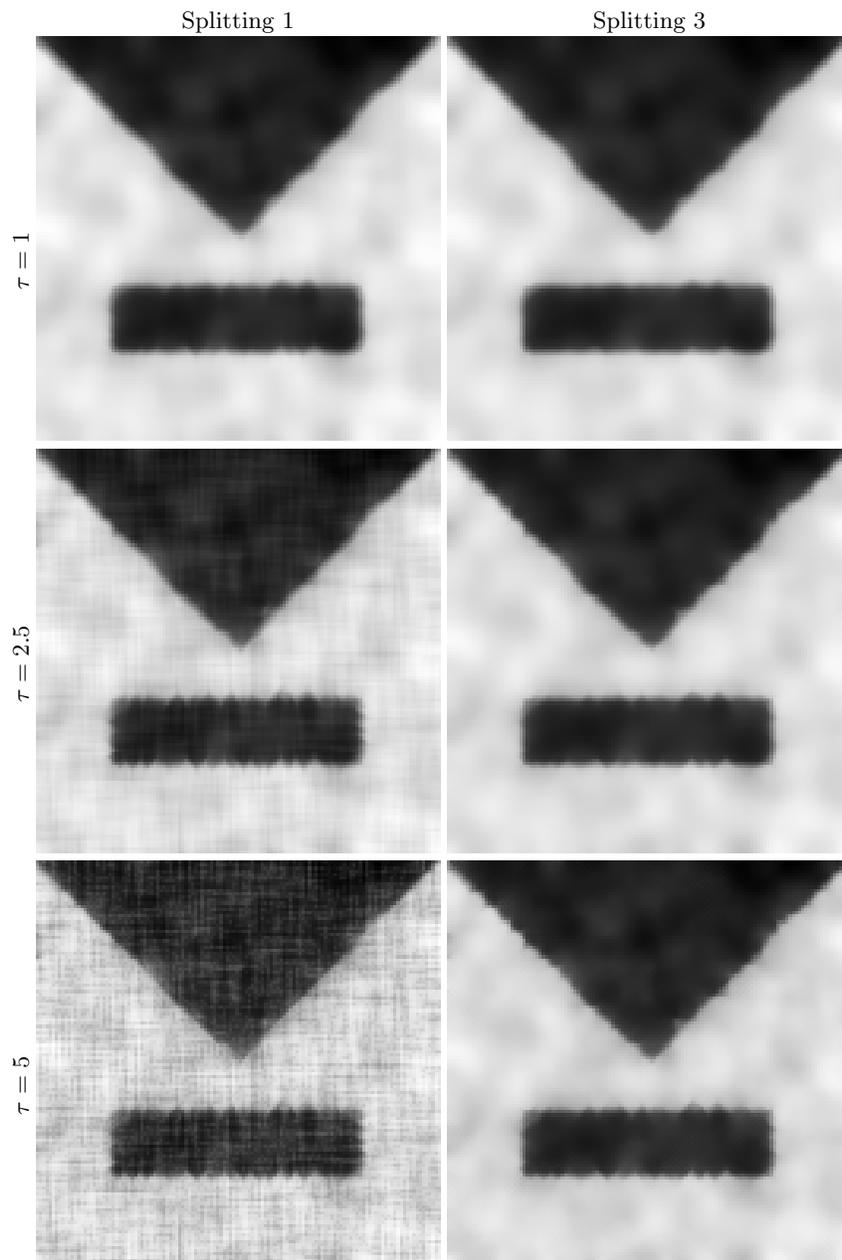


Fig. 4. Impact of the directional splitting on the results of the noise-filtering procedure: splitting 1 on the left, splitting 3 on the right. In all cases, the stopping time of the diffusion was $T = 25$, the time step τ increases from top to bottom, $\tau \in \{1, 2.5, 5\}$.

Acknowledgements

The definition of consistent directional splitting originates from an informal discussion with Joachim Weickert at the conference Algoritmy 2000 in Podbanské, Slovakia.

References

1. F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll. Image selective smoothing and edge-detection by nonlinear diffusion. *SIAM Journal on Numerical Analysis*, 29(1):182–193, 1992.
2. Tony Lindeberg. *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, 1994.
3. P. Perona and J. Malik. Scale-space and edge-detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
4. Bart M. ter Haar Romeny, editor. *Geometry-Driven Diffusion in Computer Vision*. Kluwer Academic Publishers, 1994.
5. J Weickert. Coherence-enhancing diffusion filtering. *International Journal of Computer Vision*, 31(2-3):111–127, 1999.
6. J. Weickert, B.M. ter Haar Romeny, and M. A. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Transactions on Image Processing*, 7:398–410, 1998.
7. Joachim Weickert. Theoretical foundations of anisotropic diffusion in image processing. *Computing, Suppl. 11*, pages 221–236, 1996.
8. Joachim Weickert. A review of nonlinear diffusion filtering. In B. M. ter Haar Romeny, L. Florack, J. Koenderink, and M. Viergever, editors, *Scale-Space Theory in Computer Vision*, volume 1252 of *Lecture Notes in Computer Science*, pages 3–28. Springer, 1997. Invited paper.
9. Joachim Weickert. *Anisotropic Diffusion in Image Processing*. European Consortium for Mathematics in Industry. B. G. Teubner, Stuttgart, 1998.
10. Joachim Weickert and Hanno Schar. A scheme for coherence-enhancing diffusion filtering with optimized rotation invariance. Technical Report 4/2000, Computer Vision, Graphics, and Pattern Recognition Group, Department of Mathematics and Computer Science, University of Mannheim, Germany, February 2000.
11. Joachim Weickert, K.J. Zuiderveld, B.M. ter Haar Romeny, and W.J. Niessen. Parallel implementations of AOS schemes: A fast way of nonlinear diffusion filtering. In *Proc. 1997 IEEE International Conference on Image Processing*, volume 3, pages 396–399, 1997.