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, pp. 260–271, 1997.

Recursive Separable Schemes for Nonlinear Diffusion Filters

Joachim Weickert*

Image Sciences Institute,
Utrecht University Hospital,
E.01.334, Heidelberglaan 100,
3584 CX Utrecht, The Netherlands.
E-mail: Joachim.Weickert@cv.ruu.nl

Abstract. Poor efficiency is a typical problem of nonlinear diffusion filtering, when the simple and popular explicit (Euler-forward) scheme is used: for stability reasons very small time step sizes are necessary. In order to overcome this shortcoming, a novel type of semi-implicit schemes is studied, so-called additive operator splitting (AOS) methods. They share the advantages of explicit and (semi-)implicit schemes by combining simplicity with absolute stability. They are reliable, since they satisfy recently established criteria for discrete nonlinear diffusion scale-spaces. Their efficiency is due to the fact that they can be separated into one-dimensional processes, for which a fast recursive algorithm with linear complexity is available. AOS schemes reveal good rotational invariance and they are symmetric with respect to all axes. Examples demonstrate that, under typical accuracy requirements, they are at least ten times more efficient than explicit schemes.

1 Introduction

Although nonlinear diffusion filters are quite popular, their practical applicability suffers from the fact that the widely-used explicit (Euler-forward) discretization is only stable for very small time steps.
This problem is addressed in the present paper by proposing schemes which are separable and which do not suffer from any time step size restriction since all stability-relevant terms are discretized in an implicit manner. The backbone of these schemes is the Thomas algorithm for solving a tridiagonal system of linear equations. It is fast, stable and requires only a few lines programming work. Its forward and backward substitution steps constitutes a recursive scheme with a causal and an anticausal filter. The presented schemes can be implemented in arbitrary dimensions, they create discrete nonlinear

* Current address: DIKU, Dept. of Computer Science, University of Copenhagen, Universitetsparken 1, 2100 Copenhagen East, Denmark.

diffusion scale-spaces, and their computational and memory requirements are linear in the image size.
As a prototype of a well-founded nonlinear diffusion filter[1] we focus on AOS schemes for a spatial regularization of the Perona–Malik filter [17] due to Catté, Lions, Morel and Coll [4].
The present paper is organized as follows:
Section 2 gives a brief survey on the nonlinear diffusion model (which will shall call Catté equation henceforth). In Section 3 we first investigate the limitations of explicit and semi-implicit schemes as a motivation to study AOS schemes. Then the structure of AOS is explained, and it is sketched why they create a discrete nonlinear diffusion scale-space. We shall observe that AOS schemes can be reduced to one-dimensional discrete diffusion processes leading to tridiagonal systems of linear equations. They are solved in a recursive way by the Thomas algorithm, a special version of the Gaussian elimination scheme. Algorithmic features and complexity estimates will be presented. In Section 4 we evaluate the results by checking the performance of AOS schemes with respect to rotational invariance and accuracy. This enables us to propose suitable time step sizes and to analyse the efficiency gain in comparison to the widely-used explicit scheme. We conclude the paper with a summary in Section 5.
Due to space limitations only the main ideas can be presented. Proofs and full details can be found in [21].

Related work. This work has been influenced by a number of related approaches which shall be mentioned here.
Implicit splitting-based approaches have been proposed for linear diffusion filtering in [8,3] and also in [2] where their realization as a recursive filters is suggested. Impressive results on improved efficiency by means of recursive filtering may be found in [5], and the close relation between recursive filters and linear scale-space has been clarified in [15]. Semidiscrete analogues of the linear diffusion scale-space can be found in [13].
In the nonlinear diffusion field one can find several alternatives to the conventional two-level explicit finite-difference scheme, for instance three-level methods [6], semi-implicit approaches [4], multigrid methods [1], finite element techniques [12], numerical schemes with wavelets as trial functions [6], and spectral methods [6]. Also hardware proposals for nonlinear diffusion filtering have been made, see e.g. [7].
Semi-implicit algorithms have also been proposed for other nonlinear PDEs in image processing, for instance for mean curvature motion [2].

2 The Continuous Filtering Process

In the m-dimensional case the filter of Catté, Lions, Morel and Coll [4] has the following structure:

[1] Overviews of other methods can be found in [9,20].

Let $\Omega := (0, a_1) \times .... \times (0, a_m)$ be our image domain and consider a (scalar) image $f(x)$ as a bounded mapping from Ω into the real numbers $\mathbb{R}$. Then a filtered image $u(x,t)$ of $f(x)$ is calculated by solving the diffusion equation with the original image as initial state, and reflecting boundary conditions:

$$\partial_t u = \sum_{l=1}^{m} \partial_{x_l} \Big(g(|\nabla u_\sigma|^2) \, \partial_{x_l} u \Big) \tag{1}$$

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

$$\partial_n u|_{\partial\Omega} = 0, \tag{3}$$

where n denotes the normal to the image boundary $\partial\Omega$.
The "time" t is a scale parameter: larger values lead to simpler image representations. The whole embedding of the original image into a one-parameter family of simplified images is called *scale-space*, a concept which has been introduced to image processing by Taizo Iijima more than 30 years ago [11,22]. In order to reduce smoothing at edges, the diffusivity g is chosen as a decreasing function of the edge detector $|\nabla u_\sigma|$, where ∇u_σ is the gradient of a Gaussian-smoothed version of u:

$$\nabla u_\sigma := \nabla (K_\sigma * u), \tag{4}$$

$$K_\sigma := \frac{1}{(2\pi\sigma^2)^{m/2}} \exp\left(-\frac{|x|^2}{2\sigma^2} \right). \tag{5}$$

We use the following form for the diffusivity:

$$g(s) := \begin{cases} 1 & (s \leq 0) \\ 1 - \exp\left(\frac{-3.315}{(s/\lambda)^4} \right) & (s > 0). \end{cases} \tag{6}$$

For such rapidly decreasing diffusivities smoothing on both sides of an edge is much stronger than smoothing across it. As a result, the gradient at edges may even be enhanced, see [17] for more details. λ plays the role of a contrast parameter: Structures with $|\nabla u_\sigma| > \lambda$ are regarded as edges, where the diffusivity is close to 0, while structures with $|\nabla u_\sigma| < \lambda$ are considered to belong to the interior of a region. Here the diffusivity is close to 1. Thus, we have a selective smoothing process which prefers intraregional smoothing to interregional blurring. After some time it leads to segmentation-like results which are piecewise almost constant.
The parameter $\sigma > 0$ makes the filter insensitive to noise at scales smaller than σ. It is also a regularization parameter which guarantees well-posedness of the process: Catté *et al.* [4] have shown that their filter has a unique solution which is infinitely times differentiable for $t > 0$. In earlier work [18] the author has proved that this solution depends continuously on the original image. He also established a scale-space interpretation for the continuous Catté equation and its anisotropic generalizations. In addition to invariances such as the preservation of the average grey value, it has been shown that – it

spite of its contrast-enhancing potential – these equations create smoothing scale-spaces: the obey a maximum–minimum principle, have a large class of smoothing Lyapunov functionals, and converge to a constant steady state.

3 AOS Schemes

3.1 Limitations of Explicit and Semi-Implicit Schemes

Let us now consider finite difference approximations to the m-dimensional Catté equation.
A discrete image can be regarded as a vector $f \in \mathbb{R}^N$, whose components f_i, $i \in J := \{1, ..., N\}$ display the grey values at the pixels. Pixel i represents the location x_i, and h_l is the grid size in l direction. We consider discrete times $t_k := k\tau$, where $k \in \mathbb{N}_0$ and τ is the time step size. By u_i^k and g_i^k we denote approximations to $u(x_i, t_k)$ and $g(\nabla u_\sigma(x_i, t_k))$, respectively, where the gradient is replaced by central differences.
The simplest discretization of the Catté equation with reflecting boundary conditions is given by

$$\frac{u_i^{k+1} - u_i^k}{\tau} = \sum_{l=1}^{m} \sum_{j \in \mathcal{N}_l(i)} \frac{g_j^k + g_i^k}{2h_l^2} (u_j^k - u_i^k). \tag{7}$$

where $\mathcal{N}_l(i)$ consist of the two neighbours of pixel i along the direction l (boundary pixels may have only one neighbour). In vector–matrix notation this becomes

$$\frac{u^{k+1} - u^k}{\tau} = \sum_{l=1}^{m} A_l(u^k)\, u^k \tag{8}$$

where A_l describes the diffusive interaction in l direction. It is easily seen that A_l is a symmetric, irreducible matrix with zero row sums and nonnegative off-diagonals. Due to the neighbourhood structure, at most 2 off-diagonal elements per row are nonvanishing. Writing (8) as

$$u^{k+1} = \left(I + \tau \sum_{l=1}^{m} A_l(u^k)\right) u^k, \tag{9}$$

we observe that we can calculate u^{k+1} directly (explicitly) from u^k without any matrix inversions. For this reason it is called *explicit scheme*. Each explicit iteration step can be performed very fast. Unfortunately, each step has to be very small as well: one can show [21] that in order to guarantee stability, the step size has to satisfy

$$\tau \leq \left(\frac{2}{h_1^2} + ... + \frac{2}{h_m^2}\right)^{-1}. \tag{10}$$

For most practical applications, this restriction requires to use a very high number of iterations, such that the explicit scheme is rather slow.
Thus, we consider a slightly more complicated discretization next, namely

$$\frac{u^{k+1} - u^k}{\tau} \;=\; \sum_{l=1}^{m} A_l(u^k)\, u^{k+1}. \tag{11}$$

We observe that this scheme does not give the solution u^{k+1} directly (explicitly): It requires to solve a linear system first. For this reason it is called a *linear-implicit (semi-implicit)* scheme. The solution u^{k+1} is given by

$$u^{k+1} \;=\; \Big(I - \tau \sum_{l=1}^{m} A_l(u^k)\Big)^{-1} u^k. \tag{12}$$

This scheme can be shown to be unconditionally stable [21].
However, this does not necessarily imply that it is superior to the explicit one. For dimensions ≥ 2 there appears a problem: it is not possible to order the pixels in such a way that in the i-th row all nonvanishing elements of the system matrix can be found within the positions $[i, i-m]$ to $[i, i+m]$: Usually, the matrix reveals a much larger bandwidth. Applying direct algorithms such as Gaussian elimination would destroy the zeros within the band and would lead to an immense storage and computation effort. Typical iterative algorithms such as the Jacobi, Gauß–Seidel, SOR or preconditioned CG method reveal another limitation: their convergence becomes slow for large τ, since this increases the condition number of the system matrix. Thus, in spite of its absolute stability, the semi-implicit scheme is often not so much faster than the explicit one. To take full advantage of absolute stability, it is desirable to find an efficient scheme whose effort is independent of the time step size.

3.2 AOS Schemes

In order to address the abovementioned problem let us consider a modification of the semi-implicit scheme (12), namely the *additive operator splitting (AOS) scheme*

$$u^{k+1} \;=\; \frac{1}{m} \sum_{l=1}^{m} \Big(I - m\tau A_l(u^k)\Big)^{-1} u^k. \tag{13}$$

This scheme has several interesting properties:

- **Consistency.** It has the same first-order Taylor expansion in τ as the explicit scheme (9) and the semi-implicit scheme (12): all methods are $O(\tau + h_1^2 + ... + h_m^2)$ approximations to the continuous equation. From this viewpoint, all schemes are consistent to the original equation. One should not make the mistake to regard the AOS scheme as an algebraically

incorrect reformulation of the semi-implicit scheme: The explicit scheme is also different from the semi-implicit one, but it approximates the same continuous diffusion process.

- **Equal treatment of all axes.** Since it is an *additive* splitting, all coordinate axes are treated in exactly the same manner. This is in contrast to conventional splitting techniques from the literature [14], which are *multiplicative*. In the nonlinear case, the latter ones can produce different results if the image is rotated by 90 degrees.
- **Discrete scale-space properties.** One can verify [21] that, for every (!) step size τ,

$$Q(u^k) \;:=\; \frac{1}{m} \sum_{l=1}^{m} \Big(I - m\tau A_l(u^k) \Big)^{-1} \tag{14}$$

is continuous in its argument, symmetric, all row sums are 1, all entries are nonnegative, the diagonal elements are positive, and Q is irreducible. These are the six criteria such that a scheme of type $u^{k+1} = Q(u^k)u^k$ creates a discrete diffusion scale-space with the following properties [19]:

(a) *Average grey level invariance:*
The average grey level $\mu := \frac{1}{N} \sum_{j \in J} f_j$ is not affected by the discrete diffusion filter:

$$\frac{1}{N} \sum_{j \in J} u_j^k = \mu \qquad \forall\, k \in \mathbb{N}_0. \tag{15}$$

This invariance is required in scale-space based segmentation algorithms [16], and it is useful for applications where the grey value is linked to physical properties, for instance in medical imaging.

(b) *Extremum principle:*

$$\min_{j \in J} f_j \le u_i^k \le \max_{j \in J} f_j \qquad \forall\, i \in J, \quad \forall\, k \in \mathbb{N}_0. \tag{16}$$

This property is much more than a stability result which forbids under- and overshoots. It also ensures that iso-intensity linking towards the original image is possible, an important causality property [10].

(c) *Smoothing Lyapunov sequences:*
The process is a simplifying, information-reducing transform with respect to many aspects:
The p-norms $\|u^k\|_p := (\sum_{i=1}^N |u_i^k|^p)^{1/p}$ and all even central moments $M_{2n}[u^k] := \frac{1}{N} \sum_{j=1}^N (u_j^k - \mu)^{2n}$ are decreasing in k, and the entropy $S[u^k] := -\sum_{j=1}^N u_j^k \ln u_j^k$, a measure of uncertainty and missing information, is increasing in k (if f_j is positive for all j).

(d) *Convergence to a constant steady state:*

$$\lim_{k \to \infty} u_i^k = \mu \quad \forall\, i \in J. \tag{17}$$

The discrete scale-space evolution tends to the most global image representation that is possible: a constant image with the same average grey level as f.

- **Efficiency.** The operators $B(u^k) := I - m\tau A_l(u^k)$ describe one-dimensional diffusion processes along the x_l axes. Under a consecutive pixel numbering along the direction l they come down to strictly diagonally dominant tridiagonal matrices. Then, the most efficient way to calculate $B^{-1}(u^k)u^k$ is the so-called *Thomas algorithm*, a Gaussian elimination algorithm for tridiagonal systems. This algorithm requires only $5N-4$ multiplications or divisions, and $3N-3$ subtractions.
  The Thomas algorithm can be thought of as a recursive filter consisting of three steps: determination of the filter coefficients, a causal recursion and an anticausal recursion. If we want to solve $Bu = d$ where B has diagonal $(\alpha_1, ..., \alpha_N)$, upper diagonal $(\beta_1, ..., \beta_{N-1})$ and lower diagonal $(\gamma_1, ..., \gamma_{N-1})$, then its algorithmic formulation is as follows:

```
/* step 1: filter coefficients */
m_1 := α_1
for i = 1, 2, ..., N-1:
    l_i := γ_i / m_i
    m_{i+1} := α_{i+1} - l_i β_i

/* step 2: causal filter */
y_1 := d_1
for i = 2, 3, ..., N:
    y_i := d_i - l_{i-1} y_{i-1}

/* step 3: anticausal filter */
u_N := y_N / m_N
for i = N-1, N-2, ..., 1:
    u_i := (y_i - β_i u_{i+1}) / m_i
```

- **Usable for presmoothing.** We may also calculate the presmoothing $u_\sigma = K_\sigma * u$ by means of an AOS scheme. It is well-known that Gaussian convolution with standard deviation σ is equivalent to linear diffusion filtering ($g \equiv 1$) for some time $T = \sigma^2/2$.
- **Linear complexity.** In order to assess the complexity of AOS algorithms, let us consider dimensions $m \geq 2$ and focus on terms of order N (number of pixels). Then it can be shown [21] that one entire Catté iteration has a memory requirement of $4N \times 4$ Byte, if the calculations are performed in single precision. If we use a look-up table for the diffusivity g, then the total computational effort is $11mN$ multiplications or divisions, $(10m-1)N$ additions or subtractions, and N look-ups. This is less than twice the typical effort needed for an explicit scheme, a rather low price for gaining absolute stability.

4 Evaluation

Fig. 1. Nonlinear diffusion filtering of a Gaussian-like test image ($\lambda = 8$, $\sigma = 1.5$). (a) LEFT: Original image, $\Omega = (0, 101)^2$. (b) MIDDLE: Explicit scheme, 800 iterations, $\tau = 0.25$, CPU: 8.97 s. (c) RIGHT: AOS scheme, 40 iterations, $\tau = 5$, CPU: 0.84 s.

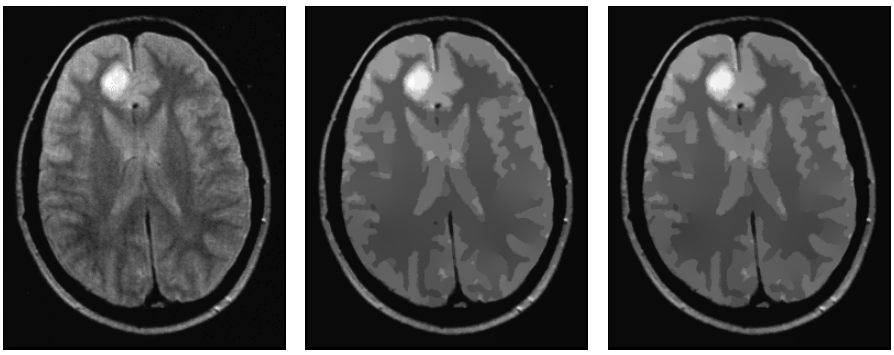

Fig. 2. Nonlinear diffusion filtering of a medical image. ($\lambda = 2$, $\sigma = 1$). (a) LEFT: Original image, $\Omega = (0, 255) \times (0, 308)$. (b) MIDDLE: Explicit scheme, 800 iterations, $\tau = 0.25$, CPU: 72.5 s. (c) RIGHT: AOS scheme, 40 iterations, $\tau = 5$, CPU: 6.75 s.

Let us now determine appropriate time step sizes for AOS schemes. For too large time steps they will still reveal average grey value invariance, stability based on an extremum principle, Lyapunov functionals, and convergence to constant steady state, but they will be less good approximations to the continuous equation. Thus, we should expect problems with those properties which a naturally linked to continuous ideas and which can only be satisfied approximately by discrete schemes: rotational invariance and accuracy.

Table 1. CPU times for one AOS iteration.

image size	CPU time	image size	CPU time
64^2	0.0086 s	16^3	0.0159 s
512^2	0.711 s	64^3	1.15 s
4096^2	145 s	256^3	237 s

Figure 1 is used as a test for rotational invariance. It depicts a Gaussian-like image and its filtered versions. We observe that both the explicit scheme with $\tau = 0.25$ and the AOS scheme with $\tau = 5$ give good results. Thus, even for 20 times larger step sizes the AOS scheme does not introduce artefacts with respect to rotational invariance. Since every AOS iteration is almost twice as expensive as an explicit one, we finally gain an increase in efficiency by a factor ten.

Next we consider Figure 2 which depicts the filtering of a brain image. The situation is similar as in Figure 1: Both the explicit scheme with $\tau = 0.25$ and the AOS scheme with $\tau = 5$ are perceptually satisfying and do not differ very much. For a more detailed study with other time step sizes the reader is referred to [21].

After these visual inspections we shall evaluate the accuracy and efficiency more quantitatively. To this end we measure both CPU time and approximation error for the explicit scheme (9), the semi-implicit scheme (12) and the AOS scheme (13). Since there is no analytical solution to the Catté equation known, we have to use a good numerical approximation to a test example as a standard for comparison. In our case we took the explicit scheme with the small step size $\tau = 0.1$ and applied it 2000 times to the test image from Fig. 2. If v denotes this reference solution, the relative l^2 error of an approximation u is given by

$$\frac{\|u - v\|_2}{\|v\|_2}. \tag{18}$$

The linear system of the 2-D semi-implicit scheme is solved iteratively by a Gauß–Seidel algorithm. Every second iteration the residue is calculated, and the process is stopped when the l^2 norm of the residue is diminished to 0.01 or 0.1 times its initial value, respectively.

Figure 3 shows the result of the comparison. The depicted curves are created by running each scheme with different time step sizes. Larger time steps allow to reach the stopping time $T = 200$ with less CPU effort, but give also rise to larger approximation errors.

We observe that for very high accuracy requirements the explicit scheme is most appropriate[2]. This is at the expense of a high overall computational

[2] One can achieve higher accuracies by methods which are of second order in time. Such high accuracies, however, are not often required in image processing.

effort. On the other hand, even relaxing the accuracy requirements to a relative l^2 error of 1 % does not permit to find a more efficient technique. For errors between 1 % and 1.7 % the semi-implicit scheme with residue accuracy $\alpha = 0.1$ is fastest, and for errors larger than 1.7 % the AOS scheme becomes rapidly superior. In our previous experiments we have observed that the accuracy of AOS with $\tau = 5$ appears to be tolerable for many applications. This corresponds to an l^2 error of about 2.2 %. In this case, AOS is almost 2.5 times more efficient than the semi-implicit scheme with $\alpha = 0.1$, more than 3.5 times faster than the semi-implicit scheme with $\alpha = 0.01$, and about 11 times more efficient than the explicit scheme. Although these relations have been illustrated by one example only, additional experiments have indicated that these basic relations between explicit, semi-implicit and AOS discretizations carry over to a large class of images: The accuracy requirements of many practical problems allow an efficiency gain by one order of magnitude. All one has to do is to replace the explicit scheme by an AOS scheme with 20 times larger time step sizes.
Table 1 shows the measured CPU times on a single R10000 processor of an SGI Challenge XL both for 2-D and 3-D images. Three-dimensional data sets from medicine with typical sizes such as $256 \times 256 \times 64$ can be processed in less than one minute per AOS iteration. In many practical applications, between 1 and 10 iterations are sufficient for the denoising of such data sets. This means that the filtering often takes less time than the image acquisition itself. We are currently testing our schemes for the filtering of 3-D ultrasound images and the preprocessing of 3-D MR data for segmentation. In both cases first results are encouraging.

5 Conclusions

We have presented absolutely stable additive operator splitting schemes for the nonlinear diffusion filter of Catté, Lions, Morel, and Coll. These schemes satisfy all criteria for discrete nonlinear diffusion scale-spaces, and they are easy to implement in any dimension. Both computational and storage effort is linear in the number of pixels. Experiments have shown that under realistic accuracy requirements one can gain an increase of efficiency by a factor 10. This makes this type of schemes attractive for applications such as medical 3D data sets. Implementations of AOS schemes on parallel architectures and generalizations to anisotropic diffusion filters with diffusion tensors will be discussed in forthcoming publications.

References

1. S.T. Acton, A.C. Bovik, M.M. Crawford, *Anisotropic diffusion pyramids for image segmentation*, Proc. IEEE Int. Conf. Image Processing (ICIP–94, Austin, Nov. 13–16, 1994), Vol. 3, IEEE Computer Society Press, Los Alamitos, 478–482, 1994.

2. L. Alvarez, *Images and PDE's*, 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, 3–14, 1996.
3. L.D. Cai, *Some notes on repeated averaging smoothing*, J. Kittler (Ed.), Pattern recognition, Lecture Notes in Comp. Science, Vol. 301, Springer, Berlin, 597–605, 1988.
4. 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.
5. R. Deriche, *Fast algorithms for low-level vision*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 12, 78–87, 1990.
6. 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.
7. T. Gijbels, P. Six, L. Van Gool, F. Catthoor, H. De Man, A. Oosterlinck, *A VLSI-architecture for parallel non-linear diffusion with applications in vision*, Proc. IEEE Workshop on VLSI Signal Processing, Vol. 7, 398–707, 1994.
8. A.R. Gourlay, *Implicit convolution*, Image Vision Comput., Vol. 3, 15–23, 1985.
9. B.M. ter Haar Romeny (Ed.), *Geometry-driven diffusion in computer vision*, Kluwer, Dordrecht, 1994.
10. R.A. Hummel, *Representations based on zero-crossings in scale space*, Proc. IEEE Comp. Soc. Conf. Computer Vision and Pattern Recognition (CVPR '86, Miami Beach, June 22–26, 1986), IEEE Computer Society Press, Washington, 204–209, 1986.
11. T. Iijima, *Basic theory of pattern normalization (for the case of a typical one-dimensional pattern)*, Bulletin of the Electrotechnical Laboratory, Vol. 26, 368–388, 1962 (in Japanese).
12. J. Kačur, K. Mikula, *Solution of nonlinear diffusion appearing in image smoothing and edge detection*, Appl. Num. Math., Vol. 17, 47–59, 1995.
13. T. Lindeberg, *Scale-space theory in computer vision*, Kluwer, Boston, 1994.
14. G.I. Marchuk, *Splitting and alternating direction methods*, P.G. Ciarlet, J.-L. Lions (Eds.), Handbook of numerical analysis, Vol. I, 197–462, 1990.
15. M. Nielsen, L. Florack, R. Deriche, *Regularization, scale-space, and edge detection filters*, B. Baxton, R. Cipolla (Eds.), Computer vision – ECCV '96, Volume II, Lecture Notes in Comp. Science, Vol. 1065, Springer, Berlin, 70–81, 1996.
16. W.J. Niessen, K.L. Vincken, J.A. Weickert, M.A. Viergever, *Nonlinear multiscale representations for image segmentation*, Computer Vision and Image Understanding, 1997, in press.
17. P. Perona, J. Malik, *Scale space and edge detection using anisotropic diffusion*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 12, 629–639, 1990.
18. J. Weickert, *Theoretical foundations of anisotropic diffusion in image processing*, W. Kropatsch, R. Klette, F. Solina (Eds.), Theoretical foundations of computer vision, Computing Suppl. 11, Springer, Wien, 221–236, 1996.
19. J. Weickert, *Nonlinear diffusion scale-spaces: From the continuous to the discrete setting*, 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, 111–118, 1996.
20. J. Weickert, *Anisotropic diffusion in image processing*, Teubner Verlag, Stuttgart, 1997, to appear.

21. J. Weickert, B.M. ter Haar Romeny, M.A. Viergever, *Efficient and reliable schemes for nonlinear diffusion filtering*, IEEE Trans. Image Proc., 1998, to appear.
22. J. Weickert, S. Ishikawa, A. Imiya, *On the history of Gaussian scale-space axiomatics*, J. Sporring, M. Nielsen, L. Florack, P. Johansen (Eds.), Gaussian scale-space theory, Kluwer, Dordrecht, 1997, in press.

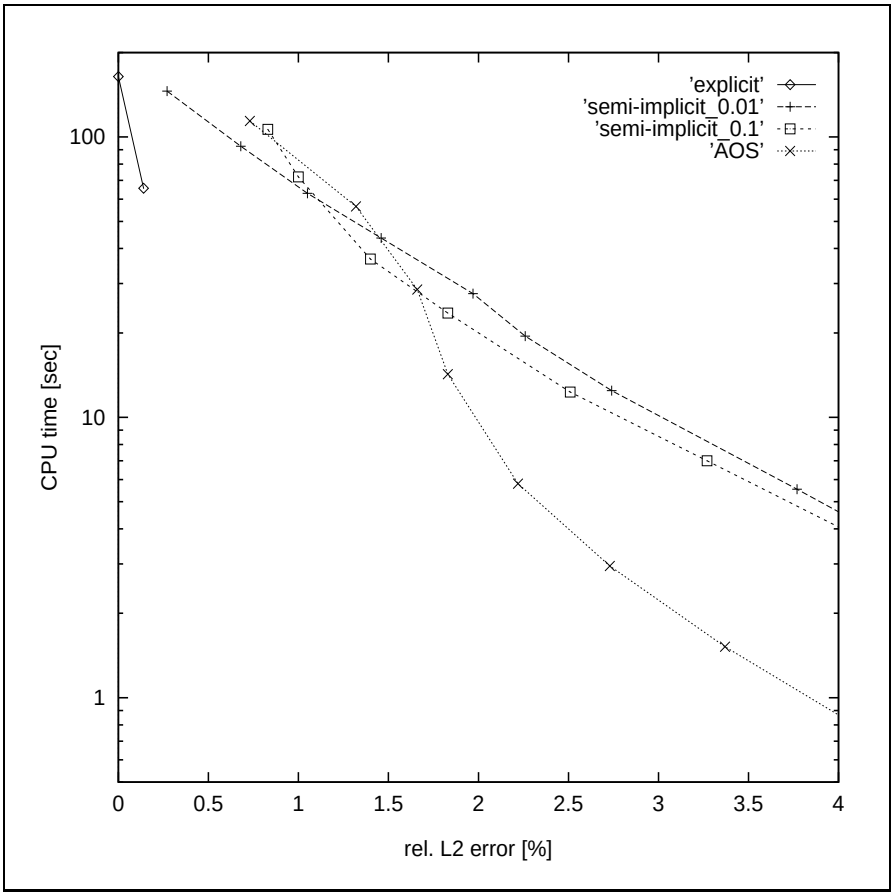

Fig. 3. Tradeoff between efficiency and accuracy of nonlinear diffusion solvers. The data were calculated on the test image from Fig. 2, size $\Omega = (0, 255) \times (0, 308)$. Filter parameters: $\lambda = 2$, $\sigma = 1$. Stopping time: $T = 200$.