Minimally Stochastic Schemes for Singular Diffusion Equations

Bernhard Burgeth[1], Joachim Weickert[1], and Sibel Tari[2]

[1]Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Bldg. 27
Saarland University, 66041 Saarbrücken, Germany
{burgeth,weickert}@mia.uni-saarland.de
http://www.mia.uni-saarland.de

[2]Department of Computer Engineering
Middle East Technical University
06531 Ankara, Turkey
stari@metu.edu.tr
http://www.ceng.metu.edu.tr

Abstract. Total variation (TV) and balanced forward-backward (BFB) diffusion are prominent examples of singular diffusion equations: Finite extinction time, the experimentally observed tendency to create piecewise constant regions, and being free of parameters makes them very interesting. However, their appropriate numerical treatment is still a challenge. In this paper a minimally stochastic approach to these singular equations is presented. It is based on analytical solutions of two-pixel signals and stochastic rounding. This introduces regularisation via integer arithmetic and does not require any limits on the diffusivity. Experiments demonstrate the favourable performance of the proposed probabilistic method.

Keywords: Randomisation, total variation, balanced forward-backward diffusion, singular diffusivity

1 Introduction

1.1 The Setting

Initiated with the work of Perona and Malik [11] nonlinear diffusion filters have become an important tool for image processing. The basic setting of diffusion filtering is as follows. An initial image $f : \Omega \longrightarrow \mathbb{R}$ given on an two-dimensional domain $\Omega \subset \mathbb{R}^2$ is subjected to an evolutionary process governed by the following partial differential equation (PDE) with Neumann boundary conditions:

$$\begin{aligned} \partial_t u &= \text{div}\ (g(|\nabla u|) \cdot \nabla u) \quad \text{on} \quad \Omega \times (0, \infty) \\ u(x, 0) &= f(x) \qquad\qquad\qquad \text{for all } x \in \Omega\,, \\ \partial_n u(x, t) &= 0 \qquad\qquad\qquad\quad\ \text{for all } x \in \partial\Omega \times (0, \infty)\,, \end{aligned} \tag{1}$$

with outward normal derivative ∂_n on the image domain boundary $\partial\Omega$. This evolution process creates more simplified versions $u(\cdot, t)$ of f the larger the time parameter t is. One can steer this process to achieve edge preservation and intraregional smoothing by specifying the diffusivity g as a nonnegative and decreasing function of $|\nabla u|$.
Many nonlinear diffusion filters rely on bounded diffusivities [6, 11]. However, in recent years unbounded diffusivities that became singular at zero have received special attention [8, 2, 10, 7]. In numerical experiments these filters create cartoon-like, piecewise constant images. In this paper we will focus on two choices for the diffusivity g, both rendering the corresponding PDE singular: The specification

$$g(|\nabla u|) = \frac{1}{|\nabla u|} \tag{2}$$

gives rise to the total variation (TV) diffusion [2, 9]. The TV diffusion filter is associated with TV regularisation if a penaliser $\Psi(|\nabla u|^2) = 2|\nabla u|$ is used [14]. Among the most interesting properties of TV diffusion are finite extinction time [3], certain shape-preserving qualities [4], and equivalence results to TV regularisation for one-dimensional signals [5, 12].

The specification

$$g(|\nabla u|) = \frac{1}{|\nabla u|^2} \tag{3}$$

generates the so-called balanced forward-backward diffusion (BFB), [10]. For this type of diffusion actual edge enhancement occurs. Note that neither TV nor BFB diffusion require any filter parameter tuning. Generalisations of this diffusion filters replacing the square by a positive exponent p also have been considered in [1, 16].

Numerical difficulties are the price to be paid for the appealing properties of TV or BFB diffusion: In order to apply classical finite difference schemes, one needs bounded diffusivities. This is achieved by replacing $|\nabla u|$ by $\sqrt{|\nabla u|^2 + \epsilon^2}$ in the denominators of (2) and (3). However, the time step size in explicit finite difference schemes is reciprocal to bounds on the diffusivity function to ensure stability, and the condition numbers of system matrices emerging from absolutely stable semi-implicit or implicit schemes are increasing functions of such bounds. This entails high computational complexity and/or potential amplification of numerical errors. Moreover the bounded diffusivity introduces the unpleasant side effect that blurring artefacts occur and theoretical considerations for singular diffusion filters are no longer applicable.

An alternative that does not require a regularised diffusivity is described in [15]: In a two-pixel setting analytic solutions of systems of ordinary differential equations associated with a spatial discretisation of the singular PDE are employed for numerical evaluation. In [18] the same idea of utilising analytical solutions of ODE-systems has been put to work successfully in the more complicated framework of four pixels. Both approaches lead to absolutely stable explicit scheme at the expense of having conditional consistency only: When the product of the time step size and the diffusivity becomes large, a linear diffusion process is approximated. This means that for increasing time step sizes, mor and more blurring artifacts arise.

The goal of the present paper is to address this problem. By introducing an approximation that allows only integers as grey values, we bound the gradient away from zero: The employed one-sided discretisation $|\nabla u|_{i,j}$ of $|\nabla u|$ in (15) entails that either $|\nabla u|_{i,j} = 0$, which can be treated separately, or $|\nabla u|_{i,j} \geq \frac{1}{\sqrt{2}h}$ with grid size h. This implies that the discrete approximations $g_{i,j}$ for the diffusivity are bounded by $\sqrt{2}\,h$ in the case of TV-diffusion and by $2\,h^2$ for BFB-diffusion. Hence we are allowed to use larger time step sizes without visual deterioration than in the conventional 2- or 4-pixel schemes.

Since diffusion is an inherently continuous process that should also be allowed to proceed in infinitesimally small steps, it is not possible to design a satisfying diffusion scheme that uses integer arithmetic in a deterministic framework by conventional rounding. As a remedy, we introduce a minimal amount of randomisation in the spirit of [13]. It is realised by a stochastic rounding procedure which introduces fluctuations that are small enough to be invisible, but large enough to have a beneficial regularising effect.

The paper is structured as follows: The two-pixel scheme based on an analytic solution of a system of ODEs is introduced in the first part of the next section. In its second part the analytic two-pixel scheme is randomised by stochastic rounding leading to the proposed minimally stochastic method. Numerical experiments in section 3 show the favourable performance of the minimally stochastic approach when compared to the purely deterministic method. Section 4 with a short summary and remarks about ongoing work completes the paper.

2 Schemes Based on Two Pixel Interaction

2.1 Deterministic Approach

We will start our investigation with the simplest possible case. We are considering a one-dimensional version of (1) discretised by two pixels with homogenous Neumann boundary conditions:

$$f = (f_1, f_2), \quad \text{resp.,} \quad u = (u_1, u_2)\,.$$

A space-discrete, but time-continuous scheme for (1) is then given by

$$\begin{aligned} \dot{u}_1 &= \frac{g_{1+\frac{1}{2}}}{h^2}\,(u_2 - u_1) \\ \dot{u}_2 &= \frac{-g_{1+\frac{1}{2}}}{h^2}\,(u_2 - u_1) \end{aligned}$$

with initial conditions $u_i(0) = f_i$, $i = 1, 2$. Here the discrete approximants g_1 and g_2 of the diffusivity g at pixel 1 and 2 are calculated using dummy pixels

$u_0 := u_1$ and $u_3 := u_2$ yielding $g_{1+\frac{1}{2}}$ by

$$g_{1+\frac{1}{2}} := \frac{g_1 + g_2}{2}\,.$$

In general, first order differences are approximated by standard central differences $\frac{1}{2h}(u_{i-1} - u_{i+1})$ with grid size h.

We assume that $g_{1+\frac{1}{2}}$ is independent of time, that is, constant with respect to t in this coupled system of ordinary differential equations. In order to de-couple this system of ODEs we introduce $w_1(t) = u_2(t) - u_1(t)$ and $v_1(t) = u_2(t) + u_1(t)$, in fact

$$\begin{pmatrix} w_1(t) \\ v_1(t) \end{pmatrix} = \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix} \cdot \begin{pmatrix} u_1(t) \\ u_2(t) \end{pmatrix}\,. \tag{4}$$

Then the function w_1 satisfies the linear first order ODE

$$\dot{w}_1 = \frac{2}{h^2}\, g_{1+\frac{1}{2}} w_1\,,$$

which is readily solved to give

$$w_1(t) = \exp\left(-\frac{2}{h^2}\, g_{1+\frac{1}{2}} \cdot t\right) w_1(0)\,.$$

For the sum $v_1(t)$ we obtain the ODE

$$\dot{v}_1(t) = 0\,,$$

yielding $v_1(t) = v_1(0) = u_2(0) + u_1(0)$ for all $t \geq 0$. With this at our disposal solving the equation system (4) gives

$$\begin{aligned} \begin{pmatrix} u_1(t) \\ u_2(t) \end{pmatrix} &= \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix}^{-1} \cdot \begin{pmatrix} w_1(t) \\ v_1(t) \end{pmatrix} \\ &= \begin{pmatrix} u_1(0) \\ u_2(0) \end{pmatrix} + \frac{1}{2}\left(1 - \exp\left(-\frac{2t}{h^2} g_{1+\frac{1}{2}}\right)\right)(u_2(0) - u_1(0)) \begin{pmatrix} 1 \\ -1 \end{pmatrix} \end{aligned} \tag{5}$$

Considering now n-pixel signals we may apply this reasoning to any pair of pixels u_i and u_{i+1}. Thus we obtain

$$\begin{aligned} u_i(t) &= u_i(0) \quad + \frac{1}{2}\left(1 - \exp\left(-\frac{2t}{h^2} g_{i+\frac{1}{2}}\right)\right)(u_{i+1}(0) - u_i(0)) \\ u_{i+1}(t) &= u_{i+1}(0) - \frac{1}{2}\left(1 - \exp\left(-\frac{2t}{h^2} g_{i+\frac{1}{2}}\right)\right)(u_{i+1}(0) - u_i(0))\,, \end{aligned}$$

or in its time discrete variant after k iterations with time step size τ

$$\begin{aligned} u_i^{k+1} &= u_i^k \quad + \frac{1}{2}\left(1 - \exp\left(-\frac{2\tau}{h^2} g_{i+\frac{1}{2}}^k\right)\right)(u_{i+1}^k - u_i^k) \\ u_{i+1}^{k+1} &= u_{i+1}^k - \frac{1}{2}\left(1 - \exp\left(-\frac{2\tau}{h^2} g_{i+\frac{1}{2}}^k\right)\right)(u_{i+1}^k - u_i^k)\,. \end{aligned} \tag{6}$$

However, this ensures interaction between the two neighbouring pixels u_i^k and u_{i+1}^k only, pixel u_{i-1}^k, say, is not involved. In order to overcome this problem we consider also a shifted version of the signal, follow the procedure indicated above and average the two signal versions in an additive operator splitting (AOS) approach [17]: We allow for diffusion between u_i^k and u_{i+1}^k with time step size 2τ yielding

$$\tilde{u}_i^{k+1} = u_i^k + \frac{1}{2}\left(1 - \exp\left(-\frac{4\tau}{h^2} g_{i+\frac{1}{2}}^k\right)\right)(u_{i+1}^k - u_i^k)\,, \tag{7}$$

and we enable diffusion between u_i^k and u_{i-1}^k with time step size 2τ by setting

$$\tilde{\tilde{u}}_i^{k+1} = u_i^k - \frac{1}{2}\left(1 - \exp\left(-\frac{4\tau}{h^2} g_{i-\frac{1}{2}}^k\right)\right)(u_i^k - u_{i-1}^k)\,. \tag{8}$$

Then averaging $u_i^{k+1} = \frac{1}{2}(\tilde{u}_i^{k+1} + \tilde{\tilde{u}}_i^{k+1})$ results in

$$\begin{aligned} u_i^{k+1} = u_i^k &+ \frac{1}{4}\left(1 - \exp\left(-\frac{4\tau}{h^2} g_{i+\frac{1}{2}}^k\right)\right)(u_{i+1}^k - u_i^k) \\ &- \frac{1}{4}\left(1 - \exp\left(-\frac{4\tau}{h^2} g_{i-\frac{1}{2}}^k\right)\right)(u_i^k - u_{i-1}^k)\,. \end{aligned} \tag{9}$$

The combination of these two steps according to the AOS-framework permits the transport of information throughout the image domain, since it provides a coupling between all pixels. Only this ensures the usefulness of the two-pixel module described in (6), res., in (7) and (8).
Note that a formal first order Taylor expansion w.r.t. τ of the exponential expressions yields the explicit scheme

$$\begin{aligned} u_i^{k+1} = u_i^k &+ \frac{\tau}{h^2} g_{i+\frac{1}{2}}^k (u_{i+1}^k - u_i^k) \\ &- \frac{\tau}{h^2} g_{i-\frac{1}{2}}^k (u_i^k - u_{i-1}^k)\,. \end{aligned} \tag{10}$$

The stability of scheme (10) will be destroyed by large diffusivity values. In contrast to that the exponential scheme (9) remains stable, However, as all unconditionally stable explicit schemes, it is only conditionally consistent: If the product of the time step size and the diffusivity becomes large the algorithm turns into simple averaging, and therefore approximates linear diffusion.

In the two-dimensional case of images an analog derivation leads to the scheme

$$\begin{aligned} u_{i,j}^{k+1} = u_{i,j}^k &+ \frac{1}{8}\left(1-\exp\left(-\frac{8\tau}{h^2}g_{i+\frac{1}{2},j}^k\right)\right)(u_{i+1,j}^k - u_{i,j}^k) \\ &+ \frac{1}{8}\left(1-\exp\left(-\frac{8\tau}{h^2}g_{i-\frac{1}{2},j}^k\right)\right)(u_{i-1,j}^k - u_{i,j}^k) \\ &+ \frac{1}{8}\left(1-\exp\left(-\frac{8\tau}{h^2}g_{i,j+\frac{1}{2}}^k\right)\right)(u_{i,j+1}^k - u_{i,j}^k) \\ &+ \frac{1}{8}\left(1-\exp\left(-\frac{8\tau}{h^2}g_{i,j-\frac{1}{2}}^k\right)\right)(u_{i,j-1}^k - u_{i,j}^k) \end{aligned}$$

Since we are averaging over twice as many neighbours as in the 1-D case, the weight 4 had been replaced by 8. This scheme is also well-suited for singular diffusivities, it is unconditionally stable and conditionally consistent.

2.2 Minimally Stochastic Approach

We want to construct an integer-valued analog to the process (6), that is, a system

$$\begin{aligned} u_m^{k+1} &= u_m^k + \omega \\ u_n^{k+1} &= u_n^k - \omega \end{aligned} \tag{11}$$

where ω can only assume integer values. This warrants that the integer grey values of the initial image remain integer valued during the whole evolution process. As already mentioned conventional rounding is not an feasible option, hence we introduce a form of randomised rounding. This amounts to the design of a randomising module that requires the data of only two pixels as input. Instead of rounding by $[x]$ = integer part of x, this module utilizes a stochastic rounding function $SR : \mathbb{R} \longrightarrow \mathbb{Z}$ defined by

$$SR(x) := \begin{cases} [x] & \text{with probability } \; 1-|x-[x]| \\ [x]+1 & \text{with probability } \; |x-[x]|. \end{cases}$$

One finds, for example,

$$SR(2.7) = \begin{cases} 2 & \text{with probability } \; 0.3 \\ 3 & \text{with probability } \; 0.7. \end{cases}$$

We employ this random variable to turn (11) with

$$\omega := SR\left[\frac{1}{2}\left(1-\exp\left(-\frac{2\tau}{h^2}\,\frac{g_n^k+g_m^k}{2}\right)\right)(u_n^k - u_m^k)\right] . \tag{12}$$

into a randomised and integer-valued variant of a 2-pixel scheme

The regularising effect of the proposed stochastic rounding allows for larger time steps. A standard deterministic rounding would not be appropriate if the image is piecewise almost flat. In this case deterministic rounding would not permit to diffuse small quantities which would entail unphysical results. Instead we allow for fluctuations of one grey level in magnitude and hereby exploit the insensitivity of the visual system to small changes in greyvalues.

So far the exchange of information between two pixels is ensured. Now the task that remains is to transport the information to other pixels. The idea close at hand would be to use an additive operator splitting like in the deterministic case. However, this would come down to averaging four integer solutions in each pixel, such that there is no guarantee that the result is an integer number again. This is the reason why we use a multiplicative operator splitting for our randomised approach. Since it leads to a sequential application of the randomised two-pixel interactions, integer results are ensured. In the 2-D setting there are 8 different ways of passing through all pixels in a regular order, as is indicated in Fig. 1. Selecting one of these cases, however, would introduce a directional bias for a

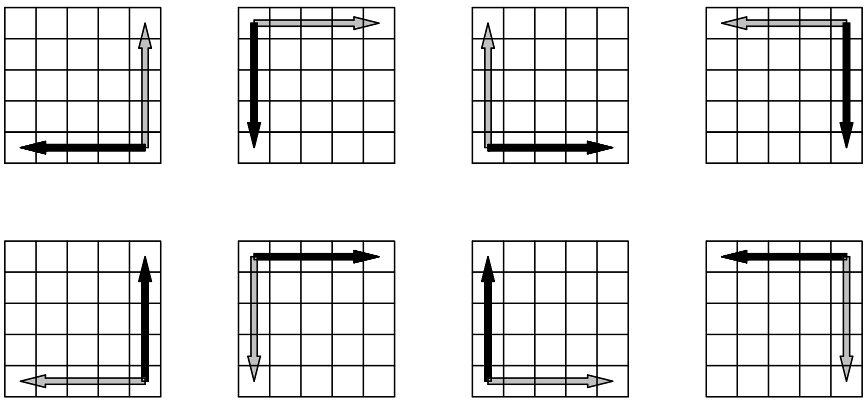

Fig. 1. Extension of the two-pixel-scheme to a 2D-image by applying it to overlapping pairs of pixels. Selection of the starting point and marching directions indicated by black and grey arrows.

nonlinear PDE such as TV flow. In order to avoid this problem, we introduce a second randomisation in our algorithm: We randomly choose one of the eight cases which are considered to be equally likely, namely of having probability $\frac{1}{8}$ each. From the numerical point of view the following issues have turned out to be beneficial:

– If the initial data (f_i) are integer valued the scheme in (11) produces integer values only.

– Since the diffusivities considered are unbounded the case that $g_m^k = \infty$ or $g_n^k = \infty$ must be accounted for by setting

$$\omega := SR\left(\frac{u_n^k - u_m^k}{2}\right) . \tag{13}$$

– From the numerical point of view it is advantageous to compute reciprocal diffusivities $\frac{1}{g_m^k}$ and use the harmonic mean for averaging:

$$\omega = \begin{cases} SR\left\{\left(1 - \exp\left(-\frac{4\tau}{h^2}\left(\frac{1}{g_n^k} + \frac{1}{g_m^k}\right)^{-1}\right)\right)\frac{u_m^k - u_n^k}{2}\right\} & \text{for } \frac{1}{g_m^k}, \frac{1}{g_n^k} > 0\,, \\ SR\left(\frac{u_m^k - u_n^k}{2}\right) & \text{for } \frac{1}{g_m^k} \text{ or } \frac{1}{g_n^k} = 0\,. \end{cases} \tag{14}$$

It is important to remark that the proposed minimally stochastic scheme produces filtered data consisting of integer values as soon as the initial data are integer valued making it suitable for simple hardware. The scheme also obeys a minimum-maximum-principle since the two-pixel process does. This is an important stability issue.

3 Numerical Experiments

In this section we display some results of numerical experiments to visualise the properties of the deterministic exponential and the minimally stochastic approach. We consider a 256×256 greyvalue medical image and a 128×128 image where in 70% of its pixels the grey value is replaced by an value randomly chosen according to a uniform distribution on $\{0, 1, \ldots, 255\}$. For the discretisation of $|\nabla u|$ we used one sided differences:

$$|\nabla u_{i,j}| = \left\{\frac{1}{2}\left(\left(\frac{u_{i+1,j} - u_{i,j}}{h}\right)^2 + \left(\frac{u_{i,j} - u_{i-1,j}}{h}\right)^2\right) + \frac{1}{2}\left(\left(\frac{u_{i,j+1} - u_{i,j}}{h}\right)^2 + \left(\frac{u_{i,j} - u_{i-1,j-1}}{h}\right)^2\right)\right\}^{\frac{1}{2}} \tag{15}$$

We subject the images to TV-diffusion based on both the deterministic and minimally stochastic two-pixel-scheme. The total diffusion time of 100 is achieved with time step sizes $\tau = 0.01,\ 0.1,\ 1$, that is, with 10000, 1000, 100 iterations. The sequence of filtered images indicates clearly the stabilising effect of the randomisation: The minimally stochastic computation allows for about 10 times larger time steps when compared with a deterministic counterpart of the same visual quality. While with a time step size of $\tau = 1$ the deterministic scheme produces an output degraded by fluctuations and blurring effects, the minimally stochastic approach still yields a satisfactory result.

The situation is similar but less pronounced in the case of BFB-diffusion. Here the total diffusion time is 3000 tackled with time step sizes $\tau = 3,\ 10,\ 30$

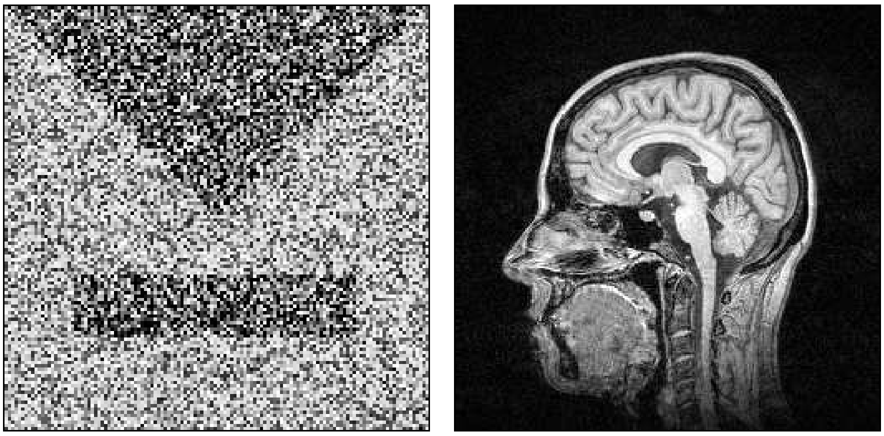

Fig. 2. Test images. *Left:* A 128×128 image polluted with 70% uniform noise. A 256×256 image without additional noise.

which entails 1000, 300, 100 iterations. Again the regularising effect of the minimally stochastic computation is clearly discernable, however, the gain is now an about three times larger time step in comparison with a qualitatively similar deterministic result.

The computational gain achieved by the minimally stochastic approach over the deterministic method is documented for both TV- and BFB-diffusion in table 1. 10000 iterations each have been performed on a Athlon XP 2.4 Ghz CPU for a grey value image of size 256×256. One can say that the deterministic and the minimally stochastic scheme are computationally equally costly.

	deterministic	minimally stochastic
TV	7 min 20.6 sec	7 min 28.3 sec
BFB	7 min 25.8 sec	7 min 26.1 sec

Tab. 1. CPU time necessary for 10000 iterations performed with the deterministic explicit or minimally stochastic sheme for TV- and BFB-diffusion.

4 Conclusion

The usage of singular diffusivities has advantages, like feature preserving qualities and the absence of tuning parameters, for instance. However, numerical intricacies turn the actual calculations into a challenging task. In this paper we introduce a minimally stochastic approach that regularises the singular diffusion filter. It is based on a time-continuous but space-discrete explicit two-pixel-scheme for which an analytical solution can be derived. This two-pixel-scheme receives a random component by employing stochastic rounding. The regularising effect of this randomisation allows for much larger time steps when compared

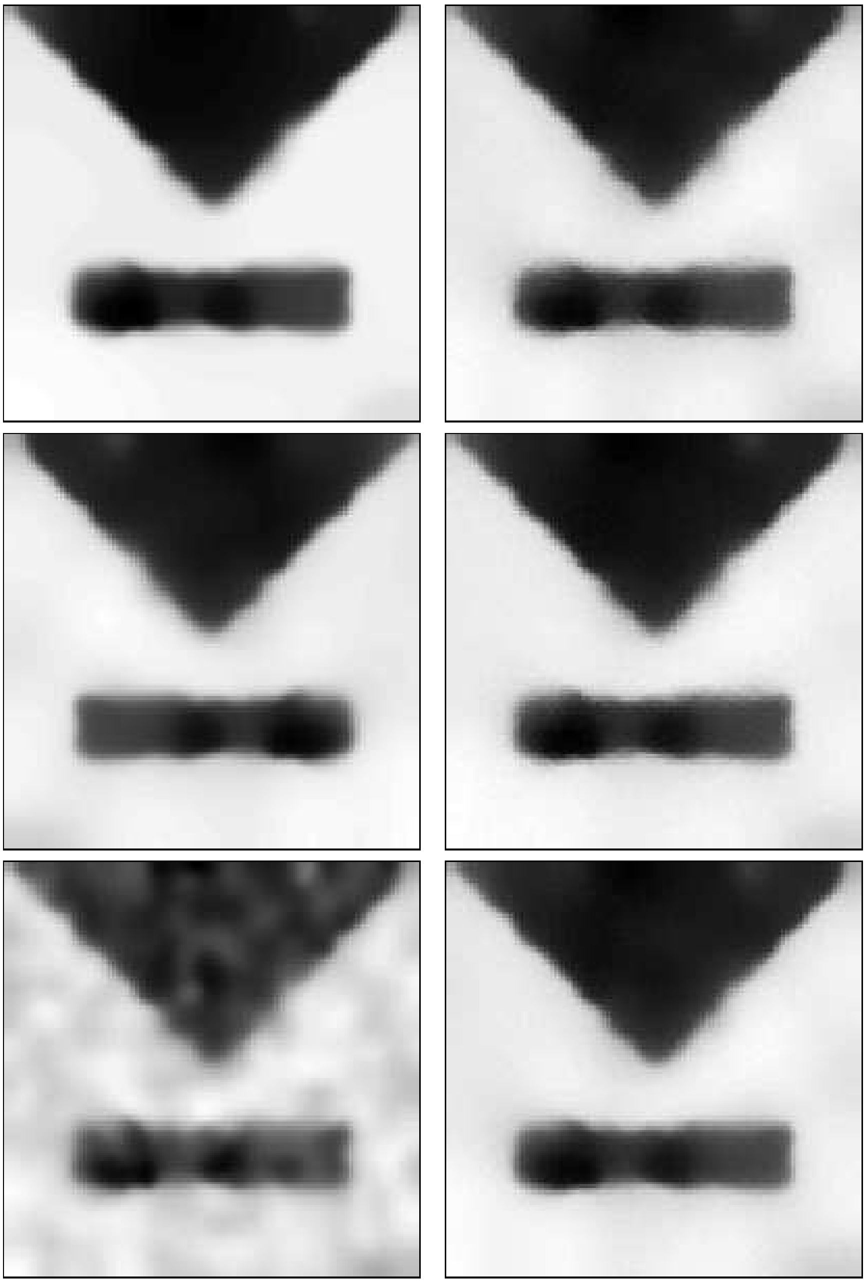

Fig. 3. Comparing deterministic and minimally stochastic computations of TV diffusion filtering with total diffusion time 100.
Left column: Deterministic calculation with explicit scheme.
Right column: Minimally stochastic calculation.
From top to bottom: Time step size $\tau = 0.01, 0.1,$ and 1 requiring 10^4, 10^3, and 10^2 iterations, respectively.

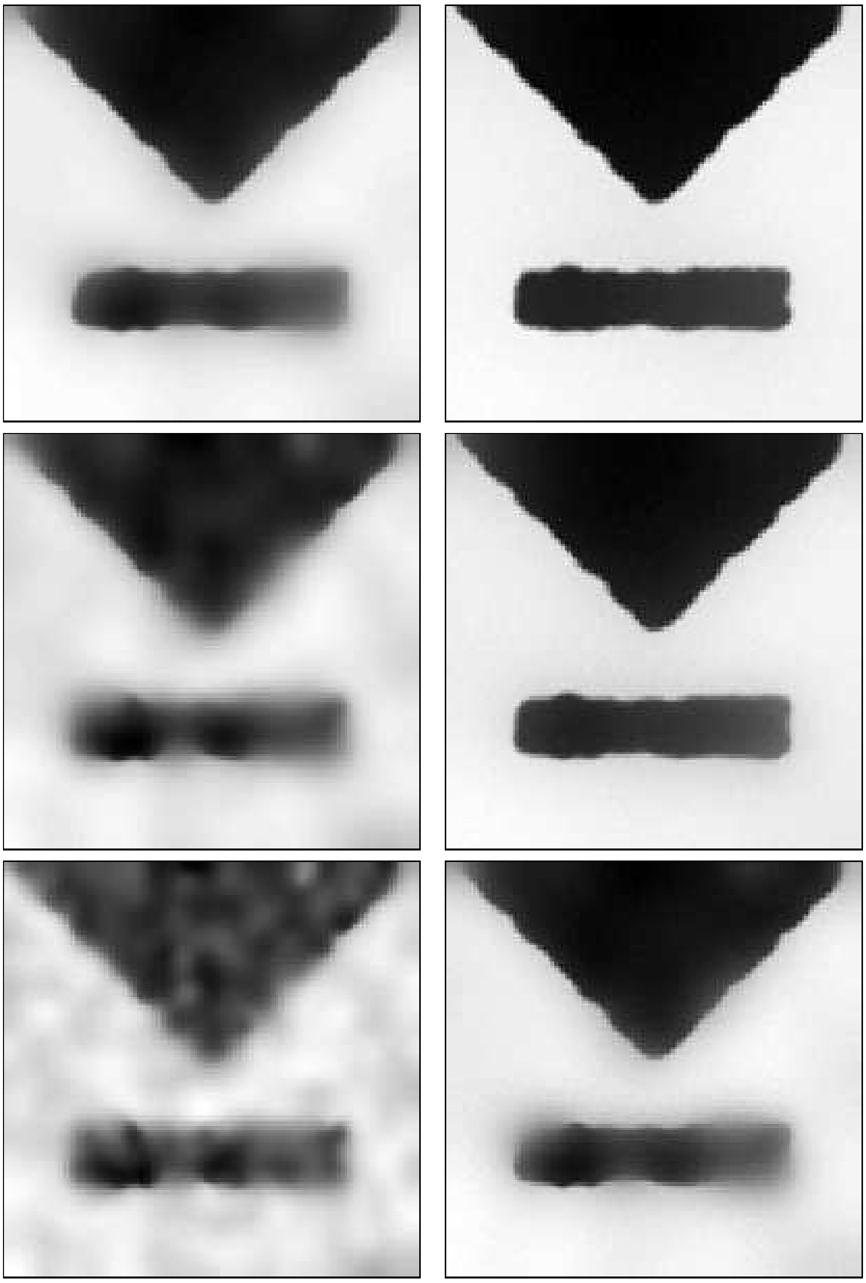

Fig. 4. Comparing deterministic and minimally stochastic computations of BFB diffusion filtering with total diffusion time 3000.
Left column: Deterministic calculation with explicit scheme.
Right column: Minimally stochastic calculation.
From top to bottom: Time step size $\tau = 3$, 10, and 30 requiring 1000, 300, and 100 iterations, respectively.

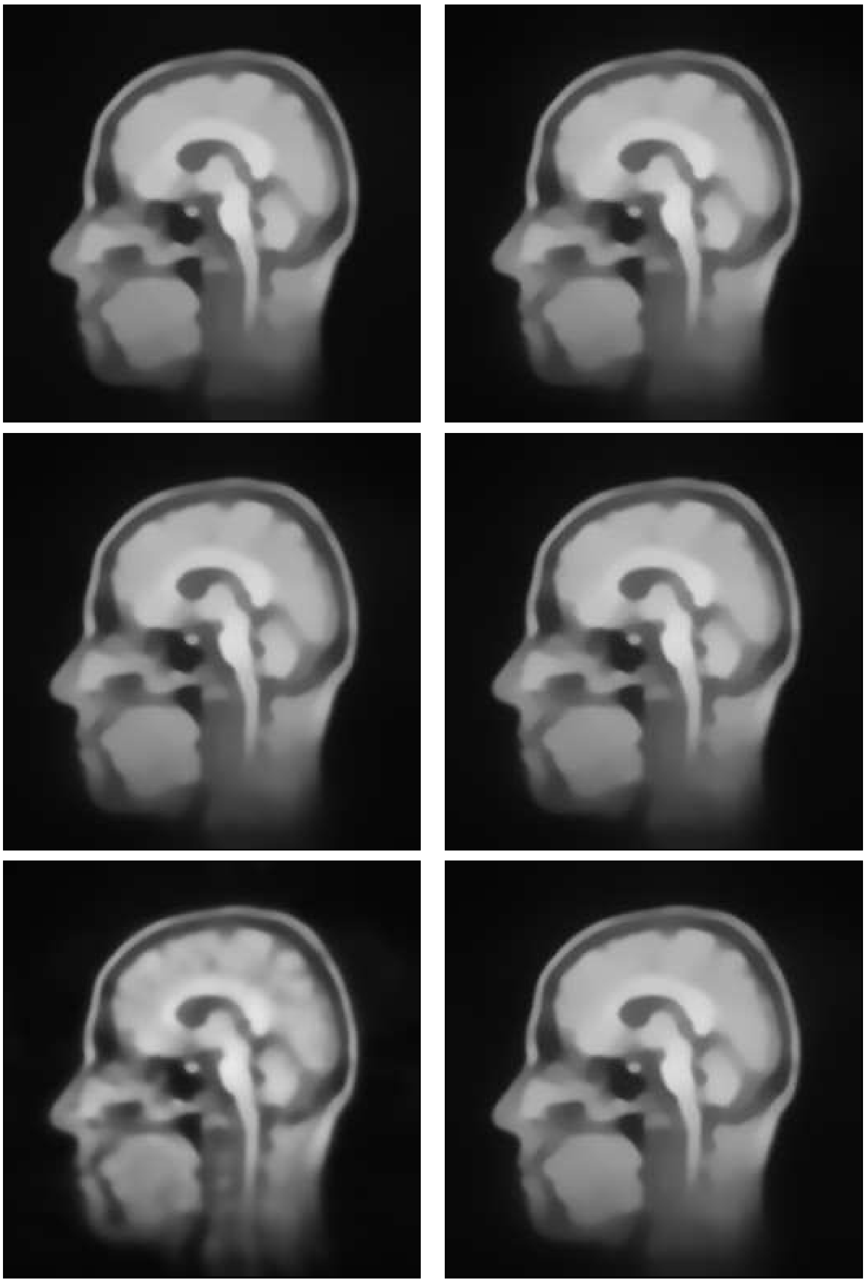

Fig. 5. Comparing deterministic and minimally stochastic computations of TV diffusion filtering with total diffusion time 100.
Left column: Deterministic calculation with explicit scheme.
Right column: Minimally stochastic calculation.
From top to bottom: Time step size $\tau = 0.01, 0.1,$ and 1 requiring 10^4, 10^3, and 10^2 iterations, respectively.

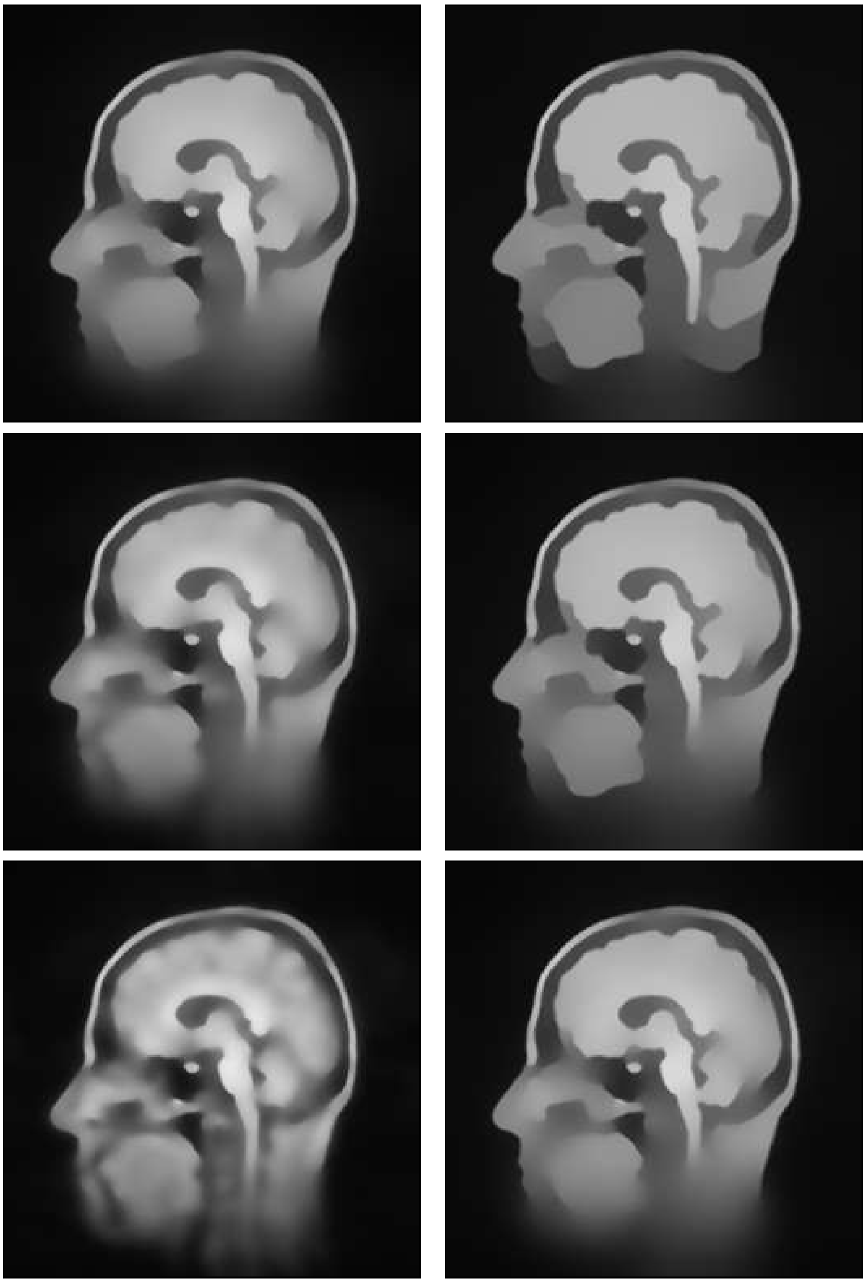

Fig. 6. Comparing deterministic and minimally stochastic computations of BFB diffusion filtering with total diffusion time 3000.
Left column: Deterministic calculation with explicit scheme.
Right column: Minimally stochastic calculation.
From top to bottom: Time step size $\tau = 3$, 10, and 30 requiring 1000, 300, and 100 iterations, respectively.

with the deterministic two-pixel-scheme, and for integer valued initial data it can be realised in such a way that only integer arithmetic is required. The numerical experiments show the favourable performance of the minimally stochastic scheme.
Ongoing research dedicated to the general class of diffusivities $g(|\nabla u|) = \frac{1}{|\nabla u|^p}$, $p > 0$, encompasses the usage of a more sophisticated four-pixel scheme and a deeper investigation of the performance.

References

1. L. Alvarez, F. Guichard, P.-L. Lions, and J.-M. Morel. Axioms and fundamental equations in image processing. *Archive for Rational Mechanics and Analysis*, 123:199–257, 1993.
2. F. Andreu, C. Ballester, V. Caselles, and J. M. Mazón. Minimizing total variation flow. *Differential and Integral Equations*, 14(3):321–360, March 2001.
3. F. Andreu, V. Caselles, J. I. Diaz, and J. M. Mazón. Qualitative properties of the total variation flow. *Journal of Functional Analysis*, 188(2):516–547, February 2002.
4. G. Bellettini, V. Caselles, and M. Novaga. The total variation flow in R^N. *Journal of Differential Equations*, 184(2):475–525, 2002.
5. T. Brox, M. Welk, G. Steidl, and J. Weickert. Equivalence results for TV diffusion and TV regularisation. In L. D. Griffin and M. Lillholm, editors, *Scale-Space Methods in Computer Vision*, volume 2695 of *Lecture Notes in Computer Science*, pages 86–100, Berlin, 2003. Springer.
6. 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*, 32:1895–1909, 1992.
7. Q. S. Chang and I. Chern. Acceleration methods for total variation based denoising problems. *SIAM Journal on Scientific Computing*, 25:982–994, 2003.
8. F. Dibos and G. Koepfler. Global total variation minimization. *SIAM Journal on Numerical Analysis*, 37(2):646–664, 2000.
9. F. Dibos and G. Koepfler. Total variation minimization by the Fast Level Sets Transform. In *Proc. First IEEE Workshop on Variational and Level Set Methods in Computer Vision*, pages 145–152, Vancouver, Canada, July 2001. IEEE Computer Society Press.
10. S. L. Keeling and R. Stollberger. Nonlinear anisotropic diffusion filters for wide range edge sharpening. *Inverse Problems*, 18:175–190, January 2002.
11. P. Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:629–639, 1990.
12. I. Pollak, A. S. Willsky, and Y. Huang. Nonlinear evolution equations as fast and exact solvers of estimation problems. *IEEE Transactions on Signal Processing*, 2004. To appear.
13. U. S. Ranjan and K. R. Ramakrishnan. A stochastic scale space for multiscale image representation. In M. Nielsen, P. Johansen, O. F. Olsen, and J. Weickert, editors, *Scale-Space Theories in Computer Vision*, volume 1682 of *Lecture Notes in Computer Science*, pages 441–446. Springer, Berlin, 1999.

14. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
15. G. Steidl and J. Weickert. Relations between soft wavelet shrinkage and total variation denoising. In L. Van Gool, editor, *Pattern Recognition*, volume 2449 of *Lecture Notes in Computer Science*, pages 198–205. Springer, Berlin, 2002.
16. V. I. Tsurkov. An analytical model of edge protection under noise suppression by anisotropic diffusion. *Journal of Computer and Systems Sciences International*, 39(3):437–440, 2000.
17. 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(3):398–410, March 1998.
18. M. Welk, J. Weickert, and G. Steidl. A four-pixel scheme for singular differential equations. In R. Kimmel, N. Sochen, and J. Weickert, editors, *Scale-Space and PDE Methods in Computer Vision*, volume 3459 of *Lecture Notes in Computer Science*, pages 585–597, Berlin, 2005. Springer.