Rotationally Invariant Wavelet Shrinkage⋆

Pavel Mrázek and Joachim Weickert

Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Building 27
Saarland University, 66123 Saarbrücken, Germany
{mrazek,weickert}@mia.uni-saarland.de
http://www.mia.uni-saarland.de

Abstract. Most two-dimensional methods for wavelet shrinkage are efficient for edge-preserving image denoising, but they suffer from poor rotation invariance. We address this problem by designing novel shrinkage rules that are derived from rotationally invariant nonlinear diffusion filters. The resulting Haar wavelet shrinkage methods are computationally inexpensive and they offer substantially improved rotation invariance.

1 Introduction

Wavelet shrinkage is a fast nonlinear method for discontinuity-preserving image denoising [1]. It is based on the idea to decompose an image in terms of a wavelet basis, to shrink all coefficients with small magnitude, and to reconstruct the filtered image from the shrunken coefficients. The success of this procedure is based on the assumption that the original image can be represented by a relatively small number of wavelet coefficients with large magnitude, while moderate Gaussian noise affects all coefficients, although to a less severe amount.

Several ways have been proposed to improve wavelet shrinkage. One of them is to make the shrinkage translation invariant [2]. Very recently it also became clear that shift-invariant wavelet denoising can be improved substantially by iterating [3,4,5]. While these ideas work well for processing 1D signals, filtering of 2D images creates an additional, but very fundamental problem: the shrinkage should be rotationally invariant. Unfortunately, this is not the case for the frequently used separable approaches. Several attempts to create wavelet transforms with improved rotation invariance have appeared in the literature, including the directional cycle spinning [6], complex wavelets [7], or the elaborated edgelet and curvelet transforms [8]. These ideas are not only relatively difficult to implement, most of them are also computationally significantly more complex than traditional shift-invariant wavelet shrinkage in 2D.

In the present paper we address this problem by proposing a novel class of shift-invariant 2D wavelet shrinkage methods for iterated image denoising with a high degree of rotation invariance. This class is computationally as simple as classical shift-invariant wavelet shrinkage. It is inspired by considering a connection between wavelet shrinkage and nonlinear diffusion filtering that has already proved fruitful in the 1D case [9]. Two-dimensional nonlinear diffusion filtering is based on a continuous differential equation that is rotationally invariant. From numerical analysis it is well-known how to discretise such equations in a consistent way such that rotation invariance is approximated well. By identifying a discrete diffusion filter with a wavelet shrinkage formulation, we can find novel shrinkage rules that lead to substantially improved rotation invariance.

The paper is organised as follows. Sections 2 and 3 provide a brief introduction to translation-invariant Haar wavelet shrinkage, and nonlinear diffusion in 2D. The connections between the two procedures are exploited in Section 4 to establish the conditions on diffusivity and shrinkage functions under which the two methods (restricted to one-step / one-scale) are equivalent. In Section 5 we evaluate the rotation invariance of the new multiscale iterated wavelet filter. The paper is concluded with a summary in Section 6.

2 Wavelet Shrinkage

The discrete wavelet transform represents a one-dimensional signal f in terms of shifted versions of a dilated lowpass scaling function, and shifted and dilated versions of a bandpass wavelet function. In this paper we restrict our attention to the discrete transform with Haar wavelets, well suited for piecewise constant signals with discontinuities. The Haar wavelet transform is described by a low-pass filter L with coefficients $(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$, and a high-pass filter H with coefficients $(\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})$ [10].

The easiest way to design a two-dimensional wavelet transform is to use separable filters [10]. The 2D wavelet transform then describes a 2D signal $\mathbf{f} = (f_{i,j})$ with $i = 0,...,N_x - 1$ and $j = 0,...,N_y - 1$ by its low-pass component at level n, $\mathbf{v}^n$, and three channels of wavelet coefficients $\mathbf{w}_x^l$, $\mathbf{w}_y^l$ and $\mathbf{w}_{xy}^l$ at levels $l = 1, \ldots, n$. This wavelet representation is created by an alternating application of the one-dimensional low-pass and high-pass filters L and H in the directions of axes x and y:

$$\mathbf{v}^{l+1} = L(x) * L(y) * \mathbf{v}^l, \qquad \mathbf{w}_y^{l+1} = L(x) * H(y) * \mathbf{v}^l, \qquad (1)$$

$$\mathbf{w}_x^{l+1} = H(x) * L(y) * \mathbf{v}^l, \qquad \mathbf{w}_{xy}^{l+1} = H(x) * H(y) * \mathbf{v}^l, \qquad (2)$$

with the initial condition $\mathbf{v}^0 = \mathbf{f}$. For image smoothing, the wavelet coefficients $\mathbf{w}_x, \mathbf{w}_y, \mathbf{w}_{xy}$ are subjected to a shrinkage function S, and the filtered image $\mathbf{u}$ is reconstructed from the shrunken coefficients using an inverse procedure to (1)–(2).

Let us now consider a *single* decomposition level, and the wavelet shrinkage steps which contribute to the output pixel $u_{i,j}$. Using the translation-invariant

$$v^\alpha = \begin{pmatrix} 0 & \tfrac12 & \tfrac12 \\ 0 & \tfrac12 & \tfrac12 \\ 0 & 0 & 0 \end{pmatrix} \quad w_y^\alpha = \begin{pmatrix} 0 & -\tfrac12 & \tfrac12 \\ 0 & \tfrac12 & \tfrac12 \\ 0 & 0 & 0 \end{pmatrix} \quad w_x^\alpha = \begin{pmatrix} 0 & \tfrac12 & -\tfrac12 \\ 0 & \tfrac12 & -\tfrac12 \\ 0 & 0 & 0 \end{pmatrix} \quad w_{xy}^\alpha = \begin{pmatrix} 0 & -\tfrac12 & \tfrac12 \\ 0 & \tfrac12 & -\tfrac12 \\ 0 & 0 & 0 \end{pmatrix}$$

$$v^\beta = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \tfrac12 & \tfrac12 \\ 0 & \tfrac12 & \tfrac12 \end{pmatrix} \quad w_y^\beta = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -\tfrac12 & -\tfrac12 \\ 0 & \tfrac12 & \tfrac12 \end{pmatrix} \quad w_x^\beta = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \tfrac12 & -\tfrac12 \\ 0 & \tfrac12 & -\tfrac12 \end{pmatrix} \quad w_{xy}^\beta = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -\tfrac12 & \tfrac12 \\ 0 & \tfrac12 & -\tfrac12 \end{pmatrix}$$

$$v^\gamma = \begin{pmatrix} \tfrac12 & \tfrac12 & 0 \\ \tfrac12 & \tfrac12 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad w_y^\gamma = \begin{pmatrix} -\tfrac12 & -\tfrac12 & 0 \\ \tfrac12 & \tfrac12 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad w_x^\gamma = \begin{pmatrix} \tfrac12 & -\tfrac12 & 0 \\ \tfrac12 & -\tfrac12 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad w_{xy}^\gamma = \begin{pmatrix} -\tfrac12 & \tfrac12 & 0 \\ \tfrac12 & -\tfrac12 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

$$v^\delta = \begin{pmatrix} 0 & 0 & 0 \\ \tfrac12 & \tfrac12 & 0 \\ \tfrac12 & \tfrac12 & 0 \end{pmatrix} \quad w_y^\delta = \begin{pmatrix} 0 & 0 & 0 \\ -\tfrac12 & -\tfrac12 & 0 \\ \tfrac12 & \tfrac12 & 0 \end{pmatrix} \quad w_x^\delta = \begin{pmatrix} 0 & 0 & 0 \\ \tfrac12 & -\tfrac12 & 0 \\ \tfrac12 & -\tfrac12 & 0 \end{pmatrix} \quad w_{xy}^\delta = \begin{pmatrix} 0 & 0 & 0 \\ -\tfrac12 & \tfrac12 & 0 \\ \tfrac12 & -\tfrac12 & 0 \end{pmatrix}$$

Fig. 1. The first-level Haar wavelet coefficients expressed using 3×3 masks centered at the pixel i, j. The masks represent multiplication of the input signal with the given coefficients, so e.g. $w_x^\alpha = \frac{1}{2} f_{i,j} + \frac{1}{2} f_{i,j+1} - \frac{1}{2} f_{i+1,j} - \frac{1}{2} f_{i+1,j+1}$.

scheme [2], we have to consider four 2×2 neighbourhoods in which the pixel (i, j) is involved. We denote by upper index α the neighbourhood $\{i, i+1\} \times \{j, j+1\}$; by β the positions $\{i, i+1\} \times \{j-1, j\}$; by γ the neighbourhood $\{i-1, i\} \times \{j, j+1\}$; and, finally, by δ the positions $\{i-1, i\} \times \{j-1, j\}$.

The input signal in neighbourhood α is first transformed into v^α, w_y^α, w_x^α and w_{xy}^α; see Fig. 1 for the definition of the corresponding masks. The wavelet coefficients $w_y^\alpha, w_x^\alpha, w_{xy}^\alpha$ are then subjected to a shrinkage function S, and the (i, j) pixel of the filtered signal belonging to the neighbourhood α is obtained using

$$u_{i,j}^\alpha = \frac{1}{2} \left(v^\alpha + S(w_y^\alpha) + S(w_x^\alpha) + S(w_{xy}^\alpha) \right). \tag{3}$$

Similar expressions can be derived for the results arising from the neighbourhoods β, γ and δ; the necessary masks are shown in Fig. 1. To obtain the final result of a shift-invariant 2D Haar wavelet shrinkage on a single level, the four intermediate results $u_{i,j}^\alpha$, $u_{i,j}^\beta$, $u_{i,j}^\gamma$ and $u_{i,j}^\delta$ have to be averaged. The complete formula for a single-level Haar wavelet shrinkage filter then reads

$$u_{i,j} = \frac{1}{8} \big(v^\alpha + v^\beta + v^\gamma + v^\delta + S(w_y^\alpha) + S(w_x^\alpha) + S(w_{xy}^\alpha) - S(w_y^\beta) + S(w_x^\beta)$$
$$- S(w_{xy}^\beta) + S(w_y^\gamma) - S(w_x^\gamma) - S(w_{xy}^\gamma) - S(w_y^\delta) - S(w_x^\delta) + S(w_{xy}^\delta) \big). \tag{4}$$

3 Nonlinear Diffusion

The basic idea behind nonlinear diffusion filtering [11] is to obtain a family $u(x, y, t)$ of filtered versions of the signal $f(x, y)$ as the solution of a suitable

diffusion process

$$u_t = \text{div}\left(g(|\nabla u|^2)\,\nabla u\right) \tag{5}$$

with f as initial condition: $u(x, y, 0) = f(x, y)$. Here subscripts denote partial derivatives, and the diffusion time t is a simplification parameter: larger values correspond to more pronounced filtering.

The divergence expression on the right hand side of (5) can be decomposed in 2D by means of two orthonormal basis vectors $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$:

$$\text{div}\left(g(|\nabla u|^2)\,\nabla u\right) = \sum_{p=1}^{2} \partial_{\boldsymbol{x}_p}\left(g(|\nabla u|^2)\,\partial_{\boldsymbol{x}_p} u\right). \tag{6}$$

Choosing the diagonal directions $\boldsymbol{x}_1 := (\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$ and $\boldsymbol{x}_2 := (\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})$, and replacing the derivatives in (5), (6) by finite differences, we can write the explicit finite difference discretisation of the nonlinear diffusion as

$$u_{i,j}^{k+1} = u_{i,j}^k + \tau \sum_{(I,J)\in\mathcal{D}(i,j)} g_{I,J}\,\frac{u_{I,J}^k - u_{i,j}^k}{2}. \tag{7}$$

Here the upper index k denotes solution at time $k\tau$ with τ standing for the time step, the set $\mathcal{D}(i, j)$ contains the diagonal neighbours of pixel (i, j), and the grid size is assumed to be 1. The term $g_{I,J} \approx g(|\nabla u(x, y)|^2)\big|_{\substack{x=(i+I)/2 \\ y=(j+J)/2}}$ represents the diffusivity belonging to the connection between pixels $u_{i,j}$ and $u_{I,J}$, where the gradient magnitude can be estimated from discrete data using 2×2 masks:

$$
\begin{aligned}
|\nabla u|^2 &= (\partial_x u)^2 + (\partial_y u)^2 \\
&\approx \left(\frac{u_{i,j} + u_{i,J} - u_{I,j} - u_{I,J}}{2}\right)^2 + \left(\frac{u_{i,j} - u_{i,J} + u_{I,j} - u_{I,J}}{2}\right)^2.
\end{aligned} \tag{8}
$$

This diagonal discretisation represents a consistent finite difference approximation to the continuous equation. It has been used successfully by Keeling and Stollberger [12]. Since its spatial consistency can be shown to be of second order, the rotation invariance of the continuous equation is approximated well.

4   Diffusion-Inspired 2D Wavelet Shrinkage

Let us now investigate the connection between a single-level wavelet shrinkage (4) and an explicit diffusion iteration (7). To this end, we consider the first diffusion iteration, starting from the initial signal $\mathbf{f} = (f_{i,j})$ and creating a solution $\mathbf{u} = (u_{i,j})$. We will express the diffusion iteration in the terms of the wavelet coefficients from Section 2.

For the first iteration, (7) becomes

$$u_{i,j} = f_{i,j} + \tau \sum_{(I,J)\in\mathcal{D}(i,j)} g_{I,J}\,\frac{f_{I,J} - f_{i,j}}{2}. \tag{9}$$

The first term on the right-hand side can be rewritten as

$$
f_{i,j} = \frac{1}{8}
\begin{array}{|c|c|c|}
\hline
\frac{1}{2} & 1 & \frac{1}{2} \\
\hline
1 & 2 & 1 \\
\hline
\frac{1}{2} & 1 & \frac{1}{2} \\
\hline
\end{array}
+ \frac{1}{8}
\begin{array}{|c|c|c|}
\hline
-\frac{1}{2} & -1 & -\frac{1}{2} \\
\hline
-1 & 6 & -1 \\
\hline
-\frac{1}{2} & -1 & -\frac{1}{2} \\
\hline
\end{array}
= \frac{1}{8}\left(v^\alpha + v^\beta + v^\gamma + v^\delta\right) + \frac{1}{8}\left(w_x^\alpha + w_y^\alpha + w_{xy}^\alpha\right.
$$

$$
\left. + w_x^\beta - w_y^\beta - w_{xy}^\beta - w_x^\gamma + w_y^\gamma - w_{xy}^\gamma - w_x^\delta - w_y^\delta + w_{xy}^\delta\right) \quad (10)
$$

where the 3×3 boxes stand for a mask multiplication with the input signal, and the v and w represent the wavelet coefficients for position (i, j); see Fig. 1. Then, the gradient magnitude for the diffusivity calculation can be estimated from the wavelet coefficients w_x and w_y:

$$
g_{I,J} = g\left((w_x^\omega)^2 + (w_y^\omega)^2\right) \quad (11)
$$

where $\omega = \alpha$ if $(I, J) = (i + 1, j + 1)$, $\omega = \beta$ for $(I, J) = (i + 1, j - 1)$, γ for $(i - 1, j + 1)$, and δ for $(i - 1, j - 1)$. Finally, the last term from (9) may be expressed using wavelet coefficients as

$$
f_{I,J} - f_{i,j} \in \{-w_x^\alpha - w_y^\alpha,\ -w_x^\beta + w_y^\beta,\ w_x^\gamma - w_y^\gamma,\ w_x^\delta + w_y^\delta\} \quad (12)
$$

where (I, J) is assigned to the expression involving α, β, γ or δ as above.

To summarise (9)–(12), we can write a single iteration of nonlinear diffusion using the wavelet decomposition components v and w in the form

$$
u_{i,j} = \frac{1}{8}\left(v^\alpha + v^\beta + v^\gamma + v^\delta + (w_x^\alpha + w_y^\alpha)(1 - 4\tau g^\alpha) + w_{xy}^\alpha + (w_x^\beta - w_y^\beta)(1 - 4\tau g^\beta)\right.
$$

$$
\left. - w_{xy}^\beta + (-w_x^\gamma + w_y^\gamma)(1 - 4\tau g^\gamma) - w_{xy}^\gamma + (-w_x^\delta - w_y^\delta)(1 - 4\tau g^\delta) + w_{xy}^\delta\right) \quad (13)
$$

where the symbol g^ω stands for $g\left((w_x^\omega)^2 + (w_y^\omega)^2\right)$, $\omega \in \{\alpha, \beta, \gamma, \delta\}$.

Comparing the diffusion iteration (13) and the single-level wavelet shrinkage (4), we observe that the two equations are equivalent under the conditions

$$
S(w_x^\omega) = w_x^\omega \left(1 - 4\tau g\left((w_x^\omega)^2 + (w_y^\omega)^2\right)\right), \quad (14)
$$

$$
S(w_y^\omega) = w_y^\omega \left(1 - 4\tau g\left((w_x^\omega)^2 + (w_y^\omega)^2\right)\right), \quad (15)
$$

$$
S(w_{xy}^\omega) = w_{xy}^\omega. \quad (16)
$$

Equations (14)–(16) connect the diffusivity function g controlling nonlinear diffusion to the shrinkage function S of wavelet shrinkage. If these conditions hold true, the two two-dimensional procedures (limited to a single scale / single iteration) are equivalent.

The equations (14), (15) are similar to the one-dimensional situation which was analysed in detail in [9]. The surprising fact in the 2D equations (14)–(16) is the use of different shrinkage rules for the different channels of wavelet coefficients, while the classical wavelet shrinkage applies the same shrinkage function S to each of them separately. In (14)–(15), the shrinkage of w_x and w_y is interconnected via the common estimation of the image gradient; the third channel, w_{xy}, is left by (16) unshrunken.
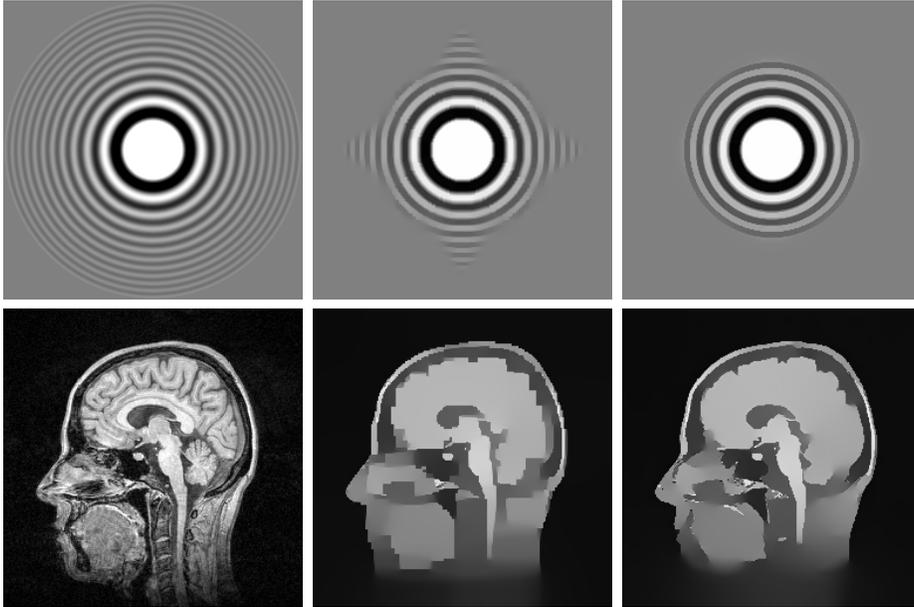
Fig. 2. Experiments on rotation invariance. Left: input image; center: filtered with classical iterated shift invariant wavelet shrinkage; right: method with channels coupled using (14)–(16). Top: ring image, 20 iterations on 8 levels of the wavelet decomposition. Bottom: head image, 100 iterations on 4 levels.

These three shrinkage rules inherit a fundamental property from their diffusion origin: the rotation invariance of the nonlinear diffusion filter. It holds exactly for the grid size tending to zero, but we shall see that this property is also approximated well in realistic discrete situations with non-vanishing grid size. Our diffusion-inspired idea to improve the invariance to rotation by coupling the shrinkage of wavelet channels as in (14)–(16) represents a very simple solution which hardly increases the computational complexity of a wavelet filter.

The diffusion–wavelet connection has been shown for a single iteration of a single-scale filter. In general, nonlinear diffusion is a single-scale iterative process, while wavelet shrinkage finds the solution using a single step on multiple scales. A hybrid multi-scale iterated filter seems a powerful and efficient alternative [5,4]. It can be understood either as a nonlinear diffusion on the Laplacian pyramid of the signal [13], or as iterated shift-invariant wavelet shrinkage [3]. The rotation invariance of such iterated multiscale filter is tested in the next section.

5 Experiments on Rotation Invariance

In this section we compare the rotation invariance of two wavelet-based filters: the classical iterated translation-invariant 2D wavelet shrinkage (4) with separate shrinkage of the coefficient channels, and the novel filter with shrinkage
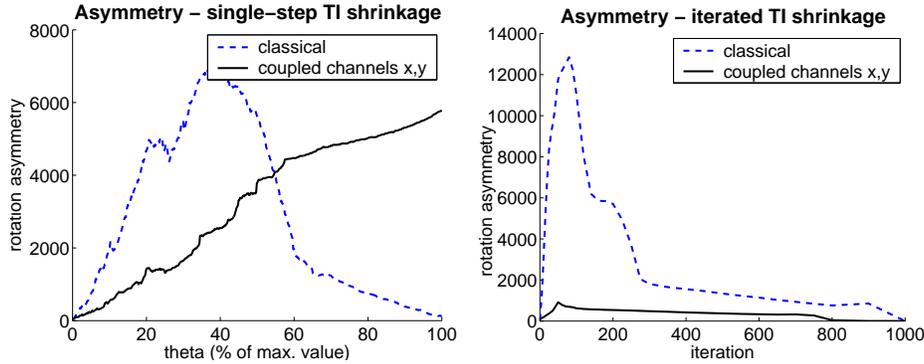
Fig. 3. Evaluation of the errors in rotation symmetry of the filtered ring image. Single-step (left), and iterated (right) shift-invariant wavelet shrinkage.

rules coupled according to (14)–(16). In all cases, we employ the Haar wavelet basis combined with hard thresholding: all coefficients with magnitude below a specified threshold θ are set to zero. The wavelet decomposition is calculated on multiple scales.

In the first experiment, we start from the rotationally symmetric ring image (Fig. 2 top left). Examples of images obtained after 5 iterations of each method are seen in Fig. 2 top. One can observe that using the coupled shrinkage (14)–(16), the filtered result reveals a much better rotational symmetry. The difference between the two methods is further visualised on a medical image at the bottom of Fig. 2. At a comparable level of image simplification, the new method is able to avoid the blocky artefacts of the classical transform.

The graphs in Fig. 3 present a numerical evaluation of the errors in rotational symmetry of the filtered ring image. The measure of asymmetry was calculated as a sum of signal variances along circles of varied diameter, centered at the center of rotation of the input image. In agreement with the design principles, the rotation symmetry of the *iterated* filter with coupled channels is very good (Fig. 3 right), but the coupled shrinkage outperforms the classical transform even using a single step of the multi-scale procedure, if the shrinkage parameter θ is not bigger than 50% of the value which flattens the image completely (Fig. 3 left).

6   Conclusions

In this paper we have addressed one of the main problems that are encountered when 2D wavelet shrinkage methods are to be used: the design of techniques with good rotation invariance. To this end we have established a connection between shift invariant Haar wavelet shrinkage on a single scale and an explicit discretisation of a nonlinear diffusion filter. Since diffusion filtering approximates a rotationally invariant continuous process, we have obtained a technique for

constructing shrinkage schemes with a high degree of rotational invariance. It turned out that all one has to do is to modify the shrinkage rules such that two of the wavelet channels are shrunken in a coupled way while the third one is left unaffected. The resulting diffusion-inspired wavelet shrinkage represents a straightforward and computationally efficient solution to the problem of designing rotationally invariant iterated multiscale wavelet shrinkage filters. Our experiments have shown that in this respect it clearly outperforms classical 2D shrinkage methods. In our future work we plan to investigate extensions of this design principle to other wavelets and other discretisations of nonlinear diffusion filters.

References

1. Donoho, D.L., Johnstone, I.M.: Ideal spatial adaptation by wavelet shrinkage. Biometrica **81** (1994) 425–455
2. Coifman, R.R., Donoho, D.: Translation invariant denoising. In Antoine, A., Oppenheim, G., eds.: Wavelets in Statistics. Springer, New York (1995) 125–150
3. Chambolle, A., Lucier, B.L.: Interpreting translationally-invariant wavelet shrinkage as a new image smoothing scale space. IEEE Transactions on Image Processing **10** (2001) 993–1000
4. Mrázek, P., Weickert, J., Steidl, G., Welk, M.: On iterations and scales of nonlinear filters. In Drbohlav, O., ed.: Computer Vision Winter Workshop 2003, Czech Pattern Recognition Society (2003) 61–66
5. Fletcher, A.K., Ramchandran, K., Goyal, V.K.: Wavelet denoising by recursive cycle spinning. In: Proc. IEEE Int. Conf. Image Proc. 2002. (2002)
6. Yu, T., Stoschek, A., Donoho, D.: Translation- and direction- invariant denoising of 2-D and 3-D images: Experience and algorithms. In Unser, M.A., Aldroubi, A., Laine, A.F., eds.: Wavelet Applications in Signal and Image Processing IV. Volume 2825 of SPIE. (1996) 608–619
7. Kingsbury, N.G.: Complex wavelets for shift invariant analysis and filtering of signals. Journal of Applied and Computational Harmonic Analysis **10** (2001) 234–253
8. Starck, J.L., Candès, E.J., Donoho, D.L.: The curvelet transform for image denoising. IEEE Transactions on Image Processing **11** (2002) 670–684
9. Mrázek, P., Weickert, J., Steidl, G.: Correspondences between wavelet shrinkage and nonlinear diffusion. In Griffin, L.D., ed.: Scale-Space Theories in Computer Vision. Lecture Notes in Computer Science. Springer, Berlin (2003)
10. Strang, G., Nguyen, T.: Wavelets and Filter Banks. Wellesley-Cambridge Press (1997)
11. Perona, P., Malik, J.: Scale space and edge detection using anisotropic diffusion. IEEE Transactions on Pattern Analysis and Machine Intelligence **12** (1990) 629–639
12. Keeling, S.L., Stollberger, R.: Nonlinear anisotropic diffusion filters for wide range edge sharpening. Inverse Problems **18** (2002) 175–190
13. Steidl, G., Weickert, J., Brox, T., Mrázek, P., Welk, M.: On the equivalence of soft wavelet shrinkage, total variation diffusion, total variation regularization, and SIDEs. Technical report, Series SPP-1114, Department of Mathematics, University of Bremen, Germany (2003)