Variational methods for denoising matrix fields

S. Setzer[1], G. Steidl[1], B. Popilka[1], and B. Burgeth[2]

[1] University of Mannheim,
Dept. of Mathematics and Computer Science
steidl@math.uni-mannheim.de

[2] Saarland University,
Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science
burgeth@mia.uni-saarland.de

Summary. The restoration of scalar-valued images via minimization of an energy functional is a well-established technique in image processing. Recently also higher-order methods have proved their advantages in edge preserving image denoising. In this paper, we transfer successful techniques like the minimization of the Rudin-Osher-Fatemi functional and the infimal convolution to matrix fields, where our functionals couple the different matrix channels. For the numerical computation we use second-order cone programming. Moreover, taking the operator structure of matrices into account, we consider a new operator-based regularization term. Using matrix differential calculus, we deduce the corresponding Euler-Lagrange equation and apply it for the numerical solution by a steepest decent method.

1 Introduction

Matrix-valued data have gained significant importance in recent years, e.g. in diffusion tensor magnetic resonance imaging (DT-MRI) and technical sciences (inertia, diffusion, stress, and permittivity tensors). As most measured data these matrix-valued data are also polluted by noise and require restoration. Regularization methods have been applied very successfully for denoising of scalar-valued images where recently higher-order methods, e.g. in connection with the infimal convolution [9] have provided impressive results. In this paper, we want to transfer these techniques to matrix fields. However, unlike vectors, matrices can be multiplied providing matrix-valued polynomials and also functions of matrices. These useful notions rely decisively on the strong interplay between the different matrix entries. Thus the corresponding regularization terms should take the relation between the different matrix channels into account.

Filtering methods for matrix fields based on matrix-valued nonlinear partial differential equations (PDEs) have been proposed in [6] for singular and in [7] for Perona-Malik-type diffusivity functions. These approaches rely on an operator-algebraic point of view on symmetric matrices as instances of self-adjoint Hilbert space operators, and are based on a basic differential calculus for matrix fields. Since the proposed techniques exploit the greater algebraic potential of matrices, if compared to vectors, they ensure appropriate matrix channel coupling, and more important, are also applicable to indefinite matrix fields.

Approaches to positive definite matrix field filtering with a differential geometric background have been suggested in [21, 10]. In their setting the set of positive definite matrices is endowed with a structure of a manifold, and the methodology is geared towards application to DT-MRI data. Comprehensive survey articles on the analysis of matrix fields utilizing a wide range of different techniques can be found in [25] and the literature cited therein.

This paper is organized as follows: In Section 2, we start by considering various variational methods for denoising scalar-valued images, in particular we adapt the infimal convolution technique to our discrete setting and introduce a corresponding simplified version. In Section 3, we turn to the matrix-valued setting. After giving the necessary preliminaries in Subsection 3.1, we consider component-based regularization terms related to the Rudin-Osher-Fatemi approach and to infimal convolution in Subsection 3.2. These functionals couple the different matrix channels as originally proposed by [22]. In Subsection 3.3 we introduce a new operator-based functional and derive the corresponding Euler-Lagrange equation which contains the Jordan product of matrices. In contrast to the ordinary matrix product the Jordan product two symmetric matrices is again a symmetric matrix. Finally, in Section 4, we present numerical examples comparing the component-based and the operator-based approach as well as first-order and infimal convolution methods.

2 Variational methods for scalar-valued images

First-order methods.

A well-established method for restoring a scalar-valued image u from a given degraded image f consists in calculating

$$\arg\min_u \int_\Omega (f-u)^2 \; + \; \alpha\,\Phi(|\nabla u|^2)\,dxdy \tag{1}$$

with a regularization parameter $\alpha > 0$ and an increasing function $\Phi : [0, \infty] \to \mathbb{R}$ in the penalizing term. The first summand encourages similarity between the restored image and the original one, while the second term rewards smoothness. The appropriate choice of the function Φ ensures that important image structures such as edges are preserved while areas with small gradients are smoothed.

The Euler-Lagrange equation of (1) is given by

$$0 = f - u + \alpha \operatorname{div}(\Phi'(|\nabla u|^2)\nabla u). \tag{2}$$

Thus, the minimizer u can be considered as the steady state ($t \to \infty$) of the reaction-diffusion equation

$$\partial_t u = f - u + \alpha \operatorname{div}(\Phi'(|\nabla u|^2)\nabla u) \tag{3}$$

with initial image $u(\cdot, 0) = f$ and homogeneous Neumann boundary conditions. For an interpretation of (2) as a fully implicit time discretization of a diffusion equation see [19, 23].

In this paper, we are mainly interested in the frequently applied ROF–model introduced by Rudin, Osher and Fatemi [18] which uses the function

$$\Phi(s^2) := \sqrt{s^2} = |s|. \tag{4}$$

For this function, the functional (1) is strictly convex and the penalizing functional $J(u) = \int_\Omega \sqrt{u_x^2 + u_y^2}\, dxdy$ is positively homogeneous, i.e. $J(\alpha u) = \alpha J(u)$ for $\alpha > 0$. Since Φ in (4) is not differentiable at zero, we have to use its modified version

$$\Phi(s^2) = \sqrt{s^2 + \varepsilon^2}, \tag{5}$$

with a small additional parameter ε if we want to apply (3).

For digital image processing, we consider a discrete version of (1). Let us introduce this discrete version in matrix-vector notation. For the sake of simplicity, we restrict our attention to quadratic images $\mathbf{f} \in \mathbb{R}^{n,n}$. We transform $\mathbf{f}$ into a vector $f \in \mathbb{R}^N$ with $N = n^2$ in the following way

$$\operatorname{vec} \mathbf{f} := \begin{pmatrix} f_0 \\ \vdots \\ f_{n-1} \end{pmatrix},$$

where f_j denotes the j-th column $\mathbf{f}$. The partial derivatives in (1) are discretized by forward differences. More precisely, we introduce the difference matrix $D_1 := \begin{pmatrix} D_x \\ D_y \end{pmatrix}$ with $D_x := I_n \otimes D$, $D_y := D \otimes I_n$ and

$$D := \begin{pmatrix} -1 & 1 & 0 & \ldots & 0 & 0 & 0 \\ 0 & -1 & 1 & \ldots & 0 & 0 & 0 \\ & & \ddots & \ddots & \ddots & & \\ 0 & 0 & 0 & \ldots & -1 & 1 & 0 \\ 0 & 0 & 0 & \ldots & 0 & -1 & 1 \\ 0 & 0 & 0 & \ldots & 0 & 0 & 0 \end{pmatrix}. \tag{6}$$

Here $A \otimes B$ is the *Kronecker product* (tensor product) of A and B. Now our discrete version of (1) reads

$$\operatorname*{arg\,min}_{u \in \mathbb{R}^N} \frac{1}{2} \|f - u\|_{\ell_2}^2 \; + \; \alpha \| \, |D_1 u| \, \|_{\ell_1}, \tag{7}$$

where $|D_1 u| \in \mathbb{R}^N$ is defined componentwise by $|D_1 u| = \left((D_x u)^2 + (D_y u)^2\right)^{1/2}$. For computations it is useful to consider the dual formulation of (7). Since we will need the dual form of various similar functionals later, let us consider more generally

$$\arg\min_{u \in \mathbb{R}^N} \frac{1}{2} ||f - u||_{\ell_2}^2 \; + \; \alpha \, || \, |Lu| \, ||_{\ell_1}, \tag{8}$$

with $L \in \mathbb{R}^{pN,N}$ and $|w|(j) = \left(\sum_{k=0}^{p-1} w(j+kN)^2\right)^{1/2}$, $j = 1, \ldots, N$. For $L := D_1$ the functional (8) coincides with (7). Since the penalizing functional $J(u) = || \, |Lu| \, ||_{\ell_1}$ is positively homogeneous, its Legendre–Fenchel dual J^* is the indicator function of the convex set

$$\mathcal{C} := \{v : \langle v, w \rangle \le J(w), \quad \forall w \in \mathbb{R}^N\}$$

and the minimizer $\hat{u}$ of (8) is given by $\hat{u} = f - \Pi_{\alpha\mathcal{C}} f$, where $\Pi_{\mathcal{C}}$ denotes the orthogonal projection of f onto $\mathcal{C}$. It can be proved, see, e.g. [], that

$$\mathcal{C} := \{L^{\mathrm{T}} V : \; \| \, |V| \, \|_\infty \le 1\} \tag{9}$$

and consequently $\hat{u} = f - L^{\mathrm{T}} V$, where V is a solution of

$$\|f - L^{\mathrm{T}} V\|_2^2 \quad \to \min \quad \text{s.t.} \quad \| \, |V| \, \|_{\ell_\infty} \le \alpha. \tag{10}$$

This minimisation problem can be solved by second-order cone programming (SOCP) [13], or alternatively by Chambolle's descent algorithm [8, 12].

Higher order methods.

For various denoising problems higher-order methods with functionals including higher-order derivatives have proved useful. In particular the drawback of so-called staircasing known from the ROF–model can be avoided in this way. An example is shown in Fig. 1.

Here we focus only on the infimal convolution method introduced by Chambolle and Lions [9]. We consider

$$\arg\min_{u \in \mathbb{R}^N} \frac{1}{2} ||f - u||_{\ell_2}^2 \; + \; (J_1 \Box J_2)(u), \tag{11}$$

where

$$(J_1 \Box J_2)(u) := \inf_{u_1, u_2 \in \mathbb{R}^N} \{J_1(u_1) + J_2(u_2) \; : \; u_1 + u_2 = u\}$$

denotes the so–called *infimal convolution* of J_1 and J_2. Note that the infimal convolution is closely related to the dilation operation in morphological image processing. In this paper, we will only deal with

$$J_1(u) := \alpha_1 \, \| \, |D_1 u| \, \|_{\ell_1}, \quad \text{and} \quad J_2(u) := \alpha_2 \, \| \, |D_2 u| \, \|_{\ell_1}, \tag{12}$$

where $D_2 := \begin{pmatrix} D_{xx} \\ D_{yy} \end{pmatrix} = \begin{pmatrix} I_n \otimes D^{\mathrm{T}} D \\ D^{\mathrm{T}} D \otimes I_n \end{pmatrix}$ is a discrete second-order derivative operator. Alternatively, we can also use $|(D_{xx}^{\mathrm{T}}, D_{yy}^{\mathrm{T}}, D_{xy}^{\mathrm{T}}, D_{yx}^{\mathrm{T}})^{\mathrm{T}}\, u|$ which is related to the Frobenius norm of the Hessian of u. It is easy to check that for (12) the infimum in (11) is attained, so that (11) can be rewritten as

$$\arg \min_{u_1, u_2 \in \mathbb{R}^N} \frac{1}{2} \| f - u_1 - u_2 \|_{\ell_2}^2 \; + \; \alpha_1 \| \, |D_1 u_1| \, \|_{\ell_1} + \alpha_2 \| \, |D_2 u_2| \, \|_{\ell_1}. \tag{13}$$

Using that $(J_1 \square J_2)^*(u) = J_1^*(u) + J_2^*(u)$ for proper convex functionals, the minimizer $\hat{u}$ of (11) is given by $\hat{u} = f - \Pi_{\alpha_1 \mathcal{C}_1 \cap \alpha_2 \mathcal{C}_2} f$, where J_i^* are the indicator functions of the convex sets $\mathcal{C}_i$ associated with J_i, $i = 1, 2$. Consequently, by (9), we obtain that $\hat{u} = f - v$, where v is the solution of

$$\begin{aligned} \|f - v\|_2^2 \quad \to \min \quad \text{s.t.} \quad & v = D_1^{\mathrm{T}} V_1 = D_2^{\mathrm{T}} V_2, \\ & \| \, |V_1| \, \|_\infty \le \alpha_1, \; \| \, |V_2| \, \|_\infty \le \alpha_2. \end{aligned} \tag{14}$$

Now we have that

$$D_2^{\mathrm{T}} = D_1^{\mathrm{T}} \begin{pmatrix} D_x & 0 \\ 0 & D_y \end{pmatrix}.$$

Consequently, assuming that $V_1 = \begin{pmatrix} D_x & 0 \\ 0 & D_y \end{pmatrix} V_2$, we may rewrite (14) as

$$\begin{aligned} \|f - D_2^{\mathrm{T}} V\|_2^2 \quad \to \min \quad \text{s.t.} \quad & \| \, \left| \begin{pmatrix} D_x V^1 \\ D_y V^2 \end{pmatrix} \right| \, \|_\infty \le \alpha_1, \\ & \| \, |V| \, \|_\infty \le \alpha_2. \end{aligned} \tag{15}$$

Note that this minimisation problem is similar but not equivalent to (14). The solution of (14), (15) or of the primal problem (13) can be computed by SOCP. In our numerical experiments we prefer the dual setting since it allows for a much faster computation with standard software for SOCP than the primal problem.

3 Variational methods for matrix-valued images

In this section, we want to transfer the regularization methods for scalar–valued images reviewed in the previous section to matrix–valued images. While it seem to be straightforward to replace the square of scalar values in the data fitting term by the squared Frobenius norm, the penalizing term may be established in different ways as we will see in the Subsections 3.2 and 3.3. First we need some notation.

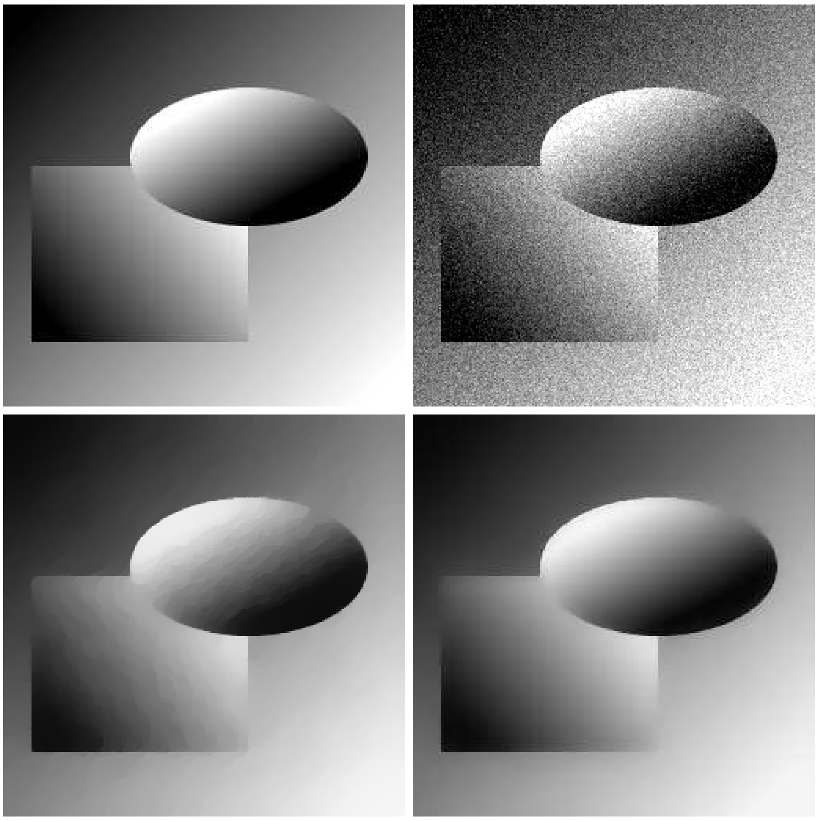

Fig. 1. Top: Original image (left), noisy image (right). **Bottom:** Denoised image by the ROF method (7) with $\alpha = 50$ and by the modified infimal convolution approach (15) with $\alpha_1 = 50$, $\alpha_2 = 180$ (right). The ROF based image shows the typical staircasing effect. The performance of the inf-conv method is remarkable.

3.1 Preliminaries.

Let $\mathrm{Sym}_m(\mathbb{R})$ be the vector space of symmetric $m \times m$ matrices. This space can be treated as a Euclidian vector space with respect to the trace inner product

$$\langle A, B \rangle := \mathrm{tr}\, AB = (\mathrm{vec}A, \mathrm{vec}B),$$

where $(\cdot, \cdot)$ on the right-hand side denotes the Euclidian inner vector product. Then

$$\langle A, A \rangle = \mathrm{tr}\, A^2 = \|A\|_F^2 = \|\mathrm{vec}A\|_{\ell_2^2}$$

is the squared *Frobenius norm* of A. In addition to this vector structure matrices are (realizations of) linear operators and carry the corresponding features. In particular they can be applied successively. Unfortunately, the original matrix multiplication does not preserve the symmetry of the matrices. The *Jordan–product* of matrices $A, B \in \mathrm{Sym}_m(\mathbb{R})$ defined by

$$A \bullet B := \frac{1}{2}(AB + BA)$$

preserves the symmetry of the matrices but not the positive semi-definiteness.

In $\mathrm{Sym}_m(\mathbb{R})$, the positive semi-definite matrices $\mathrm{Sym}^+_m(\mathbb{R})$ form a closed convex set whose interior consists of the positive definite matrices. More precisely, $\mathrm{Sym}^+_m(\mathbb{R})$ is a cone with base [1, 4, 5]. In our numerical examples we will only consider the cases $m = 2$ and $m = 3$ where the positive definite matrices can be visualized as ellipses, resp. ellipsoids, i.e. $A \in \mathrm{Sym}^+_3(\mathbb{R})$ can be visualized as the ellipsoid

$$\{x \in \mathbb{R}^3 : x^{\mathrm{T}} A^{-2} x = 1\}$$

whose axis lengths are given by the eigenvalues of A.

3.2 Component-based regularization

In the following, let $F : \mathbb{R}^2 \supset \Omega \to \mathrm{Sym}_m(\mathbb{R})$ be a matrix field. In this subsection, we transfer (1) to matrix-valued images in a way that emphasizes the individual matrix components. We will see that for this approach the denoising methods from the previous section can be translated in a direct way. However, the specific question arises whether these methods preserve positive definiteness.

Instead of (1) we are dealing with

$$\arg\min_U \int_\Omega \|F - U\|_F^2 + \alpha\, \Phi\left(\mathrm{tr}\,(U_x^2 + U_y^2)\right)\, dxdy, \tag{16}$$

where the partial derivatives are taken componentwise. The penalizing term $J(U)$ in (16) was first mentioned by Deriche and Tschumperlé [22]. Rewriting this term as

$$J(U) = \int_\Omega \Phi\big(\|U_x\|_F^2 + \|U_y\|_F^2\big)\, dxdy = \int_\Omega \Phi\big(\sum_{j,k=1}^n \nabla u_{jk}^{\mathrm{T}} \nabla u_{jk}\big)\, dxdy \tag{17}$$

we see its component-based structure implied by the Frobenius norm. However, due to the sum on the right–hand side, Φ is applied to coupled matrix coefficients and we should be careful here. By [3], the Euler–Lagrange equation of (17) is given by

$$0 = F - U + \alpha\left(\partial_x(\Phi'(\mathrm{tr}(U_x^2 + U_y^2))U_x + \partial_y(\Phi'(\mathrm{tr}(U_x^2 + U_y^2))U_y\right). \tag{18}$$

Again, we are only interested in the function Φ given by (4).

For computations we consider the discrete counterpart of (16), where we once more replace the derivative operators by simple forward difference operators

$$\arg\min_U \sum_{i,j=0}^{N-1} \frac{1}{2} \|F(i,j) - U(i,j)\|_F^2 \; + \; \alpha\, J(U), \tag{19}$$

$$J(U) := \sum_{i,j=0}^{N-1} \left(\|U(i,j) - U(i-1,j)\|_F^2 + \|U(i,j) - U(i,j-1)\|_F^2\right)^{1/2}$$

with $U(-1,j) = U(i,-1) = 0$. The functional in (19) is strictly convex and thus has a unique minimizer.

We say that the discrete matrix field $F : \mathbb{Z}_n^2 \to \mathrm{Sym}_m^+(\mathbb{R})$ has all eigenvalues in an interval $\mathcal{I}$ if all the eigenvalues of every matrix $F(i,j)$ of the field lie in $\mathcal{I}$. By the following proposition, the minimizer of (19) preserves positive definiteness.

Proposition 1. *Let all eigenvalues of* $F : \mathbb{Z}_n^2 \to \mathrm{Sym}_m^+(\mathbb{R})$ *be contained in the interval* $[\lambda_{\min}, \lambda_{\max}]$. *Then the minimizer* U *of* (19) *has all eigenvalues in* $[\lambda_{\min}, \lambda_{\max}]$.

The proof in the appendix is based on Courant's Min-Max principle and the projection theorem for convex sets.

To see how the methods from Section 2 carry over to matrix fields, we rewrite (19) in matrix-vector form. To this end, let $N = n^2$ and $M := m(m+1)/2$. We reshape $F : \mathbb{Z}_n^2 \to \mathrm{Sym}_m(\mathbb{R})$ into the vector

$$f := \begin{pmatrix} \varepsilon_{1,1} & \mathrm{vec}\,(F_{1,1}) \\ \vdots & \\ \varepsilon_{1,m} & \mathrm{vec}\,(F_{1,m}) \\ \varepsilon_{2,2} & \mathrm{vec}\,(F_{2,2}) \\ \vdots & \\ \varepsilon_{2,m} & \mathrm{vec}\,(F_{2,m}) \\ \vdots & \\ \varepsilon_{m,m} & \mathrm{vec}\,(F_{m,m}) \end{pmatrix} \in \mathbb{R}^{MN},$$

where $F_{k,l} := (F_{k,l}(i,j))_{i,j=0}^{n-1}$ and $\varepsilon_{k,l} := \begin{cases} \sqrt{2} & \text{for } k \neq l \\ 1 & \text{otherwise} \end{cases}$.

Then (19) becomes

$$\arg\min_{u \in \mathbb{R}^{MN}} \frac{1}{2} \|f - u\|_{\ell_2}^2 \; + \; \alpha \|\, |\, (I_M \otimes D_1)\, u|\, \|_{\ell_1}. \tag{20}$$

This problem has just the structure of (8) with $L := I_M \otimes D_1 \in \mathbb{R}^{2MN,MN}$ and $p = 2M$. Thus it can be solved by applying SOCP to its dual given by (10).

Similarly, we can transfer the infimal convolution approach to the matrix-valued setting. Obviously, we have to find

$$\arg \min_{u_1,u_2\in\mathbb{R}^{MN}} \frac{1}{2}\|f-u_1-u_2\|_{\ell_2}^2 + \alpha_1\|\,|(I_M\otimes D_1)\,u_1|\,\|_{\ell_1} + \alpha_2\|\,|(I_M\otimes D_2)\,u_1|\,\|_{\ell_1}.$$

In our numerical examples we solve the corresponding modified dual problem

$$\|f-(I_M\otimes D_2^{\mathrm{T}})\,V\|_2^2 \;\to \min \quad \text{s.t. } \|\,|\left(I_M\otimes\begin{pmatrix} D_x & 0\\ 0 & D_y\end{pmatrix}\right)V|\,\|_\infty \le \alpha_1, \quad \|\,|V|\,\|_\infty \le \alpha_2. \tag{21}$$

by SOCP.

3.3 Operator-based regularization

In this subsection, we introduce a regularizing term that emphasises the operator structure of matrices. For $A \in \mathrm{Sym}_m(\mathbb{R})$ with eigenvalue decomposition $A = Q\Lambda Q^{\mathrm{T}}$, let $\Phi(A) = Q\Phi(\Lambda)Q^{\mathrm{T}}$, where $\Lambda := \mathrm{diag}\,(\lambda_1,\ldots,\lambda_n)$ and $\Phi(\Lambda) := \mathrm{diag}\,(\Phi(\lambda_1),\ldots,\Phi(\lambda_n))$. We consider the following minimisation problem

$$\arg\min_U \int_\Omega \|F-U\|_F^2 + \alpha\,\mathrm{tr}\left(\Phi(U_x^2+U_y^2)\right)\,dxdy. \tag{22}$$

In contrast to (16) the trace is taken after applying Φ to the matrix $U_x^2+U_y^2$. By the next proposition we have that the functional in (22) with Φ definied by (4) is strictly convex:

Proposition 2. *For given $F : \mathbb{R}^2 \supset \Omega \to \mathrm{Sym}_m(\mathbb{R})$ and $\Phi(s^2) = \sqrt{s^2}$, the functional in* (22) *is strictly convex.*

The proof is given in the appendix.

An example in [20] shows that the solution of (22) does in general not preserve positive definiteness. The next proposition shows that the functional (22) has an interesting Gâteaux derivative.

Proposition 3. *Let Φ be a differentiable function. Then the Euler-Lagrange equations for minimizing the functional* (22) *are given by*

$$0 = F-U+\alpha\left(\partial_x\left(\Phi'(U_x^2+U_y^2)\bullet U_x\right)+\partial_y\left(\Phi'(U_x^2+U_y^2)\bullet U_y\right)\right). \tag{23}$$

The proof of the proposition is provided in the appendix and makes use of matrix differential calculus. In contrast to (18) the Jordan product of matrices appears in (23) and the function Φ' is applied to matrices.

We apply Proposition 3 to compute a minimizer of (22) by solving the corresponding reaction–diffusion equation for $t\to\infty$

$$U_t = F - U + \alpha \left(\partial_x \left(\Phi'(U_x^2 + U_y^2) \bullet U_x\right) + \partial_y \left(\Phi'(U_x^2 + U_y^2) \bullet U_y\right)\right) \tag{24}$$

with Φ as in (5), homogeneous Neumann boundary conditions and initial value F by a difference method. More precisely, we use the iterative scheme

$$U^{(k+1)} = (1-\tau)U^{(k)} + \tau F + \tau\alpha \left(\partial_x \left(G^{(k)} \bullet U_x^{(k)}\right) + \partial_y \left(G^{(k)} \bullet U_y^{(k)}\right)\right) \tag{25}$$

with sufficiently small time step size τ and $G^{(k)} := \Phi'((U_x^{(k)})^2 + (U_y^{(k)})^2)$. The inner derivatives including those in G were approximated by forward differences and the outer derivatives by backward differences so that the penalizing term becomes

$$\frac{1}{h_1}\left(G(i,j) \bullet \frac{U(i+1,j) - U(i,j)}{h_1} - G(i-1,j) \bullet \frac{U(i,j) - U(i-1,j)}{h_1}\right)$$
$$+ \frac{1}{h_2}\left(G(i,j) \bullet \frac{U(i,j+1) - U(i,j)}{h_2} - G(i,j-1) \bullet \frac{U(i,j) - U(i,j-1)}{h_2}\right),$$

where h_i, $i = 1,2$ denote the pixel distances in x and y–direction. Alternatively, we have also worked with symmetric differences for the derivatives. In this case we have to replace e.g. $G(i,j)$ in the first summand by $\tilde{G}(i+1,j) + \tilde{G}(i,j))/2$ and $\tilde{G}$ is now computed with symmetric differences.

Finally, we mention that a diffusion equation related to (24) was examined in [6]. Moreover, in [24] an anisotropic diffusion concept for matrix fields was presented where the function Φ was also applied to a matrix.

4 Numerical Results

Finally, we present numerical results demonstrating the performance of the different methods. All algorithms were implemented in MATLAB. Moreover, we have used the software package MOSEK for SOCP.

SOCP [15] amounts to minimize a linear objective function subject to the constraints that several affine functions of the variables have to lie in a *second-order cone* $\mathcal{C}^{n+1} \subset \mathbb{R}^{n+1}$ defined by the convex set

$$\mathcal{C}^{n+1} = \left\{ \begin{pmatrix} x \\ \bar{x}_{n+1} \end{pmatrix} = (x_1, \ldots, x_n, \bar{x}_{n+1})^{\mathrm{T}} : \|x\|_2 \le \bar{x}_{n+1} \right\}.$$

With this notation, the general form of a SOCP is given by

$$\inf_{x \in \mathbb{R}^n} f^{\mathrm{T}} x \quad \text{s.t.} \begin{pmatrix} A_i x + b_i \\ c_i^{\mathrm{T}} x + d_i \end{pmatrix} \in \mathcal{C}^{n+1}, \; i = 1, \ldots, r. \tag{26}$$

Alternatively, one can also use the rotated version of the standard cone:

$$\mathcal{K}^{n+2} := \left\{ \left(x, \bar{x}_{n+1}, \bar{x}_{n+2}\right)^{\mathrm{T}} \in \mathbb{R}^{n+2} : \|x\|_2^2 \le 2\, \bar{x}_{n+1} \bar{x}_{n+2} \right\}.$$

This allows us to incorporate quadratic constraints. Problem (26) is a convex program for which efficient, large scale solvers are available [17]. For rewriting our minimisation problems as a SOCP see [20].

We start by comparing the component-based regularization with the operator-based regularization. First we are interested in the 1D matrix–valued function $F : \mathbb{Z}_{16} \rightarrow \mathrm{Sym}_2^+(\mathbb{R})$ in Fig. 2. We added white Gaussian noise with standard deviation 0.1 to all components of the original data in $[0, 1]$. Then we computed the minimizer of the component-based functional (19) (left) and of the operator-based functional (**??**) (right) both by SOCP. The latter was computed using the fact that

$$\mathrm{tr}\,|U| = \max\{\big(4u_{12}^2 + (u_{11} - u_{22})^2\big)^{1/2}, |u_{11} + u_{22}|\}$$

for $U \in \mathrm{Sym}_2^+(\mathbb{R})$, cf. [20]. The middle of the figure shows the Frobenius norm of the difference between the original and the denoised signal $(\sum_{i=1}^N \|F(i) - \hat{U}(i)\|_F^2)^{1/2}$ in dependence on the regularization parameter α. We remark that the shape of the curve and its minimal point do not change if we use the error measure $\sum_{i=1}^N \|F(i) - \hat{U}(i)\|_F$ instead. The actual minima w.r.t. the Frobenius norm are given by min $= 0.2665$ at $\alpha = 0.8$ for (19) and by min $= 0.2276$ at $\alpha = 0.8$ for (**??**). The denoised signals corresponding to the smallest error in the Frobenius–norm are shown at the bottom of Fig. 2. It appears that the operator-based method performs slightly better w.r.t. these error norms. The visual results confirm this impression. The larger ellipses obtained by the first method (19) slightly overlap while there are gaps between the smaller ones. We do not have this effect for the minimizer of (**??**) on the right-hand side.

Next we consider the 2D matrix–valued function $F : \mathbb{Z}_{32}^2 \rightarrow \mathrm{Sym}_2^+(\mathbb{R})$ in Fig. 3. To all components of the original data in [0,2] we added white Gaussian noise with standard deviation 0.6. As in the previous example, we compare the minimizer of the component-based approach (16) resp. (19) with those of the operator-based approach (22). For computing the minimizer of the first functional we applied SOCP while the second one was computed via the reaction–diffusion equation (25) with time step size $\tau = 0.00025$. The iterations were stopped when the relative error in the ℓ_2-norm between two consecutive iterations became smaller than 10^{-8} (approximately 20000 iterations) although the result becomes visually static much earlier. The middle row of the figure contains the error plots for both methods. The actual minima w.r.t. the Frobenius norm are given by min $= 12.19$ at $\alpha = 1.75$ for (19) and by min $= 10.79$ at $\alpha = 1.2$ for (22). Hence, with respect to the computed errors the operator-based method outperforms the component-based one. The corresponding denoised images are shown in the bottom row of the figure.

In the following two examples, we consider bivariate matrix-valued functions which map to $\mathrm{Sym}_3(\mathbb{R})$. We use ellipsoids to visualize this form of data as described in Section 3.1. Furthermore, the color of the ellipsoid associated with a matrix A is chosen with respect to the normalized eigenvector corresponding to the largest eigenvalue of A. Fig. 4 shows a function $F : \mathbb{Z}_{12}^2 \rightarrow \mathrm{Sym}_3(\mathbb{R})$. As

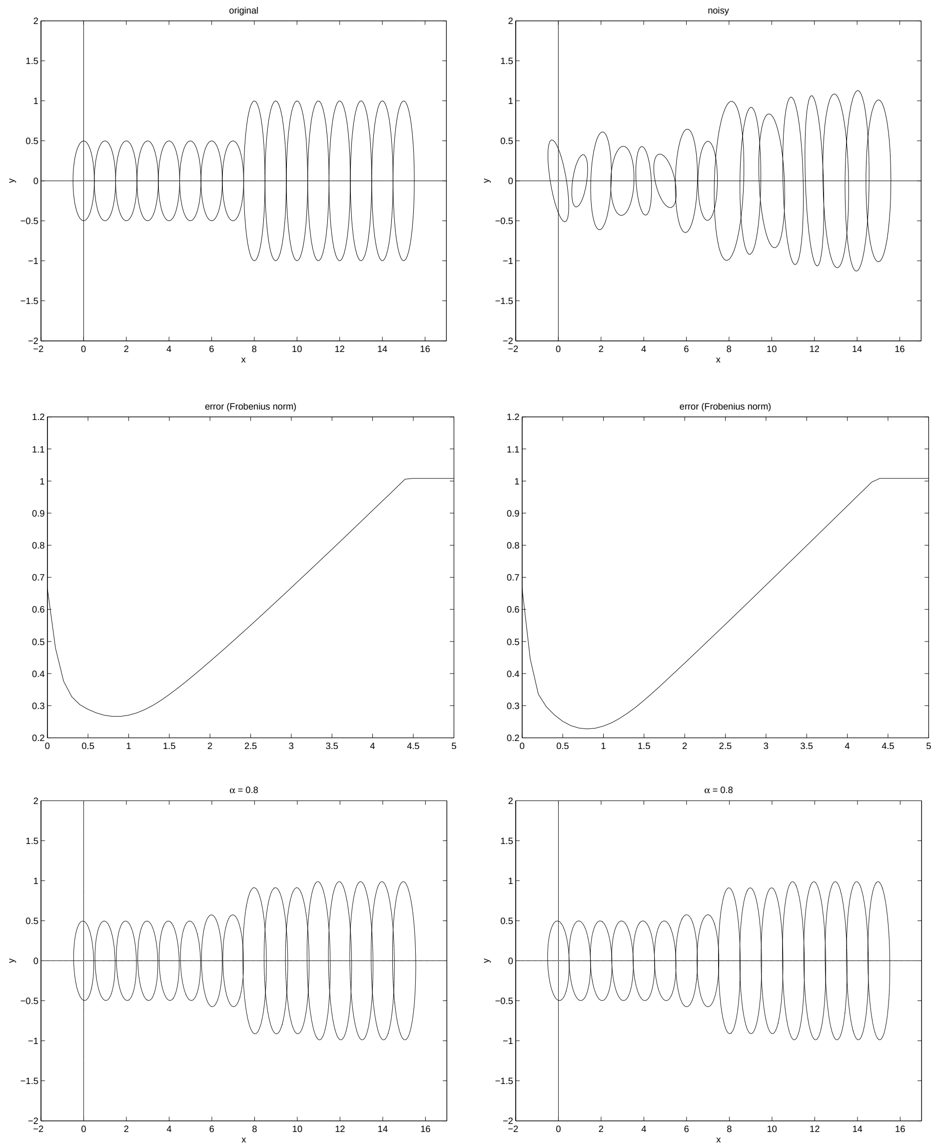

Fig. 2. Denoising of a matrix–valued signal. **Top**: Original signal (left), noisy signal (right). **Middle**: Error of the Frobenius norm in dependence on the regularization parameter α for the minimizers of the component-based functional (19) (left) and the operator-based functional (**??**) (right). **Bottom**: Denoised image for α corresponding to the smallest error in the Frobenius norm for the component-based functional (left) and the operator-based functional (right).

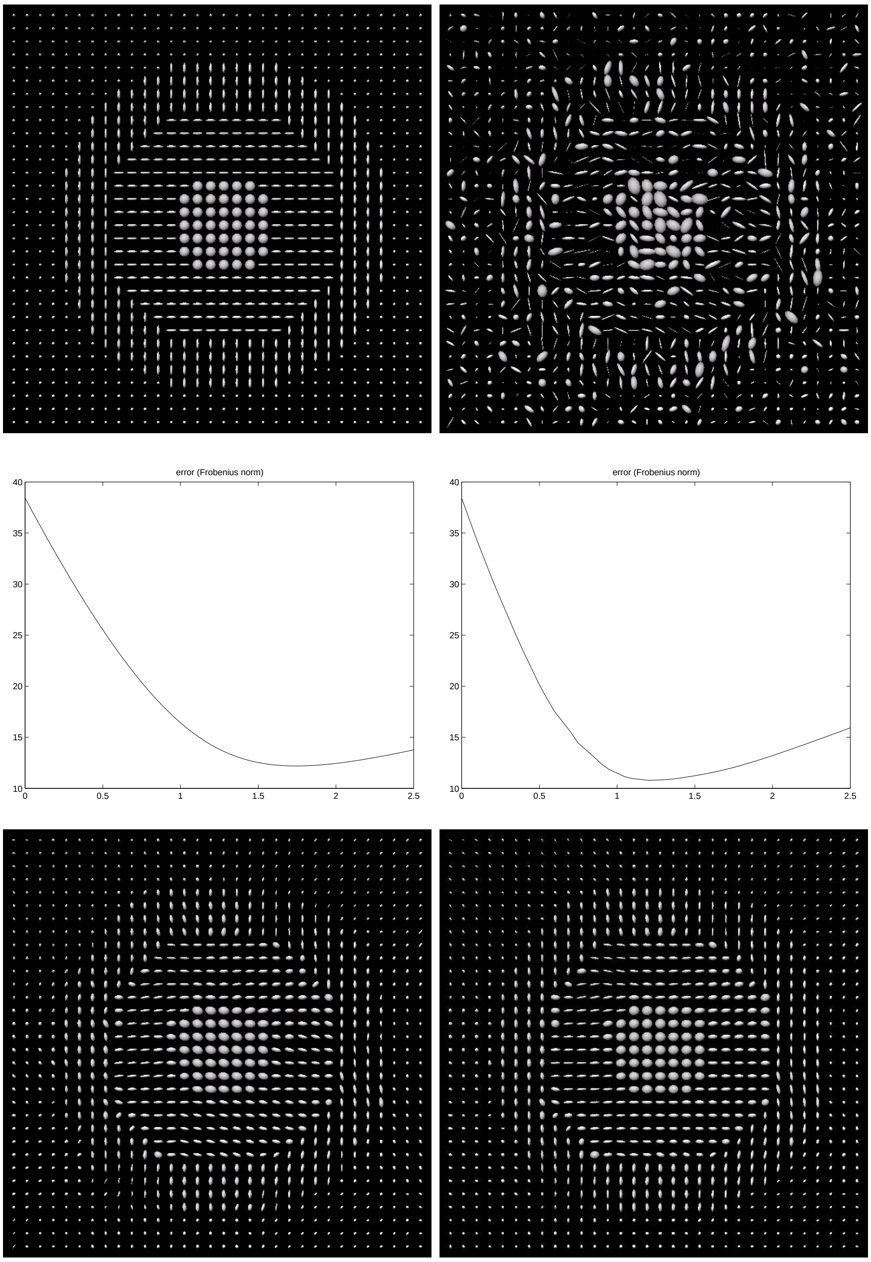

Fig. 3. Denoising of a $\mathrm{Sym}_2(\mathbb{R})$–valued image. **Top**: Original image (left), noisy image (right). **Middle**: Error of the Frobenius norm in dependence on the regularization parameter α for the minimizers of the component-based functional (19) (left) and the operator-based functional (22) (right). **Bottom**: Denoised image for α corresponding to the smallest error in the Frobenius norm for the component-based functional (left) and the operator-based functional (right).

before, we added white Gaussian noise to all components. The matrix components of the original data lie in the interval $[-0.5, 0.5]$ and the standard deviation of the Gaussian noise is 0.06. The denoising results are displayed in the last two rows of Fig. 4. We computed the minimizers of the component-based method (19) (top) by SOCP. The smallest error, measured in the Frobenius-norm, is 1.102 and was obtained for the regularization parameter $\alpha = 0.11$. In addition, we considered the minimizer of the infimal convolution approach (21) (bottom). Again we applied SOCP and found the optimal regularization parameters to be $\alpha_1 = 0.14$ and $\alpha_2 = 0.08$ for this method. The corresponding Frobenius-norm error is 0.918. We see that the infimal convolution approach is also suited for matrix-valued data.

In our final experiment, we applied the two component-based methods (19) and (21) to a larger data set. Fig. 5 shows the orginal data and the minimizers of (19) and (21). The components of the original data lie in $[-4000, 7000]$ and we used the regularization parameters $\alpha = 600$ for (19) and $\alpha_1 = 500$, $\alpha_2 = 600$ for (21), respectively.

A Proofs

Proof of Proposition 1. Using that the minimal and maximal eigenvalues $\lambda_{\min}(A), \lambda_{\max}(A)$ of a symmetric matrix A fulfill

$$\lambda_{\min}(A) = \min_{\|v\|=1} v^{\mathrm{T}} A\, v, \quad \lambda_{\max}(A) = \max_{\|v\|=1} v^{\mathrm{T}} A\, v,$$

it is easy to check that the set $\mathcal{C}$ of matrices having all eigenvalues in $[\lambda_{\min}, \lambda_{\max}]$ is convex and closed. Let $\mathcal{J}$ be the functional in (19). Assume that some matrices $\hat{U}(i,j)$ are not contained in $\mathcal{C}$. Let $P\hat{U}(i,j)$ denote the orthogonal projection (w.r.t. the Frobenius norm) of $\hat{U}(i,j)$ onto $\mathcal{C}$. Then we obtain by the projection theorem [11, p. 269] that

$$\begin{aligned} \|F(i,j) - P\hat{U}(i,j)\|_F &\le \|F(i,j) - \hat{U}(i,j)\|_F, \\ \|P\hat{U}(i,j) - P\hat{U}(k,l)\|_F &\le \|\hat{U}(i,j) - \hat{U}(k,l)\|_F. \end{aligned}$$

Consequently, $\mathcal{J}(P\hat{U}) \le \mathcal{J}(\hat{U})$ which contradicts our assumption since the minimizer is unique. This completes the proof. $\square$

Proof of Proposition 2. Since $\|F - U\|_F^2$ is strictly convex, it remains to show that the functional

$$J(U) := \mathrm{tr}\left(\sqrt{U_x^2 + U_y^2}\right)$$

is convex. Moreover, since J is positively homogeneous we only have to prove that J is subadditive, cf. [2, p. 34], i.e.,

$$J(\tilde{U} + U) \le J(\tilde{U}) + J(U).$$

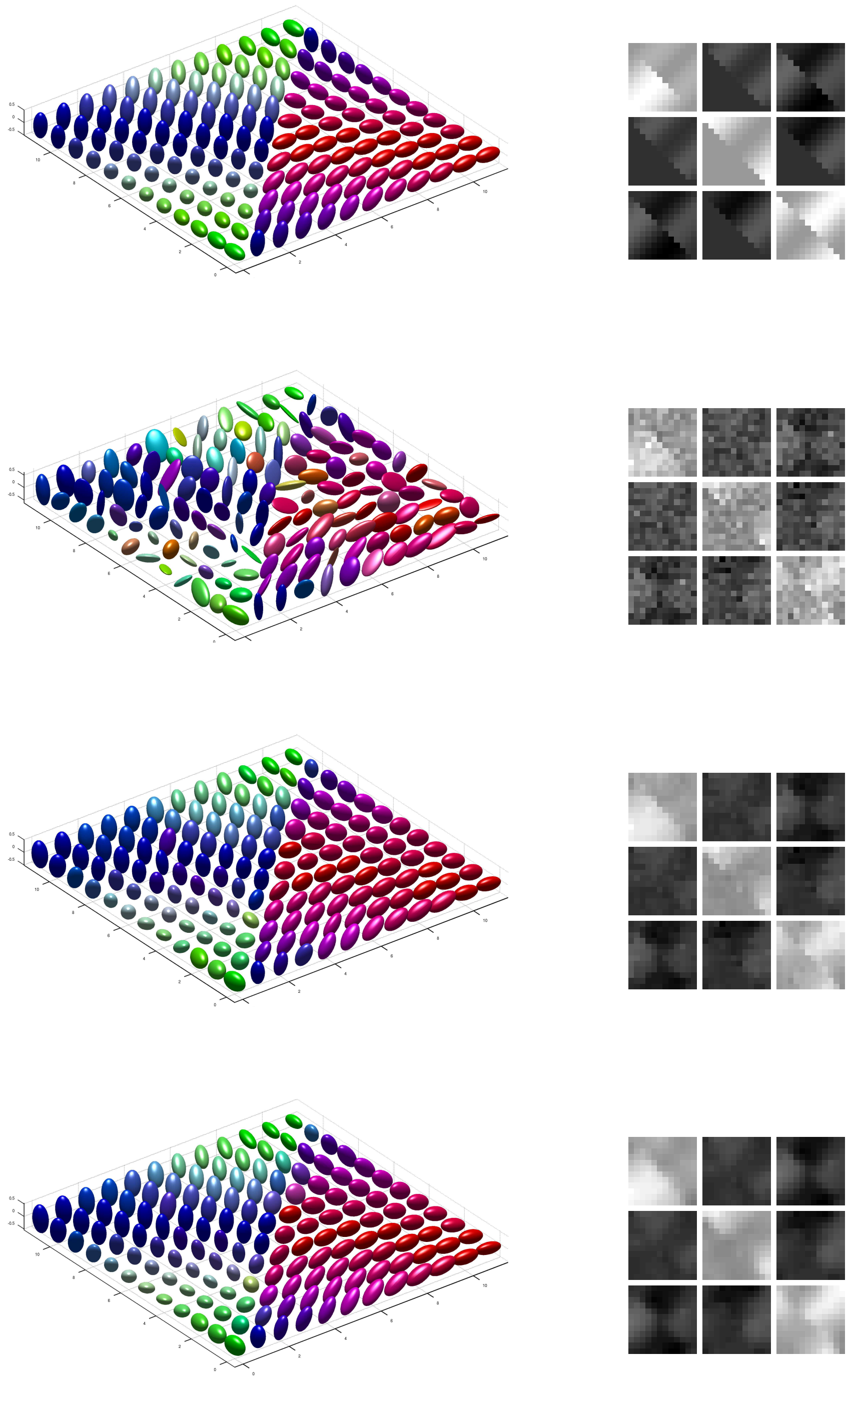

Fig. 4. Denoising of a $\mathrm{Sym}_3(\mathbb{R})$-valued image. **Top to Bottom**: Original image, noisy image, minimizer of the component-based method (19) for $\alpha = 0.11$, minimizer of the component-based infimal convolution approach (21) with parameters $\alpha_1 = 0.14$, $\alpha_2 = 0.08$. Visualization: ellipsoids (left), components of the matrix-valued data (right).

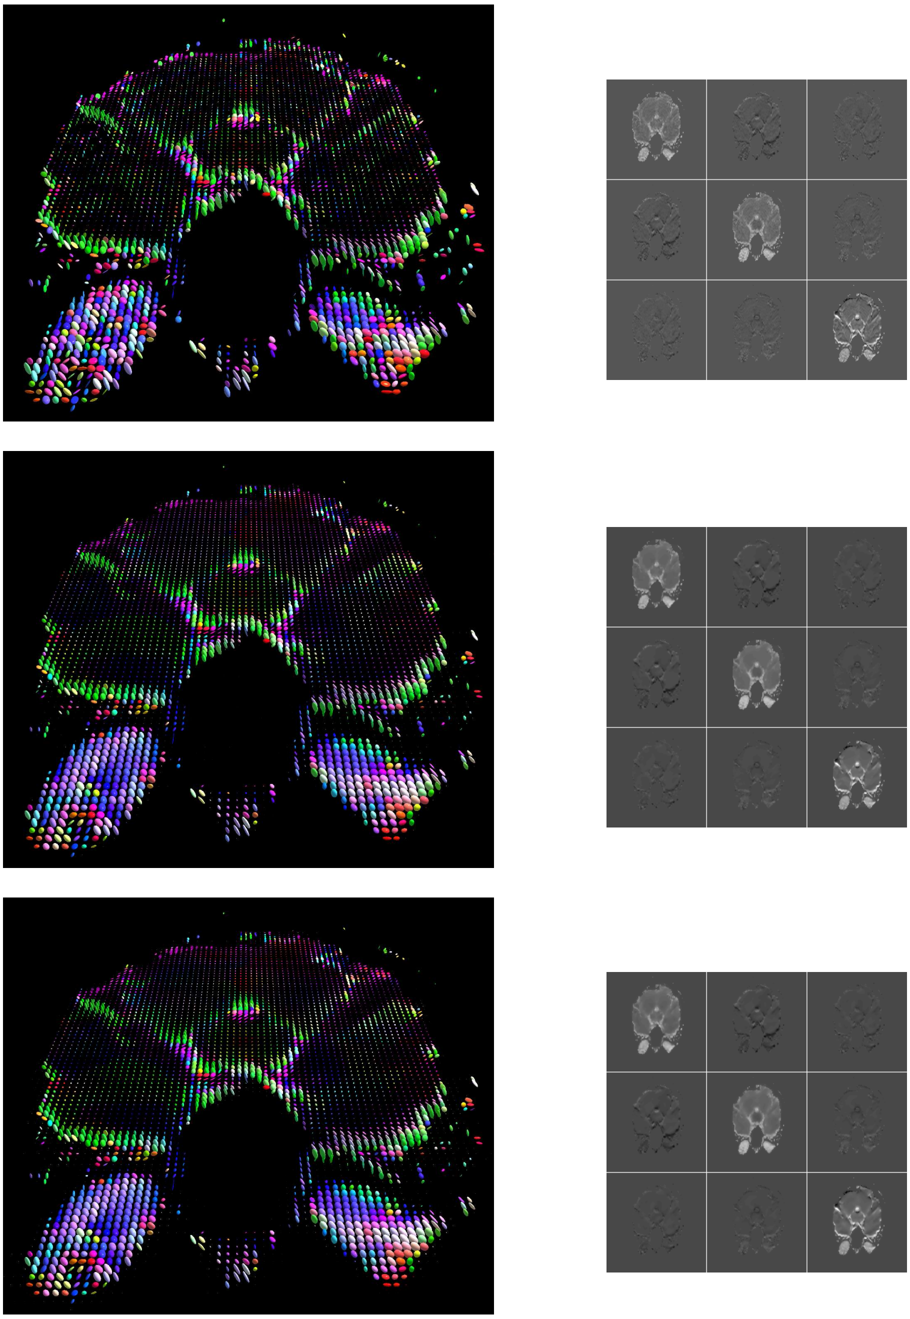

Fig. 5. Denoising of a real-world DT-MRI matrix field with values in $\mathrm{Sym}_3(\mathbb{R})$. **Top**: Original image. **Middle**: Minimizer of the component-based method (19) for $\alpha = 600$. **Bottom**: Minimizer of the infimal convolution approach (21) for $\alpha_1 = 500$, $\alpha_2 = 600$.

This can be rewritten as

$$\operatorname{tr}\left(\sqrt{(\tilde U_x+U_x)^2+(\tilde U_y+U_y)^2}\right)\le \operatorname{tr}\left(\sqrt{\tilde U_x^2+\tilde U_y^2}\right)+\operatorname{tr}\left(\sqrt{U_x^2+U_y^2}\right).$$

To prove this relation, we recall the definition of the *trace norm*, cf. [14, p. 197], which is defined as the sum of the singular values of a matrix $A\in\mathbb{R}^{s,t}$:

$$\|A\|_{\mathrm{tr}}=\operatorname{tr}(\sqrt{A^*A}).$$

Then we have for the symmetric matrices $\tilde U_x,\tilde U_y,U_x,U_y$ that

$$\|\begin{pmatrix}\tilde U_x+U_x\\ \tilde U_y+U_y\end{pmatrix}\|_{\mathrm{tr}}=\operatorname{tr}\left(\sqrt{(\tilde U_x+U_x)^2+(\tilde U_y+U_y)^2}\right)$$

Since $\|\cdot\|_{\mathrm{tr}}$ is a norm it follows that

$$\begin{aligned}\|\begin{pmatrix}\tilde U_x+U_x\\ \tilde U_y+U_y\end{pmatrix}\|_{\mathrm{tr}}&\le\|\begin{pmatrix}\tilde U_x\\ \tilde U_y\end{pmatrix}\|_{\mathrm{tr}}+\|\begin{pmatrix}U_x\\ U_y\end{pmatrix}\|_{\mathrm{tr}}\\ &=\operatorname{tr}(\sqrt{\tilde U_x^2+\tilde U_y^2})+\operatorname{tr}(\sqrt{U_x^2+U_y^2})\end{aligned}$$

and we are done. □

Proof of Proposition 3. Let $\varphi(U_x,U_y):=\operatorname{tr}\left(\Phi(U_x^2+U_y^2)\right)$. The Euler-Lagrange equations of (22) are given, for $i,j=1,...,n$ and $i\ge j$, by

$$0=\frac{\partial}{\partial u_{ij}}\|F-U\|_F^2-\alpha\left(\frac{\partial}{\partial x}\left(\frac{\partial\varphi}{\partial u_{ijx}}\right)+\frac{\partial}{\partial y}\left(\frac{\partial\varphi}{\partial u_{ijy}}\right)\right).$$

For a scalar-valued function f and an $n\times n$ matrix X, we set $\frac{\partial f(X)}{\partial X}:=\left(\frac{\partial f(X)}{\partial x_{ij}}\right)_{i,j=1}^n$. Then, by symmetry of F and U, the Euler-Lagrange equations can be rewritten in matrix-vector form as

$$W_n\circ\frac{U-F}{\alpha}=\frac{1}{2}\left(\frac{\partial}{\partial x}\left(\frac{\partial\varphi}{\partial U_x}\right)+\frac{\partial}{\partial y}\left(\frac{\partial\varphi}{\partial U_y}\right)\right),\qquad(27)$$

where W_n denotes the $n\times n$ matrix with diagonal entries 1 and other coefficients 2, and $A\circ B$ stands for the *Hadamard product* (componentwise product) of A and B.

We consider $f(X):=\operatorname{tr}\Phi(X^2)$. Then we obtain by [16, p. 178] and $\operatorname{tr}(A^{\mathrm T}B)=(\operatorname{vec}A)^{\mathrm T}\operatorname{vec}B$ that

$$\begin{aligned}\operatorname{vec}\frac{\partial f(X)}{\partial X}&=\operatorname{vec}\left(\operatorname{tr}(\Phi'(X^2)\,\frac{\partial(X^2)}{\partial x_{ij}})\right)_{i,j=1}^n\\ &=\operatorname{vec}\left((\operatorname{vec}\Psi)^{\mathrm T}\operatorname{vec}\frac{\partial(X^2)}{\partial x_{ij}}\right)_{i,j=1}^n\end{aligned}$$

where $\Psi := \Phi'(X^2)$. By [16, p. 182] and since Ψ is symmetric this can be rewritten as

$$\mathrm{vec}\frac{\partial f(X)}{\partial X} = \mathrm{vec}\, W_n \circ ((I_n \otimes X) + (X \otimes I_n))\,\mathrm{vec}\,\Psi.$$

Using that $\mathrm{vec}(ABC) = (C^{\mathrm{T}} \otimes A)\mathrm{vec}B$ we infer that

$$\mathrm{vec}\frac{\partial f(X)}{\partial X} = \mathrm{vec}W_n \circ \mathrm{vec}(X\Psi + \Psi X).$$

This implies that

$$\frac{\partial f(X)}{\partial X} = 2\, W_n \circ (\Psi \bullet X). \tag{28}$$

Applying (28) with $f(U_x) := \varphi(U_x, U_y)$ and $f(U_y) := \varphi(U_x, U_y)$, respectively, in (27) we obtain the assertion. □

Acknowledgements. The authors like to thank J. Weickert and S. Didas for fruitful discussions.

References

1. A. Barvinok. *A Course in Convexity, Graduate Studies in Mathematics*. AMS, Providence, RI, 2002.
2. J. M. Borwein and A. S. Lewis. *Convex Analysis and Nonlinear Optimization*. Springer, New York, 2000.
3. T. Brox, J. Weickert, B. Burgeth, and P. Mrázek. Nonlinear structure tensors. *Image and Vision Computing*, 24(1):41–55, 2006.
4. B. Burgeth, A. Bruhn, S. Didas, J. Weickert, and M. Welk. Morphology for matrix-data: Ordering versus PDE-based approach. *Image and Vision Computing*, 25(4):496–511, 2007.
5. B. Burgeth, A. Bruhn, N. Papenberg, M. Welk, and J. Weickert. Mathematical morphology for matrix fields induced by the Loewner ordering in higher dimensions. *Signal Processing*, 87(2):277–290, 2007.
6. B. Burgeth, S. Didas, L. Florack, and J. Weickert. Singular PDEs for the processing of matrix-valued data. Lecture Notes in Computer Science. Springer, Berlin, 2007. To appear.
7. B. Burgeth, S. Didas, and J. Weickert. A generic approach to diffusion filtering of matrix-fields. *Computing*. Submitted.
8. A. Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, (20):89–97, 2004.
9. A. Chambolle and P.-L. Lions. Image recovery via total variation minimization and related problems. *Numerische Mathematik*, 76:167–188, 1997.
10. C. Chefd'Hotel, D. Tschumperlé, R. Deriche, and O. Faugeras. Constrained flows of matrix-valued functions: Application to diffusion tensor regularization. In A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, editors, *Computer Vision – ECCV 2002*, volume 2350 of *Lecture Notes in Computer Science*, pages 251–265. Springer, Berlin, 2002.

11. P. G. Ciarlet. *Introduction to Numerical Linear Algebra and Optimisation*. Cambridge University Press, Cambridge, 1989.
12. S. Didas, G. Steidl, and S. Setzer. Combined ℓ_2 data and gradient fitting in conjunction with ℓ_1 regularization. *Adv. Comput. Math.*, 2007, accepted.
13. D. Goldfarb and W. Yin. Second-order cone programming methods for total variation-based image restoration. *SIAM J. Scientific Computing*, 2(27):622–645, 2005.
14. R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, Cambridge, U.K., 1991.
15. S. B. M. S. Lobo, L. Vandenberghe and H. Lebret. Applications of second-order cone programming. *Linear Algebra and its Applications*, 1998.
16. J. R. Magnus and H. Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. J. Wiley and Sons, Chichester, 1988.
17. H. Mittelmann. An independent bechmarking of SDP and SOCP solvers. *Mathematical Programming Series B*, 95(2):407–430, 2003.
18. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica A*, 60:259–268, 1992.
19. O. Scherzer and J. Weickert. Relations between regularization and diffusion filtering. *Journal of Mathematical Imaging and Vision*, 12(1):43–63, Feb. 2000.
20. G. Steidl, S. Setzer, B. Popilka, and B. Burgeth. Restoration of matrix fields by second-order cone programming. *Computing*, 2007, submitted.
21. D. Tschumperlé and R. Deriche. Diffusion tensor regularization with constraints preservation. In *Proc. 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 948–953, Kauai, HI, Dec. 2001. IEEE Computer Society Press.
22. D. Tschumperlé and R. Deriche. Diffusion tensor regularization with contraints preservation. In *Proc. 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 948–953, Kauai, HI, 2001. IEEE Computer Science Press.
23. J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart, 1998.
24. J. Weickert and T. Brox. Diffusion and regularization of vector- and matrix-valued images. In M. Z. Nashed and O. Scherzer, editors, *Inverse Problems, Image Analysis, and Medical Imaging*, volume 313 of *Contemporary Mathematics*, pages 251–268. AMS, Providence, 2002.
25. J. Weickert and H. Hagen, editors. *Visualization and Processing of Tensor Fields*. Springer, Berlin, 2006.