

Diffusion tensor magnetic resonance image regularization

O. Coulon^{a,*}, D.C. Alexander^b, S. Arridge^b

^a Centre National de la Recherche Scientifique, Laboratoire des Sciences de l'Information et des Systèmes, Marseille, France

^b Department of Computer Science, University College London, London, UK

Received 20 February 2003; received in revised form 11 June 2003; accepted 23 June 2003

Abstract

As multi-dimensional complex data become more common, new regularization schemes tailored to those data are needed. In this paper we present a scheme for regularising diffusion tensor magnetic resonance (DT-MR) data, and more generally multi-dimensional data defined by a direction map and one or several magnitude maps. The scheme is divided in two steps. First, a variational method is proposed to restore direction fields with preservation of discontinuities. Its theoretical aspects are presented, as well as its application to the direction field that defines the main orientation of the diffusion tensors. The second step makes use of an anisotropic diffusion process to regularize the magnitude maps. The main idea is that for a range of data it is possible to use the restored direction as a prior to drive the regularization process in a way that preserves discontinuities and respects the local coherence of the magnitude map. We show that anisotropic diffusion is a convenient framework to implement that idea, and define a regularization process for the magnitude maps from our DT-MR data. Both steps are illustrated on synthetic and real diffusion tensor magnetic resonance data.

© 2003 Elsevier B.V. All rights reserved.

Keywords: Diffusion tensor magnetic resonance images; Regularization; Tensor field; Direction field; Variational methods; Anisotropic diffusion

1. Introduction

With the development of new imaging modalities, image processing algorithms have to deal with images of various types and increasing complexity and dimensionality. The information provided is not restricted to a scalar value but can be multi-dimensional and constrained with various conditions. The field of medical imaging, in particular, has seen the appearance of various new multi-dimensional magnetic resonance (MR) acquisition techniques, a recent example of which is *diffusion tensor MR imaging* (DT-MRI). DT-MR images contain measurements of the diffusion properties of water molecules within tissues (Basser et al., 1994). These measurements provide information about the structure and physiology of tissue. A DT-MR 3D volume contains at each voxel a diffusion tensor (DT)

represented by a 3×3 symmetrical positive-definite matrix. The DT expresses a Gaussian model of the water diffusion process. It contains directional information, due to the anisotropy of the diffusion, as well as magnitude information such as the diffusivity.

Recent interest has been paid in the image processing and computer vision community to complex multi-dimensional data: multi-valued images (Whitaker and Gerig, 1994), colour images (Sapiro and Ringach, 1996; Trahanias et al., 1996; Tang et al., 1999), vector fields with various constraints (Tschumperlé and Deriche, 2001a,b) direction fields (Perona, 1998; Chan and Shen, 2000; Tang et al., 2000), or even images defined on surfaces (Kimmel, 2001; Bertalmio et al., 2000). Traditional scalar image processing techniques are often inappropriate given the multi-dimensional nature of these data and further difficulties often arise due to the constraints imposed by the space in which the data are defined. In some cases, these data can be separated into two types of information of different nature, such as direction and magnitude. Colour images for example can be represented by chromacity (directional information) and

* Corresponding author. Present address: Laboratoire LSIS, équipe LXAO, ESIL, Campus de Luminy, case 925, 13288 Marseille Cedex 09, France. Tel.: +33-4-91-82-85-27; fax: +33-4-91-82-85-51.

E-mail address: olivier.coulon@lsis.org (O. Coulon).

brightness (scalar magnitude), and methods have been proposed that process those two types of information separately (e.g. Chan and Shen, 2000; Tang et al., 2000), taking into account the specific nature of each. Nevertheless, for colour images, the direction field that defines the chromaticity and the brightness map are not directly related and can be processed independently. Other type of data exists for which the direction field and the local structure of the intensity map can be correlated. For instance with MR angiography images, the blood flow magnitude is expected to have some degree of smoothness along the flow direction. Therefore, the blood flow direction map can be used as a prior to describe the structure of the blood flow magnitude image.

The purpose of this paper is to propose a regularization method for medical data containing both direction and intensity information, and for which the directional information and the local structure of the magnitude image are related. Examples of these kinds of data are optic flow images, MRA volumes, MR flow imaging, strain tensor images, or DT-MR images. The method presented extends earlier work reported in (Coulon et al., 2001). The general idea of the scheme that we propose to regularize such data relies on two steps. First, the direction field should be restored, and we present a general method to do that. Second, the restored direction field should be used as a prior to drive the magnitude image regularization, and we show how Weickert's tensor-driven anisotropic diffusion scheme (Weickert, 1998) can be used for that purpose. The paper is organized as follows. In Section 2, we present some concepts needed to understand the complex nature of DT-MR images, as well as some background ideas related to direction field restoration and anisotropic diffusion process. A general discrete variational method for restoring direction fields is proposed in Section 3.1. Then in Section 3.2, we show how the restored direction field can be used as a prior to drive the intensity map regularization via an anisotropic diffusion process. Both steps are illustrated on DT-MR data and results are presented on synthetic and real data in Section 4.

2. Background

2.1. DT-MR images

In the following, we introduce a few concepts about DT-MR images. This acquisition technique has become increasingly popular over the past few years and has a wide range of applications from clinical diagnosis, e.g., in case of stroke, to research applications, such as brain connectivity studies (Lebihan et al., 2001). A DT-MR acquisition sequence consists of the acquisition of a number of 3D *diffusion weighted* (DW) images, each of which measures the amount of water diffusion in a

particular direction. From these DW images, a Gaussian model of the diffusion as a function of direction is fitted at every voxel. This model is expressed by a diffusion tensor, represented by a 3×3 symmetrical positive-definite matrix $\mathbf{D} = (D_{ij})$. Example images of the 6 coefficients of the matrix are shown in Fig. 1. We will use $(\lambda_i)_{i=1,2,3}$ to denote its eigenvalues, with $\lambda_1 \geq \lambda_2 \geq \lambda_3$, and $(\mathbf{v}_i)_{i=1,2,3}$ the associated eigenvectors. Therefore,

$$\mathbf{D} = (\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3) \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} (\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3)^T. \quad (1)$$

Various scalar measurements can be derived from the DT-MR images to describe tissue microstructure at each voxel (see for instance (Basser and Pierpaoli, 1996; Lebihan et al., 2001)). Amongst those measurements we mention:

- The *mean diffusivity*, which characterizes the overall amount of diffusion:

$$\langle \mathbf{D} \rangle = \frac{\text{Trace}(\mathbf{D})}{3} = \frac{\sum_i \lambda_i}{3}. \quad (2)$$

It has a low value in tissues, such as grey matter (GM) and white matter (WM), in which dense tissue structure hinders diffusion, and a high value in cerebro-spinal fluid (CSF) where water is free to diffuse unhindered.

- The *fractional anisotropy* describes the directional bias and is related to the presence of oriented structures:

$$\text{FA}(\mathbf{D}) = \left(\frac{\sum_i (\lambda_i - \langle \mathbf{D} \rangle)^2}{\sum_i \lambda_i^2} \right)^{1/2}. \quad (3)$$

Fractional anisotropy varies between 0 and 1. Low values indicate media with no preferred orientation, such as GM or CSF, whereas higher values indicate tissues with strong orientation, such as WM and some internal GM structures (thalamus), in which axonal fibers constitute an oriented medium for diffusion. There is a range of shapes that anisotropic tensors can have, which lies between two extremes: *prolate* tensors, which have one large eigenvalue and two small ones, and *oblate* tensors, which have two large eigenvalues and one small one.

The first eigenvector \mathbf{v}_1 , called *principal diffusion direction* (PDD), plays a particular role in the eigensystem since in white matter it has been shown to coincide with the local orientation of fibers, therefore being a determinant clue for studying brain connectivity (white matter fiber tracking, or *tractography*, see for example (Conturo et al., 1999)). In WM, the PDD field has a very strong coherence and defines the main orientation of the tensor. In more isotropic regions, the PDD field loses its coherence and becomes meaningless with an almost random orientation, as the anisotropy decreases (Fig. 2).

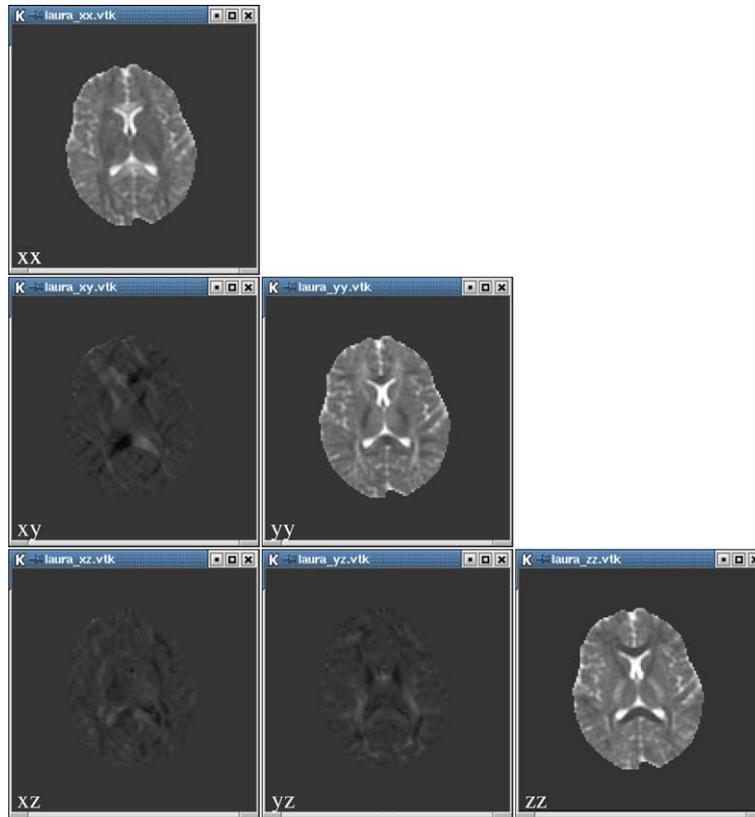


Fig. 1. A slice of the 6 coefficients that define the tensor. (This figure is available in colour, see the on-line version.)

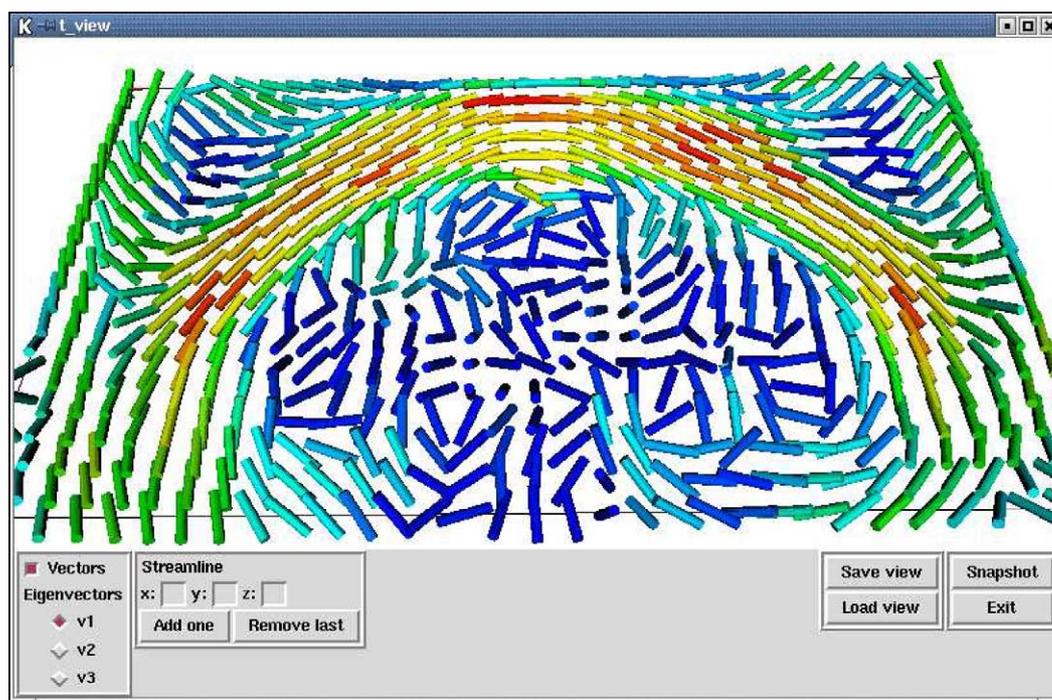


Fig. 2. PDD field, colour-mapped by anisotropy: blue represents low anisotropy, red represents high anisotropy. As the anisotropy decreases, the PDD field loses its coherence (real brain data: splenium of the corpus callosum). (This figure is available in colour, see the on-line version.)

Scalar measurements such as those in Eqs. (2) and (3), and directional information such as the PDD describe complementary aspects of the tissue microstructure and architecture at each voxel (Basser and Pierpaoli, 1996; Lebihan et al., 2001). A regularization method must aim at restoring these properties and reduce their sensitivity to noise.

A few various DT-MR image regularization can be found in the literature. They aim to regularize the whole tensor (Aldroubi and Basser, 1999), the PDD (Poupon et al., 2001), the orthonormal basis of eigenvectors (Tschumperlé and Deriche, 2001b), or the diffusion weighted images before estimation of the tensor (Parker et al., 2000; Vemuri et al., 2001), but none of them explicitly uses the structural information given by the tensor. The method we present in this paper aims to both restore the orientation of the tensor and regularize the diffusivities while taking into account the local structure of the tensor.

2.2. Direction field restoration

Direction regularization has been a recent field of interest as the problems of dealing with images containing directional information have arisen. Examples include chromaticity from colour images (Tang et al., 2000), optic flow (Weickert and Schnörr, 2001), fingerprints orientation images (Perona, 1998), PDD fields from DT-MR images (Coulon et al., 2001; Poupon et al., 2001), or direction fields from MR velocity imaging (Mohiaddin et al., 1994). A direction field is a vector field with the constraint that the vectors have a unit norm. Therefore, directions live (for the three-dimensional case) on the unit sphere S^2 , a non-linear manifold. Features living on such non-linear manifolds are said to be *non-flat* and restoration of this kind of feature has been the subject of several recent works (Chan and Shen, 2000; Tang et al., 2000; Trahanias et al., 1996; Perona, 1998).

Early work on direction field regularization was presented by Perona (1998), who proposed an angular diffusion model to regularize orientation images, using a parameterization of the unit circle. Both continuous and discrete diffusion equations are proposed with the choice of a proper metric on the unit circle. More general work was then proposed, aiming to regularize data defined between general manifolds. Tang et al. (2000) proposed a direction diffusion model using the harmonic maps framework that can be applied to direction fields in any dimension. Chan and Shen (2000) found equivalent results, using a variational model based on the \mathcal{L}_1 - or \mathcal{L}_2 -norms. In particular, the \mathcal{L}_1 model defines the *total variation* (TV) energy, whose minimization provides a discontinuity preserving regularization. Both continuous and discrete models are proposed. An extension of Chan and Shen work was presented in (Coulon et al., 2001) and applied to DT-MR data. A geometric ap-

proach to norm constraint is proposed in (Tschumperlé and Deriche, 2001a), with the integration of a constant norm constraint in the PDE that defines the regularization process. This results in a projection of the diffusion flow in the plane orthogonal to the considered vector and allows for the use of various diffusion equations in this constrained framework. As we show later, the same result is obtained by Chan and Shen in some particular cases of their variational approach. All three approaches (Tang et al., 2000; Chan and Shen, 2000; Tschumperlé and Deriche, 2001a) have in common the use of external coordinates, after embedding the considered manifold (S^2) in a Euclidean space (\mathbb{R}^3), associated with a projection term. Kimmel and Sochen (2000), on the other hand, use a local coordinate system and compute the associated Beltrami flow to perform a regularization of directions with preservation of discontinuities. Finally, let us mention the work of Poupon et al. (2001), who propose a Markovian approach to PDD field restoration from DT-MR images. Their approach could be applied to other direction fields, and allows for the integration of complex rules in the definition of the model. Beside the need for a more complex optimization method, the major drawback of Poupon's method is the necessity to discretize the feature space (S^2).

2.3. Anisotropic diffusion

Several of the above-mentioned works extend the framework of *anisotropic diffusion*, widely used for scalar images, to direction fields. Anisotropic diffusion has become a popular subject of research in the image processing community. It started with the early work of Koenderink (1984) who showed that the heat equation provides the means with which to build a scale-space and achieve isotropic smoothing with a number of useful properties. Perona and Malik (1990) then proposed an anisotropic diffusion scheme that provides an edge-preserving smoothing controlled by the norm of the gradient. Their work triggered a lot of research and, since then, many other anisotropic schemes with various formulations have been proposed for different purposes ((Catté et al., 1992; Alvarez et al., 1992; Sapiro and Tannenbaum, 1993; Krissian et al., 1997; Kornprobst, 1998; Weickert, 1999a) to mention a few), with the underlying idea that the diffusion process should be driven by the local image geometry. Theoretical properties, well-posedness, numerical aspects, and classification of those scheme have also been widely studied (see for instance (You et al., 1996; Alvarez et al., 1993; Catté et al., 1992; Niessen et al., 1994; Shah, 1996)), and developments have included extensions to vector-valued images and non-flat data ((Whitaker and Gerig, 1994) and see previous section in this paper), links with curve and surface evolution theory (Shah, 1996; Alvarez et al., 1993), and variational approaches (Shah, 1996; Teboul

et al., 1998). More generally, the anisotropic diffusion framework has grown to the more general one of partial differential equations (PDEs) for image processing. For reviews of geometry-driven anisotropic diffusion and the use of PDE's, one can refer to (Weickert, 1997b; Sapiro, 1995; Caselles et al., 1998). With variational formulations, one can establish the link between regularization as the minimization of a functional, and anisotropic diffusion that achieves that minimization (see also (Nielsen et al., 1997; Radmoser et al., 2000)). An explicit link is made between diffusion time and the regularization parameter in (Scherzer and Weickert, 2000).

The original formulation of non-linear diffusion (Perona and Malik, 1990) is defined as follows. From an original image $u^{(0)}$, one builds the one parameter family of images $u(\cdot, t)$ that satisfies the following diffusion equation:

$$\begin{cases} \frac{\partial u}{\partial t} &= \operatorname{div}(g(\|\nabla u\|)\nabla u), \\ u(\mathbf{x}, 0) &= u^{(0)}, \end{cases} \quad (4)$$

where g is the *conductance* function that defines the amount of smoothing across the image. The idea is that a high gradient norm indicates the presence of an edge, therefore if we want to preserve this edge, g should have a low value. Various conductance functions have been proposed. It has been shown that this continuous formulation defines an ill-posed problem and other formulations have been proposed for which existence and uniqueness of the solution have been proved. In particular, making the conductance a function of $\|\nabla u_\sigma\|$, where ∇u_σ is the gradient of u smoothed by a Gaussian function with parameter σ , ensures the existence and uniqueness of the solution (Catté et al., 1992).

Weickert developed the more general tensor-driven formulation of anisotropic diffusion (Weickert, 1998), defined by the following diffusion equation:

$$\frac{\partial u(\mathbf{x}, t)}{\partial t} = \operatorname{div}(\mathbf{M}\nabla u), \quad (5)$$

where \mathbf{M} is the flow tensor.¹ The flow tensor is a continuous function of the image and its derivatives, and is defined via a 3×3 symmetrical, positive-definite matrix. Scale-space and restoration properties of the scheme have been studied and can be found in (Weickert, 1998). With time t increasing, a larger smoothing is applied and the tensor \mathbf{M} is the way to define explicitly how this smoothing is driven across the image, in terms of magnitude but also direction. Different definitions have been proposed, leading to different smoothing behaviours, such as edge-enhancing (Weickert, 1998) or coherence-enhancing (Weickert, 1999a).

¹ \mathbf{M} is usually called *diffusion tensor* but we will avoid this terminology to distinguish it from the measured water diffusion tensor in the DT-MR images. For the same reason we will use the term “smoothing” instead of “diffusion”.

We also mention vector-valued diffusion, which is used to smooth all channels of a multichannel image simultaneously. Common examples are colour images (Weickert, 1999b; Tschumperlé and Deriche, 2001a) or multi-echo MR images (Whitaker and Gerig, 1994; Coulon and Arridge, 2000). A common conductance is defined for each single-channel diffusion equation, which guarantees that channels evolve in the same smoothing direction, determined by the (vector) geometry of the multichannel image.

3. Method

We present here the two steps of our method. First, a restoration scheme for direction fields is proposed, followed by a magnitude regularization using anisotropic diffusion.

3.1. A discrete variational approach to direction field restoration

We present here a variational method for restoring direction fields defined on a discrete set, followed by its application to the PDD field. Our method is a development of the total variation scheme proposed by Chan and Shen (2000) that we present below.

3.1.1. Discrete total variation model

Amongst the methods presented in Section 2.2, the Chan and Shen model (Chan and Shen, 2000) is particularly convenient, because it provides a direct discrete formulation for which no spatial derivatives are needed. Their scheme is general to non-flat data and is based on a minimization of the total variation energy. We present here the particular case of directions on \mathcal{S}^2 . We adopt the same notation as (Chan and Shen, 2000).

Let $f : \Omega_n \rightarrow \mathcal{S}^2$ be our direction distribution on \mathcal{S}^2 , where $\Omega_n \subset \mathbb{N}^n$ is the n -dimensional discrete image domain, and f_α be the direction at voxel α . Chan and Shen define the *fitted total variation (TV) energy* to be minimized:

$$\varepsilon^{\text{TV}}(f, \lambda) = \sum_{\alpha \in \Omega_n} e(f, \alpha) + \lambda \sum_{\alpha \in \Omega_n} \frac{1}{2} d_l^2(f_\alpha^{(0)}, f_\alpha), \quad (6)$$

where $d_l(f, g) = \sqrt{\|f - g\|^2}$ is the Euclidean distance between vectors, $f^{(0)}$ is the original direction map, and $e(f, \alpha)$ is the *strength function* at voxel α , that locally defines the smoothness of the direction map

$$e(f, \alpha) = \left[\sum_{\beta \in N_\alpha} d_l^2(f_\alpha, f_\beta) \right]^{1/2}, \quad (7)$$

where N_α is a neighbourhood of α . Minimising ε^{TV} therefore increases the smoothness of the map while

keeping it close to the original data. This tradeoff is controlled by the *regularization parameter* λ .

Definitions above apply to more general Riemannian manifolds. The choice of the distance d_l then depends on the metric induced by the Riemannian structure of the manifold. For a more general approach, with precise considerations on the concept of locally Riemannian distances and the choice of d_l , as well as a more general interpretation of the strength function, one should refer to the original paper (Chan and Shen, 2000).

The variational problem is solved by studying the associated Euler–Lagrange equations (see Section 3.1.2 for more details), which leads to the following differential equation to minimize ε^{TV} :

$$\frac{df_x}{dt} = \sum_{\beta \in N_x} w_{\alpha\beta} \Pi_{f_x}(f_\beta) + \lambda \Pi_{f_x}(f_x^{(0)}), \quad (8)$$

where Π_{f_x} is the orthogonal projection on the plane tangent to M at f_x , and $w_{\alpha\beta}$ is a weight defined by

$$w_{\alpha\beta}(f) = \frac{1}{e(f, \alpha)} + \frac{1}{e(f, \beta)}. \quad (9)$$

Let us point out that the use of the projection Π_{f_x} recalls the scheme proposed in (Tschumperlé and Deriche, 2001a), in which one can find a simple explanation of the role of this projection. The weights $w_{\alpha\beta}$ define the amount of interaction between voxels α and β . Notice that in the presence of a discontinuity in the neighbourhood of α the strength function tends to have a high value, therefore $w_{\alpha\beta}$ has a low value, slowing down the regularization process. That is why this scheme, like other total variation-based scheme, is known to *preserve the location of discontinuities*.

On the other hand, the behaviour observed at the location of discontinuities is not always convenient. In presence of a discontinuity in the neighbourhood of α , all the weights $w_{\alpha\beta}$, for all $\beta \in N_x$, tend to have a low value, which means that α tends to have no influence from *any* of its neighbours, leading to an *under-regularization around discontinuities* [an illustration of this behaviour is presented in (Coulon et al., 2001)]. This comes from the fact that $w_{\alpha\beta}$ is defined using all neighbours of α and all neighbours of β . Ideally, we would like the voxel α to benefit from the influence of all its neighbours that are “on the same side” of the discontinuity. To achieve that, $w_{\alpha\beta}$ needs to be defined using only α and β . We propose in next section a discrete regularising Φ -function model, directly inspired from the Chan and Shen TV model, which ameliorates the behaviour at discontinuities.

3.1.2. Generalized Φ -function model

Let $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ be a differentiable function with continuous derivative, and let $\overline{\Omega_n}$ be the graph dual to Ω_n , or the set of second-order cliques in Ω_n (if $\beta \in N_x$, then $(\alpha, \beta) \in \overline{\Omega_n}$). We define the following energy:

$$\varepsilon^\Phi(f, \lambda) = \sum_{(\alpha, \beta) \in \overline{\Omega_n}} \Phi(d_l(f_\alpha, f_\beta)) + \frac{\lambda}{2} \sum_{\alpha \in \Omega_n} d_l^2(f_x^{(0)}, f_x). \quad (10)$$

The regularising component, defined on the second-order cliques in Ω_n , recalls the Gibbs energy used in the Markovian formulations of restoration (Besag, 1974), for which the minimization is a maximum a posteriori estimation. The fact that the energy is defined on cliques reflects the idea that, on a discrete grid, common local geometry descriptors do not necessarily apply and the map is sometimes best described in terms of interaction between neighbours or distance between neighbouring features rather than, for instance, a discretized gradient. In terms of behaviour of the process at discontinuities, this is a major difference with the TV original scheme.

Let us compute the gradient of ε^Φ on \mathcal{S}^2 :

$$\frac{\partial \varepsilon^\Phi(f, \lambda)}{\partial f_x} = \sum_{\beta \in N_x} \frac{\partial}{\partial f_x} \Phi(d_l(f_\alpha, f_\beta)) + \frac{\lambda}{2} \frac{\partial}{\partial f_x} d_l^2(f_x^{(0)}, f_x). \quad (11)$$

If F is a scalar function on \mathbb{R}^3 and ∇F denotes its gradient, let \widehat{F} be its restriction on \mathcal{S}^2 with $\widehat{\partial F}/\partial f$ its gradient on \mathcal{S}^2 . Then, as stated in (Chan and Shen, 2000) and demonstrated in Appendix A:

$$\frac{\partial}{\partial f} \widehat{F}(f) = \Pi_f(\nabla F(f)). \quad (12)$$

In particular,

$$\begin{aligned} \frac{\partial}{\partial f} \Phi(d_l(f, g)) &= \Pi_f \nabla \Phi(\sqrt{\|f - g\|^2}) \\ &= \Pi_f \left(\frac{\Phi'(\sqrt{\|f - g\|^2})}{\sqrt{\|f - g\|^2}} (f - g) \right) \\ &= - \frac{\Phi'(\sqrt{\|f - g\|^2})}{\sqrt{\|f - g\|^2}} \Pi_f(g) \\ &= - \frac{\Phi'(d_l(f, g))}{d_l(f, g)} \Pi_f(g). \end{aligned} \quad (13)$$

For the same reasons

$$\frac{\partial}{\partial f} d_l^2(f, g) = -2 \Pi_f(g). \quad (14)$$

Therefore, we get the following differential equation to minimize ε^Φ :

$$\begin{aligned} \frac{\partial f_x}{\partial t} &= \sum_{\beta \in N_x} \frac{\Phi'(d_l(f_\alpha, f_\beta))}{d_l(f_\alpha, f_\beta)} \Pi_{f_x}(f_\beta) + \lambda \Pi_{f_x}(f_x^{(0)}) \\ &= \sum_{\beta \in N_x} G(d_l(f_\alpha, f_\beta)) \Pi_{f_x}(f_\beta) + \lambda \Pi_{f_x}(f_x^{(0)}), \end{aligned} \quad (15)$$

with $G(x) = \Phi'(x)/x$. Let us make a few comments about this differential equation:

- we demonstrated here discrete results that are consistent with other constrained continuous PDEs (Tschumperlé and Deriche, 2001a);
- it recalls the discrete scheme proposed by Perona and Malik in their pioneering work on anisotropic diffusion (Perona and Malik, 1990) (see Section 2.3). The function $G(x) = \Phi'(x)/x$ (a classical result) is the *conductance* function. It controls the local intensity of the smoothing, or the interaction between neighbours, as a function of the distance between neighbouring features (instead of a function of the gradient on the map for the Perona–Malik scheme). See Fig. 3 for an example of this conductance function.

We show in the next section how we apply this scheme to a particular type of directional data, the PDD field derived from DT-MRI data. The behaviour of the scheme is shown in Section 4. A comparison of our Φ -function model with the TV model proposed by Chan and Shen can be found in (Coulon et al., 2001) and shows that the clique-based energy, defined by Eq. (10), provides better management of discontinuities than the node-based energy proposed by Chan and Shen (2000).

3.1.3. Application to PDD fields

In the particular case of PDD maps, we still want discontinuities to be preserved. However, the discontinuities can arise in different ways since they can be due to transition between tracts with different directions, or transitions between tissues with different anisotropy levels. Practically speaking, because the PDD is defined as an eigenvector of a matrix, it has no sign (\mathbf{v}_1 is indistinguishable from $-\mathbf{v}_1$), and the angle between two PDDs is always in the interval $[0; \pi/2]$. To take this constraint into account, when measuring distance between f_α and any neighbour f_β , we first check that $f_\beta \cdot f_\alpha > 0$ (i.e., they are on the same hemisphere of \mathbf{S}^2) and if not, we “flip” f_β to $(-f_\beta)$.

For scalar images, the influence of the choice of function Φ has been studied and various functions have

been proposed (Perona and Malik, 1990; Green, 1990; Rudin et al., 1992; Geman and McLure, 1985; Teboul et al., 1998). Because the components of the diffusion flow in *gauge coordinates* (a local coordinate system defined by the gradient and the plane orthogonal to the gradient) is a function of the first and second derivative of Φ , it is possible to predict the behaviour of the diffusion process in every direction. Convex Φ -functions bring the most interesting theoretical properties but some non-convex functions have been shown to lead to better experimental results (see for example (Teboul et al., 1998)). However, no such theoretical study has been done for direction maps and we propose here an empirical choice with a simple geometric interpretation that proved to work well on the PDD data. Specifically we are looking for an increasing function of the distance between vectors, so we decided to use a function of the cosine of the angular difference. Therefore, we propose to define the “conductance” using

$$G(d_l(f_\alpha, f_\beta)) = (f_\alpha \cdot f_\beta)^{2m} = \cos(\theta)^{2m}, \quad (16)$$

with θ the angle between f_α and f_β and $m \in \mathbb{N}$. Because $\theta \in [0; \pi/2]$, it is easy to prove that $\cos(\theta) = (1 - d_l(f_\alpha, f_\beta)^2/2)$. Therefore, one must have

$$G(x) = \frac{\Phi'(x)}{x} = (1 - x^2/2)^{2m} \quad (17)$$

and

$$\Phi(x) = \frac{1 - (1 - x^2/2)^{2m+1}}{2m + 1}. \quad (18)$$

The functions Φ and G are shown in Fig. 3. The function G explicitly defines the discontinuity preserving behaviour as a function of the distance between features and the severity of this behaviour is tuned by the parameter m .

One important constraint of a PDD map is that it does not have any organization or meaning in isotropic regions. Regularization in isotropic regions is therefore meaningless and could actually lead to the creation of a “fake” organization. Therefore, we define

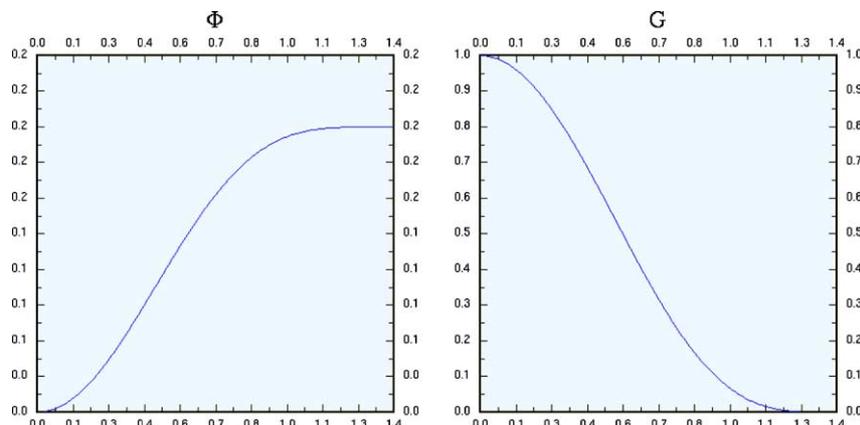


Fig. 3. The two functions $\Phi(x)$ and $G(x)$ with $m = 3$. (This figure is available in colour, see the on-line version.)

$$a_{\alpha,\beta} = \frac{\text{FA}(\alpha) + \text{FA}(\beta)}{2}, \quad (19)$$

with FA the fractional anisotropy measure defined in Eq. (3). We then alter our scheme to obtain a *weighted regularization*, defined by the functional

$$\varepsilon_w^\phi(f, \lambda) = \sum_{(\alpha,\beta) \in \Omega_n} a_{\alpha,\beta} \Phi(d_I(f_\alpha, f_\beta)) + \frac{\lambda}{2} \sum_{\alpha \in \Omega_n} d_I^2(f_\alpha^{(0)}, f_\alpha), \quad (20)$$

minimized by the differential equation

$$\frac{\partial f_\alpha}{\partial t} = \sum_{\beta \in N_\alpha} a_{\alpha,\beta} \frac{\Phi'(d_I(f_\alpha, f_\beta))}{d_I(f_\alpha, f_\beta)} \Pi_{f_\alpha}(f_\beta) + \lambda \Pi_{f_\alpha}(f_\alpha^{(0)}). \quad (21)$$

Thus, the effects of the regularization are weak where anisotropy is close to zero. Furthermore, the weighting also decreases the influence of isotropic tissues over anisotropic ones, increasing the preservation of discontinuities between GM and WM.

The minimization is then performed using a discretization scheme proposed by Chan and Shen (geodesic marching; see (Chan and Shen, 2000), Section 4.2.1).

3.1.4. Tensor reconstruction

Once the PDD has been restored, we reorient the tensor with the new restored PDD. The second and third eigenvectors must be first reoriented. We use a similar approach to the preservation of principal directions algorithm presented in (Alexander, 2001), which computes the reorientation of DTs to accompany a non-rigid transformation applied to the whole image. The second eigenvector is projected on the plane orthogonal to the regularized first eigenvector in order to compute the new eigensystem:

- Let $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ be the original set of eigenvectors, and \mathbf{v}_1^r the regularized first eigenvector.
- Define $\mathbf{v}_2^r = \mathbf{v}_2 - (\mathbf{v}_2 \cdot \mathbf{v}_1^r) \mathbf{v}_1^r$.
- Define $\mathbf{v}_3^r = \mathbf{v}_3 \times \mathbf{v}_2^r$.

The new tensor is then constructed using the new set of eigenvectors and the original eigenvalues

$$\mathbf{D}^r = (\mathbf{v}_1^r \ \mathbf{v}_2^r \ \mathbf{v}_3^r) \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} (\mathbf{v}_1^r \ \mathbf{v}_2^r \ \mathbf{v}_3^r)^T. \quad (22)$$

The re-orientated tensor then has the same shape (eigenvalues) as the original one with a different orientation.

3.2. Intensity regularization using anisotropic diffusion

3.2.1. Generalities

Direction and magnitude are separated in the regularization process, and we present here the magnitude regularization. When dealing with colour images,

brightness (i.e. the magnitude information) and chromaticity (direction) can be processed separately and independently, since the organization and coherence of the two maps are not necessarily related. In other types of image there is a correlation between the two types of information, in particular, when the direction map can be used to describe the local structure of the intensity image: DT-MR images, MR flow imaging, optic flow, deformation fields, strain tensor images. For instance in the case of flow information, intensity is expected to be smooth in the direction of the flow. The local coherence of the intensity image is therefore indicated by the direction map itself. The case of DT-MR images is similar although more complex: in WM, where the tensor is anisotropic, eigenvalue coherence is defined by the direction of the tracts, i.e. the PDD field. In isotropic tensor regions, there is no such directional coherence and the eigenvalue fields are expected to be isotropic.

The purpose of this section is to show how tensor-driven anisotropic diffusion, as formulated by Weickert (1998), can be used to perform a regularization that incorporates the prior knowledge of the image coherence expressed by the direction field. This will be illustrated by our application to DT-MR images.

As opposed to previous formulations where conductance is expressed as a scalar function (e.g. Eq. (4), sometimes designated by the term *non-linear diffusion* instead of *anisotropic diffusion*), the tensor-driven formulation defined by Eq. (5) allows us to define explicitly directions and magnitude of smoothing, under a number of conditions. This makes a perfect framework when one has a prior model of the local structure in the image and explicitly knows how and in what direction the smoothing process must be performed. Weickert's models rely on an evaluation of the local structure using the *structure tensor*: $J_\rho = K_\rho * (\nabla u_\sigma \otimes \nabla u_\sigma)$, where K_ρ is a Gaussian function and ∇u_σ is the smoothed gradient (Weickert, 1998). The eigenvectors of J_ρ indicate the directions with the most and the least variations, while the eigenvalues are descriptors of local structure, for example: constant area, edges, corners, anisotropic structures. The direction with the least variation is called the *coherence direction*. By defining the flow tensor that drives the smoothing with the same eigenvectors than J_ρ and some well-chosen functions of its eigenvalues, one can define the smoothing behaviour using the local structure information (Weickert, 1999a, 1998). Evaluating J_ρ necessitates the choice of an *integration scale*, ρ , which gives rise to a trade-off between locality and robustness.

3.2.2. Introducing a priori structural information

After having computed the local structure information, Weickert shows how to define the flow tensor. In particular, for the *coherence enhancing* scheme (Weick-

ert, 1999a), the coherence direction \mathbf{v} is used together with a local index of coherence to build a flow that smooths along \mathbf{v} only where the index is high enough.

The idea is similar in the case of multi-dimensional data such as the ones mentioned above. The prior information about the image map is partly contained into the direction map: for instance with flow images, \mathbf{v} is the flow direction. For a general approach to building an anisotropic flow tensor that takes into account the scalar map structure, we need:

- A local coordinate system to define the eigenvectors of the flow tensor. This is where the a priori directional information is introduced to define the first eigenvector.
- A local index that indicates where the direction map describes properly the scalar map structure, and where the anisotropic smoothing will be performed. For instance, in the case of flow images, we want to smooth along the flow only where there is one: the flow magnitude is the coherence index. One can then define the first eigenvalue as a function of this index, and the second and third eigenvalues with a small constant positive value. Therefore, the smoothing is performed only in the direction of coherence.

The desired behaviour can be more complex than an anisotropic smoothing in one direction. Several types of smoothing might be required depending on the scalar map characteristics, or we may want to take into account particular discontinuities. This is illustrated by DT-MR images, for which the structural information is directly available in the reorientated diffusion tensor \mathbf{D}^r . In the case of a prolate tensor (see Section 2.1), which arises in white matter fiber bundles, coherence is defined in one direction, and regularity is expected along the fibers, that is in the direction of the restored PDD. With an oblate tensor, local coherence lies within a plane defined by the first two eigenvectors. For an isotropic tensor, underlying tissues have an isotropic nature, and regularity is expected in every direction. These three extreme cases are used to define a smoothing behaviour that takes into account the expected regularity and we show how this can be achieved in the next section.

3.2.3. From diffusion tensor to flow tensor

According to the DT-MR image local coherence description presented in previous section, the smoothing applied across the image should be defined as follows:

- isotropic smoothing in isotropic tissues (GM, CSF),
- strong anisotropic smoothing in WM, along the directions of most diffusion (i.e. along fiber tracts),
- preservation of transitions between isotropic and anisotropic tissues.

The above three rules are enough to guarantee the criteria mentioned at the end of previous section. In

particular, if in WM the smoothing is performed only in the direction of anisotropy, there will not be any flow between neighbouring tracts with different directions (since the flow is directed only in the direction of the tracts).

We recall the diffusion equation (5) applied on an eigenvalue map λ

$$\frac{\partial \lambda(\mathbf{x}, t)}{\partial t} = \text{div}(\mathbf{M} \nabla \lambda), \quad (23)$$

and define the flow tensor as follows:

$$\mathbf{M} = (\mathbf{v}_1^r \quad \mathbf{v}_2^r \quad \mathbf{v}_3^r) \begin{pmatrix} \mu_1 & 0 & 0 \\ 0 & \mu_2 & 0 \\ 0 & 0 & \mu_3 \end{pmatrix} (\mathbf{v}_1^r \quad \mathbf{v}_2^r \quad \mathbf{v}_3^r)^T. \quad (24)$$

Therefore, \mathbf{M} has the same eigenvectors as \mathbf{D}^r , indicating the principal directions of anisotropy (in particular, the PDD) when there is anisotropy. The amount of smoothing is defined by the eigenvalues μ_i , functions of the diffusion tensor eigenvalues λ_i , and defined as follows:

$$\mu_i = \frac{\lambda_i^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2} H(\|\nabla \text{FA}_\sigma \cdot \mathbf{v}_i^r\|) \quad (25)$$

with

$$H(x) = \frac{1}{2}(1 - \tanh(K(x - C))). \quad (26)$$

The first part of the definition (before the function H) guarantees that if \mathbf{D}^r is isotropic then \mathbf{M} is isotropic as well, and if \mathbf{D}^r is anisotropic then \mathbf{M} is anisotropic with the flow almost entirely along the principal directions of anisotropy.

The function $H(\|\nabla \text{FA}_\sigma \cdot \mathbf{v}_i^r\|)$ behaves like a contour map that looks for anisotropy changes in the directions \mathbf{v}_i^r . At the termination of a tract, there is a tissue transition (WM/GM) in the direction of the flow. Equivalently, in GM the isotropic flow might cross a boundary if we are next to WM. Such situations are indicated by a high value of $\|\nabla \text{FA}_\sigma \cdot \mathbf{v}_i^r\|$, where FA_σ is the fractional anisotropy smoothed with a Gaussian kernel of parameter σ . Therefore, the function H provides a smooth cut-off at value C with a slope K .

Fig. 4 provides a visual illustration of the flow tensor definition, using an ellipsoid representation of tensors. In the very anisotropic bow-shaped structure (*corpus callosum*), the flow is almost mono-directional, whereas in the isotropic regions (CSF or GM, in blue on the picture) the flow is isotropic. The close-ups in Fig. 5 show the edge-preserving effect of the flow definition. In Fig. 5(a), one can see a small part of the *corpus callosum*, anisotropic (yellow and red), next to isotropic tissues (GM, small blue ellipsoids, and CSF, large blue ellipsoids). The resulting flow tensor shows a high anisotropy in the *corpus callosum*, but tensors in the CSF are “flat” (green ellipsoids), to prevent flow towards the WM. In Fig. 5(b), two tracts are next to each other, the

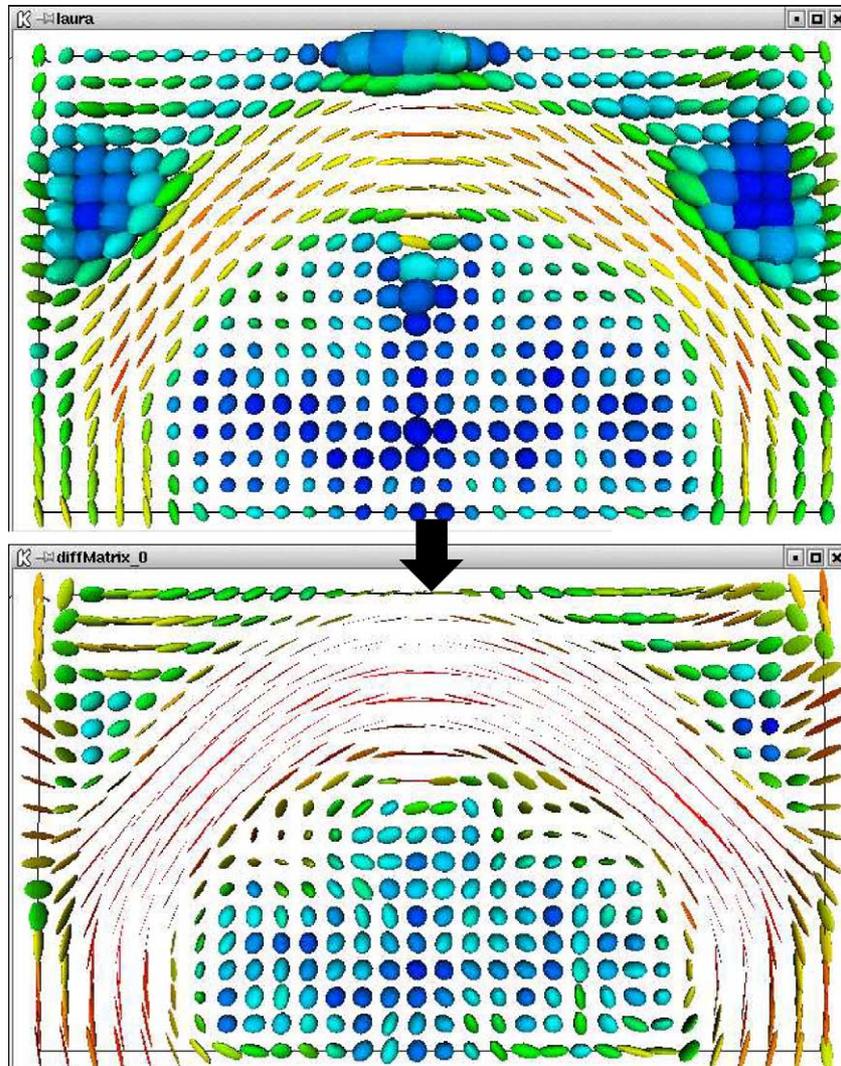


Fig. 4. Diffusion tensor (top) and corresponding flow tensor (bottom). Same colour mapping as in Fig. 2. Real brain data: splenium of the corpus callosum. (This figure is available in colour, see the on-line version.)

corpus callosum again (in plane, curved), and the *cingulum* (green ellipsoids in the foreground right corner), which runs vertically and is slightly less anisotropic. The two tracts are orthogonal and next to each other. The resulting flow tensor image shows that the flow inside each tract prevents the different orientations influencing each other, and that the flow tensor between the tracts is “flat” in the direction of the transition.

The process described by Eqs. (23)–(25) is applied on the three eigenvalue maps simultaneously, with the same flow tensor, similar to vector-valued image diffusion (Whitaker and Gerig, 1994; Weickert, 1999b). At each iteration of the diffusion process, the fractional anisotropy FA is recomputed from the new eigenvalues, which introduces an implicit coupling between the images and the contour map. The discretization of Eq. (5) is done via a simple explicit scheme with a small time step ($\Delta t = 0.2$). More stable schemes,

allowing larger time steps, can be found in (Weickert, 1997a).

4. Results

Experiments were performed on both synthetic and real DT-MR images. Synthetic images were used to assess both qualitatively and quantitatively the effects of the process. In this section, we present the results of these experiments.

4.1. Synthetic images

4.1.1. PDD restoration: robustness to noise and choice of the regularization parameter

Two noise-free tensor volumes were generated by building anisotropic structures embedded in a pure

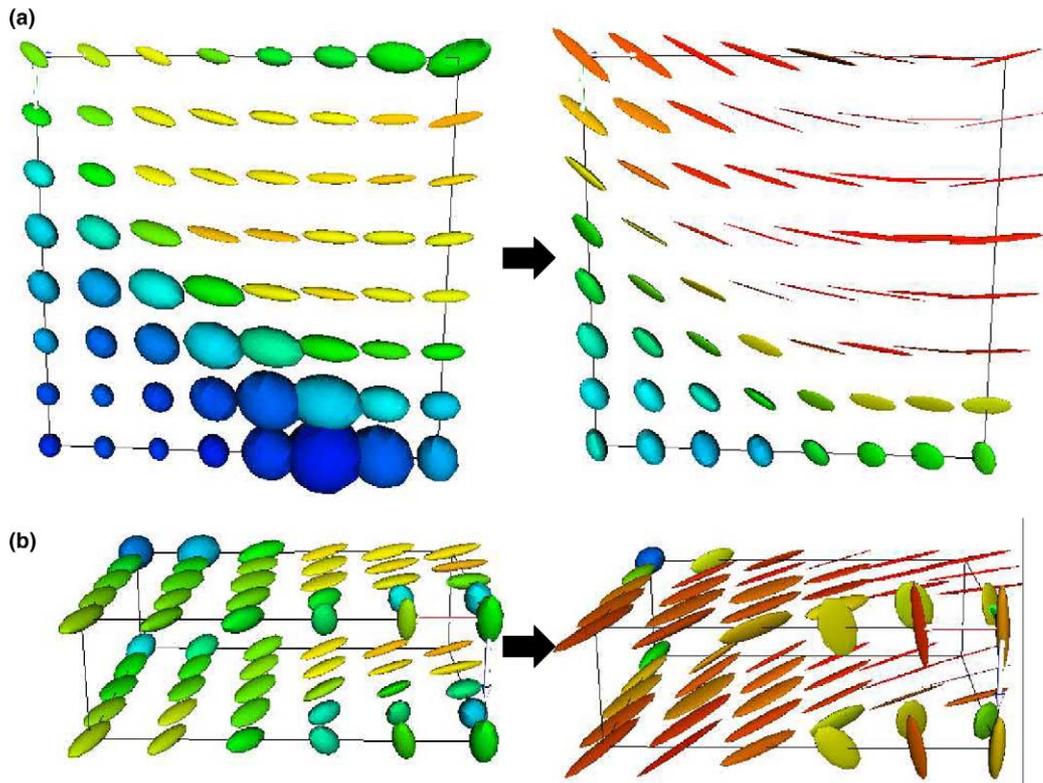


Fig. 5. Diffusion tensor (left column) and corresponding flow tensor (right column) in two sub-regions of the brain (small parts of the splenium of the corpus callosum). Same colour mapping as in Fig. 2. (This figure is available in colour, see the on-line version.)

isotropic medium. Realistic values were chosen for the eigenvalues in each type of medium (Pierpaoli and Basser, 1996; Pierpaoli et al., 1996), i.e., $\lambda_1 = 1700$ and $\lambda_2 = \lambda_3 = 200$ for the anisotropic tensors, and $\lambda_1 = \lambda_2 = \lambda_3 = 700$ for the isotropic tensors. The first volume is composed of a single anisotropic bundle in the shape of a torus. The second volume is composed of two anisotropic bundles: a “ring” around a cylinder. Once the noise-free tensor field has been constructed, we add noise to the data that reflects the noise properties of data acquired from an MR scanner using a similar approach to that described in (Alexander et al., 2002). Given each diffusion tensor, we can compute the corresponding DW measurement (see Section 2.1) for a set of 60 directions. Each of the DW images is transformed into the Fourier domain, where Gaussian noise is added with variance σ^2 . The Fourier transform is then inverted to provide noisy DW images, which are then used to reconstruct noisy versions of the DT at each voxel. A value $\sigma^2 = 25$ provides data with approximately the same level of noise as that measured in our MR scanner in the absence of any signal. A value $\sigma^2 = 400$ introduces about 10 times this level of noise, while $\sigma^2 = 1$ introduces about a tenth of this level. Figs. 10 and 11 show parts of the *torus* and *ring* synthetic noisy images.

We defined a error measure for the PDD field between a restored image and the corresponding noise-free original image. This error is defined as follows:

$$E = \sum_{\alpha \in \Omega_n} \text{FA}^{(0)}(\alpha) (1 - f_\alpha \cdot f_\alpha^{(0)}), \quad (27)$$

where $\text{FA}^{(0)}$ is the fractional anisotropy in the original noise-free image, f_α is the PDD in the restored image, and $f_\alpha^{(0)}$ the PDD in the noise-free image. On our synthetic images, this error measure effectively takes into account only the anisotropic part of the image, which is the only part we are interested in for the PDD restoration. For a given noise level, E can be plotted as a function of the regularization parameter λ in order to measure the influence of this parameter and find its optimal value. Fig. 6 shows the plots $E(\lambda)$ for noise levels $\sigma^2 = 1, 25, 100, 400$. We can see that in all situations $E(\lambda)$ reaches a single minimum λ_{opt} . For $\lambda < \lambda_{\text{opt}}$, there is too much regularization and the data end up being too smooth, whereas for $\lambda > \lambda_{\text{opt}}$ data are kept too close from the original map and the noise level is not decreased enough. Values of λ_{opt} for $\sigma^2 = 1, 25, 100, 400$, respectively, are 7.0, 1.0, 0.5, 0.25, respectively. Let us define $\lambda^*(\sigma^2)$ as the function that

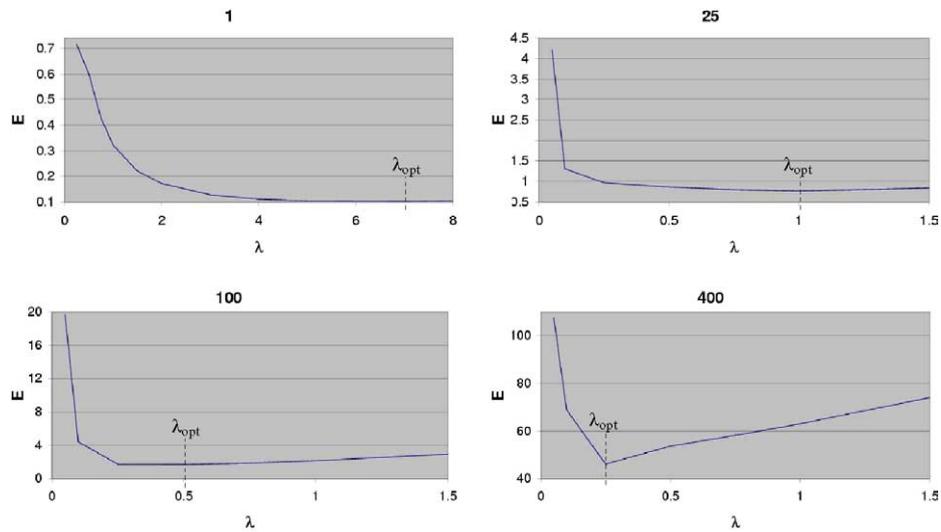


Fig. 6. Error after restoration of the *torus* PDD map as a function of the regularization parameter, for noise levels $\sigma^2 = 1, 25, 100, 400$.

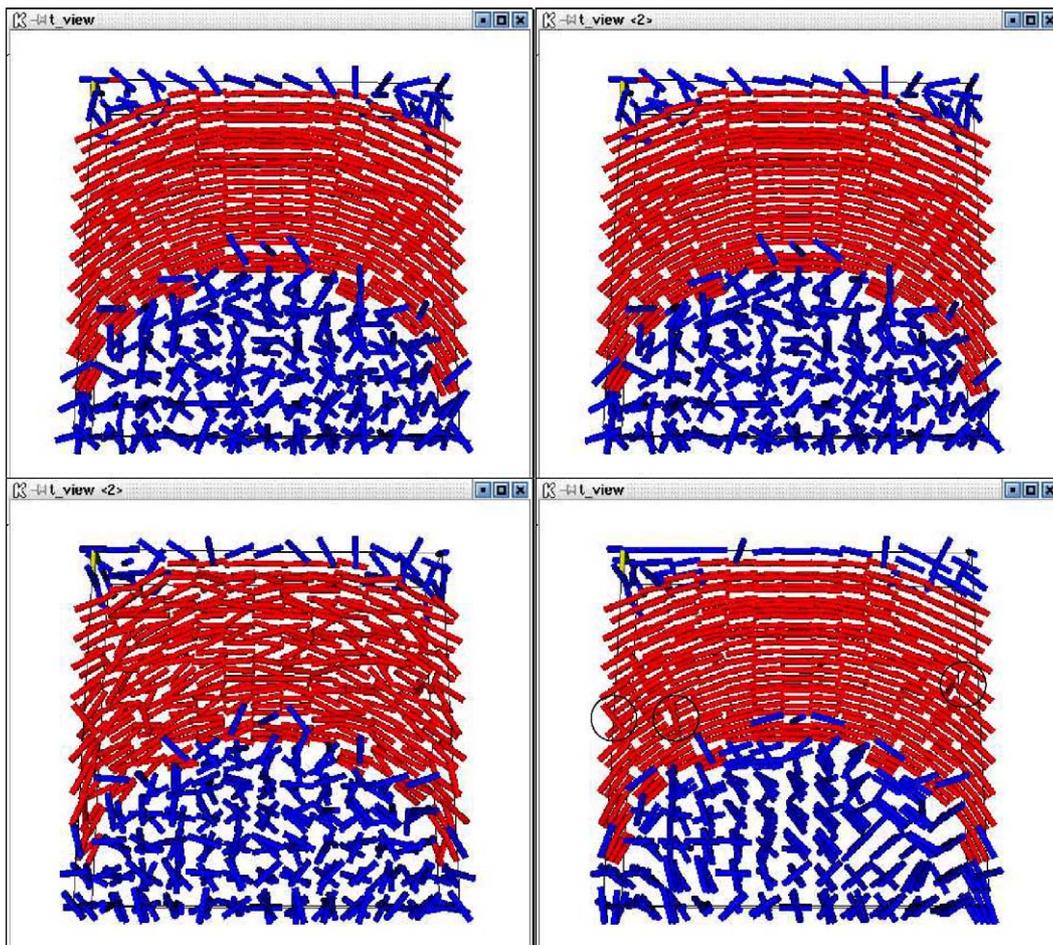


Fig. 7. PDD restoration. Left column: noisy synthetic *torus* image. Right column: the corresponding restored image. Top row: noise level $\sigma^2 = 25$. Bottom row: noise level $\sigma^2 = 400$. Blue: isotropic tensor, red: anisotropic tensor. (This figure is available in colour, see the on-line version.)

associates the noise level σ^2 with $1/\lambda_{\text{opt}}$ (the function $\lambda^*(\sigma^2)$ is sometimes called *regularization strategy* (Kirsch, 1996)). It is interesting to notice that experimentally we observe:

- when $\sigma^2 \rightarrow 0$, then $\lambda^*(\sigma^2) \rightarrow 0$;
- when $\sigma^2 \rightarrow 0$, then $E(\lambda_{\text{opt}}) \rightarrow 0$.

In other words, when the noise level tends to 0, the amount of regularization required tends to zero, and the result of the optimal restoration tends to be the original noise-free image. These two properties are the theoretical requirements for a valid (*admissible*) regularization strategy (Kirsch, 1996).

Figs. 7 and 8 show images of the result of the restoration for noise levels $\sigma^2 = 25$ (same level than real images) and $\sigma^2 = 400$, using the associated optimal regularization parameter value. The following qualitative observations can be made:

- The coherence of the field inside each structure is properly restored.
- The PDD field at the border of the anisotropic structures has not been corrupted by the isotropic part.

- In Fig. 8, the transition cylinder-ring has been properly preserved and the restoration is well behaved at the interface between those structures.
- For $\sigma^2 = 400$ one can notice a few isolated vectors that are orthogonal to their original orientation (circled in the images). This is an effect of the “sorting bias” (Basser and Pajevic, 2000): in the presence of noise λ_2 might become higher than λ_1 . In this case the sorting of the eigenvalue induces an error, and the PDD field includes v_2 instead of v_1 . At such points, the error is seen as a discontinuity by the process and the regularization is frozen. A method to correct that bias is proposed in (Basser and Pajevic, 2000).

For real images, the noise level is similar to that obtained from $\sigma^2 = 25$, therefore for the experiments on real data described in the next sections we will use the value $\lambda = 1.0$.

4.1.2. Eigenvalue regularization

Experiments were performed on the same synthetic images in order to assess the behaviour of the eigenvalue

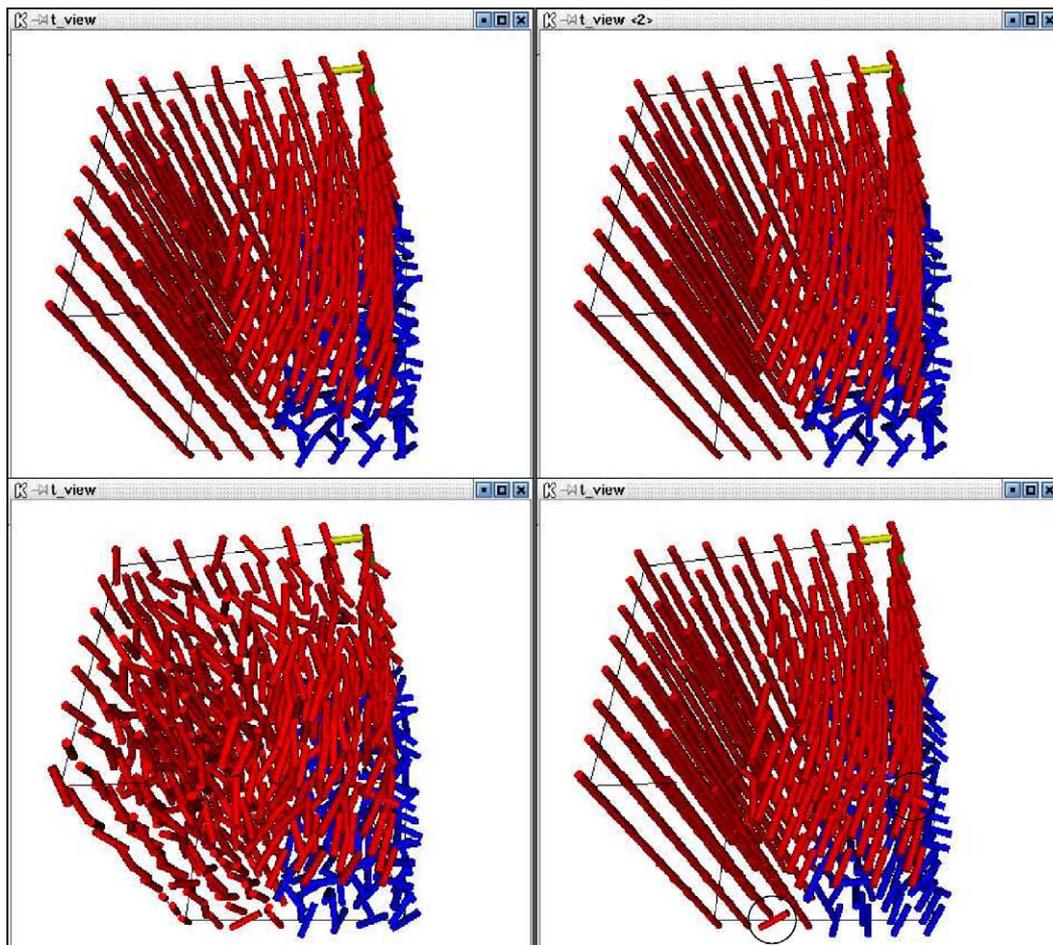


Fig. 8. PDD restoration. Left column: noisy synthetic *ring* image. Right column: the corresponding restored image. Top row: noise level $\sigma^2 = 25$. Bottom row: noise level $\sigma^2 = 400$. Blue: isotropic tensor, red: anisotropic tensor. (This figure is available in colour, see the on-line version.)

regularization process. After restoration of the PDD using the optimal parameter value for a given level of noise, the anisotropic diffusion process defined by Eqs. (5) and (24)–(26) was applied on the eigenvalue maps with a discrete time step $\Delta t = 0.2$.

We defined two regions of interest, in the isotropic and anisotropic parts of the *torus* image. For the noise level $\sigma^2 = 25$, we plotted the estimated standard deviation of each eigenvalue within each ROI against iteration number. Plots are presented in Fig. 9 and show a significant decrease of the standard deviation in every case, corresponding to an effective decrease of noise level.

After an optimal number of iterations the trend is inverted. This is due to the fact that, after some time, regions start to mix: diffusion cannot be completely stopped across boundaries (for theoretical reasons the flow cannot be zero in any direction) and after an infinite time the process converges to a constant image. This is a well-known limitation of anisotropic diffusion and is why it is necessary to determine an optimal time limit.

Note that it is also possible to make the process converge to a non-constant image by adding a data-driven term, similar to the method presented for the PDD restoration. One must then determine an optimal weight for this term. In our case, and for $\sigma^2 = 25$, 40 iterations ($t = 8$) is enough to bring down all eigenvalues to their optimal noise level both in anisotropic and isotropic media.

Figs. 10 and 11 show the results of the regularization for the two synthetic images for $\sigma^2 = 25$ and 400. It is clear that for a realistic noise level the smoothness of the field is restored within each structure. When the noise level increases, anisotropy increases in isotropic media and decreases in anisotropic bundles. Because anisotropy of the flow depends on anisotropy of the diffusion tensor, when the noise level becomes very high the flow tends to be the same everywhere. Nevertheless, for $\sigma^2 = 400$ the smoothness of the field is largely increased, while a transition between each medium can still be observed.

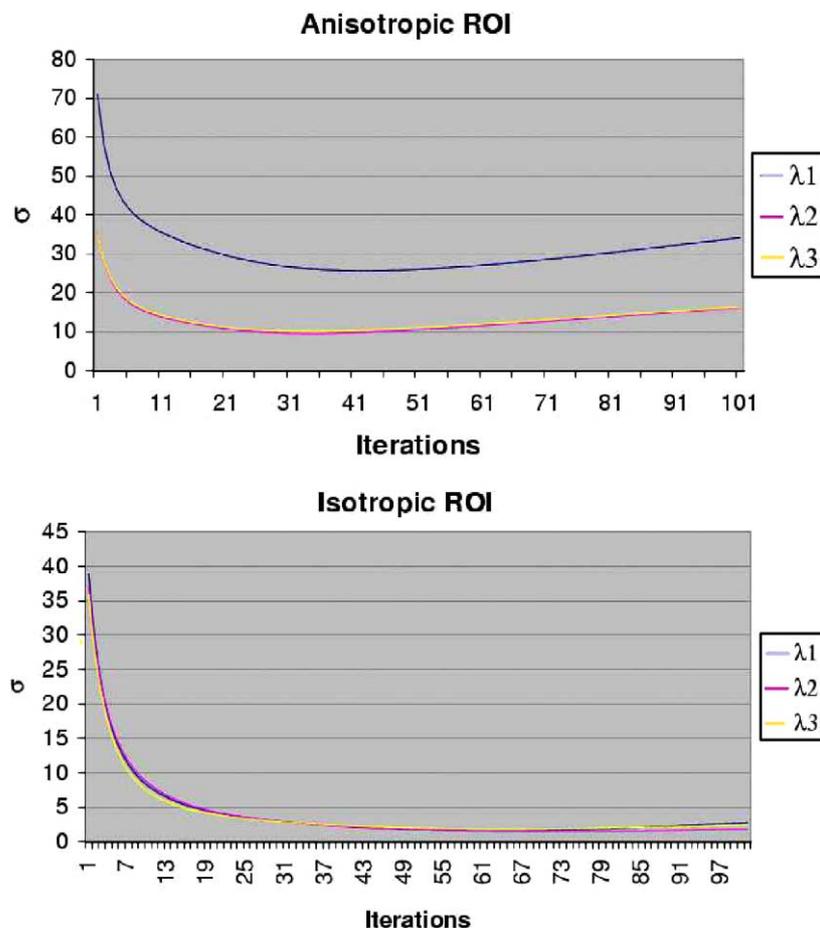


Fig. 9. Standard deviation of eigenvalues through iterations in isotropic and anisotropic media. (This figure is available in colour, see the on-line version.)

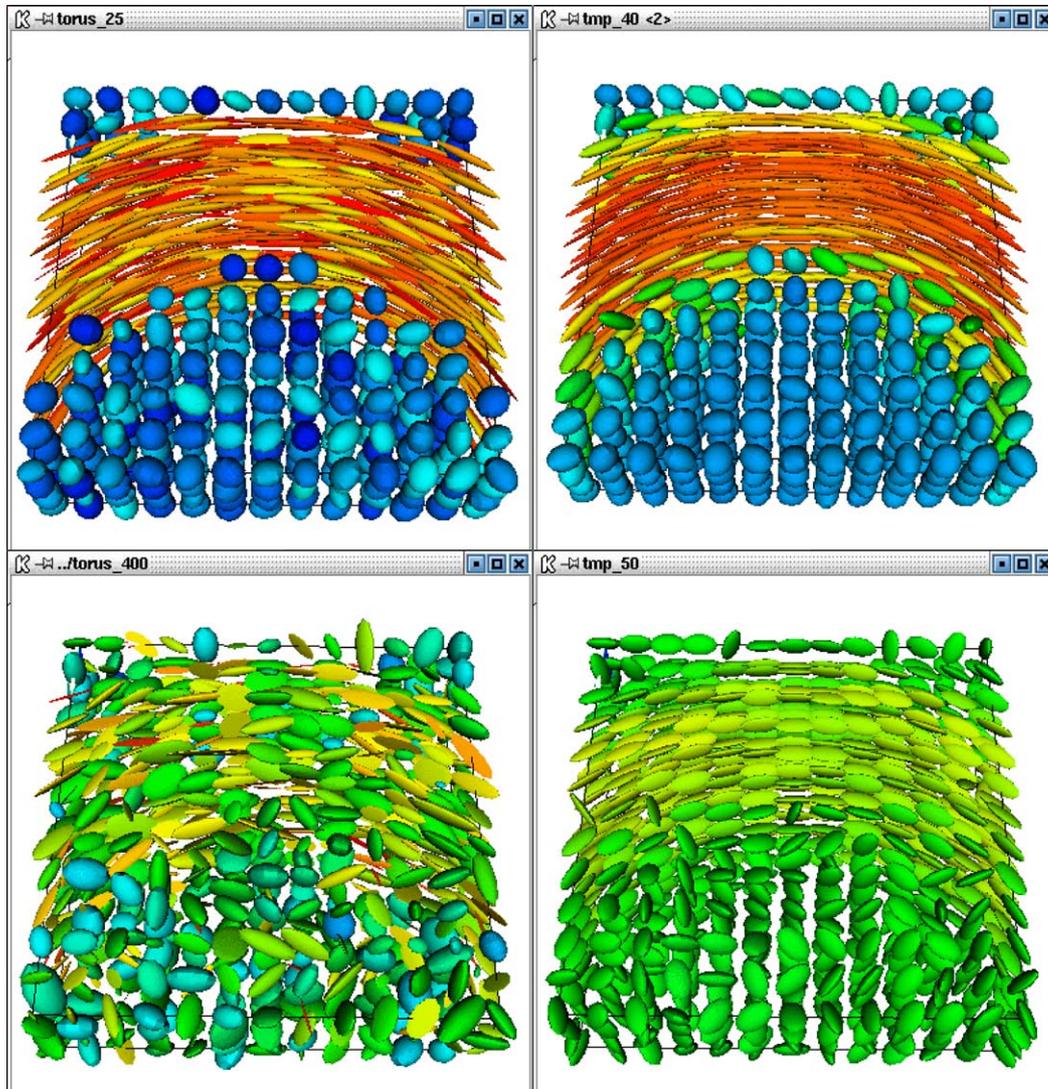


Fig. 10. Eigenvalue regularization. Left column: noisy synthetic *torus* image. Right column: the corresponding regularized image. Top row: noise level $\sigma^2 = 25$. Bottom row: noise level $\sigma^2 = 400$. Same colour mapping as in Fig. 2. (This figure is available in colour, see the on-line version.)

4.2. DT-MR images

Two types of DT-MR images were used in our experiments: spinal cord data to observe the effects of the PDD map restoration, and brain data to assess the effects of the eigenvalue regularization. We present results using those data.

4.2.1. Spinal cord data

Acquisition at the level of the spinal cord is technically more difficult and more sensitive to motion artefacts, for instance induced by breathing, therefore these data are quite noisy. In particular, the PDD field is more noisy than with brain data. The images contain a cylindrical region (the cord) inside which anisotropy is high, due to the presence of fibers, and outside which

anisotropy is low, in the CSF surrounding the cord (see Fig. 12). Discontinuities are of two types: at the interface between cord and CSF, and inside the cord, at the entrance of peripheral nerves.

Results of the PDD restoration are shown in Fig. 12 for a subsection of the cord. Fig. 12(a) shows that directions have been clearly realigned along the cord, and the smoothness has increased. At the borders of the cord, data have not been disturbed by the CSF. Fig. 12(b) shows a close-up at a discontinuity of the direction field within the cord, and we can see that the discontinuity has been preserved while the data have been smoothed. As there is no ground truth to compare results with, an essential issue of the PDD restoration is its effect on post-processing methods, in particular, white matter fiber tracking (tractography) for which the

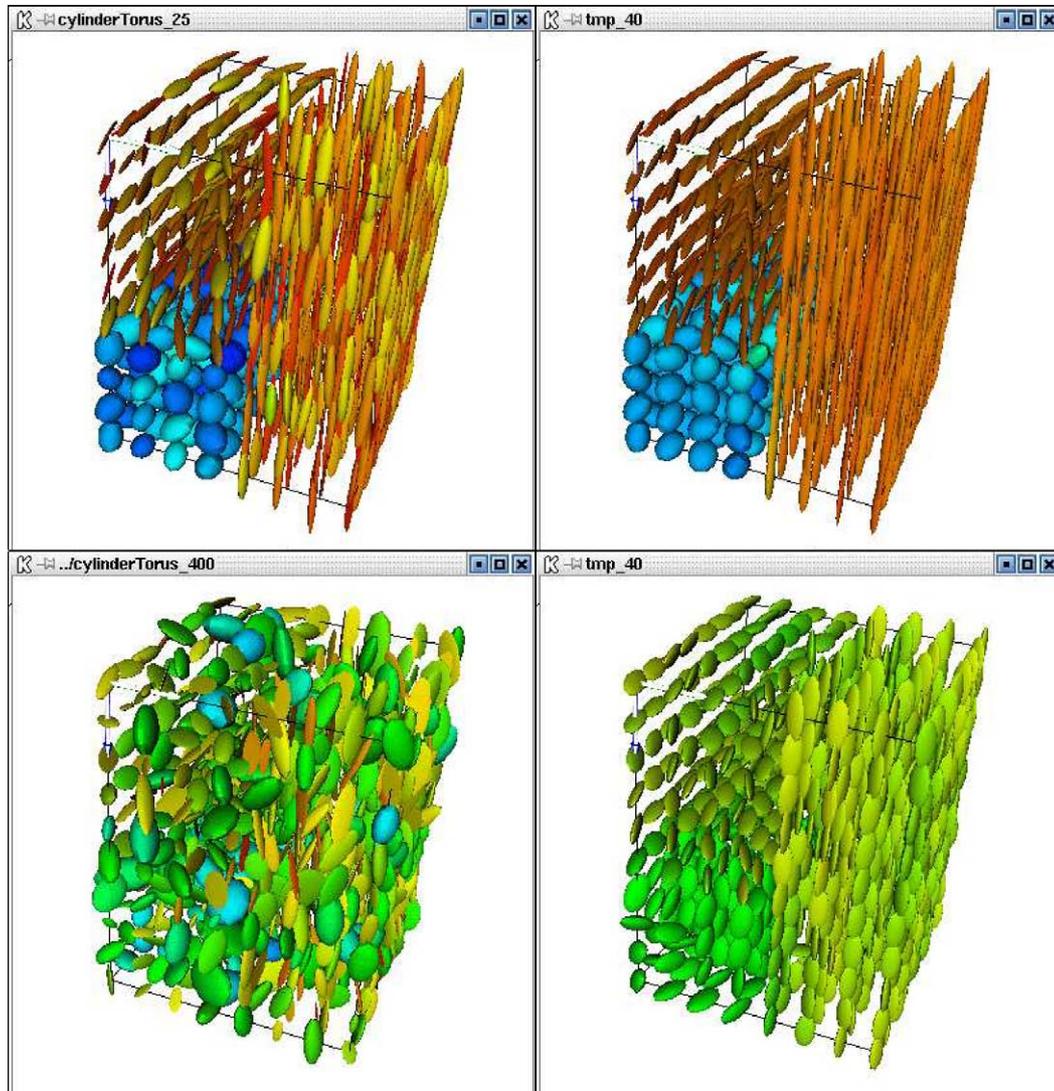


Fig. 11. Eigenvalue regularization. Left column: noisy synthetic *ring* image. Right column: the corresponding regularized image. Top row: noise level $\sigma^2 = 25$. Bottom row: noise level $\sigma^2 = 400$. Same colour mapping as in Fig. 2. (This figure is available in colour, see the on-line version.)

PDD direction is essential. Current work aims at assessing those effects using a tractography method (Parker et al., 2001) on spinal cord data.

4.2.2. Brain data

Echo-planar DT-MR brain images of size $128 \times 128 \times 42$ were acquired using an acquisition scheme similar to that described in (Jones et al., 1999), and processed using the method presented above. The tensor was first reorientated according to our direction map restoration scheme, then the eigenvalue map regularization was applied. Fig. 13 shows the results of the regularization on the three eigenvalue maps at $t = 4$. The smoothness within each structure has increased while the boundaries of those structures have been preserved. The same effect can be observed on the ellipsoid representation in

Fig. 14. As mentioned before, eigenvalues are used to define scalar measurements that can characterise tissues and assess their properties. Therefore, it is interesting to observe the effect of the regularization on those measurements. Fig. 15 shows the evolution of a fractional anisotropy map through the regularization process with increasing time. Obviously, as the smoothness of the eigenvalue maps increases, so does the smoothness of the fractional anisotropy, following a scale space-like evolution. Discontinuities are preserved and the local coherence of the anisotropy map is largely enhanced by the process. As time increases, thin anisotropic structures in the white matter disappear. This highlights the issue of the optimal amount of regularization to apply to data like brain images, that contain a lot of different structures at various scales. There is a trade-off

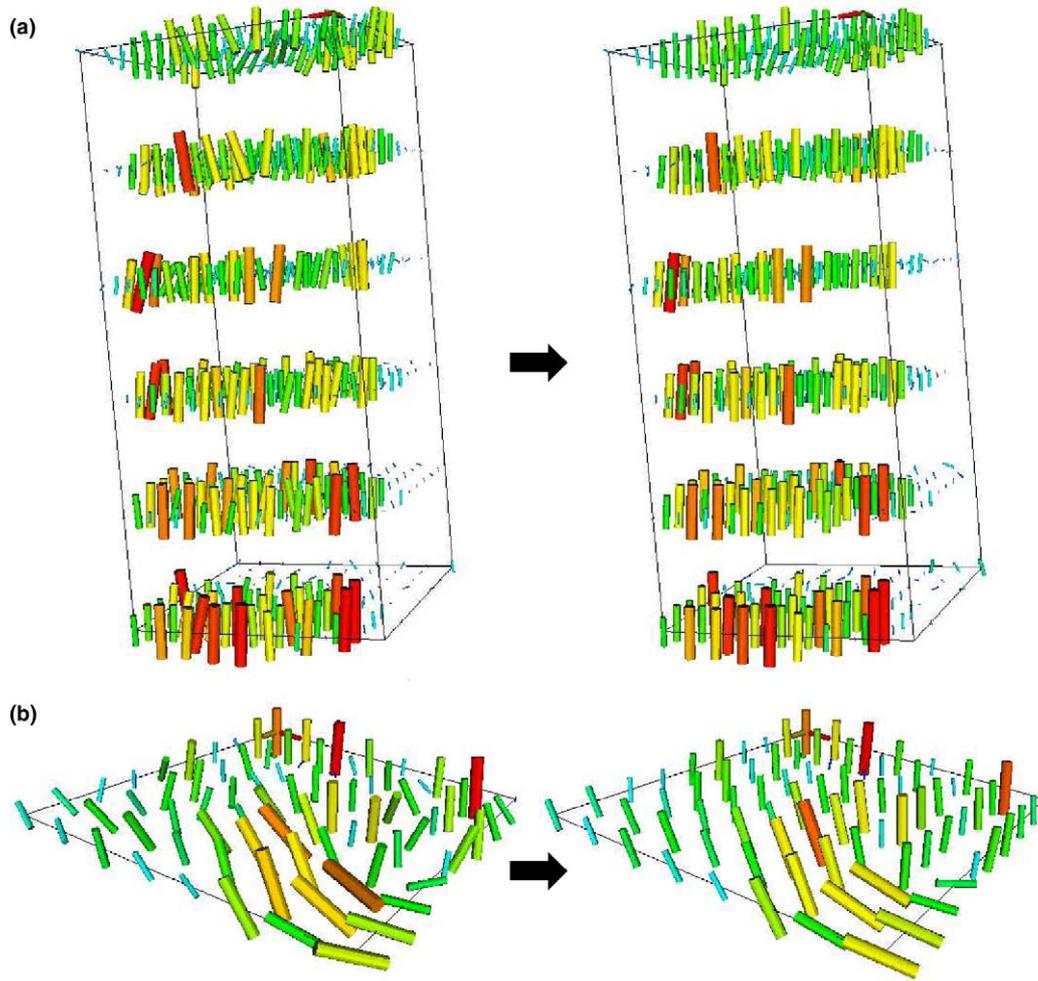


Fig. 12. Two sub-region of the spinal cord PDD field before (left) and after (right) regularization. Directions are scaled and colour-mapped by fractional anisotropy (same colour mapping as in Fig. 2). (This figure is available in colour, see the on-line version.)

between preserving fine structures and providing enough regularization, although the discontinuity-preserving strategy and the anisotropic nature of the smoothing in white matter aim at relaxing this trade-off as much as possible.

5. Conclusion

We presented a method to regularise diffusion tensor magnetic resonance images. More generally, the method defines a framework for the regularization of multi-dimensional data containing both directional information and one or several magnitude scalar components whose structure is related to the direction map. On the theoretical side, we proposed a general variational method for restoration of direction fields on a discrete domain, inspired from the approach of Chan and Shen (2000). The method can be used with

various regularising Φ -functions, and the influence of this choice on the behaviour of the scheme, for instance in terms of anisotropy or edge enhancement, is still to be done. The second part of our method proposes the use of anisotropic diffusion to drive a regularization process using the restored direction field as a prior to describe the local magnitude image structure, the main idea being to drive the smoothing along the direction field, modulated by a measure of coherence, similar to the coherence enhancing scheme proposed by Weickert (1999a). On the practical side we proposed the use of the diffusion tensor itself to define the flow tensor that drives the anisotropic diffusion tensor, as a way to take into account the structural information in the image and drive the regularization consistently with the underlying tissues. For the PDD field restoration, we proposed the use of realistic synthetic data to evaluate the adequate value of the regularization parameter.

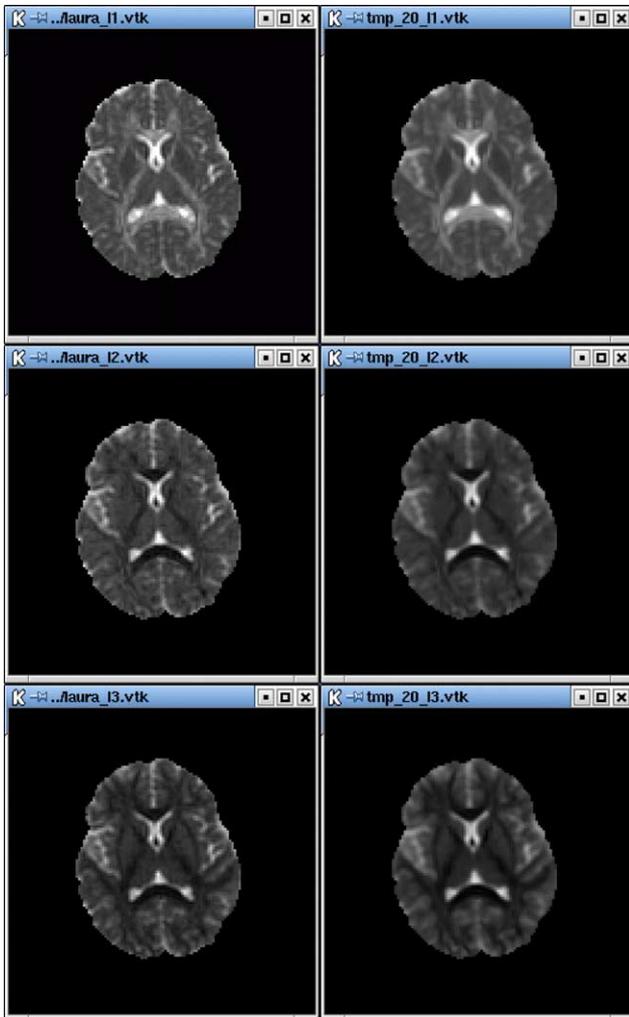


Fig. 13. Eigenvalue map regularization. Left column: λ_1 , λ_2 and λ_3 . Right column: the corresponding regularized images at time $t = 4$. (This figure is available in colour, see the on-line version.)

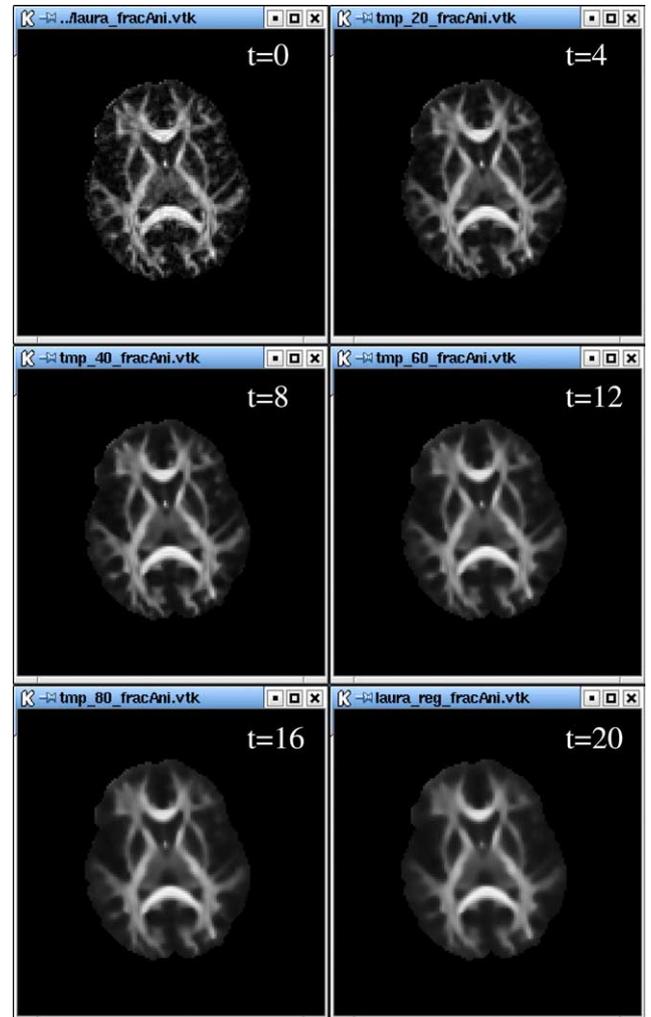


Fig. 15. Fractional anisotropy map at time $t = 0, 4, 8, 16, 12, 16$ and 20 . (This figure is available in colour, see the on-line version.)

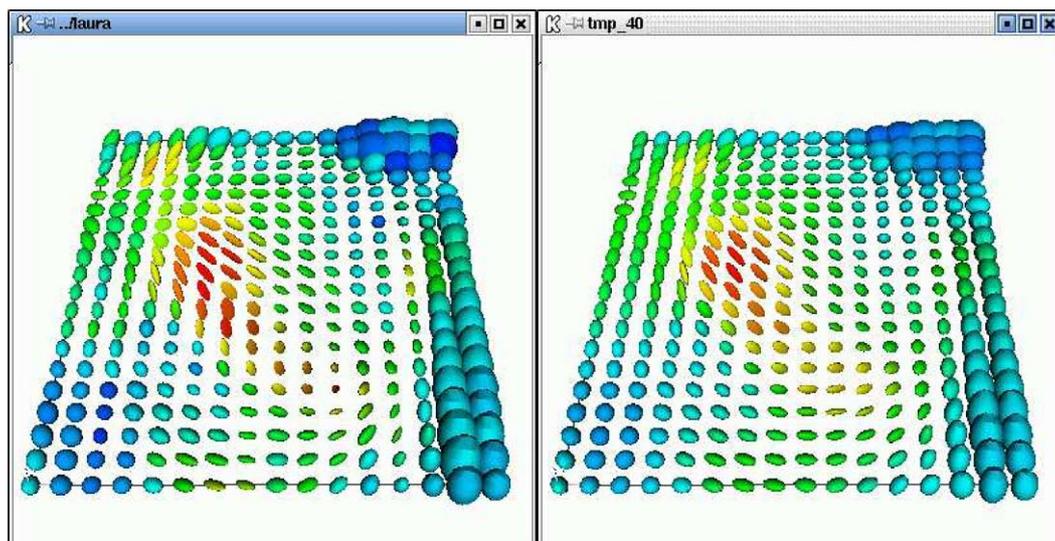


Fig. 14. Ellipsoid representation of brain data (sub-region between ventricule and caudate nucleus). Left: original image, right: after regularization at time $t = 8$ (same colour mapping as in Fig. 2). (This figure is available in colour, see the on-line version.)

The scheme was shown to work well on diffusion tensor magnetic resonance data. Validation of the pro-

cedure where $\mathbf{J}(f)$ stands for the Jacobian of the mapping to the sphere:

$$\mathbf{J}(f) = \begin{pmatrix} \frac{1}{\sqrt{x^2 + y^2 + z^2}} - \frac{x^2}{(x^2 + y^2 + z^2)^{3/2}} & -\frac{xy}{(x^2 + y^2 + z^2)^{3/2}} & -\frac{xz}{(x^2 + y^2 + z^2)^{3/2}} \\ -\frac{xy}{(x^2 + y^2 + z^2)^{3/2}} & \frac{1}{\sqrt{x^2 + y^2 + z^2}} - \frac{y^2}{(x^2 + y^2 + z^2)^{3/2}} & -\frac{yz}{(x^2 + y^2 + z^2)^{3/2}} \\ -\frac{xz}{(x^2 + y^2 + z^2)^{3/2}} & -\frac{yz}{(x^2 + y^2 + z^2)^{3/2}} & \frac{1}{\sqrt{x^2 + y^2 + z^2}} - \frac{z^2}{(x^2 + y^2 + z^2)^{3/2}} \end{pmatrix}. \quad (\text{A.3})$$

cess using applications of DT-MRI is the subject of current work. It includes statistics on scalar measures derived from the tensor in various regions of the brain, as well as tractography studies on spinal cord and brain images to assess the effects of the PDD regularization. An alternative to post-reconstruction regularization is to integrate a spatial regularization constraint in the tensor estimation process, and it is also the subject of current work. Such methods should allow acquisition constraints to be reduced. In particular, we aim to provide a way to acquire images with a higher spatial resolution while maintaining a reasonable acquisition time. When keeping a constant acquisition time, increasing spatial resolution will decrease the signal-to-noise ratio. Spatial regularization aims at fixing this.

Acknowledgements

This work is funded by the Wellcome Trust. Images were kindly provided by the NMR research unit of the Institute of Neurology, London, UK. The authors acknowledge, in particular, the contributions of Dr. C. Wheeler-Kingshott and Dr. G. Barker, Institute of Neurology, London, as well as from Dr. G. Parker, University of Manchester, for very helpful discussions on the subject.

Appendix A. Gradient restricted to \mathcal{S}^2

We define the mapping from \mathbb{R}^3 to \mathcal{S}^2 :

$$\forall f \in \mathbb{R}^3, \quad \hat{f} = \frac{f}{\|f\|}. \quad (\text{A.1})$$

For a scalar function $G: \mathbb{R}^3 \rightarrow \mathbb{R}$, we define its restriction \hat{G} on \mathcal{S}^2 with $\hat{G}(f) = G(\hat{f})$. Let $\nabla G = ((\nabla G)_x, (\nabla G)_y, (\nabla G)_z)^T$ the gradient on \mathbb{R}^3 , and define $\partial \hat{G} / \partial f$ the gradient on \mathcal{S}^2 :

$$\forall f = (x, y, z) \in \mathbb{R}^3, \quad \frac{\partial \hat{G}(f)}{\partial f} = \mathbf{J}(f) \cdot \nabla G(\hat{f}), \quad (\text{A.2})$$

Therefore,

$$\forall f \in \mathbb{R}^3,$$

$$\begin{aligned} \frac{\partial \hat{G}(f)}{\partial f} &= \begin{pmatrix} (\nabla G)_x \frac{1}{\|f\|} - \frac{(\nabla G)_x x^2 - (\nabla G)_y xy - (\nabla G)_z xz}{\|f\|^3} \\ (\nabla G)_y \frac{1}{\|f\|} - \frac{(\nabla G)_x xy - (\nabla G)_y y^2 - (\nabla G)_z yz}{\|f\|^3} \\ (\nabla G)_z \frac{1}{\|f\|} - \frac{(\nabla G)_x xz - (\nabla G)_y yz - (\nabla G)_z z^2}{\|f\|^3} \end{pmatrix} \\ &= \frac{1}{\|f\|} \left(\nabla G - (\nabla G \cdot \frac{f}{\|f\|}) \frac{f}{\|f\|} \right), \end{aligned} \quad (\text{A.4})$$

In particular, for $f \in \mathcal{S}^2$,

$$\frac{\partial \hat{G}(f)}{\partial f} = \nabla G - (\nabla G \cdot f)f = \Pi_f(\nabla G), \quad (\text{A.5})$$

where Π_f is the projection on the plane orthogonal to f .

References

- Aldroubi, A., Basser, P., 1999. Reconstruction of vector and tensor field from sampled discrete data. *Contemp. Math.* 247, 1–15.
- Alexander, D., 2001. Notes on indices of shapes and similarity for diffusion tensors. Technical Report RN/01/14, UCL, Department of Computer Science. <http://www.cs.ucl.ac.uk/staff/D.Alexander/Papers/rn14.pdf>.
- Alexander, D., Barker, G., Arridge, S., 2002. Detection and modelling of non-gaussian apparent diffusion coefficient profiles in clinical human brain data. *Magn. Reson. Med.* 48, 331–340.
- Alvarez, L., Lions, P., Morel, J., 1992. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.* 29 (3), 845–866.
- Alvarez, L., Guichard, F., Lions, P., Morel, J., 1993. Axioms and fundamental equations of image processing. *Arch. Rational Mech. Anal.* 123, 3.
- Basser, P., Pajevic, S., 2000. Statistical artefacts in diffusion tensor MRI (DT-MRI) caused by background noise. *Magn. Reson. Med.* 44, 41–50.
- Basser, P., Pierpaoli, C., 1996. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J. Magn. Reson. B* 111, 209–219.
- Basser, P., Mattiello, J., LeBihan, D., 1994. MR diffusion tensor spectroscopy and imaging. *Biophys. J.* 66, 259–267.
- Bertalmio, M., Cheng, L.-T., Osher, S., Sapiro, G., 2000. Variational problems and partial differential equations on implicit surfaces. Technical Report CAM-TR 00-23, UCLA.

- Besag, J., 1974. Spatial interaction and the statistical analysis of lattice systems. *J. R. Statist. Soc. B* 36, 192–236.
- Caselles, V., Morel, J., Sapiro, G., Tannenbaum, A., 1998. Special issue on PDEs and geometry-driven diffusion in image processing. *IEEE Trans. Image Process.* 7 (3).
- Catte, F., Coll, T., Lions, P., Morel, J., 1992. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.* 29 (1), 182–193.
- Chan, T., Shen, J., 2000. Variational restoration of non-flat image features: model and algorithm. *SIAM J. Appl. Math.* 61 (4), 1338–1361.
- Conturo, T., Lori, N., Cull, T., Akbudak, E., Snyder, A., Shimony, J., McKinstry, R., Burton, H., Raichle, M., 1999. Tracking neuronal fiber pathways in the living human brain. *Proc. Natl. Acad. Sci. USA* 96, 10422–10427.
- Coulon, O., Arridge, S., 2000. Dual echo MR image processing using multi-spectral probabilistic diffusion coupled with shock filters. In: *British Conference on Medical Image Understanding and Analysis – MIUA'2000*, London, pp. 141–144.
- Coulon, O., Alexander, D., Arridge, S., 2001. A regularisation scheme for diffusion tensor magnetic resonance images. In: *IPMI'01, International Conference on Information Processing in Medical Imaging Lecture Notes in Computer Science*, Vol. 2082. Springer, Davis, pp. 92–105.
- Geman, S., McLure, D., 1985. Bayesian image analysis: an application to single photon emission tomography. *Am. Statist. Assoc.*, 12–18.
- Green, P., 1990. On use of the EM algorithm for penalized likelihood estimation. *J. R. Statist. Soc.* 52 (3), 443–452.
- Jones, D., Horsfield, M., Simmons, A., 1999. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn. Reson. Med.* 42, 515–525.
- Kimmel, R., 2001. Intrinsic scale space for images on surfaces: the geodesic curvature flow. In: *ScaleSpace'97, First International Conference on Scale-Space Theory in Computer Vision Lecture Notes in Computer Science*, 1252. Springer, Davis, pp. 212–223.
- Kimmel, R., Sochen, N., 2000. How to comb a porcupine? Technical Report CIS 2000-02, Technion.
- Kirsch, A., 1996. An introduction to the mathematical theory of inverse problems. In: *Applied Mathematical Sciences*, Vol. 120. Springer, Berlin.
- Koenderink, J., 1984. The structure of images. *Biol. Cybernet.* 50, 363–370.
- Kornprobst, P., 1998. Contributions à la restauration d'images et à l'analyse de séquences: approches variationnelles et équations aux dérivées partielles. Ph.D. Thesis, Université de Nice-Sophia Antipolis.
- Krissian, K., Malandain, G., Ayache, N., 1997. Directional anisotropic diffusion applied to segmentation of vessels in 3D images. In: *Scale-Space'97. Lecture Notes in Computer Science*, Vol. 1251. Springer, Berlin, pp. 345–348.
- Lebihan, D., mangin, J.-F., Poupon, C., Clark, C., Pappata, S., Molko, N., Chabriat, H., 2001. Diffusion Tensor Imaging: concepts and applications. *J. Magn. Reson. Imaging* 13, 534–546.
- Mohiaddin, R., Yang, G., Kilner, P., 1994. Visualization of flow by vector analysis of multidirectional cine MR velocity mapping. *J. Comput. Assist. Tomogr.* 18, 383–392.
- Nielsen, M., Florack, L., Deriche, R., 1997. Regularization, scale-space and edge detection filters. *J. Math. Imaging Vision* 7 (4), 291–308.
- Niessen, W., ter Haar Romeny, B., Viergever, M., 1994. Numerical analysis of geometry-driven diffusion equations. In: ter Haar Romeny, B.M. (Ed.), *Geometry-Driven Diffusion in Computer Vision*. Kluwer Academic Publishers, Dordrecht, pp. 393–410.
- Parker, G., Schnabel, J., Symms, M., Werring, D., Barker, G., 2000. Nonlinear smoothing for reduction of systematic and random errors in diffusion tensor imaging. *J. Magn. Reson. Imaging* 11, 702–710.
- Parker, G., Wheeler-Kingshott, C., Barker, G., 2001. Distributed anatomical brain connectivity derived from diffusion tensor imaging. In: *IPMI'01, International Conference on Information Processing in Medical Imaging Lecture Notes in Computer Science*, Vol. 2082. Springer, Davis, pp. 106–120.
- Perona, P., 1998. Orientation diffusion. *IEEE Trans. Image Process.* 7 (3), 457–467.
- Perona, P., Malik, J., 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* 12 (7), 629–639.
- Pierpaoli, C., Basser, P., 1996. Toward a quantitative assessment of diffusion anisotropy. *Magn. Reson. Med.* 36, 893–906.
- Pierpaoli, C., Jezzard, P., Basser, P., Barnett, A., 1996. Diffusion tensor imaging of the human brain. *Radiology* 201, 637–648.
- Poupon, C., Mangin, J., Clark, C., Frouin, V., Régis, J., LeBihan, D., Bloch, I., 2001. Towards inference of human brain connectivity from MR diffusion tensor data. *Med. Image Anal.* 5, 1–15.
- Radmoser, E., Scherzer, O., Weickert, J., 2000. Scale-space properties of nonstationary iterative regularisation methods. *J. Visual Commun. Image Representation* 11, 96–114.
- Rudin, L., Osher, S., Fatemi, E., 1992. Nonlinear total variation based noise removal algorithms. *Physica D* 60, 259–268.
- Sapiro, G., 1995. Geometric partial differential equations in image analysis: past present and future. In: *IEEE Conference on Image Processing*, Washington, DC.
- Sapiro, G., Ringach, D., 1996. Anisotropic diffusion of multivalued images with applications to colour filtering. *IEEE Trans. Image Process.* 5 (11), 1582–1586.
- Sapiro, G., Tannenbaum, A., 1993. Affine invariant scale-space. *Int. J. Comput. Vision* 11, 25–44.
- Scherzer, O., Weickert, J., 2000. Relations between regularization and diffusion filtering. *J. Math. Imaging Vision* 12, 43–63.
- Shah, J., 1996. A common framework for curve evolution, segmentation, and anisotropic diffusion. In: *IEEE Conference on Computer Vision and Pattern Recognition*, San Francisco, CA, pp. 136–142.
- Tang, B., Sapiro, G., Caselles, V., 1999. Color image enhancement via chromaticity diffusion. Technical Report, ECE, University of Minnesota.
- Tang, B., Sapiro, G., Caselles, V., 2000. Diffusion of general data on non-flat manifolds via harmonic maps theory: The direction diffusion case. *Int. J. Comput. Vision* 36 (2), 149–161.
- Teboul, S., Blanc-Féraud, L., Aubert, G., Barlaud, M., 1998. Variational approach for edge preserving regularization using coupled PDE's. *IEEE Trans. Image Process.* 7 (3), 387–397.
- Trahanias, P., Karako, D., Venetsanopoulos, A., 1996. Directional processing of colour images: theory and experimental results. *IEEE Trans. Image Process.* 5 (6), 868–880.
- Tschumperlé, D., Deriche, R., 2001a. Constrained and unconstrained PDEs for vector image restoration. In: *12th Scandinavian Conference on Image Analysis*. Bergen, Norway.
- Tschumperlé, D., Deriche, R., 2001b. Regularization of orthonormal vector sets using coupled PDE's. In: *IEEE Workshop on Variational and Level Set Methods in Computer Vision*, Vancouver, Canada.
- Vemuri, B., Chen, Y., Rao, M., McGraw, T., Wang, Z., Mareci, T., 2001. Fiber tract mapping from diffusion tensor MRI. In: *IEEE Workshop on Variational and Level Set Methods in Computer Vision*. Vancouver, Canada.
- Weickert, J., 1997a. Recursive separable schemes for nonlinear diffusion filters. In: *First International Conference on Scale-Space Theory in Computer Vision – ScaleSpace'97 Lecture Notes in Computer Science*, Vol. 1252. Springer, Utrecht, pp. 260–271.

- Weickert, J., 1997b. A review of non-linear diffusion filtering. In: First International Conference on Scale-Space Theory in Computer Vision – ScaleSpace'97 Lecture Notes in Computer Science, Vol. 1252. Springer, Utrecht, pp. 3–28.
- Weickert, J., 1998. Anisotropic Diffusion in Image Processing. Teubner, Stuttgart.
- Weickert, J., 1999a. Coherence-enhancing diffusion filtering. *Int. J. Comput. Vision* 31 (2/3), 111–127.
- Weickert, J., 1999b. Coherence-enhancing diffusion of colour images. *Image Vision Comput.* 17, 201–212.
- Weickert, J., Schnörr, C., 2001. Variational optic flow computation with a spatio-temporal smoothness constraint. *J. Math. Imaging Vision* 14 (3), 245–255.
- Whitaker, R., Gerig, G., 1994. Vector-valued diffusion. In: ter Haar Romeny, B. (Ed.), *Geometry-driven Diffusion in Computer Vision*. Kluwer Academic Publishers, Dordrecht, pp. 93–134.
- You, Y.-L., Xu, W., Tannenbaum, A., Kaveh, M., 1996. Behavioral analysis of anisotropic diffusion in image processing. *J. Visual Commun. Image Representation* 11 (5), 1539–1553.