



Structure-adaptive sparse denoising for diffusion-tensor MRI

Lijun Bao^{a,*}, Marc Robini^b, Wanyu Liu^{a,b,*}, Yuemin Zhu^b

^aHIT-INSa Sino French Research Centre for Biomedical Imaging, Harbin Institute of Technology, Harbin 150006, China

^bCREATIS, CNRS UMR 5220, INSERM U1044, INSA Lyon, University of Lyon, Villeurbanne Cedex 69621, France

ARTICLE INFO

Article history:

Received 21 October 2011

Received in revised form 23 January 2013

Accepted 28 January 2013

Available online 17 February 2013

Keywords:

Diffusion tensor MRI

Denoising

Block matching

Structure-adaptive grouping

Sparsity

ABSTRACT

Diffusion tensor magnetic resonance imaging (DT-MRI) is becoming a prospective imaging technique in clinical applications because of its potential for in vivo and non-invasive characterization of tissue organization. However, the acquisition of diffusion-weighted images (DWIs) is often corrupted by noise and artifacts, and the intensity of diffusion-weighted signals is weaker than that of classical magnetic resonance signals. In this paper, we propose a new denoising method for DT-MRI, called structure-adaptive sparse denoising (SASD), which exploits self-similarity in DWIs. We define a similarity measure based on the local mean and on a modified structure-similarity index to find sets of similar patches that are arranged into three-dimensional arrays, and we propose a simple and efficient structure-adaptive window pursuit method to achieve sparse representation of these arrays. The noise component of the resulting structure-adaptive arrays is attenuated by Wiener shrinkage in a transform domain defined by two-dimensional principal component decomposition and Haar transformation. Experiments on both synthetic and real cardiac DT-MRI data show that the proposed SASD algorithm outperforms state-of-the-art methods for denoising images with structural redundancy. Moreover, SASD achieves a good trade-off between image contrast and image smoothness, and our experiments on synthetic data demonstrate that it produces more accurate tensor fields from which biologically relevant metrics can then be computed.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Magnetic resonance imaging (MRI) has benefited from many technological developments since its advent in the 1970s. However, the fundamental trade-offs between image resolution and signal-to-noise ratio (SNR) on the one hand, and between physiological and clinical constraints on acquisition speed on the other hand, often translate to spurious artifacts such as noise, partial volume, and bias field (Basser and Pajevic, 2000; Farrell et al., 2007). An eloquent example is diffusion-tensor MRI (DT-MRI), which has become quite popular over the last decade because of its potential for in vivo and non-invasive characterization of the three-dimensional (3D) fiber architecture of anatomical organs (Behrens et al., 2003; Bondiau et al., 2008; Clatz et al., 2005; Delingette et al., 2012; Deriche et al., 2009; Descoteaux et al., 2009; Durrleman et al., 2011; Fillard et al., 2011; Galanaud et al., 2010; Galban et al., 2005; Guevara et al., 2011; Le Bihan, 2003; Lenglet et al., 2009; Messe et al., 2011; Mori et al., 2009; Rohmer et al., 2007; Smith et al., 2006; Wakana et al., 2004; Wu et al., 2006). The effects of Rician noise on DT-MRI are severe because of the inherent nature of the imaging process—the higher the tissue anisotropy, the lower the intensity in the

diffusion-weighted images (DWIs), and hence the higher the sensitivity to noise (Awate and Whitaker, 2007).

Denoising magnetic resonance images is an important problem; the most popular approaches are Bayesian statistical approaches (Awate and Whitaker, 2007; Basser and Pajevic, 2003; Lu et al., 2006), PDE-based approaches (Chen and Edward, 2005; Fillard et al., 2007), wavelet-based methods (Pižurica et al., 2006; Yu et al., 2009), methods based on spatial correlation (Barash and Comaniciu, 2004; Kervrann and Boulanger, 2006; Manjón et al., 2008), and sparse representation denoising (Chatterjee and Milanfar, 2009; Donoho et al., 2006; Elad and Aharon, 2006; Varshney et al., 2008). No particular method shows good performance for all relevant aspects of MRI, including bias reduction, artifacts removal, structure- and edge-preservation, and good generality/reliability trade-off.

Buades et al. (2005) surveyed spatial correlation techniques and proposed the so-called non-local means (NL-means) method which is based on the assumption that natural images usually have structural redundancy. The underlying idea is that any pixel has similar pixels that are distributed not only in its local neighborhood but also in the whole spatial domain of the image, which allows to account for the global information associated with large structures, and hence to overcome the limitations of local search. The NL-means approach has been applied to denoising (Boulanger et al., 2010; Brox et al., 2008; Katkovnik et al., 2010; Katkovnik and Spokoiny, 2008), super-resolution restoration (Manjón et al.,

* Corresponding authors. Address: HIT-INSa Sino French Research Centre for Biomedical Imaging, Harbin Institute of Technology, Harbin 150006, China.

E-mail address: baolijun@xmu.edu.cn (L. Bao).

2010; Rousseau, 2010), compressed sensing (Danielyan et al., 2010; Marim et al., 2010) and inpainting (Elad et al., 2005). NL-means detects pixel similarity by exploring patch similarity for better robustness to noise, and similar patches can be grouped together to achieve sparse representation in a suitable transform domain in which denoising reduces to a shrinkage operation. Therefore, non-local self-similarity and sparse representations are key elements of state-of-the-art filtering algorithms.

The similarity measure used in NL-means evaluates the correlation between patches globally via the standard Euclidean distance and neglects the local consistency between a patch and its center pixel. In particular, in the case where different patches are close in terms of Euclidean distance but have center pixels with quite different values, the standard similarity measure may cause a loss of fine structures and edge blurring – the same behavior is observed when the number of patches similar to the reference patch is too small. Alternatively, the block-matching algorithm (BM3D) proposed by Dabov et al. (2007), which exploits a specific non-local image model by grouping similar patches and by collaborative filtering, is another interesting denoising approach. The BM3D filter stacks similar two-dimensional (2D) image patches into 3D arrays which are processed in a 3D transform domain to attenuate the noise component. However, patches containing fine image details or singularities or sharp edges are examples where a non-adaptive transform is not able to produce good sparse representations (Foi et al., 2007; Hammond and Simoncelli, 2008; Sikora and Makai, 1995); for such patches, the filter may introduce artifacts around discontinuities where the visual attention is often mainly focused.

We propose a 3D structure-adaptive sparse denoising (SASD) algorithm for DT-MRI which we apply to human cardiac DWIs. Our method can be decomposed into four steps. First, we form groups of similar patches in the DWI to be denoised using a modified structure-similarity index. Second, each set of similar patches is arranged into a 3D array, and a structure-adaptive window pursuit method is used to adapt to the local image features. Third, the resulting structure-adaptive 3D arrays are denoised by Wiener shrinkage in the transform domain defined by 2D principal component analysis in the image domain and Haar transformation in the third dimension. Finally, the noise-free DWI estimate is obtained by weighted averaging of the denoised structure-adaptive patches.

The paper is organized as follows. Section 2 describes the SASD algorithm and its main components, including the definition of the proposed similarity measure, the design of the structure-adaptive window pursuit method, and the Wiener filtering operation in the transform domain. The experimental setup is discussed in Section 3, and our results are presented in Section 4, where our approach is compared to state-of-the-art denoising methods, namely, Bayesian least-squares Gaussian scale mixture (BLS-GSM) (Portilla et al., 2003) and Field of experts (FOE) (Roth and Black, 2005). Concluding remarks are given in Section 5.

2. Methodology

In a nutshell, the proposed approach is inspired by the NL-means method to filter sparse local representations of the image to be denoised. Therefore, we describe the standard NL-means method in Section 2.1 prior to our SASD algorithm in Section 2.2.

2.1. Standard NL-means

The NL-means method is based on pointwise estimation: a given pixel is denoised by weighted averaging of the pixels with similar neighborhoods in the noisy image. Let $\mathbf{Y} = \{\mathbf{Y}(i); i \in \Omega\}$ be an image to be denoised, where i is the pixel index (Ω is the spatial

image domain) and $\mathbf{Y}(i)$ is the gray level intensity at pixel i . The NL-means estimate of \mathbf{Y} is defined by the weighted average

$$\text{NL}(\mathbf{Y}(i)) = \sum_{j \in \Omega} w(i, j) \mathbf{Y}(j), \quad (1)$$

where $w(i, j)$ measures the similarity between pixels i and j and between their neighbors (the $w(i, j)$'s are in $(0, 1)$ and satisfy $\sum_j w(i, j) = 1$ for all i). Given a neighborhood system $\{N_i; i \in \Omega\}$ (N_i denotes the subset of Ω that contains the neighbors of pixel i), the weights $w(i, j)$ are computed as follows:

$$w(i, j) = \frac{1}{Z(i)} \exp \left\{ -\frac{\|\mathbf{Y}(N_i) - \mathbf{Y}(N_j)\|_2^2}{2h^2} \right\}, \quad (2)$$

where $\mathbf{Y}(N_k)$ denotes the patch centered in pixel k (N_k is a rectangular pixel region centered on k), $Z(i)$ is a normalization factor, and h controls the decay of the exponential function (this smoothing parameter is set equal to σn , where σ is the standard deviation of the noise and n is the patch size).

The standard similarity measure (2) only uses the standard Euclidean distance between patches and hence does not account for the local consistency between pixels and their neighbors. Consequently, two patches $\mathbf{Y}(N_i)$ and $\mathbf{Y}(N_j)$ can be similar in terms of Euclidean distance even though $\mathbf{Y}(i) - \mathbf{Y}(j)$ is large, and thus it can happen that $w(i, j)$ is large for patches centered on pixels with very different values. Furthermore, in the case of a reference patch having a very small number of similar patches, the constraint that the sum of the $w(i, j)$'s over j must be equal to one gives non-negligible weights to dissimilar patches, which produces distortion in the denoised image. Consequently, NL-means denoising based on the similarity measure (2) does not properly recover image details such as fine structures and tends to blur edges.

2.2. Structure-adaptive sparse denoising

The above discussion on NL-means suggests two immediate levers for improving patch-based denoising: first, the patch-similarity measure should ensure local consistency, and second, the weights used for averaging should preclude dissimilar patches. Besides, when using patches of fixed shape and size, the performance is limited by the lack of sparsity around singularities or edges, which translates to ringing artifacts in the denoised estimates. Hence the idea to replace the usual fixed-size rectangular patches with structure-adaptive patches for more effective sparse representation and thus better quality estimates.

This leads to the proposed SASD algorithm schematized in Fig. 1. This algorithm relies on a new constrained similarity measure for patch matching – the structural similarity (SSIM) index – and on a new structure-adaptive window (SAW) pursuit method, both of which use the non-stationary degree (NSD) operator of Liu et al. (1995). The SSIM index and the SAW pursuit method are the building blocks to form structure-adaptive 3D arrays of similar patches, which are then denoised in the domain of a 3D separable transform that promotes sparsity. (This transform is defined as follows: the structure-adaptive patches of a given group are projected onto principal components, and a 1D Haar transformation is performed in the dimension along which the patches are stacked.) We proceed with the formal presentation of the SASD approach, leaving the descriptions of the similarity measure and the SAW pursuit method to Sections 2.2.2 and 2.2.3, respectively. The full algorithm is given in Section 2.2.4.

2.2.1. Formal description of SASD

The core of the SASD method consists in performing the following three steps for each pixel i . We assume that the patches

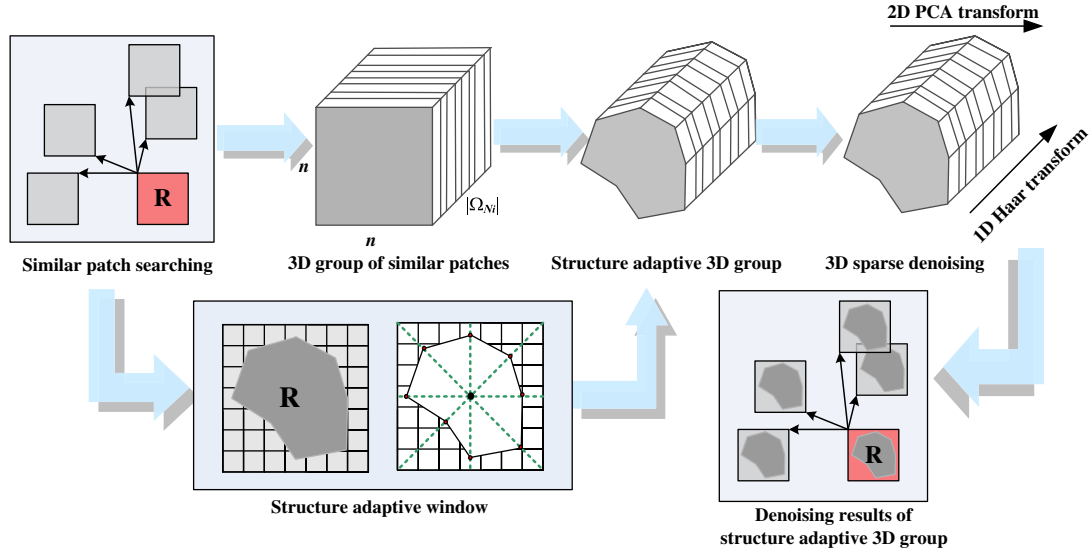


Fig. 1. Schematic representation of the SASD algorithm (R stands for the current reference patch).

are defined on square regions of size n , and we denote by N_k the set of pixels of the patch centered on pixel k .

- First, the set Ω_{N_i} of patches similar to the reference patch $\mathbf{Y}(N_i)$ is formed using the proposed similarity measure, and the elements of Ω_{N_i} are stacked into a 3D rectangular array $\mathbf{S}_{\mathbf{Y}(N_i)} = \{\mathbf{Y}(N_j) : j \in \Omega_{N_i}\}$. This array is of size $n \times n \times |\Omega_{N_i}|$, and Ω_{N_i} is defined by

$$\Omega_{N_i} = \left\{ j \in \Omega \mid \mu < \frac{\overline{\mathbf{Y}(N_i)}}{\mathbf{Y}(N_j)} < \frac{1}{\mu} \text{ and } v < \text{SSIM}(\mathbf{D}(N_i), \mathbf{D}(N_j)) \right\}, \quad (3)$$

where $\overline{\mathbf{Y}(N_i)}$ is the mean of \mathbf{Y} in N_i , \mathbf{D} is the NSD operator, and $\mu, v \in (0, 1)$ are given thresholds (the rationale for this choice is given in Section 2.2.2, where we also discuss the setting of μ and v).

- Second, the structure-adaptive neighborhood N_i^{SA} of pixel i (which is a subset of N_i) is computed by SAW pursuit and is used as a template for creating a structure-adaptive set of patches $\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}} = \{\mathbf{Y}(N_j^{\text{SA}}) : j \in \Omega_{N_i}\}$, where $N_j^{\text{SA}} \subseteq N_j$ is the translation of N_i^{SA} by $j - i$ (in pixel coordinates).
- Third, the noise component in $\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}}$ is attenuated after a 3D transformation which consists in first decomposing $\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}}$ in a 2D principal component (PC) basis and then performing a 1D Haar wavelet transform along the dimension perpendicular to the stacked PCs (we use PC bases because they are more adaptive to image features than a fixed dictionary and because they do not require its construction). The actual denoising is performed by Wiener filtering of the PC projection coefficients in the wavelet domain, and a subsequent 3D inverse transformation gives a structure-adaptive set of patch estimates $\mathbf{S}_{\mathbf{X}(N_i)}^{\text{SA}} = \{\hat{\mathbf{X}}(N_j^{\text{SA}}) : j \in \Omega_{N_i}\}$. Formally, letting $\Gamma_{3\text{D}}$ and $\Gamma_{3\text{D}}^{-1}$ be the direct and inverse 3D transforms, we have

$$\mathbf{S}_{\mathbf{X}(N_i)}^{\text{SA}} = \Gamma_{3\text{D}}^{-1} \left(W \left(\Gamma_{3\text{D}} \left(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}} \right) \right) \right), \quad (4)$$

where W denotes the Wiener filtering operator – Wiener denoising is an element-by-element multiplication of the Fourier spectrum of $\Gamma_{3\text{D}} \left(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}} \right)$, say $\overline{\Gamma_{3\text{D}} \left(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}} \right)}$, by

$$\frac{|\overline{\Gamma_{3\text{D}} \left(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}} \right)}|^2}{|\overline{\Gamma_{3\text{D}} \left(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}} \right)}|^2 + \sigma}, \quad (5)$$

where σ is the noise standard deviation.

Once the above steps have been performed for all pixels, the denoised DWI, say $\hat{\mathbf{X}}$, is obtained by cumulative weighted average of the structure-adaptive denoised patches, that is,

$$\hat{\mathbf{X}}(k) = \frac{\sum_{i \in \Omega} \sum_{j \in \Omega_{N_i}} w(i) \hat{\mathbf{X}}_k(N_j^{\text{SA}})}{\sum_{i \in \Omega} \sum_{j \in \Omega_{N_i}} w(i) \chi_j(k)}, \quad (6)$$

where $\hat{\mathbf{X}}_k(N_j^{\text{SA}})$ denotes the value of the pixel in $\hat{\mathbf{X}}(N_j^{\text{SA}})$ that corresponds to pixel k in the whole image (by convention, we set $\hat{\mathbf{X}}_k(N_j^{\text{SA}}) = 0$ if such a pixel does not exist), and where the indicator function χ_j is defined by $\chi_j(k) = 1$ if pixel k is in N_j^{SA} , 0 otherwise. The weights $w(i)$ are given by

$$w(i) = \begin{cases} \frac{1}{\sigma^2 N_C(i)} & \text{if } N_C(i) \geq 1, \\ 1 & \text{otherwise,} \end{cases} \quad (7)$$

where $N_C(i)$ denotes the number of “effective” PCs representing $\mathbf{Y}(N_i)$; more precisely, letting $\mathbf{V}(i) = \{V_1(i), \dots, V_K(i)\}$ be the PC basis associated with $\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}}$ and $(\alpha_1(i), \dots, \alpha_K(i))$ be the coordinate vector of $\mathbf{Y}(N_i^{\text{SA}})$ in $\mathbf{V}(i)$, $N_C(i)$ is the smallest integer p such that

$$\left\| \mathbf{Y}(N_i^{\text{SA}}) - \sum_{k=1}^p \alpha_k(i) V_k(i) \right\|_2 \leq n\sigma. \quad (8)$$

2.2.2. A new similarity measure

We propose to improve the robustness of patch-similarity measurement in two ways: first, we increase the performance of the SSIM index of Wang et al. (2004) by using the NSD operator of Liu et al. (1995), and second, we select similar patches by thresholding this new index and by placing bounds on the mean of the candidate patches. Another advantage of this patch selection scheme is that it limits the sizes of the groups of similar patches, thereby reducing the computational burden.

The SSIM index measures the similarity between two patches \mathbf{x} and \mathbf{y} as follows:

$$\text{SSIM}(\mathbf{x}, \mathbf{y}) = \frac{(2\mu_{\mathbf{x}}\mu_{\mathbf{y}} + C_1)(2\sigma_{\mathbf{xy}} + C_2)}{(\mu_{\mathbf{x}}^2 + \mu_{\mathbf{y}}^2 + C_1)(\sigma_{\mathbf{x}}^2 + \sigma_{\mathbf{y}}^2 + C_2)}, \quad (9)$$

where $\mu_{\mathbf{x}}$ and $\mu_{\mathbf{y}}$ are the means of \mathbf{x} and \mathbf{y} , $\sigma_{\mathbf{x}}^2$ and $\sigma_{\mathbf{y}}^2$ are the variances of \mathbf{x} and \mathbf{y} , and $\sigma_{\mathbf{xy}}$ is the covariance between \mathbf{x} and \mathbf{y} . The role of the constants C_1 and C_2 is to avoid instability when the means $\mu_{\mathbf{x}}$ and $\mu_{\mathbf{y}}$ or the variances $\sigma_{\mathbf{x}}^2$ and $\sigma_{\mathbf{y}}^2$ are close to zero. The problem with this index is that its sensitivity to noise is too high for our application (we made this observation by comparing SSIM index maps obtained from simulated cardiac DWIs with and without noise). Therefore, since structural image characteristics usually take the form of non-stationary points (i.e., edges), we propose to apply the SSIM index to the NSDs $\mathbf{D}(\mathbf{x})$ and $\mathbf{D}(\mathbf{y})$ of \mathbf{x} and \mathbf{y} , that is,

$$\text{SSIM}(\mathbf{D}(\mathbf{x}), \mathbf{D}(\mathbf{y})) = \frac{(2\mu_{\mathbf{D}(\mathbf{x})}\mu_{\mathbf{D}(\mathbf{y})} + C_1)(2\sigma_{\mathbf{D}(\mathbf{x})\mathbf{D}(\mathbf{y})} + C_2)}{(\mu_{\mathbf{D}(\mathbf{x})}^2 + \mu_{\mathbf{D}(\mathbf{y})}^2 + C_1)(\sigma_{\mathbf{D}(\mathbf{x})}^2 + \sigma_{\mathbf{D}(\mathbf{y})}^2 + C_2)}. \quad (10)$$

The NSD operator \mathbf{D} is defined by

$$\mathbf{D}(\mathbf{x}) = (\mathbf{x} * h)^2 * h - (\mathbf{x} * h * h)^2, \quad (11)$$

where $*$ denotes the 2D convolution operator and h is a rectangular filter of size $M \times M$:

$$h(i, j) = g(i)g(j), \quad \text{with} \quad g(i) = \frac{1}{M} \text{rect}\left(\frac{i}{M}\right). \quad (12)$$

The NSD of a second-order stationary signal is zero, whereas if the input signal is not second-order stationary, the NSD detector gives high output values at discontinuity locations. In other words, NSD detection gives prominence to the structural characteristics of the input image.

In Coupé et al. (2008), the selection of the patches $\mathbf{Y}(N_j)$ similar to a reference patch $\mathbf{Y}(N_i)$ is performed by upper and lower bounding the mean and the variance of $\mathbf{Y}(N_j)$. As the example in Fig. 2 shows, the local mean and variance are simple estimators that allow to discriminate different tissue classes and to detect edges, but the local variance is very sensitive to noise. Compared to the local variance, the NSD detector has better edge-enhancement capabilities,

and it is more immune to noise because it filters the signal before computing the variance (the NSD detector that produced the images in Fig. 2d and h is defined by the 3×3 rectangular filter). Consequently, we propose the following new similarity measure:

$$w(i, j) = \begin{cases} \frac{1}{z(i)} & \text{if } \mu < \frac{\mathbf{Y}(N_i)}{\mathbf{Y}(N_j)} < \frac{1}{\mu} \quad \text{and} \quad v < \text{SSIM}(\mathbf{D}(N_i), \mathbf{D}(N_j)), \\ 0 & \text{otherwise,} \end{cases} \quad (13)$$

where the parameter $\mu \in (0, 1)$ sets the bounds on the local mean ratio, and $v \in (0, 1)$ is the lower threshold for the NSD SSIM index. (Note that setting $Z(i) = \sigma^2 N_c(i)$ if $N_c(i) \geq 1$ and $Z(i) = 1$ otherwise gives the weights defined in (7).) The parameters μ and v constrain the range of acceptable variations of the local mean and the range of possible similarity values. When μ and v tend to zero, the constraints vanish and all the neighborhoods N_j equally contribute to $w(i, j)$.

Conversely, the closer μ and v are to one, the more restrictive the constraints, and the less numerous the number of selected patches (and hence the less the amount of filtering). To fix ideas, we set $\mu = 0.7$ and $v = 0.1$ in our experiments, which is a good trade-off between denoising quality and computation time.

2.2.3. Structure-adaptive patches

A natural idea to improve the performance of patch-based denoising is to use structure-adaptive patches rather than fixed-size rectangular ones. Indeed, structure-adaptive patches allow local adaptation to image features, so that the true signal is nearly homogeneous in such regions, which improves the sparsity of the 3D arrays and eases the denoising process. To efficiently construct such patches, we propose a structure-adaptive window pursuit method that uses NSD detection to find the boundary points. This method, which we call NSD-SAW for non-stationary degree structure-adaptive window, is schematized in Fig. 3. For each considered direction away from the center pixel i_0 , we take the first pixel whose NSD value is greater than some given threshold τ as a boundary point of the structure-adaptive window. The different steps of NSD-SAW are given below (the 2-D NSD map of the whole image to be denoised is precomputed, and given an ordered set $\Theta = (i_0, i_1, \dots, i_m)$ of pixel indexes, we let $\text{NSD}(\Theta)$ be the ordered set $(\text{NSD}(i_0), \text{NSD}(i_1), \dots, \text{NSD}(i_m))$ of associated NSD values).




Fig. 2. Local mean, local variance and NSD of a simulated cardiac DWI: (a and e) the noise-free and noisy DWIs (Rician noise, $\sigma = 30$); (b and f) the local means of (a and e), respectively; (c and g) the local variances; and (d and h) the NSD maps.




Fig. 3. Schematization of the NSD-SAW method.

- (1) Fix an $n \times n$ square patch (with n odd) centered on i_0 . Define the set of search directions as $\{\theta_k = (k - 1)\pi/4; k = 1, \dots, 8\}$.
- (2) For each k , extract the set of pixels in the direction of θ_k starting from i_0 , and use the successive NSD values of these pixels to form the 1D signal $\text{NSD}(\Theta(k))$, $\Theta(k) = (i_0, i_1(-k), \dots, i_{(n-1)/2}(k))$.
- (3) Select the boundary points $i(k)$ in each direction as follows: if all the values in $\text{NSD}(\Theta(k))$ are smaller or equal to τ , then the boundary point in the direction of θ_k is the boundary of the square patch; otherwise, pick the closest pixel i to i_0 such that $i \in \Theta^k$ and $\text{NSD}(i) \geq \tau$.

Fig. 4 illustrates NSD-SAW at work in a cardiac DWI. We can see from Fig. 4a and b that the NSD map of the 21×21 patch shown in

Fig. 4c has a peak-shaped bulge along the edge. The boundary point in the $\pi/2$ direction is $i_6(3)$; this is the non-stationary points with NSD greater than $\tau = 300$ that is the closest to i_0 . In the $3\pi/2$ direction, there is no non-stationary point, and thus the boundary point $i_{10}(7)$ is the boundary of the patch. As shown in Fig. 4d, no matter where the center pixel is located, NSD-SAW can find a homogeneous structure-adaptive neighborhood in coherence with image edges and structures.

2.2.4. Implementation of the SASD algorithm

The SASD algorithm is given below. To reduce the computation time, we restrict the locations of the centers of the reference patches to a set $\Omega_s \subset \Omega$ obtained by downsampling the regular grid supporting the image to be denoised. Therefore, if the image size is $N \times N$ and the downsampling factor in the horizontal and vertical directions is s , the number of processed patches is approximately N^2/s^2 . Furthermore, the search for similar patches can be restricted to a region centered on the center pixel of the reference patch and whose shape and size are fixed in advance.

- (1) For each pixel $i \in \Omega_s$, perform the following steps.
 - (a) Search for the set of square patches similar to $\mathbf{Y}(N_i)$ in the search region centered on pixel i , that is, construct the set Ω_{N_i} defined in (3), and form the 3D rectangular array $\mathbf{S}_{\mathbf{Y}(N_i)} = \{\mathbf{Y}(N_j); j \in \Omega_{N_i}\}$.
 - (b) Compute the structure-adaptive neighborhood N_i^{SA} of pixel i using the NSD-SAW method and extract the structure-adaptive 3D array $\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}} = \{\mathbf{Y}(N_j^{\text{SA}}); j \in \Omega_{N_i}\}$ from $\mathbf{S}_{\mathbf{Y}(N_i)}$.

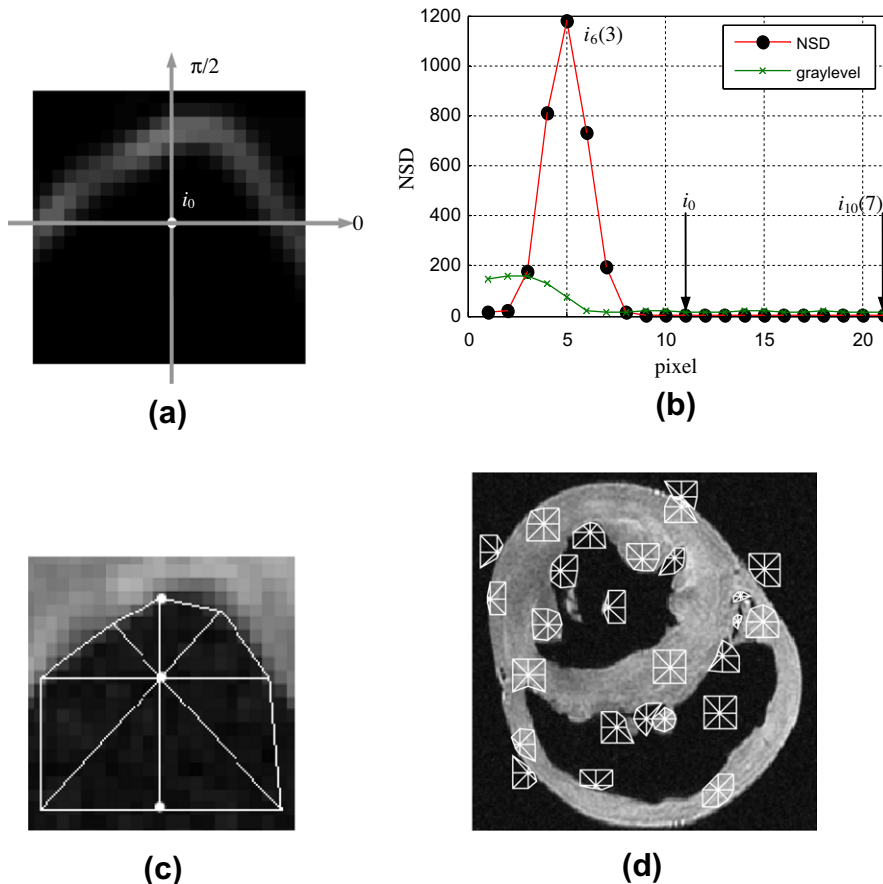


Fig. 4. Structure-adaptive window search in a cardiac DWI: (a) NSD map of the 21×21 patch shown in (c); (b) NSD profiles in the vertical direction (red curve) and associated patch intensity profiles (green curve); (c) computed structure-adaptive window; (d) examples of structure-adaptive neighborhoods. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(c) Decompose $\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}}$ into a PC basis $\mathbf{V}(i) = \{V_1(i), \dots, V_K(i)\}$:

$$\text{PCA}\left(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}}\right) = \{(\alpha_1(j), \dots, \alpha_K(j)); j \in \Omega_{N_i}\},$$

where $(\alpha_1(j), \dots, \alpha_K(j))$ is the coordinate vector of $\mathbf{Y}(N_i^{\text{SA}})$ in $\mathbf{V}(i)$.

(d) Compute the weight $w(i)$ defined in (7).

(e) Compute $\Gamma_{3\text{D}}(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}})$ by Haar transformation of $\text{PCA}(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}})$ in the dimension indexed by j :

$$\Gamma_{3\text{D}}\left(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}}\right) = \{(\tilde{\alpha}_1(j), \dots, \tilde{\alpha}_K(j)); j \in \Omega_{N_i}\},$$

where, for any $k \in \{1, \dots, K\}$, the signal $(\tilde{\alpha}_k(1), \dots, \tilde{\alpha}_k(|\Omega_{N_i}|))$ is the 1D discrete Haar wavelet transform of $(\alpha_k(1), \dots, \alpha_k(|\Omega_{N_i}|))$.

(f) Denoise $\Gamma_{3\text{D}}(\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}})$ by Wiener filtering and obtain the denoised estimate $\mathbf{S}_{\mathbf{X}(N_i)}^{\text{SA}}$ of $\mathbf{S}_{\mathbf{Y}(N_i)}^{\text{SA}}$ by 3D inverse transformation (see (4) and (5)).

(2) Compute the denoised DWI estimate $\hat{\mathbf{X}}$ by weighted averaging of the denoised structure-adaptive patches in $\bigcup_{i \in \Omega_s} \mathbf{S}_{\mathbf{X}(N_i)}^{\text{SA}}$ with the $w(i)$'s (see (6)).

3. Experimental setup

3.1. Simulated data

Since real noise-free cardiac DWIs cannot be obtained by real DT-MRI acquisitions, we simulated biologically informed DWIs based on physical measures from polarized light imaging (we used the simulation method proposed in Wang et al. (2012)). The simulated DWIs – an example of which is displayed in Fig. 5a – reveal the left and right ventricles of a human heart and contain various textures and fine structures. They form a set of realistic short-axis cardiac images (size 256×256 , intensity range

$[0, 205]$) corresponding to 12 diffusion gradient directions. This set will be used to assess the performance of our SASD algorithm both qualitatively and quantitatively.

3.2. Real data

Our real data comes from an ex vivo human heart; they were acquired in the Hospital of Neuro-Cardiology in Lyon (France). The heart was placed in a plastic container and fixed by hydrophilic gel to maintain a diastolic shape. The data were acquired with a Siemens Avanto 1.5T Scanner, using a diffusion spin-echo EPI sequence with 30 diffusion gradient directions and a single excitation. We used a large number of sensitizing directions because previous studies (Frindel et al., 2009; Kingsley, 2006a) have shown that it improves the precision of DT-MRI measurements by reducing the rotational variance due to noise propagation. Each diffusion-weighted volume consists of 52 contiguous axial slices of size 128×128 (the spatial resolution is $2 \times 2 \times 2 \text{ mm}^3$).

3.3. Performance evaluation

In the case of the simulated data, the images to be denoised are obtained by adding Rician noise to the simulated noise-free DWIs, and the quality of the denoised estimates is evaluated via the PSNR and the mean SSIM (MSSIM) defined by

$$\text{PSNR} = 20 \cdot \log_{10} \left(\frac{L}{\|\mathbf{X} - \hat{\mathbf{X}}\|_2} \right) \quad \text{and}$$

$$\text{MSSIM} = \frac{1}{|\Omega_s|} \sum_{i \in \Omega_s} \text{SSIM}(\mathbf{X}(N_i), \hat{\mathbf{X}}(N_i)), \quad (14)$$

where L is the width of the pixel intensity range, and \mathbf{X} and $\hat{\mathbf{X}}$ denote the noise-free DWI and its denoised estimate, respectively. In




Fig. 5. Denoising of a simulated DWI: (a) noise-free image; (b) noisy image (noise-free + Rician noise with standard deviation $\sigma = 20$); (c–e) images denoised with the BLS-GSM, FOE and SASD algorithms.

the case of the real cardiac data, we assess the denoised DWIs with the SNR and the contrast-to-noise ratio (CNR) defined by

$$\text{SNR} = 20 \cdot \log_{10} \left(\frac{\mu_M}{\sigma_A} \right) \quad \text{and} \quad \text{CNR} = \frac{\mu_M - \mu_A}{\sigma_A}, \quad (15)$$

where μ_M and μ_A are the mean intensities in the region of the myocardium and in the background air region, respectively, and σ_A is the standard deviation of the noise in the air region.

When denoising DWIs, image details of clinical importance are liable to be damaged. Inappropriate denoising can discard or alter these details, which leads to erroneous calculation of diffusion tensors that may affect myocardial fiber reconstruction and patho-

logical analysis (Assemlal et al., 2011; Peyrat et al., 2007). Consequently, to get further insights into the performance of the denoising methods, we consider: (i) common metrics such as the fractional anisotropy (FA), the mean diffusivity (MD), and the coherence index (CI); (ii) the mean frobenius distance between tensor-fields associated with noise-free and denoised DWIs (the diffusion tensors are computed from the DWIs using unweighted linear least square (Kingsley, 2006b)); (iii) fiber angles; and (iv) fiber architectures obtained by tractography. FA measures the confidence in the diffusion direction of water molecules (Basser and Pierpaoli, 1996) – it ranges from 0 (perfectly isotropic diffusion) to 1 (fully anisotropic diffusion) – MD measures average molecular

Table 1
Denoising performance in terms of PSNR(dB) and MSSIM for different noise levels (the data is the DWI shown in Fig. 5a corrupted by Rician noise with standard deviation σ).

σ	5		10		15		20		30	
	PSNR	MSSIM	PSNR	MSSIM	PSNR	MSSIM	PSNR	MSSIM	PSNR	MSSIM
Noisy image	30.00	0.501	23.98	0.387	20.48	0.315	18.04	0.261	14.48	0.182
BLS-GSM	38.24	0.974	31.76	0.942	29.46	0.910	28.15	0.888	25.64	0.823
FOE	38.48	0.979	34.02	0.951	31.88	0.926	29.86	0.907	25.03	0.813
SASD	38.81	0.986	34.83	0.966	32.85	0.947	30.73	0.926	27.87	0.882

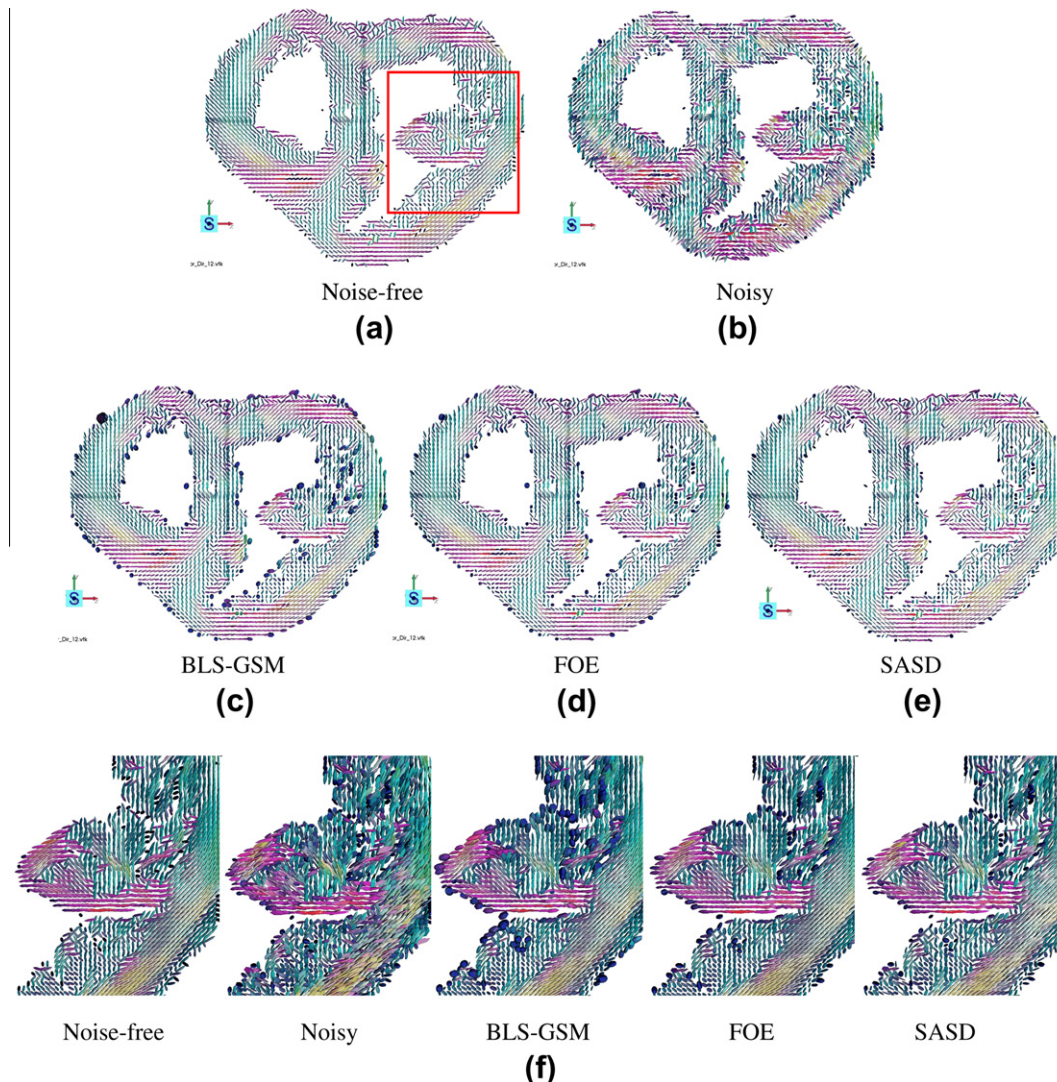


Fig. 6. Diffusion-tensor fields computed from (a) the noise-free simulated DWIs, (b) the noisy DWIs, and (c–e) the DWIs denoised by BLS-GSM, FOE and SASD. f displays enlargements of the region delimited by the red rectangle in Fig. 6a.

motion (Kingsley, 2006a), and CI estimates fiber coherence by comparing the principal eigenvector directions of nearby voxels (Basser and Pajevic, 2000). The fiber angles are evaluated in terms of azimuth and elevation angles (Jouk et al., 2000) and displayed using a directionally encoded colormap (DEC). The elevation angle

– denoted by EI – is that between the fiber and the section plane (that is, the x - y plane), and the azimuth angle – denoted by Az – is that between the projection of the fiber on the x - y plane and the x -axis. We use the DEC that assigns colors to the principal eigenvectors of a tensor field as follows: the components of the

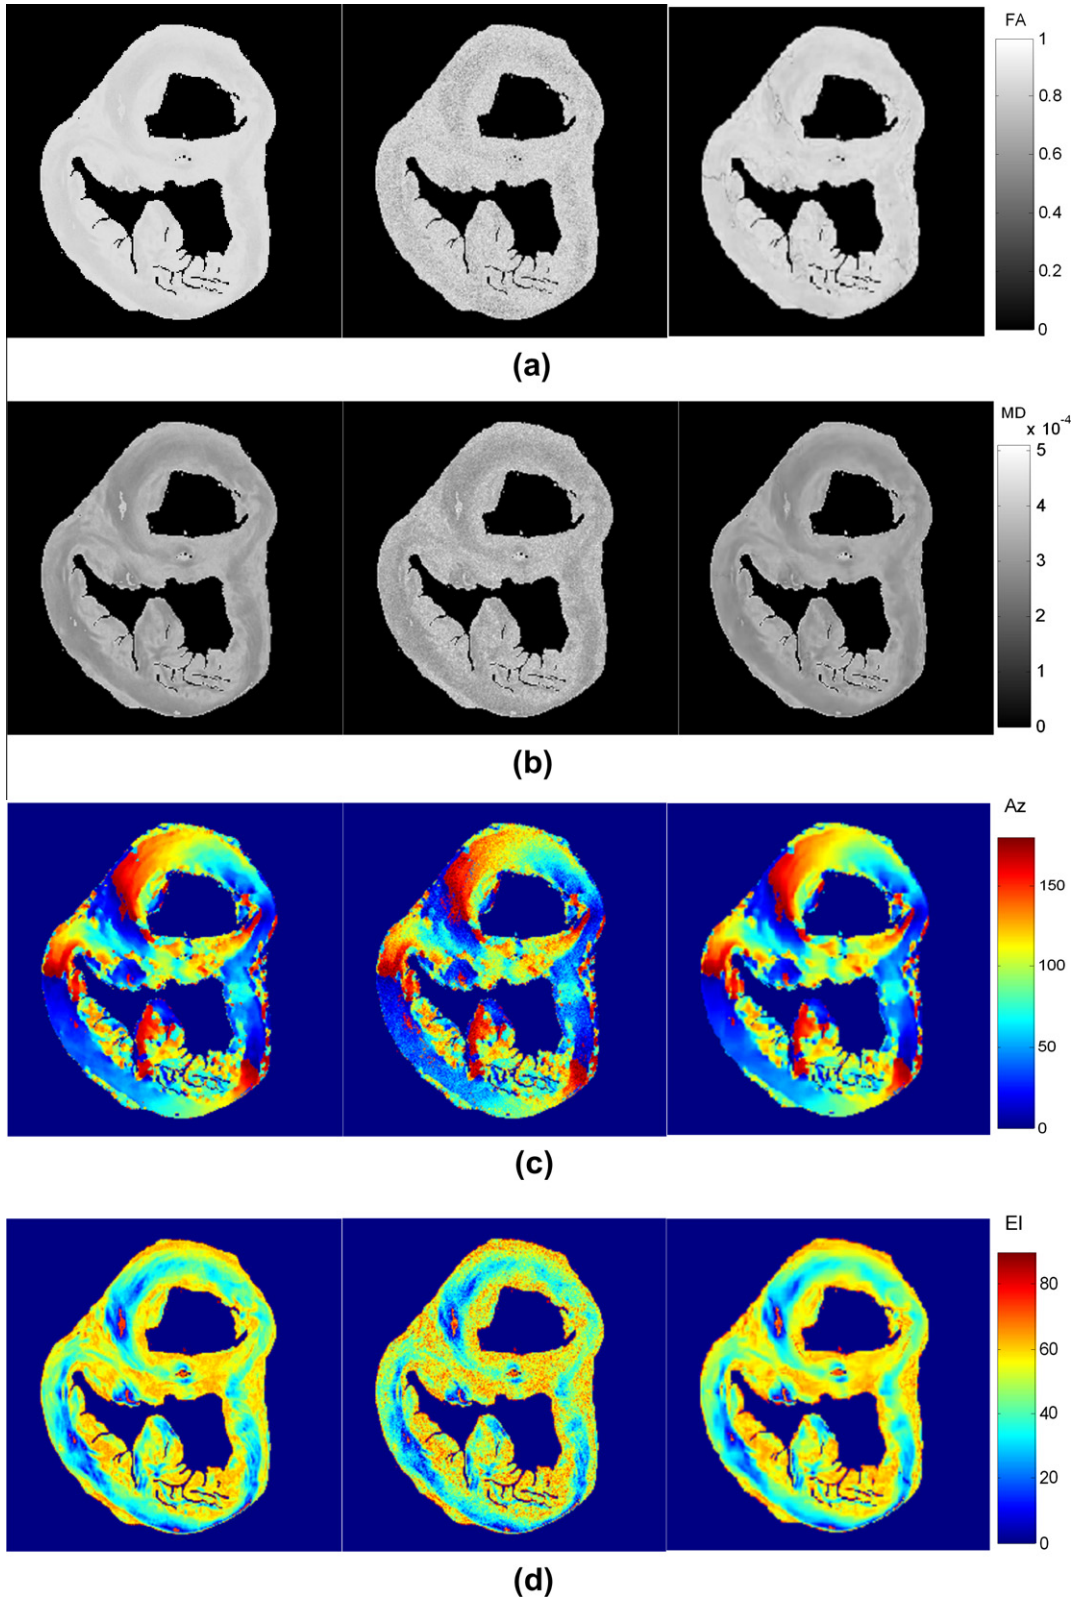


Fig. 7. Index maps obtained from the noise-free simulated DWIs (left), the noisy DWIs (middle), and the DWIs denoised by SASD (right): (a) FA, (b) MD, (c) Az, and (d) EI.

Table 2
Denosing performance in terms of mean Frobenius distance and MSE on standard indexes (the data is the DWI shown in Fig. 5a corrupted by Rician noise with standard deviation $\sigma = 20$).

	Frob. dist.	FA	MD (10^{-3} mm ² /s)	Azimuth ($^{\circ}$)	Elevation ($^{\circ}$)	CI
Noisy	0.116	0.083	0.034	18.891	5.261	0.185
BLS-GSM	0.078	0.078	0.028	15.214	5.783	0.087
FOE	0.067	0.053	0.021	14.650	3.837	0.102
SASD	0.063	0.048	0.019	13.502	3.742	0.064




Fig. 8. Denoising of a simulated DWI for different noise levels: (a) images corrupted by Rician noise with standard deviation $\sigma = 10, 30, 40$ and 50 and (b) corresponding denoised estimates computed by SASD. The case $\sigma = 20$ is displayed in Fig. 5.

eigenvectors are encoded into the red, green, and blue channels of the RGB color system, and the saturation is proportional to FA.

The performance of our SASD algorithm is compared to that of two advanced denoising methods, namely, Bayesian Least-Squares Gaussian Scale Mixture (BLS-GSM) (Portilla et al., 2003) and Fields of Experts (FOE) (Roth and Black, 2005), which produce state-of-the-art denoising performances. For suitable comparison, each method is run with the values of the parameters that produce the best solution in terms of the performance criteria described above, and these parameters are kept the same for all the experiments presented in the next section. The SASD algorithm uses patches of size 11×11 , search regions of size 41×41 (the search region of a reference patch $\mathbf{Y}(N_i)$ is the set of possible locations of the center pixels of the patches that are examined to construct Ω_{N_i}), and a downsampling factor $s = 2$. The NSD-SAW method uses a rectangular filter h of size 3×3 for computing the NSD maps and a threshold $\tau = 400$ for detecting the boundaries along the radial lines. The BLS-GSM algorithm uses Daubechies wavelets with a 7×7 window size, 4 scales and 8 directions, and the FOE algorithm uses 24 filters with a window size of 5×5 and a learning rate of 0.2.

4. Results and discussion

4.1. Denoising of the simulated DWIs

4.1.1. Comparison of BLS-GSM, FOE and SASD

Fig. 5 shows the denoising results for the simulated DWI associated with the first diffusion direction. The image to be denoised is displayed in Fig. 5b; it was obtained by adding Rician noise with standard-deviation $\sigma = 20$ to the noise-free image shown in Fig. 5a

(this noise component significantly affects the fine structures). The solution obtained by BLS-GSM is over-smoothed (Fig. 5c), and the output of FOE has chessboard artifacts pointed out by the arrows in Fig. 5d. By contrast, SASD shows good edge- and smoothness-preservation properties (the arrows in Fig. 5e indicate smooth regions that are well-recovered by SASD but that BLS-GSM and FOE fail to denoise properly), and it does not introduce spurious artifacts. Table 1 compares the results of the three methods in terms of PSNR and MSSIM for different values of the noise standard deviation. The SASD algorithm always produces the best results, and the associated performance improvement increases with increasing noise level.

Fig. 6 shows the diffusion tensor fields computed from (a) the noise-free simulated DWIs, (b) the DWIs with Rician noise with standard-deviation $\sigma = 20$, and (c–e) the DWIs denoised by the competing algorithms. To facilitate visualization and comparison, the tensors are rendered as ellipsoids with colors specified by the DEC, and enlargements of the region delimited by the red rectangle in Fig. 6a are displayed in Fig. 6f (we used a downsampling factor of 3 for Fig. 6a–e and a downsampling factor of 2 for Fig. 6f). Comparing the tensor fields computed from the noise-free and noisy DWIs, we observe that the noise induces a swelling effect and perturbs the diffusion directions, which in turn translates to errors when estimating diffusivity parameters (this latter point is illustrated in Fig. 7 and discussed next). The three denoising algorithms allow to reduce the swelling effect and to recover the original diffusion directions, but BLS-GSM and FOE tend to amplify the swelling effect at edge locations. Looking at the enlargements in Fig. 6f, we see that the tensors computed from the denoised DWIs obtained by SASD are more faithful to the original tensors in terms of both anisotropy and direction.

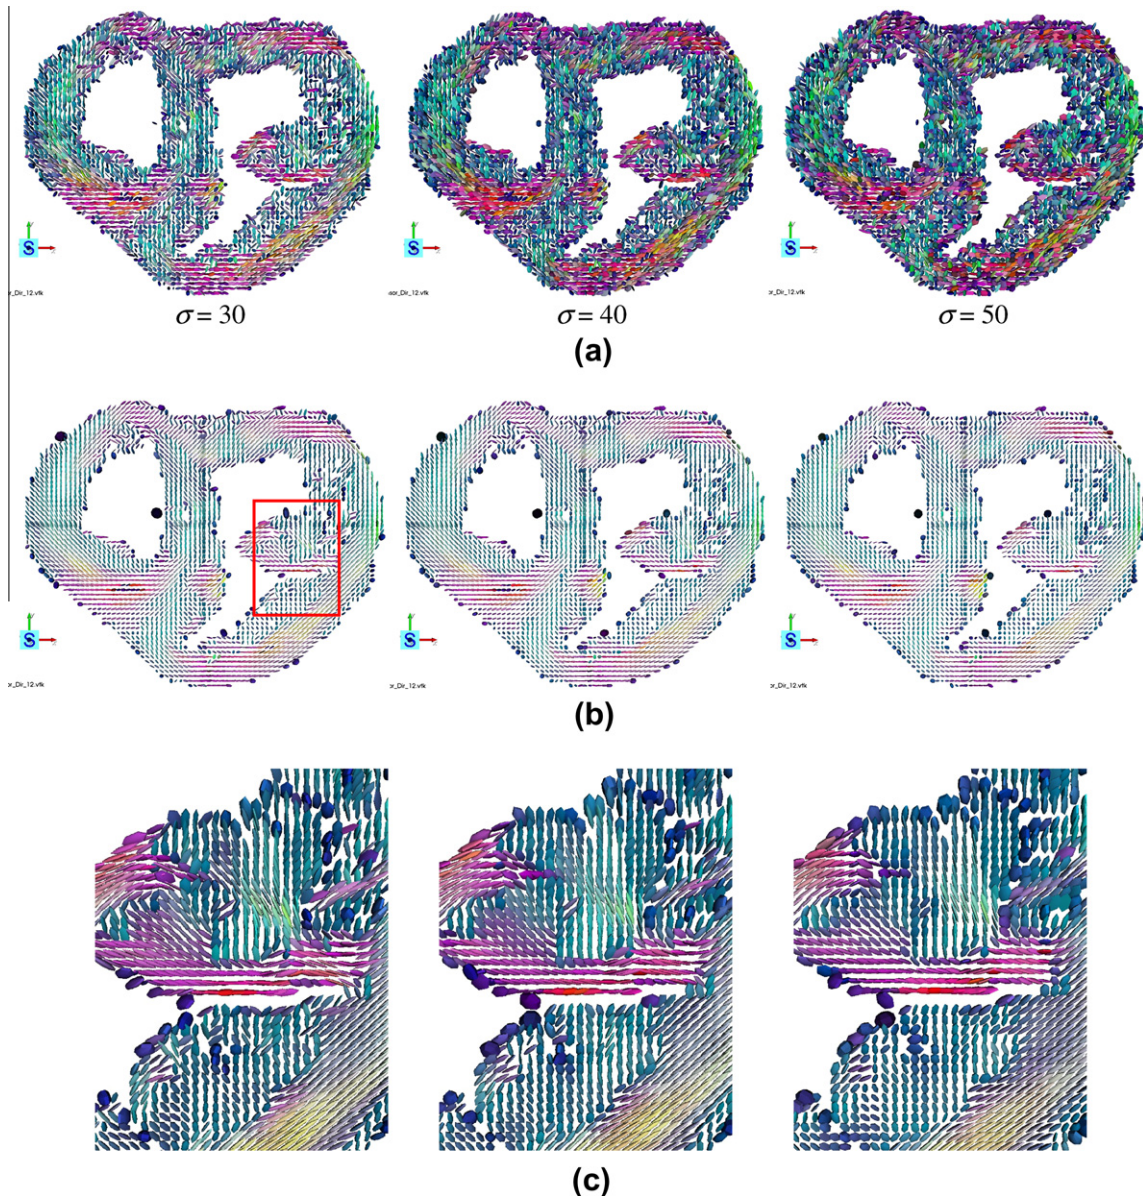


Fig. 9. Diffusion-tensor fields computed from (a) the noisy DWIs and (b) the DWIs denoised by SASD for $\sigma = 30$ (left), $\sigma = 40$ (middle) and $\sigma = 50$ (right). (c) Enlargements of the region delimited by the red rectangle in (b). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We also computed different index maps (FA, MD, Az and EI) from the three tensor fields associated with (i) the noise-free simulated DWIs, (ii) the noisy DWIs ($\sigma = 20$), and (iii) the DWIs denoised by SASD. These maps are displayed in Fig. 7. Note that the FA and MD maps obtained from the noise-free DWIs have low intensity regions corresponding to small elevation angles in the middle myocardium; this is due to the fact that diffusion is weak in the middle myocardium of the human heart. In the noisy case, the index maps have a granular aspect that corrupts fine structures and impedes the delineation of homogeneous regions. This drawback is eliminated by denoising the DWIs by SASD: the index maps obtained from the denoised DWIs have very few artifacts and are very close to that in the noise-free case.

In addition to the PSNR and the MSSIM, we can quantify the effects of denoising by means of the mean Frobenius distance – the Frobenius distance between two tensors \mathbf{T}_1 and \mathbf{T}_2 is given by $d_F(\mathbf{T}_1, \mathbf{T}_2) = (\text{Trace}((\mathbf{T}_1 - \mathbf{T}_2)^2))^{1/2}$ – and of the mean-square error (MSE) on FA, MD, Az, EI, and CI. Table 2 gives the results for the

case of a Rician noise with standard deviation $\sigma = 20$. The denoising algorithms reduce the effect of the noise in terms of all the performance metrics considered, and SASD gives the best results in all cases. In other words, of the three denoising algorithms, SASD produces the best DWI in terms of fidelity to the noise-free tensor-field.

4.1.2. Performance of SASD versus noise level

To assess the performance of our SASD method as a function of the noise level, we consider the denoising of DWIs corrupted by Rician noise with standard deviation $\sigma \in \{10, 20, 30, 40, 50\}$. Fig. 8 shows the estimates produced by the same SASD algorithm as in Section 4.1.1 for $\sigma = 10, 30, 40$ and 50 (the estimate corresponding to $\sigma = 20$ is displayed in Fig. 5e). Predictably, fine structures and texture information progressively disappear from the denoised estimates as σ increases (see, for instance, the regions pointed out by the arrows). The details are very well preserved for $\sigma = 10$ and $\sigma = 20$. When $\sigma = 30$ and $\sigma = 40$, some of the detail information




Fig. 10. DT-MRI index maps computed before (a–d) and after (e–h) denoising by SASD for $\sigma = 50$: (a and e) FA maps; (b and f) MD maps; (c and g) Az maps; (d and h) El maps.

Table 3
Denoising performance of SASD – evaluated in terms of mean Frobenius distance and MSE on standard indexes – as a function of the noise level (the results for $\sigma = 20$ are given in Table 2).

	Frob. dist.	FA	MD (10^{-3} mm ² /s)	Azimuth ($^{\circ}$)	Elevation ($^{\circ}$)	CI
$\sigma = 30$	0.194	0.163	0.075	24.390	7.853	0.279
SASD	0.091	0.070	0.031	16.018	5.354	0.079
$\sigma = 40$	0.275	0.236	0.119	31.325	10.689	0.404
SASD	0.126	0.091	0.044	18.830	6.576	0.091
$\sigma = 50$	0.315	0.270	0.135	38.238	13.347	0.517
SASD	0.166	0.114	0.061	21.035	7.475	0.101

is lost, but the denoised estimates remain satisfactory, as the principal structures and some textures are preserved. On the other hand, for $\sigma = 50$, the information loss caused by noise is too important to properly recover fine structures or textured regions. This simulation shows that SASD can achieve accurate recovery up to $\sigma = L/5$ (L is the width of the intensity range of the noise-free DWIs, which is about 200 here).

The diffusion tensor fields computed from the noisy DWIs for $\sigma = 30, 40$ and 50 and from the corresponding SASD estimates are shown in Fig. 9 (the results associated with $\sigma = 20$ are displayed in Fig. 6). We can see from Fig. 9a that the importance of the swelling effect and the disorder of the principal diffusion directions increase with the noise level. By contrast, the tensor fields computed from the denoised DWIs are well-organized, with a swelling effect apparent only at the boundaries of the myocardium. For $\sigma = 30$, most of the original diffusion directions are recovered and tensor anisotropy is well-preserved. When the noise standard deviation increases to 40 and then 50, the tensor fields computed from the denoised DWIs get smoother, but the diffusion directions become less faithful to the original ones. This is visible by comparing the enlargements in Fig. 9c with the leftmost one in Fig. 6f.

Fig. 10 gives an idea of the performance of SASD under severe noise ($\sigma = 50$). The index maps computed from the non-processed DWIs are very noisy: the fine structures and the smooth regions are hardly visible. The maps obtained from the denoised DWIs are way easier to interpret, although, compared to the case where $\sigma = 20$ (see Fig. 7), some fine structures are lost. Table 3 gives the mean Frobenius distance and the MSEs (FA, MD, Az, El and CI) for $\sigma = 30, 40$ and 50 (the results corresponding to $\sigma = 20$ are given in Table 2). The mean Frobenius distance and the MSEs computed

from the noisy DWIs increases as σ increases. Denoising the DWIs using SASD reduces the mean Frobenius distance by a factor of about 2, the MSEs on FA and MD by a factor between 1.7 and 2.7, the MSEs on the angles by a factor of about 1.6, and the MSE on CI by a factor between 2.9 and 5.1.

4.2. Denoising of real cardiac DWIs

4.2.1. Comparison of BLS-GSM, FOE and SASD in terms of SNR and CNR

Examples of denoising real cardiac DWIs using the SASD, BLS-GSM and FOE algorithms are illustrated in Fig. 11. The raw image in Fig. 11a is a DWI of an ex vivo healthy human heart; the noise component has an estimated standard deviation of $\sigma = 20.3$ (to be compared with the width of the unknown noise-free image, which is about 200). The denoised estimates produced by the different algorithms are shown in Fig. 11b–d, and the performances in terms of SNR and CNR are illustrated in Fig. 11e and f. Similarly to what we observed for the simulated DWIs, BLS-GSM blurs the edges and smoothes the fine structures, while FOE produces undesirable artifacts, especially near the edges. Our SASD algorithm produces the best result, both qualitatively and quantitatively. Note that the CNR in the denoised slices #11, 23, 26, and 29 is higher than for the other slices because the associated original images (i.e., before denoising) have higher CNR and higher structural redundancy.

4.2.2. Diffusion-tensor analysis

Fig. 12 shows the diffusion-tensor fields of a slice of the left ventricle of a human heart computed before and after denoising the DWIs. The directions of the tensors computed from the raw DWIs

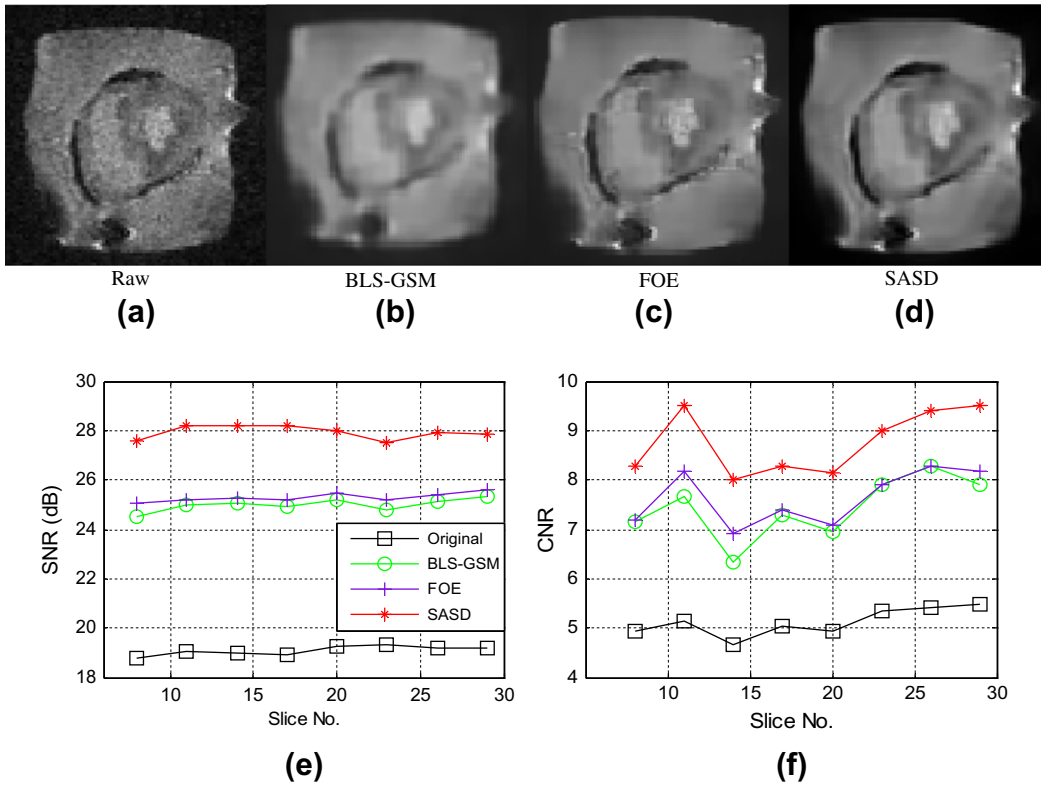


Fig. 11. Denoising of a real cardiac DWI: (a) raw image (SNR = 19.05 dB, CNR = 5.13) and images denoised by (b) BLS-GSM (SNR = 24.97 dB, CNR = 7.68), (c) FOE (SNR = 25.18 dB, CNR = 8.19), and (d) SASD (SNR = 28.22 dB, CNR = 9.49). (e and f) The SNR and CNR values achieved for different DWIs of a cardiac diffusion-weighted volume.

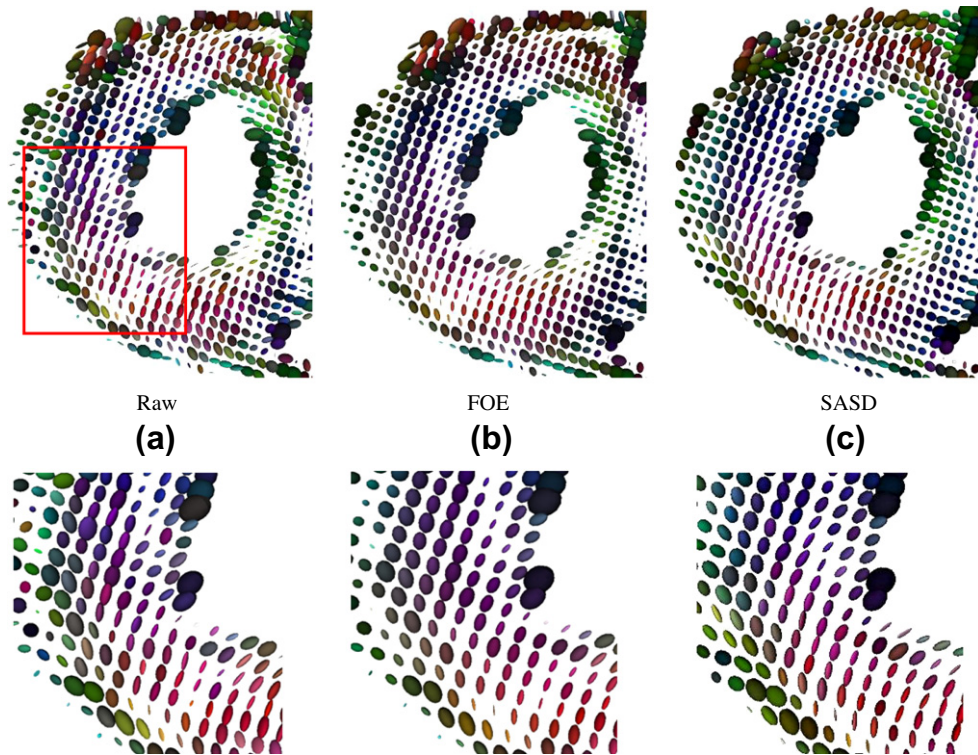


Fig. 12. Diffusion-tensor fields of a slice of the left ventricle of a human heart computed from (a) raw DWIs, (b) DWIs denoised by FOE, and (c) DWIs denoised by SASD. The associated enlargements correspond to the region delimited by the red rectangle. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 13. DT-MRI index maps of a healthy heart: (a) FA, (b) MD, and (c) representations of the fiber directions using the DEC. From left to right, the maps correspond to raw DWIs and DWIs denoised by BLS-GSM, FOE and SASD.

Table 4
Number of well-defined tensors and statistics of FA, MD and CI computed from real cardiac DWIs and from the denoised estimates obtained by the three competing algorithms.

	# of well-defined tensors	FA		MD (10^{-3} mm ² /s)		CI	
		Mean	Var	Mean	Var	Mean	Var
Raw DWIs	786	0.353	0.021	0.551	0.035	1.043	0.114
BLS-GSM	786	0.305	0.017	0.537	0.035	1.304	0.064
FOE	788	0.311	0.018	0.538	0.036	1.312	0.072
SASD	792	0.294	0.018	0.529	0.035	1.321	0.083

are noisy, and these tensors are not well organized in terms of shape. Fig. 12a illustrates the swelling effect caused by noise, which translates to an over-estimation of FA. Denoising the DWIs by SASD improves the consistency of the direction and shape of the tensors, and FOE does not perform as well as SASD in terms of the organization of the principal diffusion directions and of anisotropy preservation. Note that the diffusion-tensor field computed from the DWIs denoised by BLS-GSM is not represented because it is even worse than that obtained from the FOE estimates.

Fig. 13 shows the FA, MD and directionality maps of the heart slice considered in Fig. 12, before and after denoising by BLS-GSM, FOE and SASD. The FA map computed from the raw DWIs has a noticeable granular aspect, and the black spots in the myocardium region correspond to tensors having a negative eigenvalue (these spots also appear in the MD and directionality maps). Compared to BLS-GSM and FOE, SASD not only reduces the number of tensors having a negative eigenvalue, but also leads to more regular index maps with less artifacts. This qualitative analysis is confirmed by the quantitative results given in Table 4. We observe

that denoising increases the number of well-defined tensors (that is, the tensors with positive eigenvalues) and the mean value of CI, and that it decreases the mean values of FA and MD. Again, SASD yields the best improvement.

4.2.3. Fiber tracking

Fig. 14 shows the reconstructed fibers – using the DEC – of a whole human heart before and after denoising with the SASD algorithm; the associated quantitative measurements are given in Table 5. Fiber tracking was performed using the following thresholds: minimum FA of 0.05 (this threshold was set empirically to obtain an appropriate fiber density), minimum fiber length of 5 mm, and maximum fiber angle of 60°. The fibers obtained by tracking the raw tensors (that is, the tensors computed from the unprocessed DWIs) are not well organized and are somewhat irregular, whereas fiber tracking from the denoised tensors (that is, the tensors computed from the denoised DWIs) results in smooth and well-organized trajectories that are more consistent with the fiber architecture of a healthy human heart (Lombaert

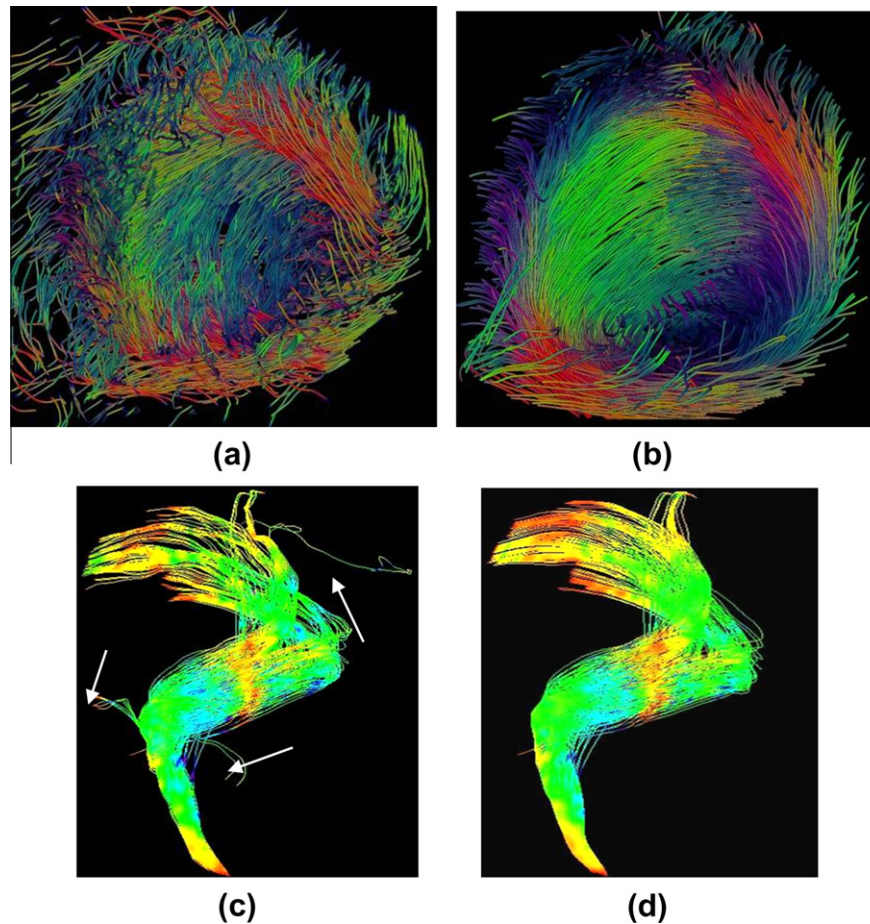


Fig. 14. Myocardium fiber reconstruction. (a and c) Fibers tracked from raw tensors. (b and d) Fibers tracked from denoised tensors (i.e., computed from the DWIs denoised by SASD).

Table 5
Quantitative analysis of the myocardium fiber architectures.

	# of fibers	Well-def. tensors	FA		MD ($10^{-3} \text{ mm}^2/\text{s}$)		Fiber length (mm)		Helix fiber angle ($^\circ$)	
			Mean	Var	Mean	Var	Mean	Var	Mean	Var
Raw DWIs	2551	33,741	0.237	0.006	0.649	0.045	36.116	304.020	7.638	3.823
SASD	2673	33,982	0.209	0.005	0.641	0.045	38.067	300.847	6.590	1.649

et al., 2012; Rohmer et al., 2007; Seemann et al., 2006). Fig. 14c and d show a local fiber bundle in the middle myocardium – the arrows point out erroneous trajectories resulting from the noise in the raw DWIs. The tracking obtained from the denoised tensors is globally much smoother than that performed from the raw tensors. Quantitatively, denoising the DWIs by SASD increases the number of fibers by 4.8% and the mean fiber length by 5.4%, and it decreases the variance of the helix fiber angle by 56.8% (this latter effect is due to the fact that denoising attenuates the disturbance in the principal diffusion direction field).

5. Conclusions

We proposed a new DWI denoising method, namely, structure-adaptive sparse denoising (SASD), that collects adaptive similar neighborhoods to increase redundancy and hence to facilitate the removal of the noise. To avoid the drawbacks of standard similarity measures, we have defined a new constrained similarity measure based on the local mean and the SSIM index. We also proposed

to search for structure-adaptive neighborhoods in the reference patches, which allows local adaptation to image features. The core of SASD is the following: similar structure-adaptive neighborhoods are grouped into 3D arrays that are denoised in a 3D transform domain using Wiener filtering. The 3D transformation consists in a 2D principal component (PC) decomposition in the image domain – which concentrates the energy in a few coefficients – followed by a 1D Haar wavelet transform in the third dimension.

The proposed SASD method is competitive and outperforms state-of-the-art denoising methods. It presents good detail-preservation properties and produces very few artifacts. Compared to BLS-GSM and FOE, SASD effectively removes the noise component in cardiac DWIs while preserving image contrast and fine structures. Our experiments show that SASD performs well for images with sufficient structural redundancy and that it achieves a good trade-off between image contrast and image smoothness. Furthermore, it is potentially very useful for cleaning highly corrupted DWIs, which facilitates subsequent operations such as tensor field analysis, fiber tracking, and clinical diagnosis in the cardiac domain.

A pertinent and persistent question about denoising is that of the minimum detectable feature size. The answer is not simple, as it depends on the precise definition of the features of interest and on their redundancy. For instance, in pathological cases involving tissue lesions, the lack of redundancy in the neighborhoods of the lesions may translate to lower denoising performance in these injured regions (indeed, if it happens that a patch has no similar patch, the denoising reduces to Wiener filtering). Nevertheless, our experiments with simulated (biologically informed) DWIs show that SASD achieves accurate recovery – the accuracy of the estimates was measured in the image domain, in the tensor-field domain and in parametric domains – up to a noise standard-deviation equal to 20% of the width of the true intensity range. In addition, our experiments on real cardiac data show that tracking fibers from the denoised DWIs leads to myocardium fibers that are more regular and more conform to anatomical knowledge.

It should be stressed that we focused on 2D denoising because we considered either 2D data or 3D data with a lower sampling rate in the dimension perpendicular to the slices (in other words, the spatial correlation is stronger in the 2D-slice domain than in the remaining dimension). However, the performance can be improved by extending the search for similar patches to adjacent slices, and the proposed approach can also be extended to fully-3D denoising (that is, using 3D patches) in the case of fully-3D data acquisition. Besides, our SASD algorithm was initially motivated by the denoising of cardiac DWIs, which is the reason why we only considered additive Rician noise. Nevertheless, we did not make any specific assumption on the organ to be imaged or on the noise statistics, and thus SASD can be applied to other denoising problems involving additive noise. Future directions of research include the study of the behavior of SASD in the presence of structured noise (as in ultrasound images) and scan-rescan reproducibility. We did not address the latter because our concern was with ex vivo hearts (contrary to in vivo imaging, ex vivo scans do not change over time except for the noise component); nevertheless, although there is as yet no clinical application of in vivo cardiac DT-MRI, some data are available for experimenting clinical cases of healthy and failing human hearts (Lombaert et al., 2012, 2011a, 2011b).

Acknowledgments

This work was supported by the NNSF of China (61271092), the International S&T Cooperation Project of China (2007DFB30320), and the French ANR (ANR-09-BLAN-0372-01). The authors would like to thank L.H. Wang for supplying the simulated cardiac DWI data used in our experiments in Section 4.1.

References

- Assemlal, H.E., Tschumperle, D., Brun, L., Siddiqi, K., 2011. Recent advances in diffusion MRI modeling: angular and radial reconstruction. *Med. Image Anal.* 15, 369–396.
- Awate, P.S., Whitaker, R.T., 2007. Feature-preserving MRI denoising: a nonparametric empirical bayes approach. *IEEE Trans. Med. Imaging* 26, 1242–1255.
- Barash, D., Comaniciu, D., 2004. A common framework for nonlinear diffusion, adaptive smoothing, bilateral filtering and mean shift. *Image Video Comput.* 22, 73–81.
- Basser, P.J., Pajevic, S., 2000. Statistical artifacts in diffusion tensor MRI caused by background noise. *Magn. Reson. Med.* 44, 41–50.
- Basser, P.J., Pajevic, S., 2003. A normal distribution for tensor-valued random variables: application to diffusion tensor MRI. *IEEE Trans. Med. Imaging* 22, 785–794.
- Basser, P.J., Pierpaoli, C., 1996. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J. Magn. Reson. B* 111, 209–219.
- Behrens, T.E., Johansen-Berg, H., Woolrich, M.W., Smith, S.M., Wheeler-Kingshott, C.A., Boulby, P.A., Barker, G.J., Sillery, E.L., Sheehan, K., Ciccarelli, O., Thompson, A.J., Brady, J.M., Matthews, P.M., 2003. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nat. Neurosci.* 6, 750–757.
- Bondiau, P.Y., Clatz, O., Sermesant, M., Marcy, P.Y., Delingette, H., Frenay, M., Ayache, N., 2008. Biocomputing: numerical simulation of glioblastoma growth using diffusion tensor imaging. *Phys. Med. Biol.* 53, 879–893.
- Boulanger, J., Kervrann, C., Bouthemy, P., Elbau, P., Sibarita, J.B., Salameo, J., 2010. Patch-based nonlocal functional for denoising fluorescence microscopy image sequences. *IEEE Trans. Med. Imaging* 29, 442–454.
- Brox, T., Kleinschmidt, O., Cremers, D., 2008. Efficient nonlocal means for denoising of textural patterns. *IEEE Trans. Image Process.* 17, 1083–1092.
- Buades, A., Coll, B., Morel, J.M., 2005. A review of image denoising algorithms with a new one. *Multiscale Model. Simul.* 4, 490–530.
- Chatterjee, P., Milanfar, P., 2009. Clustering-based denoising with locally learned dictionaries. *IEEE Trans. Image Process.* 18, 1438–1451.
- Chen, B., Edward, W., 2005. Noise removal in magnetic resonance diffusion tensor imaging. *Magn. Reson. Med.* 54, 393–407.
- Clatz, O., Sermesant, M., Bondiau, P.Y., Delingette, H., Warfield, S.K., Malandain, G., Ayache, N., 2005. Realistic simulation of the 3-D growth of brain tumors in MR images coupling diffusion with biomechanical deformation. *IEEE Trans. Med. Imaging* 24, 1334–1346.
- Coupé, P., Yger, P., Prima, S., Hellier, P., Kervrann, C., Barillot, C., 2008. An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images. *IEEE Trans. Med. Imaging* 27, 425–436.
- Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K., 2007. Image denoising by sparse 3D transform domain collaborative filtering. *IEEE Trans. Image Process.* 16, 2080–2095.
- Danielyan, A., Foi, A., Katkovnik, V., Egiazarian, K., 2010. Spatially adaptive filtering as regularization in inverse imaging: compressive sensing, upsampling and super-resolution. In: Milanfar, P. (Ed.), *Super-Resolution Imaging*. CRC Press, USA.
- Delingette, H., Billet, F., Wong, K.C., Sermesant, M., Rhode, K., Ginks, M., Rinaldi, C.A., Razavi, R., Ayache, N., 2012. Personalization of cardiac motion and contractility from images using variational data assimilation. *IEEE Trans. Biomed. Eng.* 59, 20–24.
- Deriche, R., Calder, J., Descoteaux, M., 2009. Optimal real-time Q-ball imaging using regularized Kalman filtering with incremental orientation sets. *Med. Image Anal.* 13, 564–579.
- Descoteaux, M., Deriche, R., Knosche, T.R., Anwander, A., 2009. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Trans. Med. Imaging* 28, 269–286.
- Donoho, L.D., Elad, M., Temlyakov, V.N., 2006. Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Trans. Inf. Theory* 52, 6–18.
- Durrleman, S., Fillard, P., Pennec, X., Trounev, A., Ayache, N., 2011. Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents. *Neuroimage* 55, 1073–1090.
- Elad, M., Aharon, M., 2006. Denoising via sparse and redundant representations over learned dictionaries. *IEEE Trans. Image Process.* 15, 3736–3745.
- Elad, M., Starck, J.L., Querreb, P., Donoho, D.L., 2005. Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA). *Appl. Comput. Harmon. Anal.* 19, 340–358.
- Farrell, J.A., Landman, B.A., Jones, C.K., Smith, S.A., Prince, J.L., van Zijl, P.C., Mori, S., 2007. Effects of signal-to-noise ratio on the accuracy and reproducibility of diffusion tensor imaging-derived fractional anisotropy, mean diffusivity, and principal eigenvector measurements at 1.5 T. *J. Magn. Reson. Imaging* 26, 756–767.
- Fillard, P., Arsigny, V., Pennec, X., Ayache, N., 2007. Clinical DT-MRI estimation, smoothing and fiber tracking with Log-Euclidean metrics. *IEEE Trans. Med. Imaging* 26, 1472–1482.
- Fillard, P., Descoteaux, M., Goh, A., Gouttard, S., Jeurissen, B., Malcolm, J., Ramirez-Manzanares, A., Reisert, M., Sakaie, K., Tensaouti, F., Yo, T., Mangin, J.F., Poupon, C., 2011. Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom. *Neuroimage* 56, 220–234.
- Foi, A., Katkovnik, V., Egiazarian, K., 2007. Pointwise shape-adaptive DCT for high-quality denoising and deblocking of grayscale and color images. *IEEE Trans. Med. Imaging* 16, 1395–1411.
- Frimel, C., Robini, M., Croisille, P., Zhu, Y.M., 2009. Comparison of regularization methods for human cardiac diffusion tensor MRI. *Med. Image Anal.* 13, 405–418.
- Galanaud, D., Haik, S., Linguraru, M.G., Ranjeva, J.P., Faucheux, B., Kaphan, E., Ayache, N., Chiras, J., Cozzone, P., Dormont, D., Brandel, J.P., 2010. Combined diffusion imaging and MR spectroscopy in the diagnosis of human prion diseases. *Am J. Neuroradiol.* 31, 1311–1318.
- Galban, C.J., Maderwald, S., Uffmann, K., Ladd, M.E., 2005. A diffusion tensor imaging analysis of gender differences in water diffusivity within human skeletal muscle. *NMR Biomed.* 18, 489–498.
- Guevara, P., Poupon, C., Riviere, D., Cointepas, Y., Descoteaux, M., Thirion, B., Mangin, J.F., 2011. Robust clustering of massive tractography datasets. *Neuroimage* 54, 1975–1993.
- Hammond, D.K., Simoncelli, E.P., 2008. Image modeling and denoising with orientation-adapted Gaussian scale mixtures. *IEEE Trans. Image Process.* 17, 2089–2101.
- Jouk, P.S., Usson, Y., Michalowicz, G., Grossi, L., 2000. Three-dimensional cartography of the pattern of the myofibres in the second trimester fetal human heart. *Anat. Embryol. (Berl.)* 202, 103–118.
- Katkovnik, V., Spokoiny, V., 2008. Spatially adaptive estimation via fitted local likelihood techniques. *IEEE Trans. Image Process.* 17, 873–886.

- Katkovnik, V., Foi, A., Egiazarian, K., Astola, J., 2010. From local kernel to nonlocal multiple-model image denoising. *Int. J. Comput. Vision* 86, 1–32.
- Kervrann, C., Boulanger, J., 2006. Optimal spatial adaptation for patch-based image denoising. *IEEE Trans. Image Process.* 15, 2866–2878.
- Kingsley, P.B., 2006a. Introduction to diffusion tensor imaging mathematics: Part II. Anisotropy, diffusion-weighting factors, and gradient encoding schemes. *Concepts Magn. Reson.* 28A, 123–154.
- Kingsley, P.B., 2006b. Introduction to diffusion tensor imaging mathematics: Part III. Tensor calculation, noise, simulations, and optimization. *Concepts Magn. Reson.* 28A, 155–179.
- Le Bihan, D., 2003. Looking into the functional architecture of the brain with diffusion MRI. *Nat. Rev. Neurosci.* 4, 469–480.
- Lenglet, C., Prados, E., Pons, J.P., Deriche, R., Faugeras, O., 2009. Brain connectivity mapping using riemannian geometry, control theory, and PDEs. *SIAM J. Imaging Sci.* 2, 285–322.
- Liu, W.Y., Magnin, I.E., Gimenez, G., 1995. Un nouvel opérateur pour la détection de rupture dans des signaux bruités. *Traitement du Signal* 12, 226–236.
- Lombaert, H., Peyrat, J.M., Fanton, L., Cheriet, F., Delingette, H., Ayache, N., Clarysse, P., Magnin, I., Croisille, P., 2011a. Variability of the human cardiac laminar structure. In: *Proceedings of the 14th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'11)*, pp. 160–167.
- Lombaert, H., Peyrat, J.M., Fanton, L., Cheriet, F., Delingette, H., Ayache, N., Clarysse, P., Magnin, I., Croisille, P., 2011b. Statistical atlas of human cardiac fibers: comparison with abnormal hearts. In: *Proceedings of the 14th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'11)*, pp. 207–213.
- Lombaert, H., Peyrat, J.M., Croisille, P., Rapacchi, S., Fanton, L., Cheriet, F., Clarysse, P., Magnin, I., Delingette, H., Ayache, N., 2012. Human atlas of the cardiac fiber architecture: study on a healthy population. *IEEE Trans. Med. Imaging* 31, 1436–1447.
- Lu, Y., Aldroubi, A., Gore, J.C., Anderson, A.W., Ding, Z., 2006. Improved fiber tractography with Bayesian tensor regularization. *Neuroimage* 31, 1061–1074.
- Manjón, J.V., Carbonell-Caballero, J., Lull, J.J., Garcia-Marti, G., Marti-Bonmati, L., Robles, M., 2008. MRI denoising using non-local means. *Med. Image Anal.* 12, 514–523.
- Manjón, J.V., Coupé, P., Buades, A., Fonov, V., Collins, D.L., Robles, M., 2010. Non-local MRI upsampling. *Med. Image Anal.* 14, 784–792.
- Marim, M., Angelini, E., Olivon, J., 2010. Denoising in fluorescence microscopy using compressed sensing with multiple reconstructions and non-local merging. In: *International Conference on IEEE EMBS'2010*, pp. 3394–3397.
- Messe, A., Caplain, S., Parodot, G., Garrigue, D., Mineo, J.F., Soto Ares, G., Ducreux, D., Vignaud, F., Rozec, G., Desal, H., Pelegrini-Issac, M., Montreuil, M., Benali, H., Lehericy, S., 2011. Diffusion tensor imaging and white matter lesions at the subacute stage in mild traumatic brain injury with persistent neurobehavioral impairment. *Hum. Brain Mapp.* 32, 999–1011.
- Mori, S., Oishi, K., Faria, A.V., 2009. White matter atlases based on diffusion tensor imaging. *Curr. Opin. Neurol.* 22, 362–369.
- Peyrat, J.M., Sermesant, M., Pennec, X., Delingette, H., Xu, C., McVeigh, E.R., Ayache, N., 2007. A computational framework for the statistical analysis of cardiac diffusion tensors: application to a small database of canine hearts. *IEEE Trans. Med. Imaging* 26, 1500–1514.
- Pižurica, A., Wink, M., Vansteenkiste, E., Philips, W., Roerdink, J., 2006. A review of wavelet denoising in MRI and ultrasound brain imaging. *Curr. Med. Imaging Rev.* 2, 247–260.
- Portilla, J., Strela, V., Wainwright, M.J., Simoncelli, E.P., 2003. Image denoising using scale mixtures of gaussians in the wavelet domain. *IEEE Trans. Image Process.* 12, 1338–1350.
- Rohmer, D., Sitek, A., Gullberg, G.T., 2007. Reconstruction and visualization of fiber and sheet structure with regularized tensor diffusion MRI in the human heart. *Invest. Radiol.* 42, 777–789.
- Roth, S., Black, M.J., 2005. Fields of Experts: a framework for learning image priors. In: *International Conference on IEEE Computer Vision and Pattern Recognition (CVPR'2005)*, pp. 860–867.
- Rousseau, F., 2010. A non-local approach for image super-resolution using intermodality priors. *Med. Image Anal.* 14, 594–605.
- Seemann, G., Keller, D.U.J., Weiss, D.L., Dossel, O., 2006. Modeling human ventricular geometry and fiber orientation based on diffusion tensor MRI. In: *Proceedings of the International Conference on Computers in Cardiology 2006*, pp. 801–804.
- Sikora, T., Makai, B., 1995. Shape-adaptive DCT for generic coding of video. *IEEE Trans. Circuits Syst. Video Technol.* 5, 59–62.
- Smith, S.M., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T.E., Mackay, C.E., Watkins, K.E., Ciccarelli, O., Cader, M.Z., Matthews, P.M., Behrens, T.E., 2006. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage* 31, 1487–1505.
- Varshney, K.R., Çetin, M., Fisher, J.W., 2008. Sparse representation in structured dictionaries with application to synthetic aperture radar. *IEEE Trans. Signal Process.* 56, 3548–3561.
- Wakana, S., Jiang, H., Nagae-Poetscher, L.M., van Zijl, P.C., Mori, S., 2004. Fiber tract-based atlas of human white matter anatomy. *Radiology* 230, 77–87.
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13, 1–14.
- Wang, L., Zhu, Y., Li, H., Liu, W., Magnin, I.E., 2012. Multiscale modeling and simulation of the cardiac fiber architecture for DMRI. *IEEE Trans. Biomed. Eng.* 59, 16–19.
- Wu, M.T., Tseng, W.Y., Su, M.Y., Liu, C.P., Chiou, K.R., Wedeen, V.J., Reese, T.G., Yang, C.F., 2006. Diffusion tensor magnetic resonance imaging mapping the fiber architecture remodeling in human myocardium after Infarction. *Circulation* 114, 1036–1045.
- Yu, H., Zhao, L., Wang, H., 2009. Image denoising using trivariate shrinkage filter in the wavelet domain and joint bilateral filter in the spatial domain. *IEEE Trans. Image Process.* 18, 2364–2369.