Non Local Spatial and Angular Matching : a new denoising technique for diffusion MRI

Samuel St-Jean[1], Pierrick Coupé[2], and Maxime Descoteaux[1]

[1]Sherbrooke Connectivity Imaging Lab (SCIL), Université de Sherbrooke, Sherbrooke, Québec, Canada, [2]Unité Mixte de Recherche CNRS (UMR 5800), Laboratoire Bordelais de Recherche en Informatique, Bordeaux, France

INTRODUCTION: Diffusion Weighted Images (DWIs) datasets suffer from low Signal-to-Noise Ratio (SNR), especially at high b-values. Acquiring data at high b-values contains relevant information and is now of great interest for connectomics studies [1]. High noise levels bias the measurements because of the non-Gaussian nature of the noise, which in turn can lead to a false and biased estimation of the diffusion parameters. Therefore, high SNR DWIs is important in order to draw meaningful conclusions in subsequent data or group analyses [2]. The acquired DWIs differ between themselves, but still share the same underlying structure. It is also known that natural images are redundant and can be sparsified [3]. We thus propose to use the redundancy of DWIs as a sparse representation to reduce the noise level and achieve a higher SNR using dictionary learning and sparse coding, without the need for additional acquisition time. We show quantitative results on the ISBI 2013 HARDI challenge phantom [4].

METHOD: Denoising methods are often applied on each DWIs separately [5], without taking into account the structure they share. As in [5, 6], we will use the common structure amongst DWIs to further improve the denoising. In contrast to [6], our method also uses the full information of all DWIs in the same fashion as [7] during the first step of the process in addition to the information of angular neighboring DWIs in the second step. We improve upon [7] by not enforcing orthogonality in the first step and exploiting sparsity of the DWIs in the second step. We first apply [8] to account for Rician noise and then perform brain extraction on the DWIs. We next normalize each DWIs independently (referred as $\mathbf{X}_i$). The first step is a sparse principal component analysis (SPCA) decomposition [9]. This initial denoising is performed across all the DWIs at once. We only keep the most meaningful principal components (PC) as in [10] and use this first denoised version as the input for the second step. This second step consists in doing a local 3D neighborhood denoising based on dictionary learning and sparse coding [11, 12]. In this part of the algorithm, we find the angular neighbors of each DWIs (excluding B0s). For all $\mathbf{X}_i$, we create a 4D stack made of the B0, $\mathbf{X}_i$ and the p closest angular neighbors to $\mathbf{X}_i$. With this formulation, each $\mathbf{X}_i$ can be selected as a valid neighbor multiple times, and as such will be denoised multiple times (referred as j times). This results in a 4D stack of size (n, n, n, P), where n is the size of $\mathbf{X}_i$ (in 3D) and P = (2+p). We then extract all 4D patches of size (m, m, m, P) from our current stack, which contains spatial information as well as angular neighboring information. Each 4D patches is then stored as a column vector. We constrain each column to have a unit L2 norm as required by **Equation1** [12]. The goal is to find a dictionary **D** in which $\mathbf{X}_i$ will be well represented by its coefficients α. This minimization aims to find a sparse representation of each neighborhood, while discarding noise. Learning **D** from the noisy patches also ensures the reconstruction will be tailored to the neighborhood currently being denoised [11]. The second step is repeated for each $\mathbf{X}_i$. The final denoised volume is made by computing the mean of those j representations of $\mathbf{X}_i$, i.e. $\mathbf{X}_{i\ \text{denoised}} = \textit{mean}(\mathbf{X}_{ij})$. We finally apply the inverse transformation of [8].

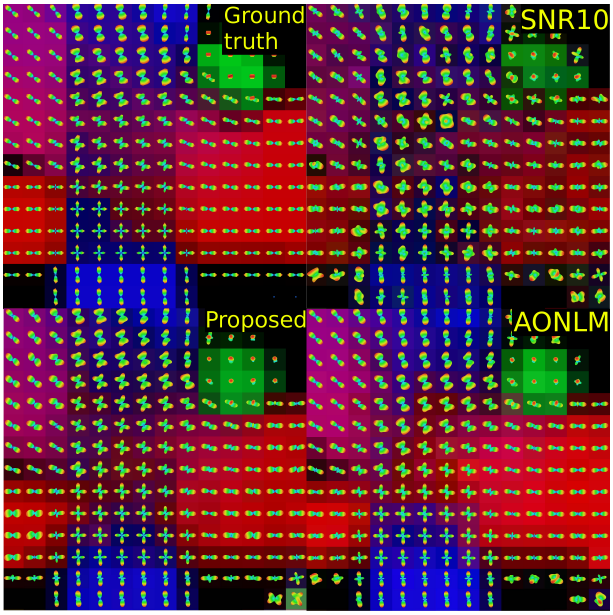

Figure 1 : ISBI Challenge CSA-ODFs reconstructed from the ISBI HARDI 2013 dataset, overlaid on the CFA map. Top left is the ground truth, top right the SNR10 data, bottom left our method and bottom right the AONLM method. The ODFs in our method more closely resemble the ground truth while having a more accurate CFA map, as confirmed by results in **Table 1**.

$$\min_{\mathbf{D},\alpha} \frac{1}{n} \sum_{i=1}^{n} \left(\frac{1}{2} ||\mathbf{x_i} - \mathbf{D}\boldsymbol{\alpha}_i||_2^2 + \lambda ||\boldsymbol{\alpha}_i||_1 \right)$$

Equation 1: The main dictionary learning equation.

DATASET: We used for our synthetic experiments the ISBI 2013 HARDI testing dataset of SNR10, corrupted with rician noise [4]. The phantom is made of 64 b=3000 s/mm^2 images with a constant b=0 image.

RESULTS: **Figure 1** shows constant solid angle (CSA) q-ball ODFs of order 4 [13] reconstructed from the ISBI phantom [4] overlaid onto the colored fractional anisotropy (CFA). We denoised the SNR10 data with AONLM [5], LPCA [7] and our method. We used AONLM as a comparison since it was the denoising algorithm used in the winning method [14]. For our method, we used a 3D block size of (3, 3, 3) and 2 angular neighbors with λ=0.01. All other methods were ran with their default parameters. Our method reconstructs ODFs that are closer to the ground truth than the AONLM denoising, **Table 1** shows two perceptual similarity metrics, the mean PSNR in dB and mean SSIM computed on the 4D volume. We also report three diffusion metrics : the root mean squared fractional anisotropy (RMSE FA), root mean squared general fractional anisotropy (RMSE GFA) and the mean angular error (mean AE) with the noiseless data for the SNR10, SNR30, AONLM denoising [5], LPCA denoising [7] (not shown on **Figure 1**) and our method. In most cases, we achieve a higher similarity than the other methods, even higher than the SNR30 dataset.

	SNR10	SNR30	AONLM	LPCA	Proposed
PSNR	19.597	29.320	23.869	28.370	**30.708**
SSIM	0.876	0.987	0.950	0.982	**0.990**
RMSE FA	0.179	**0.040**	0.132	0.128	0.132
RMSE GFA	0.152	**0.047**	0.105	0.103	0.108
Mean AE	14.047	13.578	13.111	12.905	**11.587**

Table 1 : Reconstruction metrics for the tested methods. Our method performs best for the perceptual similarity (PSNR in dB and SSIM) metrics and the mean angular error. A clear gain is achieved from the noisy input SNR10 data and our method even achieves better reconstruction than the non-denoised SNR30 data in most cases. The best result in each category is shown in bold.

DISCUSSION: The benefits of denoising the data prior to the reconstruction are shown on **Figure 1** and **Table 1**. Our method reconstructs sharper, more uniform ODFs than the SNR10 data provides without denoising. The produced ODF and the underlying CFA map are also closer to the ground truth with our method than AONLM and are better than the non-denoised SNR10 data. Using post-processing methods does not add to the clinical acquisition burden, while reaching a higher SNR. Having cleaner data also means that reconstruction algorithms will estimate less biased diffusion parameters. This in turn leads to more accurate and robust data analyses. We believe that denoising the data should be a pre-processing step part of every pipeline, just like motion and eddy current corrections that are commonly applied to correct for artifacts.

REFERENCES: [1] Van Essen & al. (2013) NeuroImage. [2] Jones & al. (2013) NeuroImage. [3] Olshausen & al. (1997) Vision Research. [4] http://hardi.epfl.ch/static/events/2013_ISBI/ [5] Manjón & al. (2010) Journal of magnetic resonance imaging. [6] Tristán-Vega & al. (2010) Medical image analysis. [7] Manjón & al. (2013) PloS one. [8] Foi, A. ISBI 2011. [9] Zou & al. (2006) Journal of Computational and Graphical Statistics. [10] Dabov & al. SPARS 2009. [11] Elad & al. (2006) IEEE transactions on image processing. [12] Mairal & al. ICML 2009. [13] Aganj & al. (2010) Magnetic resonance in medicine. [14] Garyfallidis & al. (2013) ISBI HARDI challenge.