Published in final edited form as:

Med Image Comput Comput Assist Interv. 2016 October ; 9902: 587–595. doi:
10.1007/978-3-319-46726-9_68.

XQ-NLM: Denoising Diffusion MRI Data via x - q Space Non-Local Patch Matching

Geng Chen^{1,2}, Yafeng Wu¹, Dinggang Shen², and Pew-Thian Yap²

¹Data Processing Center, Northwestern Polytechnical University, Xi'an, China

²Department of Radiology and BRIC, University of North Carolina, Chapel Hill, U.S.A.

Abstract

Noise is a major issue influencing quantitative analysis in diffusion MRI. The effects of noise can be reduced by repeated acquisitions, but this leads to long acquisition times that can be unrealistic in clinical settings. For this reason, post-acquisition denoising methods have been widely used to improve SNR. Among existing methods, non-local means (NLM) has been shown to produce good image quality with edge preservation. However, currently the application of NLM to diffusion MRI has been mostly focused on the spatial space (i.e., the x -space), despite the fact that diffusion data live in a combined space consisting of the x -space and the q -space (i.e., the space of wavevectors). In this paper, we propose to extend NLM to both x -space and q -space. We show how patch-matching, as required in NLM, can be performed concurrently in x - q space with the help of azimuthal equidistant projection and rotation invariant features. Extensive experiments on both synthetic and real data confirm that the proposed x - q space NLM (XQ-NLM) outperforms the classic NLM.

1 Introduction

Diffusion MRI suffers from low signal-to-noise ratio (SNR), especially when the diffusion weighting (i.e., b -value) is high. To improve SNR, a common approach is to repeat the acquisition multiple times so that repeated measurements can be averaged to reduce noise and enhance SNR. However, this approach is time consuming and is oftentimes prohibitive in clinical settings.

Post-acquisition denoising methods [1–3] have been shown to be a viable alternative in improving SNR without resorting to repeating acquisitions. Among existing methods, the non-local means (NLM) [4] algorithm has been shown to offer particularly good performance in edge-preserving denoising. The key feature that distinguishes NLM from the other methods is its ability to increase significantly the information available for denoising by going beyond the local neighborhood and allowing non-local or spatially distant information to be involved in denoising. Self-similar information is gathered via a patch matching mechanism where voxels with similar neighborhoods are given greater weights and those with dissimilar neighborhoods are given lesser weights in denoising. The

effectiveness of NLM stems from the fact that the non-local nature of the algorithm allows significantly more information to be available for denoising than local methods.

In particular, NLM has been applied to removing noise in diffusion-weighted (DW) images [1, 2]. Using NLM, the DW images can be denoised as individual images, as a multi-spectral vector image, or as parametric maps given by a diffusion model [1]. However, existing methods are mainly focused on neighborhood matching in the spatial domain (i.e., x -space). While effective, this approach might not perform as well in regions with highly curved white matter structures, resulting in averaging over disparate axonal directions and failing to make full use of information from differentially oriented structures.

In this paper, we show that improved denoising performance can be attained by extending the NLM algorithm beyond x -space to include q -space. The advantage afforded by this extension is twofold: (i) Non-local information can now be harnessed not only across x -space, but also across measurements in q -space; and (ii) In white matter regions with high curvature, q -space neighborhood matching corrects for such non-linearity so that information from structures oriented in different directions can be used more effectively for denoising without introducing artifacts. To allow NLM to be carried out in x - q space, the patch for neighborhood similarity evaluation is defined in q -space. This involves mapping a spherical neighborhood in q -space to a disc and then computing a set of rotation invariant features from the disc for patch matching.

2 Methods

2.1 x - q Space Non-Local Means

Our method utilizes neighborhood matching in both x -space and q -space for effective denoising. For each voxel at location $\mathbf{x}_i \in \mathbb{R}^3$, the diffusion-attenuated signal measurement $S(\mathbf{x}_i, \mathbf{q}_k)$ corresponding to wavevector $\mathbf{q}_k \in \mathbb{R}^3$ is denoised by averaging over non-local measurements that have similar q -neighborhoods. We estimate the denoised signal $NLM(S)(\mathbf{x}_i, \mathbf{q}_k)$ as

$$NLM(S)(\mathbf{x}_i, \mathbf{q}_k) = \sum_{(\mathbf{x}_j, \mathbf{q}_l) \in \mathcal{V}_{i,k}} w_{i,k;j,l} [S(\mathbf{x}_j, \mathbf{q}_l) + c_{i,k;j,l}], \quad (1)$$

where $\mathcal{V}_{i,k}$ is the search neighborhood in x - q space associated with $S(\mathbf{x}_i, \mathbf{q}_k)$, $c_{i,k;j,l}$ is a variable used to compensate for differences in signal levels due to spatial intensity inhomogeneity and signal decay in q -space, $w_{i,k;j,l}$ is the weight indicating the similarity between $S(\mathbf{x}_i, \mathbf{q}_k)$ and $S(\mathbf{x}_j, \mathbf{q}_l)$.

Instead of restricting neighborhood matching to x -space, we introduce x - q space neighborhood matching by defining a patch in q -space. For each point in x - q space, $(\mathbf{x}_i, \mathbf{q}_k)$, we define a spherical patch, $\mathcal{N}_{i,k}$, centered at $\mathbf{q}_k$ with fixed $q_k = |\mathbf{q}_k|$ and subject to a neighborhood angle α_p . The samples on this spherical patch are mapped to a disc using azimuthal equidistant projection (AEP, Section 2.2) before computing rotation invariant features via polar complex exponential transform (PCET, Section 2.3) for patch matching.

Fig. 1 illustrates the concept of patch matching in x - q space. The search radius in x -space is s and the search angle in q -space is α_s . Matching between different shells is allowed.

To deal with the fact that q -space is not always sampled in a shell-like manner, we project the samples in the neighbor of the spherical patch radially onto the spherical patch and weight them accordingly. For instance, we project measurement $S(\mathbf{x}_i, \mathbf{q}_k)$ to the shell, i.e.,

$$S(\mathbf{x}_i, \frac{|\mathbf{q}_k|}{|\mathbf{q}_{k'}|} \mathbf{q}_{k'}) = S(\mathbf{x}_i, \mathbf{q}_{k'}) \exp \left\{ -\frac{(\sqrt{b_{k'}} - \sqrt{b_k})^2}{h_{\text{projection}}^2} \right\}. \quad (2)$$

where $b_k = t|\mathbf{q}_k|^2$ and $b_{k'} = t|\mathbf{q}_{k'}|^2$ are the respective b -values and t is the diffusion time. $\mathcal{N}_{i,k}$ is then constructed using the projected measurements.

The diffusion-weighted signal can vary due to inhomogeneity or natural signal decay in q -space when the b -value is increased. This difference in signal level is compensated using $c_{i,k;j,l}$ defined as the mean difference between two patches, i.e., @ @ @

$c_{i,k;j,l} = \overline{\mathbf{S}(\mathcal{N}_{i,k})} - \overline{\mathbf{S}(\mathcal{N}_{j,l})}$, where $\overline{\cdot}$ denotes the mean and $\mathbf{S}(\mathcal{N}_{i,k})$ is a vector containing the values of all diffusion signals in $\mathcal{N}_{i,k}$. Before feature computation for patch matching, we centralize the signal vector as @ @ @ $\hat{\mathbf{S}}(\mathcal{N}_{i,k}) = \mathbf{S}(\mathcal{N}_{i,k}) - \overline{\mathbf{S}(\mathcal{N}_{i,k})}$.

2.2 Azimuthal Equidistant Projection (AEP)

AEP [5] maps the coordinates on a sphere to a plane where the distances and azimuths of points on the sphere are preserved with respect to a reference point [5]. This provides a good basis for subsequent computation of invariant features for matching. The reference point, which in our case corresponds to the center of the spherical patch, will project to the center of a circular projection. Viewing the reference point as the ‘North pole’, all points along a given azimuth, θ , will project along a straight line from the center. In the projection plane, this line subtends an angle θ with the vertical. The distance from the center to another

projected point is given as ρ . We represent the reference point @ @ @ $\hat{\mathbf{q}}_0 = \frac{\mathbf{q}_0}{|\mathbf{q}_0|}$ as spherical coordinates (ϕ_0, λ_0) , with ϕ referring to latitude and λ referring to longitude. We project (ϕ, λ) to a corresponding point (ρ, θ) in a 2D polar coordinate system, where ρ is the radius and θ is the angle. Based on [5], the relationship between (ϕ, λ) and (ρ, θ) is as follows:

$$\cos \rho = \sin \phi_0 \sin \phi + \cos \phi_0 \cos \phi \cos(\lambda - \lambda_0), \quad (3a)$$

$$\tan \theta = \frac{\cos \phi \sin(\lambda - \lambda_0)}{\cos \phi_0 \sin \phi - \sin \phi_0 \cos \phi \cos(\lambda - \lambda_0)}. \quad (3b)$$

The projection can be described as a two-step mapping:

$$\mathbf{q} \rightarrow (q, \phi, \lambda) \rightarrow (q, \rho, \theta). \quad (4)$$

Note that extra care needs to be taken when using the above equations to take into consideration the fact that diffusion signals are antipodal symmetric. Prior to performing AEP, we map antipodally all the points on the sphere to the hemisphere where the reference point is located. AEP maps a q -space spherical patch $\mathcal{N}$ to a 2D circular patch $\hat{\mathcal{N}}$. Note that AEP changes only the coordinates but not the actual values of the signal vector, i.e., $\mathbf{S}(\hat{\mathcal{N}}) = \mathbf{S}(\mathcal{N})$.

2.3 Polar Complex Exponential Transform (PCET)

After AEP, we proceed with the computation of rotation invariant features. We choose to use the polar complex exponential transform (PCET) [6] for its computation efficiency. Denoting an element of $\hat{\mathcal{N}}$ as $\hat{\mathcal{S}}(\mathbf{x}, q, \rho, \theta)$, the PCET of order n , $|n| = 0, 1, 2, \dots, \infty$, and repetition l , $|l| = 0, 1, 2, \dots, \infty$, is defined as

$$M_{n,l}(\hat{\mathcal{N}}) = \frac{1}{\pi} \int_{(\mathbf{x}, q, \rho, \theta) \in \hat{\mathcal{N}}} [H_{n,l}(\rho, \theta)]^* \hat{\mathcal{S}}(\mathbf{x}, q, \rho, \theta) \rho d\rho d\theta, \quad (5)$$

where $[\cdot]^*$ denotes the complex conjugate and $H_{n,l}(\rho, \theta)$ is the basis function defined as $H_{n,l}(\rho, \theta) = e^{i2\pi n\rho^2} e^{il\theta}$. It can be easily verified that $|M_{n,l}(\hat{\mathcal{N}})|$ is invariant to rotation [6]. $|M_{n,l}(\hat{\mathcal{N}})|$'s up to maximum order m , i.e., $0 \leq |n| \leq m$, are concatenated into a feature vector $\mathbf{M}(\hat{\mathcal{N}})$.

2.4 Patch Matching

If $\mathbf{M}(\hat{\mathcal{N}}_{i,k})$ is the feature vector of the projected patch $\hat{\mathcal{N}}_{i,k}$, the matching weight $w_{i,k;j,l}$ is defined as

$$w_{i,k;j,l} = \frac{1}{Z_{i,k}} \exp \left\{ -\frac{\|\mathbf{M}(\hat{\mathcal{N}}_{i,k}) - \mathbf{M}(\hat{\mathcal{N}}_{j,l})\|_2^2}{h_{\mathbf{M}}^2(i, k)} \right\} \exp \left\{ -\frac{(\sqrt{b_k} - \sqrt{b_l})^2}{h_b^2} \right\}, \quad (6)$$

where $Z_{i,k}$ is a normalization constant to ensure that the weights sum to one:

$$Z_{i,k} = \sum_{(\mathbf{x}_j, \mathbf{q}_l) \in \mathcal{V}_{i,k}} \exp \left\{ -\frac{\|\mathbf{M}(\hat{\mathcal{N}}_{i,k}) - \mathbf{M}(\hat{\mathcal{N}}_{j,l})\|_2^2}{h_{\mathbf{M}}^2(i, k)} \right\} \exp \left\{ -\frac{(\sqrt{b_k} - \sqrt{b_l})^2}{h_b^2} \right\}. \quad (7)$$

Here $h_{\mathbf{M}}(i, k)$ is a parameter controlling the attenuation of the first exponential function. As in [7], we set $h_{\mathbf{M}}(i, k) = \sqrt{2\beta\hat{\sigma}_{i,k}^2 |\mathbf{M}(\hat{\mathcal{N}}_{i,k})|}$, where β is a constant [7] and $\hat{\sigma}_{i,k}^2$ is the estimated noise standard deviation, which can be computed globally as shown in [8] or spatial-adaptively as shown in [7]. The former is used in this paper. Similarly, h_b

$h_b = \sqrt{2}\sigma_b$ controls the attenuation of the second exponential function, where σ_b is a scale parameter. $|\mathbf{M}(\hat{\mathcal{N}}_{i,k})|$ denotes the length of the vector $\mathbf{M}(\hat{\mathcal{N}}_{i,k})$.

2.5 Adaption to Noncentral Chi Noise

The classic NLM is designed to remove Gaussian noise and needs to be modified for the noncentral Chi noise distribution typical in parallel MRI, which uses multiple receiver coils. Based on [9], we define the unbiased denoised signal $\text{UNLM}(S)(\mathbf{x}_i, \mathbf{q}_k)$ as

$$\text{UNLM}(S)(\mathbf{x}_i, \mathbf{q}_k) = \sqrt{\sum_{(\mathbf{x}_j, \mathbf{q}_l) \in \mathcal{V}_{i,k}} w_{i,k;j,l} [S(\mathbf{x}_j, \mathbf{q}_l) + c_{i,k;j,l}]^2 - 2N\sigma^2}, \quad (8)$$

where σ is the Gaussian noise standard deviation that can be estimated from the image background [8], N is the number of receiver coils. When there is only one coil (i.e., $N=1$), the noncentral Chi distribution reduces to a Rician distribution.

3 Experiments

3.1 Datasets Synthetic Data

A synthetic multi-shell dataset of a spiral was generated for the quantitative evaluation of the proposed method. The parameters used in synthetic data simulation were consistent with the real data described next: $b = 1000, 2000, 4000, 6000$ s/mm², 81 gradient directions, 128×128 voxels with resolution 2×2 mm². Four levels of Rician noise (3%, 9%, 15%, and 21%) were added to the resulting ground truth data. Rician noise was generated by adding Gaussian noise with distribution $\mathcal{N}(0, \mathbf{v}(p/100))$ in the complex domain of the signal with noise variance determined based on noise-level percentage p and maximum signal value $\mathbf{v}$ (150 in our case) [7].

Real Data—The real dataset was acquired using the same gradient directions and b -values as the synthetic dataset. A Siemens 3T TRIO MR scanner was used for data acquisition. The imaging protocol is as follows: 96×128 imaging matrix, $2 \times 2 \times 2$ mm³ resolution, TE=145 ms, TR=12,200 ms, 32 receiver coils.

3.2 Results

In all experiments, we set $s = 2$ voxels, $\beta = 0.1$, $\sigma_b = 5$, $\alpha_p = 30^\circ$, $\alpha_s = 30^\circ$ and $m = 4$. The dataset is acquired using spherical sampling, therefore we set $h_{\text{projection}}$ to a small value to disable the signal projection described in (2). We use peak-to-signal-ratio (PSNR) as the metric for performance evaluation.

Figure 2 indicates that the new patch matching scheme is robust to the variation of fiber orientations and b -values. This allows our method, XQ-NLM, to use information from differentially oriented signal profiles for effective denoising. For different noise levels, the PSNR results shown in Fig. 3 indicate that XQ-NLM significantly improves the PSNR. The largest improvement over NLM is 5.19 dB when the noise level is 15%. The denoised DW images, shown in Fig. 4, indicate that XQ-NLM is able to preserve sharp edges and

effectively remove noise, thanks to the robust q -space patch matching mechanism, as demonstrated in Fig. 2.

For the real data, the results shown in Fig. 5 indicate that XQ-NLM yields markedly improved edge-preserving results in cortical regions. The influence of denoising on the white matter fiber orientation distribution functions (ODFs) is shown in Fig. 6. Visible improvements are marked by white arrows and rectangles, indicating that XQ-NLM yields cleaner results with less spurious peaks.

4 Conclusion

We have extended NLM beyond x -space to include q -space, allowing denoising using non-local information not only in the spatial domain but also in the wavevector domain. Our method allows information from highly curved white matter structures to be used effectively for denoising. The synthetic data experiments show that our method increases the PSNR markedly and yields improved results, especially at boundaries. The real data experiments further demonstrate that our method gives cleaner white matter fiber ODF with less spurious peaks.

Acknowledgments

This work was supported in part by NIH grants (NS093842, EB006733, EB008374, EB009634, AG041721, and MH100217).

References

1. Wiest-Daesslé, N.; Prima, S.; Coupé, P.; Morrissey, SP.; Barillot, C. Non-local means variants for denoising of diffusion-weighted and diffusion tensor MRI. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*; 2007. p. 344-351.
2. Descoteaux, M.; Wiest-Daesslé, N.; Prima, S.; Barillot, C.; Deriche, R. Impact of Rician adapted non-local means filtering on HARDI. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*; 2008. p. 122-130.
3. Chen G, Zhang P, Wu Y, Shen D, Yap PT. Denoising magnetic resonance images using collaborative non-local means. *Neurocomputing*. 2016; 177(215-227)
4. Buades A, Coll B, Morel JM. A review of image denoising algorithms, with a new one. *Multiscale Modeling & Simulation*. 2005; 4(2):490-530.
5. Wessel P, Smith WH. The generic mapping tools. <http://gmt.soest.hawaii.edu>.
6. Yap PT, Jiang X, Kot AC. Two-dimensional polar harmonic transforms for invariant image representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2010; 32(7):1259-1270. [PubMed: 20489228]
7. Coupé P, Yger P, Prima S, Hellier P, Kervrann C, Barillot C. An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images. *IEEE Transactions on Medical Imaging*. 2008; 27(4):425-441. [PubMed: 18390341]
8. Manjón JV, Carbonell-Caballero J, Lull JJ, García-Martí G, Martí-Bonmatí L, Robles M. MRI denoising using non-local means. *Medical Image Analysis*. 2008; 12(4):514-523. [PubMed: 18381247]
9. Koay CG, Özarslan E, Basser PJ. A signal transformational framework for breaking the noise floor and its applications in MRI. *Journal of Magnetic Resonance*. 2009; 197(2):108-119. [PubMed: 19138540]

Fig. 1. Patch Matching in x - q Space

Patch matching involving 4 shells in q -space and a search radius of 2 voxels in x -space. The search neighborhood $\mathcal{V}$ is a combination of the sub-neighborhoods $\{\mathcal{V}(j, r)\}$ (blue) associated with different locations, $\{\mathbf{x}_j\}$, and b -values, $\{b_r\}$, i.e., $\mathcal{V} = \bigcup_{j,r} \mathcal{V}(j, r)$, where $\mathcal{V}(j, r) \equiv \mathcal{V}(\mathbf{x}_j, b_r)$. The total number of sub-neighborhoods is $4 \times (2 \times 2 + 1)^3 = 500$. The reference patch is marked by red and the search patches are marked by orange.

Fig. 2. q-Space Patch Matching

Two cases of q -space patch matching: (middle row) variation of fiber orientations and (bottom row) variation of b -values. The top row shows the profiles of the diffusion signals. Patch matching is performed using the point marked by the red arrow as the reference. The middle row shows the matching results of signal profiles in different orientations. Warm colors indicate greater agreement, cool colors indicate otherwise. The bottom row shows the matching results for different b -values (i.e., 1000, 2000, 3000, 4000 s/mm²), based on the leftmost signal profile.

Fig. 3. PSNR Comparison

Performance comparison between NLM and XQ-NLM using the synthetic dataset.

Fig. 4. DW Images – Synthetic Data

(A) Ground truth DW image; (B) Image with 3% noise; (C) NLM-denoised image; (D) XQ-NLM-denoised image.

Fig. 5. DW Images – Real Data

The reference DW image is shown on the far left. Close-up views of (left) noisy DW image, (middle) NLM-denoised image, and (left) XQ-NLM-denoised image.

Fig. 6. Fiber ODFs

White matter fiber ODFs shown in order identical to Fig 5.