

Impact of Rician Adapted Non-Local Means Filtering on HARDI

Maxime Descoteaux¹, Nicolas Wiest-Daesslé^{3,4,5}, Sylvain Prima^{3,4,5},
Christian Barillot^{3,4,5}, and Rachid Deriche²

¹ NMR Lab, NeuroSpin, CEA Saclay, France

² Project Team Odysse, INRIA Sophia Antipolis – Méditerranée, France

³ INRIA, VisAGeS Project-Team, Rennes, France

⁴ INSERM, U746, Rennes, France

⁵ University of Rennes I, CNRS, UMR 6074, IRISA, Rennes, France

Abstract. In this paper we study the impact of denoising the raw high angular resolution diffusion imaging (HARDI) data with the Non-Local Means filter adapted to Rician noise (NLMr). We first show that NLMr filtering improves robustness of apparent diffusion coefficient (ADC) and orientation distribution function (ODF) reconstructions from synthetic HARDI datasets. Our results suggest that the NLMr filtering improve the quality of anisotropy maps computed from ADC and ODF and improve the coherence of q-ball ODFs with the underlying anatomy while not degrading angular resolution. These results are shown on a biological phantom with known ground truth and on a real human brain dataset. Most importantly, we show that multiple measurements of diffusion-weighted (DW) images and averaging these images along each direction can be avoided because NLMr filtering of the individual DW images produces better quality generalized fractional anisotropy maps and more accurate ODF fields than when computed from the averaged DW datasets.

1 Introduction

In diffusion-weighted MR imaging (DWI), increasing b-value and decreasing voxel size both improve the separability of the different components at the voxel level, and allow better characterization of fiber crossing. However, these modified MR sequences also drastically reduce the signal-to-noise ratio (SNR) in the acquired images, where some intensities are close to the background noise level. This noise is severely detrimental to the analysis of diffusion-weighted (DW) data, even in the simple, cylindrically symmetrical, single-fiber model case [1]. One possible harmful consequence is that one may infer more complexity in the tissue under study than what really exists. In magnitude MR images, noise is often Rician-distributed with standard reconstruction methods, and is often conveniently approximated as Gaussian in high SNR regions. This hypothesis no longer holds for low SNR DWI. In the context of diffusion tensor MRI (DTI),

recent attempts have been made to account for the Rician noise to regularize the DW data [2], to estimate the diffusion tensor [3], or to perform both tasks simultaneously [4]. However, among the existing methods to estimate and/or regularize orientation distribution function (ODF) reconstructions from high angular resolution diffusion imaging (HARDI) [5,6,7,8], the Rician noise bias has just started to be addressed. In [5], local geometries of 3D curves are used in a relaxation labeling framework to regularize fields of DT and ODFs and show good results without modeling the Rician noise explicitly. On the other hand, [6] uses a variational flow formulation to estimate a robust field of ODF that incorporates a possibly more complex noise distribution to account for the Rician bias. Other very recent attempts, similar in spirit to our approach because working directly on the raw DWI, includes the correction of the raw DWI signal [8] and inject a Rician statistics term in the iterative algorithm that reconstructs the fibre orientation density [8] or use a linear mean square error estimator that restores DWI using a Rician noise modeling [7].

To remove noise artifacts, DTI and HARDI acquisitions are often repeated multiple times along each gradient direction (typically, 3 to 10 measurements) and each DW image is averaged to create the dataset to be processed. This averaging increases acquisition time considerably. In this paper, we propose to filter with a specific Rician-based, structure-preserving, NLMean method, which recently showed to outperform other state-of-the-art methods and in particular the classical Gaussian version of the NLMean to denoise DTI at low b-values [9]. The filtering is applied on the raw DWI before estimating the ADC and diffusion ODF estimated from q-ball imaging (QBI). We want to see if robustness of the ADC and ODF estimation is improved by the filtering of the raw DWI and if angular resolution of QBI is affected by this filtering. Most importantly, we want to see if multiple measurements of DW images and averaging these images along each direction can be avoided to potentially reduce acquisition time.

2 Method

2.1 Non-Local Means Filtering Adapted to Rician Noise

The Non Local Means (NLM) filter restores the intensity x of a voxel v by computing a weighted average of all the voxels intensity in the neighborhood V of v using carefully chosen weights w_i such that $\sum w_i = 1$. The original version of the NLM filter [10] is well adapted for Gaussian distributed noise, but must be extended to account for the Rician noise present in DWI. Two main problems need to be tackled. First, how to remove the noise, and second, how to remove the statistical *bias*. These two problems are solved using a slight modification of the NLM filter. A second order moment of the Rician distribution is incorporated, which is $E(X^2) = \mu^2 + 2\sigma^2$, where μ is the true signal value and σ is the noise variance. This value σ is estimated from the data with the pseudo-residual method [11]. The key idea is to restore the squared value of the signal instead

of the signal itself and then correct for the *bias* using an estimation of the noise variance. The Rician NLM filter (NLMr) is then:

$$NLM_r(x) = \sqrt{\sum_{x_i \in V} w_i x_i^2 - 2\sigma^2}, \quad (1)$$

where x is the voxel intensity to restore and V is a neighborhood around x . Note that Eq. 1 may yield negative values. In such cases the resulting restoration is set to zero. In practice, in our experiments on real data, negative values were only found in the background of the images.

The weights w_i are simply computed using the classical L2 norm between two neighborhoods, N and N_i of size S , around x and x_i :

$$w_i = \frac{1}{Z(i)} e^{-\frac{\sqrt{\frac{1}{S} \sum_k \|y_k - z_k\|^2}}{h}}$$

where $Z(i)$ is a normalization constant with $Z(i) = \sum_i w_i$, y_k , z_k are the intensities of the k -th voxels in neighborhood N around N_i and N_j , and h is the filtering parameter that is set to σ in practice (more details are found in [12,9]).

2.2 ADC and ODF Reconstruction

From HARDI measurements, several functions such as the ADC and the ODF can be reconstructed to recover information about the diffusion displacement probability of water molecules. Recently, the spherical harmonic basis was shown to be a powerful tool to estimate both functions [13,14,15,16,17,18]. ADC estimation on HARDI data is the generalization of diffusion tensor estimation in a higher order space and several high order anisotropy measures can be extracted from this representation [17]. However, to extract fiber orientation, the diffusion ODF is more appropriate and agrees with the underlying principal directions of fiber bundles. In this paper, we compute ADC and ODF estimation using an even order spherical harmonic basis of order L , as done in [13,14,17] and [18] respectively. We also compute the generalized fractional anisotropy (GFA) maps from the estimated q-ball ODFs.

2.3 Synthetic Datasets and Real Datasets

Synthetic Data Experiment. We perform synthetic data experiments quantifying the error on the sphere for the filtered DW images, the ADC and ODF reconstructions. We also quantify the angular error in the fiber directions detected from the ODF and the angular resolution limitation before and after filtering. We use the multi-tensor model with Rician noise to generate synthetic DWI datasets consisting of 6, 30x30 slices with fibers of width 6 voxels with crossing regions. This synthetic data generation is illustrated in [17,18]. It has the advantage of producing known ground truth ADC and ODF profiles as well as ground truth fiber orientations. We use profiles with eigenvalues $[1390,355,355] \times 10^{-6} \text{mm}^2/\text{s}$

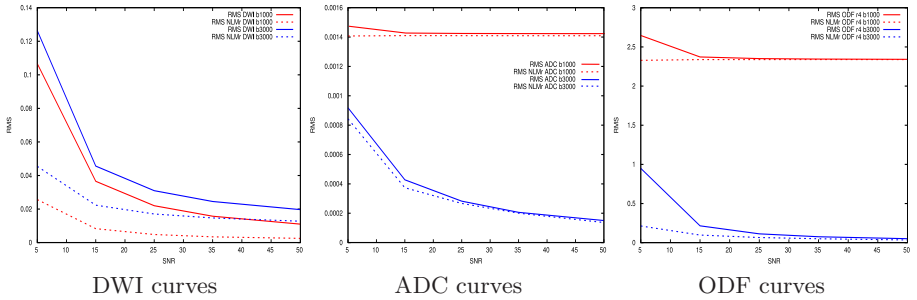


Fig. 1. Root mean squared (RMS) error curves between ground truth (solid lines) DWI signal, ADC, ODF and the denoised (dashed lines) DWI signal, ADC and ODF with the NLMr filter for reconstruction of order $L = 4$ while varying SNR between 5 and 50 and b -value 1000 (red curves) and 3000 s/mm^2 (blue curves)

for single fiber voxels ($\text{FA} = 0.7$) and crossing voxels are composed with equal volume fraction. DW images were generated along 81 gradient directions corresponding to a 3rd order tessellation of the sphere, b -value 1000 and 3000 s/mm^2 , SNR varying between 5 and 50 and crossing fiber angle between 30° and 90° . Finally, each noisy DW image was denoised with NLMr and the ADC and ODF profiles were estimated before and after NLMr filtering with a spherical harmonic basis of order $L = 4, 6$, and 8.

Real Datasets. The biological phantom was created from two excised rat spinal cords embedded in 2% agar (see [5]). The acquisition was done on a 1.5T scanner using 90 encoding directions, with $b = 3000 \text{ s}/\text{mm}^2$, $\text{TR} = 6.4 \text{ s}$, $\text{TE} = 110 \text{ ms}$, 2.8 mm isotropic voxels and four signal averages per direction. The SNR of a single DW image in the spinal cord was estimated to SNR 5 and the corresponding averaged DW image with SNR 10. We perform NLMr on each of the four DW datasets and compare the GFA maps and ODF reconstructions with the averaged dataset.

The human brain dataset was acquired on a 3T scanner using 80 encoding directions, with $b = 3000 \text{ s}/\text{mm}^2$, $\text{TR} = 4 \text{ s}$, $\text{TE} = 122 \text{ ms}$, 1.25 mm x 1.25 mm x 3 mm voxels and three signal averages per direction. The SNR of a single T2 image and single DW image in the white matter was estimated to SNR 6 and 2 respectively. The corresponding average T2 and DW image have a SNR of 11 and 4 respectively. We perform NLMr on each of the three raw DWI datasets and compare the GFA maps and ODF reconstructions with the averaged dataset. Estimation order was fixed to $L = 4$ in the ODF reconstructions as well as for the biological phantom dataset.

3 Results

Synthetic Data Experiment. First, Fig. 1 illustrates the root mean squared (RMS) error made on the sphere when comparing ground truth DWI, ADC

Table 1. Average angular error and resolution limit on synthetic data with SNR 35. We show angular resolution in the noise-free case (true θ_{sep}), average % of crossing voxels correctly detected for that true angle in the NLMr filtered noisy case and in the NLMr filtered noisy case (θ_{sep}), and the angular error from the detected ODF peaks.

order L	true θ_{sep}	% of crossing detected	θ_{sep}	angular error
4	63°	16%	65°	$1.1 \pm 3.4^\circ$
6	59°	67%	60°	$2.9 \pm 10^\circ$
8	58°	64%	60°	$8.5 \pm 12^\circ$

and ODF profiles with their respective noisy and NLMr filtered profiles on the sphere. The error was computed from the projection of each function of the sphere tessellated with a fine mesh corresponding to a 16th order tessellation.

As one would expect, DW images are closer to ground truth DW images after NLMr filtering. As a consequence, ADC and ODF estimation are more robust to noise for all SNR experiments. In particular, the error gap between non-filtered and filtered DWI, ADC and ODF reconstructions on the sphere is more important to SNR values lower than 15. This is the range in which most real HARDI acquisitions stand (especially at high b -values).

Second, Table 1 illustrates the angular resolution limit of the q-ball ODF reconstructions as well as the angular error made on the detected ODF peak. We first report the angular resolution in the noise-free case (true θ_{sep}). That is the angle under which the two fiber compartments cannot be discriminated and only a single ODF peak is detected. Then, we report the average percentage of crossing voxels correctly detected for that true separation angle when the ODF is reconstructed from the noisy DW images filtered with NLMr. We see that not all crossing voxels are correctly discriminated. However, more than one half of crossing voxels are identified as crossing for estimation order 6 and 8. In the next column, we report the angular resolution limitation of the NLMr filtered reconstructions, θ_{sep} . That is, the angle at which all crossing voxels are correctly discriminated with two maxima. We see that this separation angle limit is less than 2 degrees from the true angular resolution. Finally, we see that the average angular error made on the detected ODF peaks is less than 9 degrees for all order reconstructions. As expected, the error increases as estimation order does because high order frequencies in the estimation are more noisy. Although not shown here, the angular error on the ODF maxima is several degrees larger when no filtering is applied before ODF reconstruction. Overall, these synthetic simulation results suggest that NLMr filtering of the raw HARDI data preserves angular resolution of QBI and that angular error is kept small for small order estimations. We thus choose $L = 4$ reconstructions for the real data experiments.

Real Data Experiment. Fig. 2 illustrates the added value of NLMr filtering on the GFA maps and ODF reconstructions for the biological phantom. We see that NLMr filtering on each DWI dataset removes most of the noise artifacts and enhances the signal in the cords. Since there is major signal loss observed




Fig. 2. NLMr filtering improve ODF reconstructions on a biological phantom

in some of the individual DWI datasets, one cannot hope to obtain a better denoised individual GFA map than the averaged GFA map. However, the average of all four GFA with NLMr filtering qualitatively looks smoother than the averaged GFA without filtering. Hence, most of filtered ODF profiles are more coherent with the underlying structure and we see that the crossing information is preserved. Moreover, the NLM filtering enhances some fiber compartments, indicated in the white box, due to the neighborhood support of each separate principal orientation.

Fig. 3 illustrates the added value of NLMr filtering on the real human brain dataset. We see that each GFA image reconstructed after NLMr filtering of the respective DWI datasets is better denoised than the averaged GFA image. The ODF images show a crossing region in a coronal slice between the corpus callosum (CC) and corticospinal tract (CST). ODF profiles reconstructed after filtering have less spurious peaks and are better aligned with the known anatomy. We also see that ODF profiles remain sharp and multiple peaked. Hence, the Rician noise bias is removed while producing a qualitatively more coherent field of ODF that seems to preserve crossing fiber information, as suggested by the synthetic simulations.




Fig. 3. NLMr filtering improves ODF reconstructions and estimated GFA maps on a human brain HARDI dataset. The ODFs are overlaid on the coronal filtered GFA slice.

4 Discussion and Conclusions

Results on synthetic datasets with low SNR (< 15) and real noisy datasets with SNR between 2 and 5 and b -value 3000 s/mm^2 illustrate the added value of NLMr for HARDI. This suggests that NLMr filtering could be added as a restoration step in diffusion MRI processing pipeline before DTI and HARDI reconstructions. This restoration will become even more important for higher b -value HARDI acquisitions, where the signal is largely attenuated and Rician noise bias dominates the raw DW signal.

Overall, we have presented the impact of non-local means filtering adapted to Rician noise on HARDI datasets with a special focus on the estimation of the ADC and ODF profiles as well as the fiber detection and angular limitation of the ODF reconstruction on controlled synthetic datasets with varying SNR and b -value. The NLMr filtering was evaluated on a biological phantom with known underlying fiber bundle orientation. The added value of NLMr filtering was shown on a real human brain dataset with crossing and poor SNR. Our results suggest that the NLMr filtering improve the quality of anisotropy maps from ADC/ODF estimation and improve the coherence of q-ball ODFs with the underlying anatomy while preserving angular resolution. Therefore, the need

of multiple measurement and averaging along each gradient directions seems unnecessary for human brain acquisitions, which can potentially greatly reduce acquisition time. It is now important to further analyze the effect of NLMr filtering on the b -value used and on the choice of local HARDI reconstruction algorithm and on the resulting tractography output.

References

1. Jones, D.K., Basser, P.J.: Squashing peanuts and smashing pumpkins: how noise distorts diffusion-weighted MR data. *Magn. Reson. Med.* 52(5), 979–993 (2004)
2. Basu, S., Fletcher, P.T., Whitaker, R.T.: Rician Noise Removal in Diffusion Tensor MRI. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 117–125. Springer, Heidelberg (2006)
3. Landman, B., Bazin, P.L., Prince, J.: Diffusion Tensor Estimation by Maximizing Rician Likelihood. In: Lew, M., Sebe, N., Huang, T., Bakker, E. (eds.) ICCV 2007, Rio de Janeiro, Brazil, pp. 1–8. IEEE Computer Society, Los Alamitos (2007)
4. Fillard, P., Arsigny, V., Pennec, X., Ayache, N.: Clinical dt-mri estimation, smoothing and fiber tracking with log-euclidean metrics. *IEEE Transactions on Medical Imaging* 26(11), 1472–1482 (2007)
5. Savadjiev, P., Campbell, J.S.W., Pike, G.B., Siddiqi, K.: 3d curve inference for diffusion mri regularization and fibre tractography. *Medical Image Analysis* 10, 799–813 (2006)
6. Assemlal, H.-E., Tschumperlé, D., L.B.: Fiber tracking on hardi data using robust odf fields. In: IEEE International Conference on Image Processing, pp. 344–351 (2007)
7. Aja-Fernandez, S., Niethammer, M., Kubicki, M., Shenton, M.E., Westin, C.F.: Restoration of dwi data using a rician lmmse estimator. *IEEE Transactions Medical Imaging* (in press, 2008)
8. Clarke, R.A., Scifo, P., Rizzo, G., Dell’Acqua, F., Scotti, G., Fazio, F.: Noise correction on rician distributed data fibre orientation estimators. *IEEE Transactions Medical Imaging* (in press, 2008)
9. Wiest-Daesslé, N., Prima, S., Coupé, P., Morrissey, S.P., Barillot, C.: Rician noise removal by non-local means filtering for low signal-to-noise ratio MRI: Application to DT-MRI. In: MICCAI (in press, 2008)
10. Buades, A., Coll, B., Morel, J.M.: A non local algorithm for image denoising. In: CVPR 2005, San Diego, USA, pp. 60–65. IEEE Computer Society, Los Alamitos (2005)
11. Gasser, T., Sroka, L., Steinmetz, C.: Residual variance and residual pattern in non linear regression. *Biometrika* 73(3), 625–633 (1986)
12. Wiest-Daesslé, N., Prima, S., Coupé, P., Morrissey, S.P., Barillot, C.: Non-local means variants for denoising of diffusion-weighted and diffusion tensor MRI. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 344–351. Springer, Heidelberg (2007)
13. Frank, L.: Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine* 47(6), 1083–1099 (2002)
14. Alexander, D., Barker, G., Arridge, S.: Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine* 48(2), 331–340 (2002)

15. Ozarslan, E., Mareci, T.: Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution imaging. *Magnetic Resonance in Medicine* 50, 955–965 (2003)
16. Chen, Y., Guo, W., Zeng, Q., Yan, X., Huang, F., Zhang, H., He, G., Vemuri, B., Liu, Y.: Estimation, smoothing, and characterization of apparent diffusion coefficient profiles from high angular resolution dwi. *Computer Vision and Pattern Recognition* 1, 588–593 (2004)
17. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Apparent diffusion coefficients from high angular resolution diffusion imaging: Estimation and applications. *Magnetic Resonance in Medicine* 56, 395–410 (2006)
18. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical q-ball imaging. *Magnetic Resonance in Medicine* 58(3), 497–510 (2007)