Tina Memo No. 2007-004
Published in Proc. MIUA 2007

Multispectral MRI De-noising Using Non-Local Means.

J. V. Manjon , M. Robles and N.A.Thacker

Last updated
2 / 04 / 2007

WWW.TINA-VISION.NET

Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

Multispectral MRI De-noising Using Non-Local Means.

J. V. Manjón[1] , M. Robles[1] and N.A. Thacker[2]
1: IBIME Group, ITACA Institute, Universidad Politécnica de Valencia, Valencia, Spain.
2: Imaging Science and Biomedical Engineering, University of Manchester.

Abstract

Clinical MRI data is normally corrupted by random noise from the measurement process which reduces the accuracy and reliability of any automatic analysis. For this reason, de-noising methods are often applied to increase the SNR and improve image quality. Most of these methods work on single channel images by correcting each grey level using an implicit model of the surrounding region, but without taking into consideration the potential multispectral nature of MR images. In this paper we present an extension of a recently proposed filter to reduce random noise in multispectral MR images and test it on synthetic and real images. We compare performance to a multispectral approach based upon the imaging physics and published previously at this conference using real data. We conclude from our results that these methods can be used for de-noising of MR data.

Introduction

MRI de-noising is a common pre-processing step in many MR image processing and analysis tasks, such as segmentation or registration. Many filtering methods are based on the signal averaging principle which uses the spatial redundancy in the image. In this sense, Gaussian filters have been largely used in some applications such as fMRI but they have the disadvantage of blurring edges due to averaging of non similar patterns. In order to avoid this problem many edge preserving filters have been proposed. Probably the best known is the Anisotropic Diffusion Filter (ADF) (Guerig et al,1992; A. Samsonov and C. Johnson, 2004). Such filters respect edges by averaging pixels in the orthogonal direction of the local gradient. However such filtering can erase small features and may change image statistics. Wavelet based filters have also been applied to MRI de-noising (R. Nowak, 1999; J. C. Wood and K. M. Johnson,1999) but such filters tend to introduce characteristic artifacts that can be very problematic for the clinicians.

Most existing filters work on single channel images without taking into consideration the potential multispectral nature of MR images. There are few methods that use the multispectral information as basis for the de-noising process. One of the first attempts to use this information was the multispectral ADF proposed by Gerig (Gerig et al,1992). In their work they propose using the gradient information of the different channels to conduct the diffusion process. A wavelet base de-noising technique for multispectral images exploiting interscale and interchannel correlations has also been proposed (Scheunders, 2004). This technique is demonstrated to outperform single channel wavelet thresholding techniques. Finally, a partial volume segmentation based approach has been recently proposed in Thacker (Thacker et al, 2004) where the filtering is performed using multidimensional data and a partial volume data density model (MPVM). This approach abandons altogether local smoothness constraints and achieves noise filtering by enforcing agreement between measured data using underlying tissue proportions computed from a physics based image formation model. This work also introduced a new way of characterising noise filters, which permits direct comparison of filtering techniques on real (as opposed to simulated) data, thereby obtaining more meaningful measures without the need for a gold standard. This multi-spectral filtering method was previously evaluated in comparison to single image filtering approaches. At the time a state of the art multispectral noise filter was not implemented for comparison.

In the present work we extend the application of a new filter recently proposed by Buades (Buades et al 2005) know as Non-local means (NL-means) to de-noise multispectral MR images. The technique is first evaluated in the conventional way using simulated data sets, and then compared to the MPVM noise filter using our previously suggested evaluation method on real data.

Methods

The NL-means filter is an evolution of the Yaroslavsky filter (Yaroslavsky, 1985) which averages similar image pixels defined according to their local intensity similarity. The main difference between the NL-means and this filter is that the similarity between pixels has been made more robust to noise by using a region comparison, rather than pixel comparison and also that matching patterns are not restricted to be local. That is, pixels far from the pixel being filtered are not penalised.

Given an image Y the filtered value at a point p using the NL-means method is computed as a weighted average of neighbouring pixels N_p in the image following this formula:

$$NL(Y(p)) \;=\; \sum_{\forall q \in N_p} w(p,q)Y(q) \quad with \quad 0 \leq w(p,q) \leq 1 \quad and \quad \sum_{\forall q \in N_p} w(p,q) \;=\; 1 \qquad (1)$$

where p is point being filtering and q represents any other image pixel[1]. The weights $w(p,q)$ are based on the similarity between the neighbourhoods N_p and N_q of pixels p and q. N_i is defined as a square neighbourhood window centred around pixel i with a user-defined radius R_{sim}. Theoretically, noise filtering performed must be considered as an estimation task. Therefore, the process of linear weighting and the weight factors $w(p,q)$ can be regarded as the calculation which computes the most likely noise free grey-level value of the selected pixel, on the basis of the measured evidence. A simple example which leads to this form of solution can be derived using Likelihood, on the assumption that the measurements form each image ($Y(q)$ and $Y(p)$) can be taken as an independent estimate of the noise free value, drawn from a Gaussian distribution with variance $1/w(p,q)$. The process of selecting the most effective noise filtering algorithm can be considered as a way of optimising the match between the assumed computational form and the statistical distributions in the data.

In the case of NL-means, $w(p,q)$ is calculated as:

$$w(p,q) \;=\; \frac{1}{z(p)} exp(-\frac{\sum_i \sum_j G_\rho(i,j)(Y(p_{ij}) - Y(q_{ij}))^2}{h^2}) \;=\; \frac{1}{z(p)} exp(-d(p,q)/h^2) \qquad (2)$$

$z(p)$ is the normalising constant, h is a exponential decay control parameter and G_{rho} is a normalised Gaussian weighting function with zero mean and standard deviation ρ (generally set to 1), so that d is a Gaussian weighted squared Euclidian distance of all the pixels of the neighbourhood. The normalisation factor is defined by:

$$z(p) \;=\; \sum^{N_p} exp(-d(p,q)/h^2) \qquad (3)$$

The weighting process penalises pixels far from the centre of the neighbourhood window giving more weight to pixels near the centre[2]. To eliminate over-weighting in (1), the original NLM method w(p,p) was calculated as:

$$w(p,p) = max(w(p,q), \forall q \neq p) \qquad (4)$$

However, this correction can have the disadvantage of blurring singular points (i.e. pixels with no similar patches, like image corners and peaks or valleys) by averaging them with non similar patches[3].

As the acquisition of multispectral sequences is common in clinical practice the above method can be extended to be used on a multichannel framework. Effectively, the similarity measure can be better obtained by combining information of various channels. Therefore, we propose to use a multispectral similarity function described as follows:

$$w(p,q) \;=\; \frac{1}{z(p)} exp(\frac{-\sum_i^C d(p_i,q_i)/h_i^2}{C}) \qquad (5)$$

where $z(p)$ is the appropriate normalisation factor, C is the number of channels and h_i parameter is closely related with the noise standard deviation of each channel. We will refer to this multispectral method as MNLM.

The NLM algorithm has three parameters and the filter results depend highly on their setting. The first parameter is the radius of the search window enclosing N_P. The second parameter, R_{sim}, is the radius of the neighbourhood window used to find the similarity between two pixels. If the value of R_{sim} is increased the similarity measure will be more robust but fewer similar neighbourhoods will be found. The third parameter, h, is related to the decay of the exponential curve and controls the degree of smoothing. If h is too small, little noise will be removed while if h is set too high, the image will be blurred. In our experiments we use an 11x11 search window, which seems a reasonable value for medical images. The best setting for R_{sim} and h under different noise levels was estimated to be $R_{sim} = 1$ and h is set to the estimated standard deviation of image noise.

[1] In principle N_p can include the whole image, though efficiency requirements normally prevent this.

[2] The center pixel of the Gaussian weighting window is set to the same value that the pixels at a distance 1 to avoid over-weighting effects.

[3] To overcome this situation, we apply equation 4 only if the maximum w is above to a fixed threshold. We have fixed this value to 0.002 which corresponds to pixels far more than 2.5 times the image noise standard deviation. This modification is especially effective on low noise conditions.

To conduct the experiments we used a simulated T1-weighted, PD-weighted and T2-weighted, 1mm3 voxel resolution images (8 bit quantization) from the Brainweb phantom (C. Cocosco et al.,1997) (fig 1) [4]. All experiments were performed using MATLAB 7.0 (Mathworks Inc.).

Experiments and Results

The accuracy of the filter over typical levels of MR image noise (1 % to 9 %) was evaluated with the mono and multispectral version of the filter to show the effect of adding multiple channels in the image de-noising process.

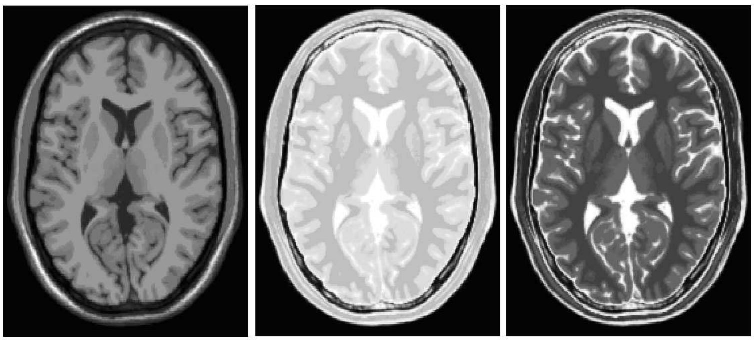

Figure 1: From left to right. T1-Weighted synthetic noise free MR image, PD and T2.

Four different methods of the NLM were evaluated, one mono-spectral and three multispectral. The Root Mean Square Error (RMSE) was used to evaluate the differences between the de-noised image (T1) and the original without noise. Results are summarised on fig 2. As can be seen, to include multiple channels improves filter

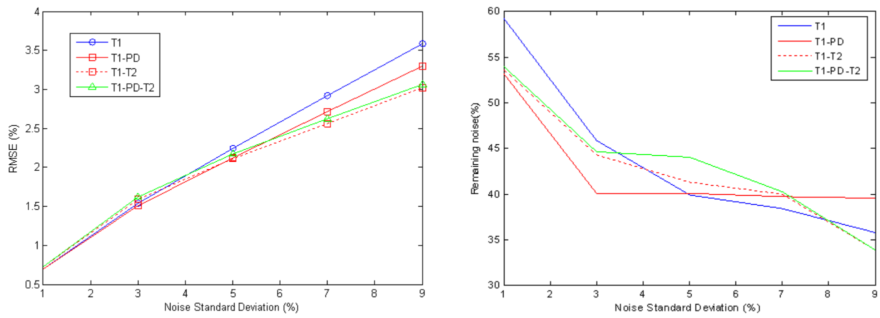

(a) (b)

Figure 2: Fractional removal of RMSE for the different methods evaluated.

performance. The best results were obtained with the combination T1-T2 probably due to a better tissue contrast.

Estimates of noise removal using simulated data have a number of restrictions. The main one of these is that we can never be sure that a simulation accurately matches the problems seen in real data. How do we know that the the filter is doing something appropriate at all locations? Secondly, performance figures at low noise levels are diffciult to interpret, as most simulation data is derived at some point from real data and contains residual noise which will inevitably be removed by the noise filter, preventing a low RMSE. In figure 2 we might argue that results below 3% noise are underestimating the true capabilites of the filter, and beyond 3 % the performance can be considered as effectively constant, as expected for an averaging based scheme.

For this reason it has been suggested that noise filters can be assessed using real image data using two complementary performance measures. The first, the residual outlier measure (ROM) simply measures the number of image values which have been changed unreasonably by the application of the filter (ie: by more than 3 standard deviations of the estimated image noise). The second evaluates the efficiency of the noise filter by measuring the proportion of additional noise which passes the filter. These measures are now given for the multispectral NLM filter.

[4]Although it is well known that magnitude MRI data is Rician distributed, for signals away from zero such as in the region of the brain, it can be well approximated by a Gaussian.

To evaluate effectiveness of the multispectral approach, four real images were used (IRTSE, PD,T2 and FLAIR top row in Figure 3). The noise level in each image was estimated using the LNE technique (Thacker et al, 2003). The proposed method has been compared with the multispectral partial volume modelling (MPVM) method proposed by Thacker (Thacker et al, 2003). Although the proposed method does not remove as much noise as the MPVM method it is less destructive.

	LNE	MPVM	MNLM	Gaussian Filter
IRTSE	58.76	0.22 (1689)	0.50 (125)	0.27 (2405)
PD	64.06	0.20 (1804)	0.55 (45)	0.26 (3127)
T2	58.2	0.17 (938)	0.44 (43)	0.26 (1909)
FLAIR	52.4	0.13 (4971)	0.56 (194)	0.27 (3827)

Table 1: Table 1: Monte-Carlo estimate of fraction of remaining noise following Filtering and data lying beyond 3 S.D. of original value following Filtering (brackets). The results for standard Gaussian smoothing with a kernal width of 1 pixel are provided for comparison.

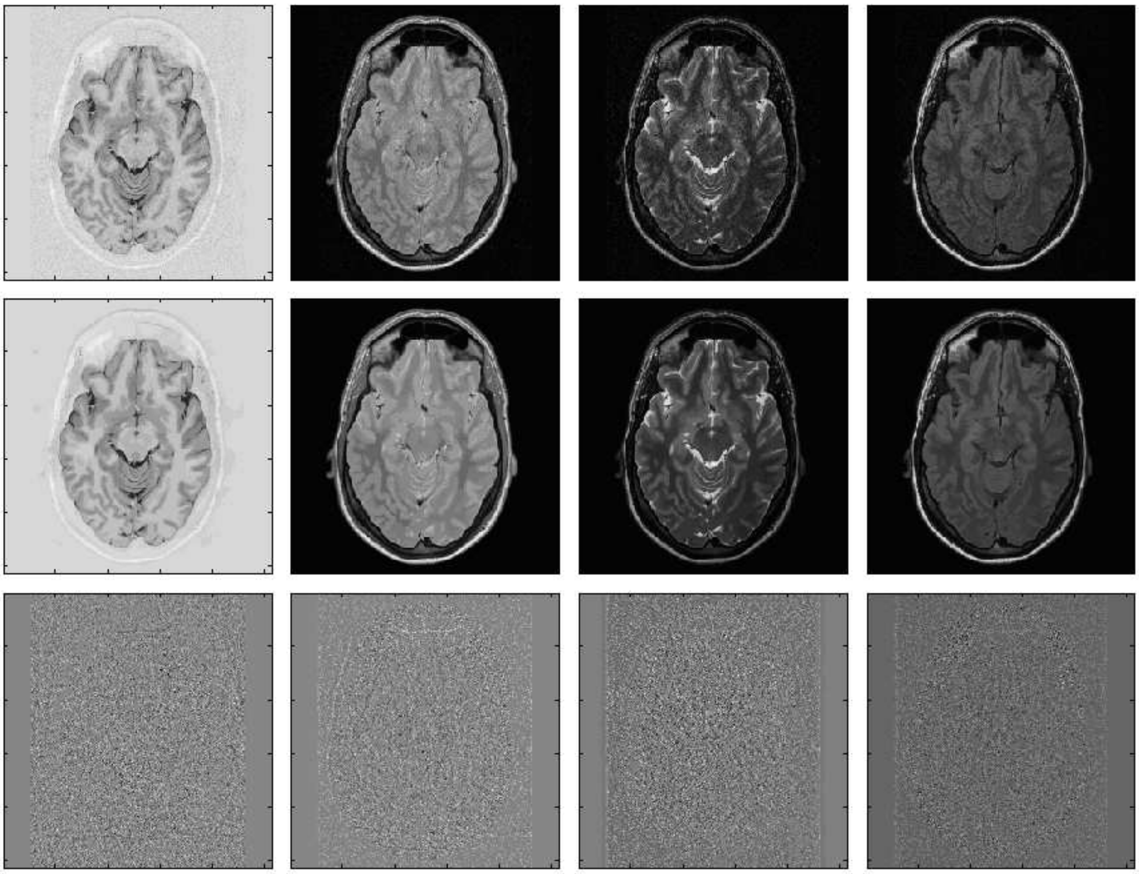

Figure 3: MNLM results (original, de-noised and residuals)

Images of the filtered images are shwn in Figure 3. Although the MPVM method is seen to have structural patterns in the estimated noise which are due to the constraints imposed on the estimated tissue proportions across the 4 input images, the MNLM method shows little evidence for tissue dependant changes.

Conclusion

In this paper we have explained how, for a noise filter to be statistically valid, its algorithmic form should be consistent with the statistical distributions present in real data. In practice, these distributions are often not measured but instead algorithms are tuned by directly adjusting for best performance. This process has been applied here for the application of the method of an NLM noise filter to multispectral data.

The new method of multispectral filtering has been evaluated in the standard way using simulated data. In addition, given the limitations of such an approach, we have compared the performance of the filter to other approaches on the basis of measures computed from real data. There appears to be sufficient evidence to conclude that empirically based filters, such as NLM, have superior performance to simple noise filters such as median and Gaussian filtering, and that also this performance is enhanced if data are combined across multiple MR sequences, in the form of a multispectral filter. The method can be thought of as a bootstrap filtering scheme which filters on the basis of recurring stucture in a given image, as opposed to say Bayesian (or any model based) methods which pre-suppose the same prior expectation for all images. The performance of this filter also compares well with a physics based approach (MPVM) which reconstructs the most likely image data from estimates of tissue proportions. Though it is slightly less efficient at removing noise, it seems to be less destructive, particularly for FLAIR images, in which the MPVM filter also corrects image processes which are not described by a linear model of image formation, such as flow artefacts. For the chosen control parameters, the level of data modification by the NLM is less than that predicted by the measured noise level of the image for a perfect noise filter (ie: for Gaussian noise we expect a value which is the integral of a Gaussian beyond 3 S.D. $\approx 1,200$) . Such under-correction errs on the side of caution for medical data where image details play an important role in medical diagnosis. As the MNLM filter does not transform the remaining noise on correlated noise (noise-to-noise principle (Buades,2005)) the filter can potentially be applied iteratively to remove further noise without introducing a significant number of outliers. However, this claim must be moderated by the observation that any averaging process is likely to suppress the appearance of unique patterns in the data.

The main advantages of NLM are probably the simplicity of this model free approach (certainly in comparison to MPVM) and its non destructive nature, though execution time may be an issue for some implementations and setting up the filter is still a process of trial and error. For situations where the image formation process matches a linear process and unique locations need to be preserved, the physics based approach (MPVM) has built in parameter estimation and superior noise filtering characteristics. These issues may yet be addressed with more research.

References

[1] Paul Scheunders. Wavelet Thresholding of Multivalued Images, IEEE TRANSACTIONS ON IMAGE PROCESSING, 13(4), 475-483, 2004.

[2] G. Gerig, O. Kubler, R. Kikinis, and F. A Jolesz. Nonlinear Anisotropic Filtering of MRI Data. IEEE Trans. Med Imag., vol. 11, pp. 221-232, 1992.

[3] M.Pokric, N.A.Thacker and D.C. Williamson, Noise Filtering and Testing for MR Using a Multi-Dimensional Partial Volume Model. Proc. MIUA, 21-24, London, 2004.

[4] A. Buades, B. Coll, J.M Morel, "A review of image denoising algorithms, with a new one" , Multiscale Modeling and Simulation (SIAM interdisciplinary journal), Vol 4 (2), pp 490-530, 2005.

[5] C.A. Cocosco, V. Kollokian, R.K.-S. Kwan, A.C. Evans : "BrainWeb: Online Interface to a 3D MRI Simulated Brain Database". NeuroImage, vol.5, no.4, part 2/4, S425, 1997 – Proceedings of 3-rd International Conference on Functional Mapping of the Human Brain, Copenhagen, May 1997.

[6] A. Samsonov and C. Johnson. Noise-Adaptive Nonlinear Diffusion Filtering of MR Images With Spatially Varying Noise Levels. Magnetic Resonance in Medicine 52:798-806, 2004.

[7] J. C. Wood and K. M. Johnson. Wavelet Packet Denoising of Magnetic Resonance Images: Importance of Rician Noise at Low SNR. Magnetic Resonance in Medicine 41:631-635, 1999.

[8] R. D. Nowak, Wavelet-based Rician noise removal for magnetic resonance imaging, IEEE Transactions on Image Processing, 8:10, 1408 -1419, 1999.

[9] L.P. Yaroslavsky, Digital Picture Processing - An Introduction, Springer Verlag, 1985.