

Classification and performance of denoising algorithms for low signal to noise ratio magnetic resonance images

Wilfred L. Rosenbaum^a, M. Stella Atkins^a, Gordon E. Sarty^b

^a Dept. of Computing Science, Simon Fraser University, Burnaby, BC, Canada, V5A 1S6

^b Dept. of Medical Imaging, University of Saskatchewan, Saskatoon, Sask.

ABSTRACT

The generation of magnitude magnetic resonance images comprises a sequence of data encodings or transformations, from detection of an analog electrical signal to a digital phase/frequency k -space to a complex image space via an inverse Fourier transform and finally to a magnitude image space via a magnitude transformation and rescaling. Noise present in the original signal is transformed at each step of this sequence. Denoising MR images from low field strength scanners is important because such images exhibit low signal to noise ratio.

Algorithms that perform denoising of magnetic resonance images may be usefully classified according to the data domain on which they operate (ie at which step of the sequence of transformations they are applied) and the underlying statistical distribution of the noise they assume. This latter dimension is important because the noise distribution for low SNR images may be decidedly non-Gaussian.

Examples of denoising algorithms include 2D wavelet thresholding (operates on the wavelet transform of the magnitude image; assumes Gaussian noise), Nowak's 2D wavelet filter (operates on the squared wavelet transform of the magnitude image; assumes Rician noise), Alexander et. al.'s complex 2D filters (operates on the wavelet transform of the complex image space; assumes Gaussian noise), wavelet packet denoising (wavelet packet transformation of magnitude image; assumes Rician noise) and anisotropic diffusion filtering (operates directly on magnitude image; no assumptions on noise distribution).

Effective denoising of MR images must take into account both the availability of the underlying data, and the distribution of the noise to be removed. We classify a number of recently published denoising algorithms and compare their performance on images from a 0.35T permanent magnet MR scanner.

Keywords: Denoising, MRI, Wavelets, Rician

1. INTRODUCTION

Removal of noise is an important part of magnetic resonance (MR) imaging, especially in low- and mid-field systems which exhibit relatively low signal to noise ratios. Wavelet-based denoising methods have proven useful in signal processing in general and in MR imaging in particular. However, the production of MR images involves several transformations of the originally sampled data, and it is not clear at which transformational stage the data should be denoised.

1.1. k -space Data Acquisition

For a single slice MR image the raw data consists of an $N_p \times N_F$ array of complex (k -space) entries, where N_p and N_F denote the number of phase-encoding and frequency-encoding steps, respectively. The observed (noisy) k -space data \mathbf{Y} may be written as $\mathbf{Y} = \mathbf{S} + \mathbf{E}$, where \mathbf{S} and \mathbf{E} denote the true signal and the noise respectively. Elements of \mathbf{E} are assumed to be complex Gaussian random variables with mean $0 + 0i$. Specifically

$$\mathbf{E}_{\phi\nu} = \mathbf{E}_{R\phi\nu} + i\mathbf{E}_{I\phi\nu}$$

where $\mathbf{E}_{R\phi\nu}$ and $\mathbf{E}_{I\phi\nu}$ are independent identically distributed $N(0, \sigma)$ variables.

Further author information:

W.L.R.: E-mail: rosen@sfu.ca

M.S.A.: E-mail: stella@cs.sfu.ca

G.E.S.: E-mail sarty@maya.usask.ca

The $N_P \times N_F$ matrix of K-space data is zero padded to form an $N \times N$ complex array ($N = 256$). The distribution of the zero padded data matrix (which we still call \mathbf{Y}) is now given by

$$\begin{aligned} \mathbf{Y}_{\phi\nu} &\approx N(\mathbf{S}_{\phi\nu}, \sigma^2) & \frac{N - N_P}{2} \leq \phi < \frac{N + N_P}{2} & \quad 0 \leq \nu < N \\ \mathbf{Y}_{\phi\nu} &= 0 & 0 \leq \phi < \frac{N - N_P}{2}, \quad \frac{N + N_P}{2} \leq \phi < N, & \quad 0 \leq \nu < N \end{aligned}$$

1.2. Inverse Fourier Transform

The inverse Fourier transform of the zero padded $N \times N$ matrix is computed. In our particular implementation the transform is not scaled, so that

$$\begin{aligned} y_{mn} &= \sum_{\phi=0}^N \sum_{\nu=0}^N Y_{\phi\nu} e^{-2\pi i(\phi m + \nu n)/N} \\ &= \sum_{\phi=0}^N \sum_{\nu=0}^N S_{\phi\nu} e^{-2\pi i(\phi m + \nu n)/N} + \sum_{\phi=0}^N \sum_{\nu=0}^N E_{\phi\nu} e^{-2\pi i(\phi m + \nu n)/N} \\ &= s_{mn} + e_{mn} \end{aligned}$$

Because the Fourier transform is orthogonal, the e_{mn} are complex Gaussian variates with mean $0 + 0i$, and common variance $N_P N_F \sigma^2$ (the zero padding entries in \mathbf{Y} do not contribute to the variance of e_{mn}). Hence $y_{mn} \approx N(s_{mn}, N_P N_F \sigma^2)$.

1.3. Magnitude Transformation

Next the magnitude of the Fourier transformed data is taken

$$x_{mn} = |y_{mn}| = \sqrt{y_{Rmn}^2 + y_{Imn}^2}$$

The nonlinear magnitude transformation changes the distribution of the noise from Gaussian to Rician. In particular, the distribution of x_{mn} is given by

$$p_x(x_{mn}; s_{mn}, N_P N_F \sigma^2) = \frac{x_{mn}}{2\pi N_P N_F \sigma^2} e^{-\frac{x_{mn}^2 + |s_{mn}|^2}{2N_P N_F \sigma^2}} I_0\left(\frac{|s_{mn}| x_{mn}}{N_P N_F \sigma^2}\right)$$

where I_0 denotes the modified zero-order Bessel function of the first kind.³

Typically, an affine transformation

$$z_{mn} = U x_{mn} + V$$

for some appropriately chosen constants U and V , which scales the magnitude data to lie in an appropriate range for display. The distribution for the scaled pixel data is

$$p_z(z_{mn}; s_{mn}, N_P N_F \sigma^2) = \frac{z_{mn} - V}{2\pi N_P N_F U^2 \sigma^2} e^{-\frac{|(z_{mn} - V)/U|^2 + |s_{mn}|^2}{2N_P N_F \sigma^2}} I_0\left(\frac{|s_{mn}|(z_{mn} - V)}{N_P N_F \sigma^2}\right) \quad (1)$$

2. DENOISING METHODS

The objective of wavelet based denoising methods is to remove whatever noise is present in an image, and leave the underlying signal unaffected, regardless of the frequency content of the signal. Typically such algorithms are based on the following outline:

1. transformation of the original noisy image to a wavelet domain
2. nonlinear processing of the wavelet transform coefficients

3. inversion of the processed wavelet transform coefficients

Wavelet shrinkage denoising refers to a broad class of algorithms that follow this approach. In wavelet shrinkage, wavelet coefficients whose absolute magnitude is below a prescribed threshold τ are set to zero. Coefficients whose absolute value are greater than τ may be left untouched (“hard thresholding”) or may be shrunk towards zero by τ (“soft-thresholding”). Methods for automatic selection of a suitable threshold parameter τ have been the subject of considerable research. One commonly used approach is to use $\tau = \sigma\sqrt{2\log(n)}$ where n is the number of image pixels, and σ is the standard deviation of the underlying Gaussian noise. This “universal threshold” is the basis of Dononho and Johnstone’s *VisuShrink* algorithm, which attempts to visual image quality by removing all of the wavelet coefficients that could statistically be attributed solely to noise.¹

2.1. k -space Denoising

As described in Section 1.1, data in k -space is complex valued with assumed complex Gaussian noise. To apply wavelet shrinkage denoising to this data, we split the data into real and imaginary components and apply shrinkage denoising independently to each component. The real and imaginary denoised transforms are inverted, reassembled to form a single complex matrix which is then inverse Fourier transformed and used to compute the (k space denoised) magnitude image.

2.2. Post Fourier Transform Denoising

After the Fourier transform is performed, the data remain complex with complex Gaussian noise. Wavelet denoising is again applied separately to the real and imaginary components, which are then combined to form a single complex matrix and the magnitude image.

2.3. Magnitude Image Denoising

Although wavelet shrinkage denoising enjoys several asymptotic properties when removing additive Gaussian noise from an image, the noise distribution in low SNR magnetic resonance images is Rician. Nonetheless, we can apply wavelet shrinkage denoising to the magnitude image, although we expect the results to be suboptimal.

2.4. Nowak’s Algorithm

In order to account for signal-dependent Rician noise, Nowak has proposed an alternative to wavelet shrinkage denoising.² Nowak’s method multiplies each wavelet coefficient w by

$$\alpha_w = \left(\frac{w^2 - 3\sigma^2}{w^2} \right)_+$$

where

$$(x)_+ = \begin{cases} x & \text{if } x \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

It is straightforward to derive a maximum likelihood estimator for σ using the magnitude image data only, using Equation (1) and the observation that for background image pixels, the true signal magnitude $|s_{mn}|$ is zero.³

3. RESULTS

In order to identify the transformational stage at which wavelet based denoising is most useful, we applied a variety of algorithms to a set of 32 spin echo MR images. The test images were comprised of 16 axial T2 weighted slices and 16 sagittal proton-density weighted slices from a single patient. Images were acquired on Millennium Technology’s 0.35 Tesla scanner, using $N_P = 160$ phase encoding steps and $N_F = 256$ frequency encoding steps. The wavelet shrinkage denoising algorithms and Nowak’s algorithm were applied to each image using a periodized orthogonal wavelet transformation, varying the choice of wavelet (Haar, Daubechies order 4 and order 10, Symmlet order 4) and the index of the coarsest scale (2,3,6). In addition the wavelet shrinkage denoising methods were performed using both hard and soft thresholding. Gaussian noise variance was estimated using median filtering; Rician noise variance was estimated from the magnitude images using the maximum likelihood estimator described above. Magnitude image SNR was estimated using the formula

$$\widehat{\text{SNR}} = 0.655 \frac{\widehat{\sigma}_F}{\widehat{\sigma}_B}$$

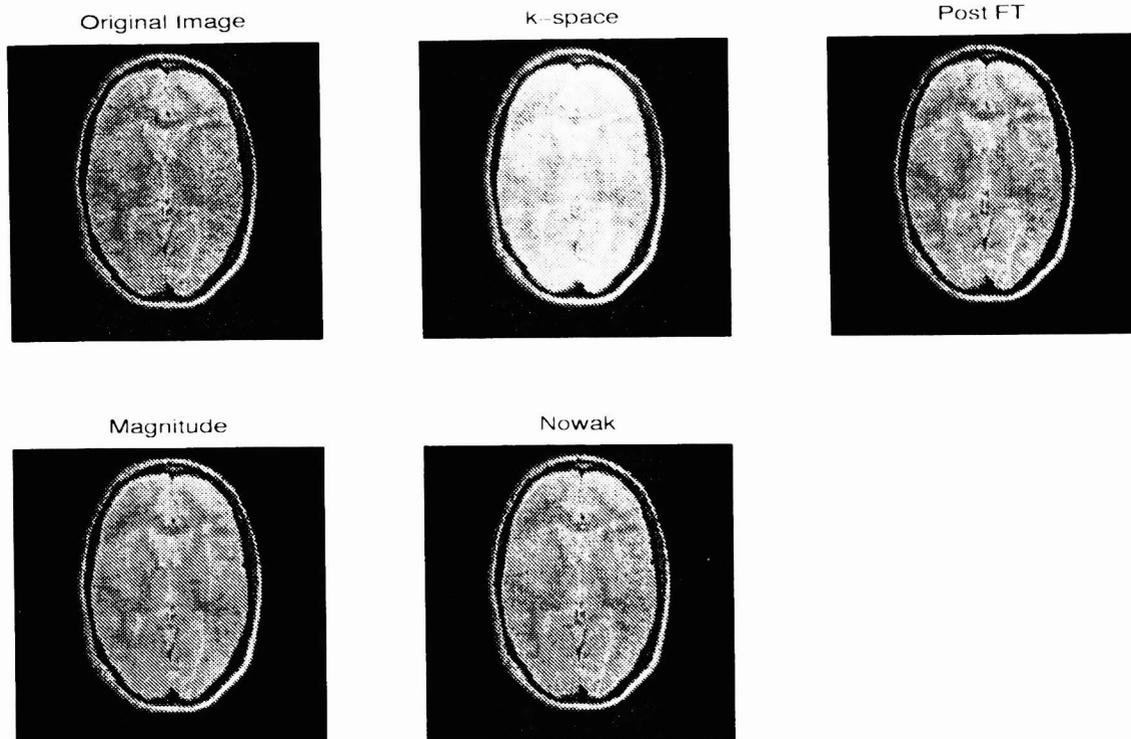


Figure 1. An original axial T2 weighted image, and the results of applying wavelet shrinkage denoising to the k -space data, to the Fourier transformed data and to the magnitude image data. The final image shows the results of using Nowak's data-adaptive wavelet filter.

where $\widehat{\sigma}_B$ is an estimate of the standard deviation of a background region, and $\widehat{\sigma}_F$ is an estimate of the standard deviation of a homogeneous foreground region.⁴ The estimated SNRs of the original test images were all in the range 3 to 9 db.

All of the denoising methods increased the estimated SNRs. To determine the effect of choice of denoising parameters on SNR we performed four separate analyses of variance using as response variable the relative increase in SNR (separate ANOVAs were performed for k space denoising, post Fourier transform denoising, magnitude image denoising and Nowak's algorithm). Independent main effects were included in each model for type of thresholding (soft or hard), coarsest scale of analysis and type of wavelet. The ANOVA model for Nowak's algorithm of course had no term for threshold type. For each wavelet shrinkage method, changes in mean relative SNR were statistically significantly associated with type of threshold and type of wavelet. Soft thresholding was associated with larger mean relative SNR. Symmlets of order 4 were associated with the largest increase in SNR, followed in order by Daubechies order 10, order 4 and order 2 (Haar). The choice of coarsest scale for analysis was not statistically associated with any change in mean relative SNR.

However, SNR is lacking as a measure of image quality. It is always possible to increase the estimated image SNR simply by increasing the image smoothing. SNR cannot distinguish between images corrupted by noise and images corrupted by artifact. Finally, SNR does not correlate well with human perception of image quality. We are ultimately left with a subjective evaluation of image quality, and hence of the denoising approaches we have examined.

Figure 1 presents the results of denoising in each of the different MRI data domains for a typical image. Figure 2 show the same images, but at a much smaller window width so as to highlight the denoising effects on the background noise. Figure 3 presents the absolute value of the difference between each denoised image and the original.

Denoising the k -space data is not recommended. Small errors in the wavelet reconstruction are magnified when

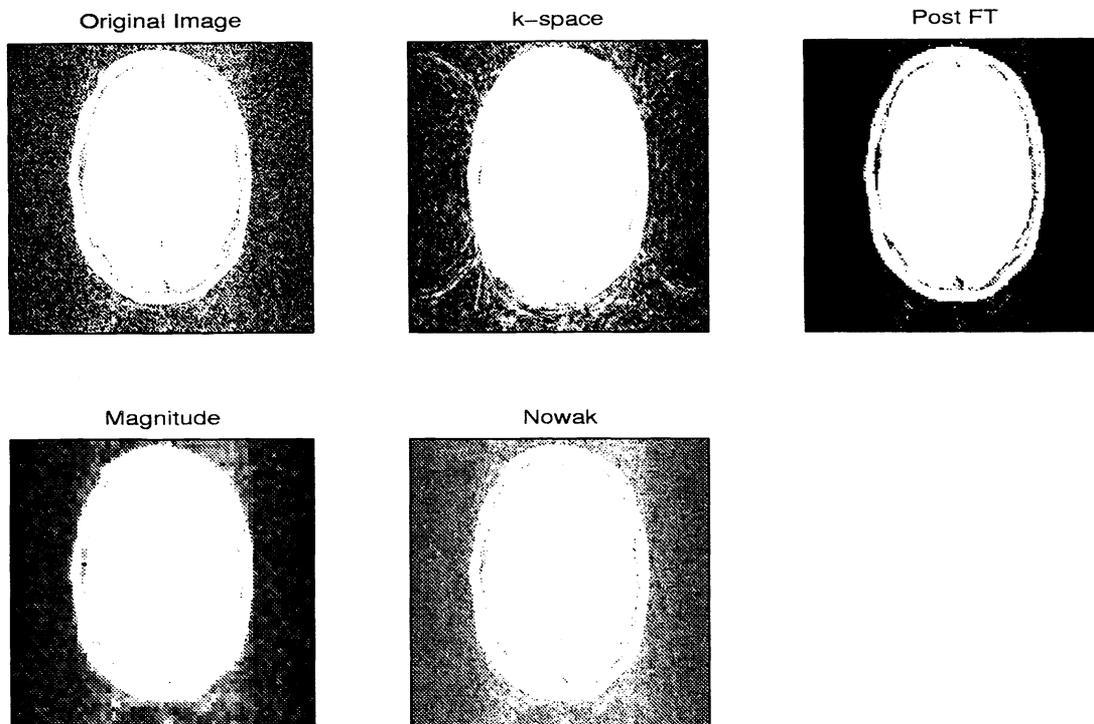


Figure 2. The same images as in Figure 1, but displayed with only 16 gray levels.

the data is Fourier transformed, occasionally resulting in image ghosting. This phenomenon is visible in Figure 2. The k -space data may be inherently too “spiky” to denoise effectively.

Denosing the complex data after it has been Fourier transformed shows promise, and was the approach that demonstrated the largest increase in relative SNR. However, this method often produces results very similar to simple intensity thresholding of the magnitude image, with the attendant loss of image detail and contrast. Figure 4 presents the cumulative distribution of pixel intensities (scaled to 8 bits for display) for a typical original image, and the same image subject to post-Fourier transform denoising. This phenomenon was not observed when denoising in the other data domains.

Direct wavelet shrinkage denoising of the magnitude image proved unsatisfactory: too much detail was lost in the process. Nowak’s data-adaptive wavelet filtering provided the best overall performance. However, Nowak’s method often does create small reconstruction artifacts, similar to those observed in JPEG compressed images. In our own informal user assessment studies, professional radiologists found these artifacts annoying, but admitted they were unlikely to result in any misdiagnosis. It may be possible to use further post-processing techniques to reduce these artifacts.⁵

ACKNOWLEDGMENTS

The research described in this manuscript was made possible by a grant from the Science Council of British Columbia.

REFERENCES

1. D. L. Donoho and I. M. Johnstone, “Ideal spatial adaptation via wavelet shrinkage,” *Biometrika* **81**, pp. 425–455, 1994.
2. R. D. Nowak, “Wavelet-based Rician noise removal for magnetic resonance images,” *IEEE Transactions on Image Processing* **8**(10), pp. 1408–1419, 1999.

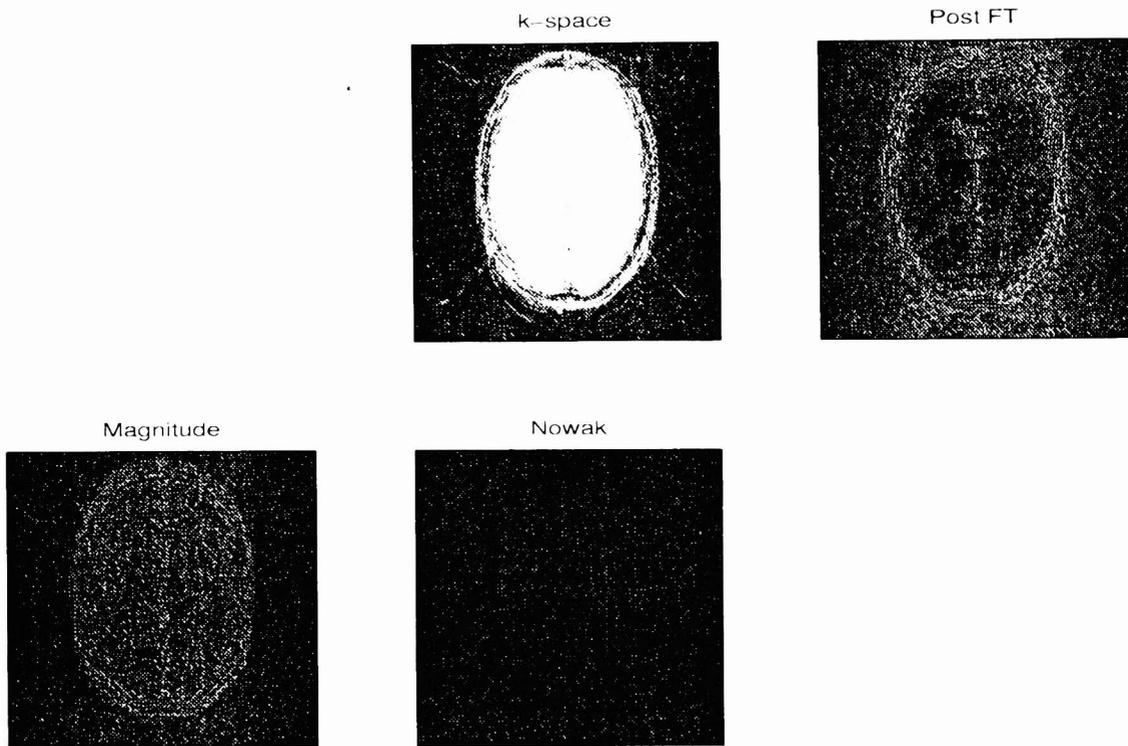


Figure 3. The absolute values of the differences between the denoised and the original images.

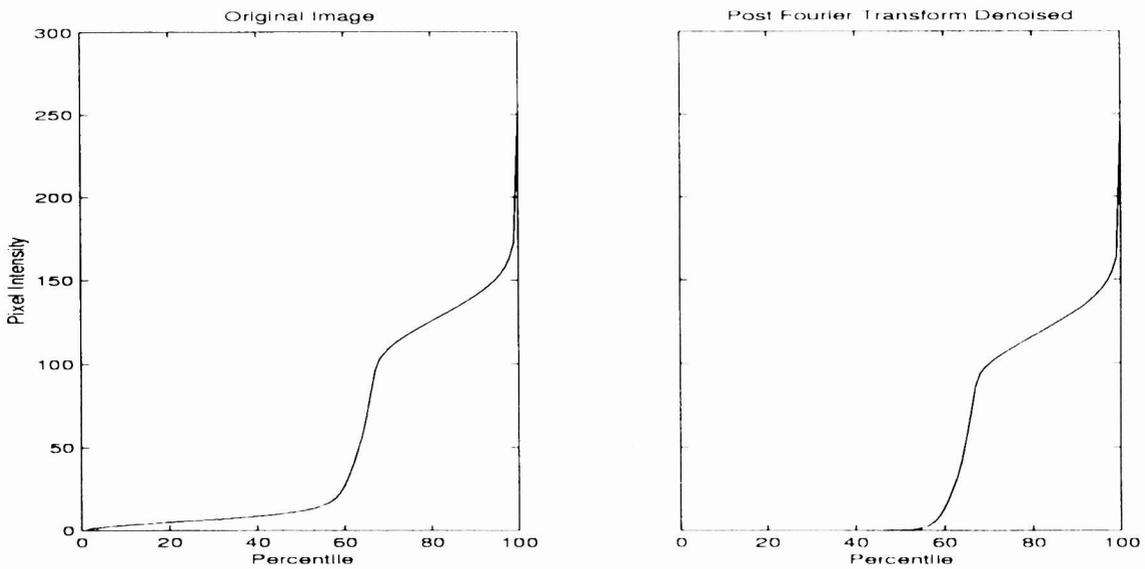


Figure 4. Percentile distribution of pixel intensities in original (left) and post Fourier transform denoised images. Note that more than 50 percent of the pixels in the denoised image have been set to 0.

3. J. Sijbers, A. J. Dekker, P. Scheunders, and D. van Dyke, "Maximum-likelihood estimation of Rician distribution parameters," *IEEE Transactions on Medical Imaging* **17**, pp. 357–361, June 1998.
4. M. Firbank, A. Coulthard, R. Harrison, and E. Williams, "A comparison of two methods for measuring the signal to noise ratio on mr images," *Physics in Medicine and Biology* **44**(12), pp. 261–264, 1999.
5. S. Comes, B. Macq, and M. Mattavelli, "Postprocessing of images by filtering the unmasked coding noise," *IEEE Transactions on Image Processing* **8**, pp. 1050–1062, 1999.