Denoising for Magnetic Resonance Imaging

Akshay Chaudhari
Stanford University
akshaysc@stanford.edu

Abstract

The utility of magnetic resonance imaging can often be diminished in regions and tissues that suffer from a low signal to noise ratio (SNR). This is especially the case in quantitative MRI where the quantitative parameters being measured can be easily biased by low SNR. In this study, the utility of different denoising algorithms, namely non-local means, bilateral filtering, and block-match and 3D filtering (BM3D) were evaluated. BM3D offered the best denoising results while maintaining visual signal quality. This algorithm was further enhanced by making it locally adaptive to the region of interest being measured. It was also shown that genetic algorithms could be easily used to find robust solutions to the multi-dimensional parameter space that the inputs of BM3D encompasses.

1. Introduction

In magnetic resonance imaging (MRI), the raw data that is needed to generate images is directly sampled in the spatial frequency domain. There exist various paradigms to sample the 3D spatial frequency volume, one such being a trajectory of several 3D radial spokes (1), akin to a kooshball. The benefit of such a strategy is that there can be significant undersampling that can be performed without aliasing artifacts, compared to the minimum sampling density prescribed by the Shannon-Nyquist sampling theorem. However, the undersampling does incur a signal-to-noise-ratio (SNR) penalty as a function of the undersampling ratio (2). In addition to visually corrupting the recovered images, noise can be a large hindrance in performing quantitative imaging where voxels with low SNR can significantly bias the quantitative parameter being measured (3).Thus, there is a need to for an efficient image reconstruction process for MRI where denoising methods can be applied to low SNR images in order to improve quantitative measures of MRI as well as increase the visual quality of the images.

1.1. Denoising

Algorithms that denoise images with a priori measures of noise have been quite popular the image processing fields, but have yet to make a mark in clinical MRI (4-6). One potential reason for this might be inefficiency of applying denoising methods to MRI examinations that can typically have between 20 to 200 slices that need to be denoised. In addition, denoising algorithms can increase the SNR of images, but often, this comes at the price locally averaging regions of interest (ROI). To make matters at hand even more challenging, quantitative MRI often times uses images with varying contrasts and signals, as can be seen in Figure 1. Denoising processes essentially downmsample high-resolution MRI data and then discards the high-resolution information by low pass filtering it. However, recent improvements in the creation of denoising methods such as the block matching and 3D filtering (BM3D) might be useful in clinical MRI (7). In order to address some of the common concerns in using denoising for MRI, this work explores the application of different denoising techniques for use in quantitative clinical MRI.

2. Denoising Algorithms

For the scope of this study, different denoising algorithms will be chosen and evaluated on three broad categories, namely; SNR improvement in ROI's, implementation efficiency, and visual quality of denoised images. The first of these (SNR improvement in ROI's) is somewhat self-explanatory for a denoising algorithm. By using this algorithm, there has to be a measurable increase in the SNR of a certain region (of a user's choosing) in the images. This SNR increase also has to be balanced with the time taken for the implementation of this algorithm because unless these methods are efficient, it is challenging to incorporate them in a clinical pipeline. In addition to the aforementioned quantitative outcome measures, visual quality of the denoised images is also a vital metric to balance so that overall clinical utility from the MR image is not diminished due to blocking and low-pass filtering artifacts. With these performance metrics, three denoising algorithms were in-

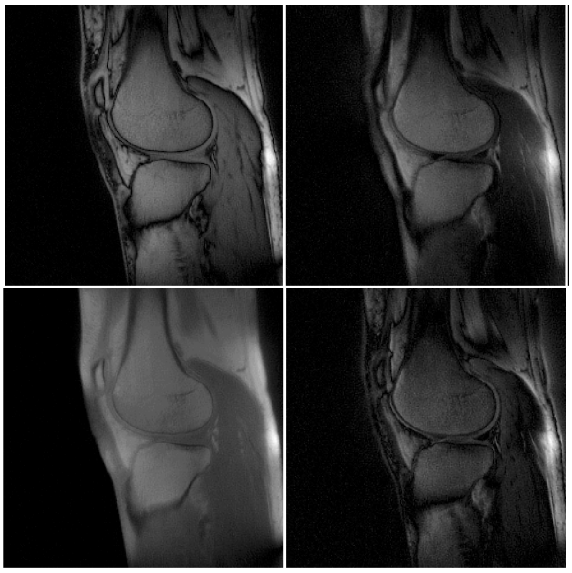

Figure 1: Images used in quantitative MRI can have varying contrasts and varying SNR's in different ROI's.

cluded in the study - BM3D, Non-Local Means (NLM), and Bilateral Filtering (7-9). In order to create a roughly linear quantitative metric with the previous outcomes, the following performance metric would be used to classify the performance of the three algorithms:

$$\rho = \frac{10^{(\lambda)}}{\sqrt{\tau}}\kappa$$

where ρ is the overall performance metric, λ is the gain factor (ratio of new denoised SNR to original SNR), τ is the algorithm run time, and κ is a visual quality. The SNR gain was implemented as an inverse logarithm because SNR improvement is the main goal of the denoising algorithm such that the performance metric ρ can be sensitized to typical SNR gains of 5-20%. The visual quality, κ between 1 and 3, simply ranks the denoised images based on their clinical appearance and is a multiplicative scalar. The denoising algorithms are penalized by the square root of their run time since a time penalty is useful for robustness but not of tantamount importance.

2.1. Bilateral Filter

The bilateral filter implemented by Tomasi and Manduchi is a non-iterative method for denoising images while preserving edges. This filter is edge preserving because it represents filtering an image with a low-pass denoising kernel, while adjusting the kernel to the spatial distributions of pixel-values from the original image. It relies on the creation of a non-linear weighting function which is which is convolved with intensity values of the image. In this manner, if there are any sharp transitions in image intensity i.e. edges, the kernel which encompasses the specific region is weighted with respect to the presence of the edge. This family of filters can produce adequate denoising parameters and the kernel weighting functions have relatively few parameters that need to be optimized on a case-by-case basis.

2.2. Non-Local Means

The Non-local means (NLM) filter, first proposed by Buades *et al*. (9) utilizes the principle that most natural images have self spatial similarities. The conceptual framework for this filter is refreshingly simplistic. Instead of using a local kernel to denoise specific parts of an image, NLM takes advantage of the self similarity of the images and tries to find patches in the rest of the image that are similar to the current patch that is being denoised. This process is referred to as neighborhood filtering and was first proposed b Yarolavsky (10). The advantage of using such neighborhood filters is that there is an increased amount of information regarding the spatial intensities. Different patches can be assigned confidence weightings as functions of similarity to the original patch as well as the Euclidean distance from the center of the original patch. In addition to this, total variation regularization constraints as well as anisotropic filtering (convolving the image in a direction only orthogonal to the maximal gradient) can be implemented to increase performance. Overall, NLM uses inherent image information more efficiently, but since it relies on a large search space, its implementation can be computational bottleneck.

2.3. Block Matching and 3D Filtering

BM3D is a natural extension of the NLM method that also utilizes the notion that natural images have self spatial similarities. The mechanisms for BM3D are similar to that of NLM. First, the algorithm searches for patches in the image that are similar to the intensities of the patch that has to be denoised. This method builds a 3D matrix where the first two dimensions are the size of the patch while the third dimension is the aggregated patches. Unlike NLM which applies a 1D filter on the 3rd dimension of this aggregated array, BM3D applies a 3D unitary sparsifying transform. This transform space is then thresholded and Weiner filtered in order to achieve coefficient shrinkage to remove high-frequency noises. Following the sparsifying transform, the inverse transformation is applied to yield a 3D block that has been denoised. The original patch is then recovered from the 3D array by assigning weights to each patch in the array as a function of the distance of the patch and its variance.

BM3D is a robust algorithm as there are several parameters that can be modified in order to achieve the best denoising for a given image. Depending on parameters, especially size of the 3D patch stack and the patch search radius, the

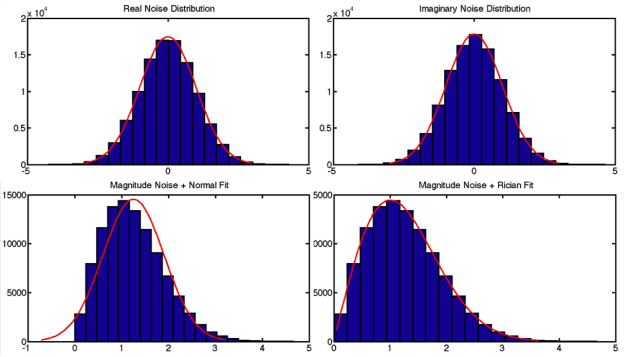

Figure 2: Variation between a Rician and Gaussian noise model when faced with magnitude noise stemming from zero-mean complex Gaussian noise.

method has the ability to function with high SNR improvements in under 1 second on most standard computing platforms.

3. Denoising Implementation

MR images are reconstructed from complex spatial frequency data, thus, the images themselves are complex in nature. Almost all images are stored and viewed in their magnitude form though. While this is not an inconvenience for common utility of the images, this does change the noise distributions that the images have. Nominally, the noise in each of the complex and imaginary channels is a zero-mean Gaussian. However, when a magnitude image is formed, the complex bi-channel zero-mean Gaussian noise turns into Rician noise for low SNR regions in the image while the noise turns into Rayleigh noise for high SNR regions in the image (11). The difference between the noise patterns can be seen in the illustrative example in Figure 2 where zero-mean complex Gaussian noise is combined and fit to a Rician and Gaussian model. It can be clearly seen that a Gaussian cannot approximate the magnitude noise as well as the Rician model.

As a starting point, all three methods were tested on magnitude as well as the complex images. All three algorithms used were implemented in Matlab 2012a. The image that was used for denoising was an MRI scan that shows a sagittal view of the human knee in Figure 3. The goal of the algorithms was to improve the SNR in the low-SNR meniscus region. SNR was computed by averaging the signal of the meniscus ROI and dividing by the standard deviation of the a 40x40 pixel block in the bottom left of the image where there is no signal present. This method is standard in MRI for measuring SNR (12).

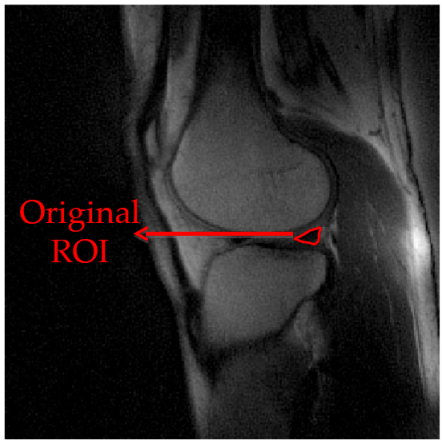

Figure 3: The image being denoised where the the meniscus was the ROI where the SNR was being measured.

Table 1: Results for complex denoising

Method	*Time*	*SNR Gain*	*Quality*	*Performance*
Bilateral	12.0s	0.99	2	5.64
NLM	66.4s	1.71	1	6.27
BM3D	1.2s	1.20	3	43.40

Table 2: Results for magnitude denoising

Method	*Time*	*SNR Gain*	*Quality*	*Performance*
Bilateral	7.3s	0.82	2	4.89
NLM	33.7s	1.14	1	2.38
BM3D	0.6s	1.21	3	62.81

3.1. Results and Discussion

The results for the complex denoising can be seen in Table 1 while the results for the magnitude denoising can be seen in Table 2. The resultant images can be seen in Figure 4.

The non-local means takes significantly higher time to process the images which results in a low performance score. In addition, NLM also produces images that look very blurred out and excessively filtered. It is curious that in the case of complex denoising that the SNR for the NLM experiences an increase of 70 % but upon further examination, this is because the accompany signal from the hyperintense cartilage leaks into the ROI causing the signal increase.

The bilateral filter creates images that look relatively good however, in both the cases of the complex and magnitude denoising, the SNR actually worsens in the low SNR meniscus ROI. Since SNR is weighted strongly in our performance metric, the bilateral filter achieves a low score here.

BM3D seems to perform the best job here because it increases the SNR the most from all three methods, but it also

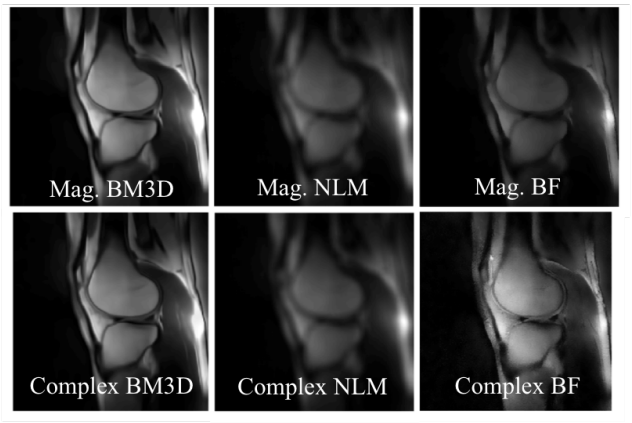

Figure 4: Resultant images from the magnitude (top row) and complex (bottom row) denoising for all 3 algorithms

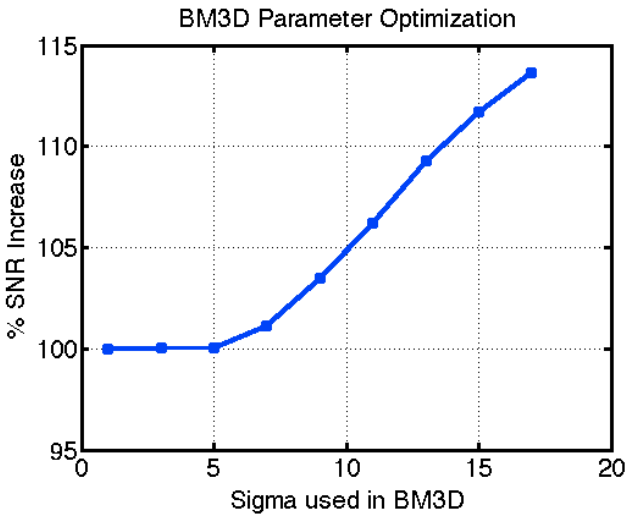

Figure 5: Impact of noise sigma on the SNR of the meniscus

performs its computations in under 1 second per slice. In addition to this, the visual quality of the images is maintained the best amongst all three filters. Due to these effects, it scores the highest in our performance metric. It is curious to note that while there is not a significant difference in SNR improvement between the complex and magnitude denoising, but we choose to proceed with the complex denoising so that the noise statistics are not biased. Biasing the noise has shown to affect quantitative MRI measurements that are generated from low SNR sub-regions.

4. Increasing BM3D Performance

Improvements in SNR of MR images by denoising methods can directly translate to needing fewer data samples acquired during the MRI process. As a rule of thumb, in order to increase the SNR of any given image by $\sqrt{2}$, the scan time for that scan should be doubled (12). Analogous to that argument, if BM3D can give us an acceptable SNR improvement of 1.2, that would mean that we can now effectively undersample the MRI data by a factor of $1.2^2 = 1.4$, which is a significant improvement. However, any additional increase in SNR is beneficial also.

4.1. Parameter Optimization

The BM3D parameter space is multi-dimensional which makes it challenging to establish a set of parameters that are globally optimal. Towards this end however, I first tried to evaluate an approximate sigma for the denoising. Figure 5 shows that as we increase the approximation of noise, the SNR of the ROI goes up. However, with this increased SNR comes the oversmoothing artifact which is not ideal. Thus, as a balance between both, a sigma of 9 (in the image intensity range of 0 to 255) was chosen.

4.2. Locally Adaptive BM3D

One method to make the BM3D algorithm more robust is by replacing global parameters used in the algorithm with case specific parameters. By running simulations on the sample MRI scan, it was seen that the performance of BM3D stayed relatively constant when the number of patches was over 60. Till this point, the SNR kept increasing. This notion could be utilized to dynamically pick how many blocks are needed for optimal denoising. In addition, the search radius for a particular block also influences how many similar blocks are found and can be denoised.

To implement a dynamic locally adaptive BM3D algorithm, the user would first have to manually mark out the foreground of the image. This image would be thresholded and a centroid would be found. The maximum search radius for the BM3D implementation would be the maximum distance between the centroid and any given point in the foreground. A histogram of all the signal intensities in the foreground would also be used in order to compare the signal distribution of the foreground to that of the ROI in question. If the mean ROI SNR is lower than 90% of the other pixels, then the 3D block size would be $(0.90)(60)blocks = 54$ blocks. Similarly, if 80% of the pixels have higher signal than the ROI, then $(0.8)(60) = 48$ blocks would be used. An illustration of this method can be seen in Figure 6 where the image values in the denoised ROI are compared to the global foreground distribution. Similarly, the size of the patch (first two dimensions) was chosen automatically based on the outline of the ROI so that appropriate sized blocks were chosen by the algorithm to compare to the ROI.

Such a method would reduce computational overhead of the ROI being. Not only that, but if the foreground of the image is small, then the algorithm does not waste time looking for blocks away from the radius.

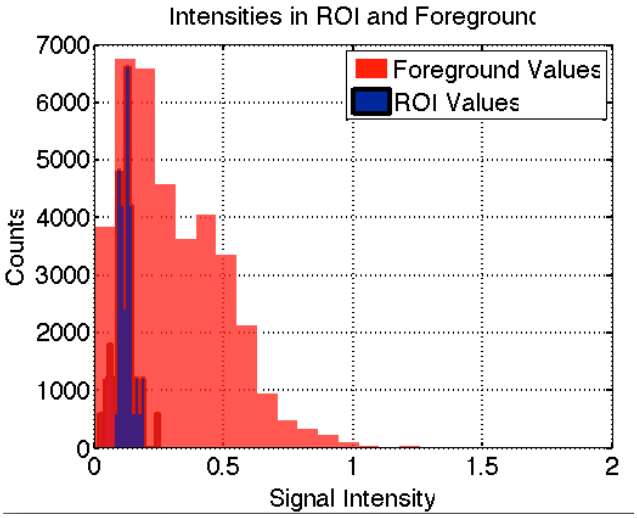

Figure 6: Comparing signal intensities in the denoised ROI to the global intensities (note that the counts for the ROI have been artifically increased to be viewable on this plot).

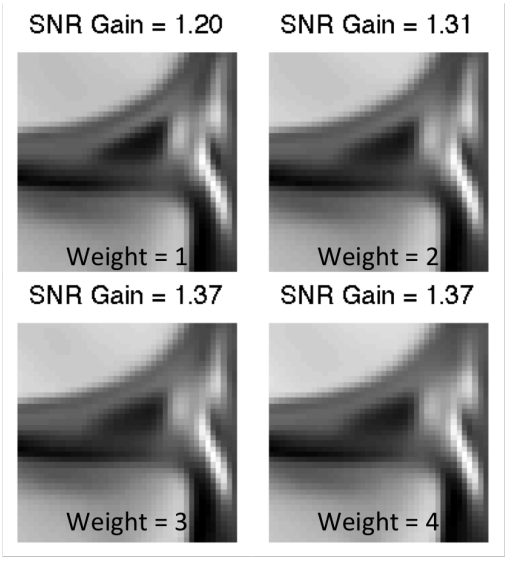

Figure 7: Resultant images as the values for the wavelet thresholding are increased linearly from 1 to 4

4.3. Thresholding Variations

Instead of setting the thresholds for the sparsifying transforms to a constant value, these thresholds can be set differently such that the finer sub-bands of the wavelet transform carry a lower weight after the transformation. Since the noise is mostly a function of the higher frequency information from the image, providing a higher weight to the coarser bands of the wavelet transform may increase SNR gains. However, this comes with downside of the blurring as this is analogous to low pass filtering the data excessively.

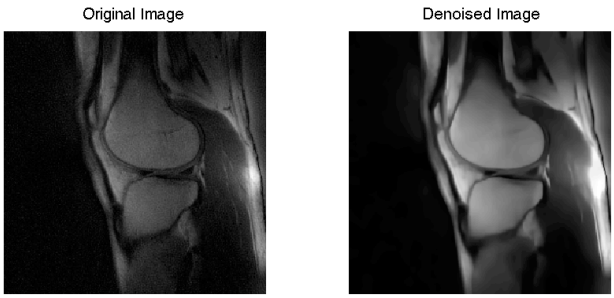

Figure 8: Before and after images for the original image and its optimally denoised variant.

In Figure 7 one can see that as the weighting for the lowest bands of the wavelet is increased from 1 to 4, the perceived SNR increases. However, a closer look at this shows that there is in fact blurring of the signal from the cartilage (white bands around the triangular meniscus) that is leaking into the meniscus. Thus, a happy medium between noise suppression and blurring was chosen to be for the weighting to be 2. This provided an SNR gain of 31% while keeping blurring minimal.

4.4. Combining All Parameter Optimizations

Combining all the parameters that were optimized previously, as well as changing some of the internal BM3D paramters (provided in the code) resulted in an SNR improvement of 25.7%. While this may not seem very significant compared to the initial 20% SNR improvements reported, the image quality is significantly better and there are fewer blocking or over-smoothing artifacts. This might be a result of the locally adaptive parameters. The before and after images can be seen in Figure 8. A SNR increase of around 26% represents having to scan a patient for $(gain^2)$ times less a=which corresponds to a scan time decrease of 58% which in the scope of MRI scans is **extremely** significant.

5. Genetic Algorithms

Genetic algorithms are a class of optimization procedures that are used to find an optimal set of parameters for an unconstrained solution space. I propose that BM3D is effective for the use of genetic algorithms due the various variables that can be optimized, but at the same time, the solution space is unconstrained because the algorithm would have different results for images with differing signal statistics.

Genetic algorithms work not towards find a global maximum for the function being optimized, but rather, a local maximum that is robust to minor changes in the input parameters. The genotype for the algorithm is the set of input

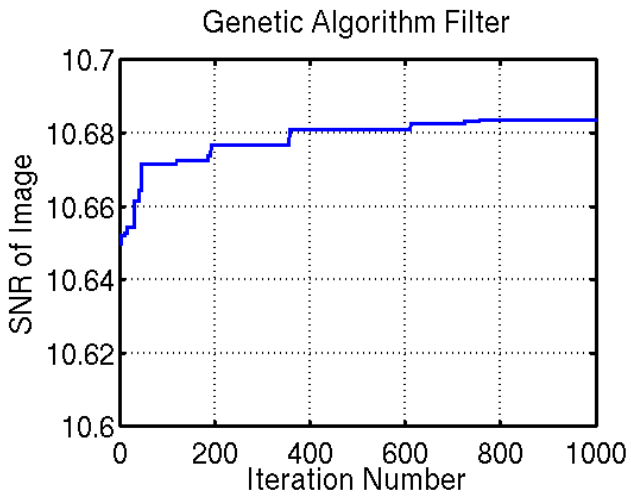

Figure 9: Genetic algorithm for implementing a thresholding matrix for Wavelet transforms

parameters that are used to generate the phenotype output. In the case of BM3D, the genotype would include parameters such as the size of the 3D array, thresholding parameters, search radius, etc. The phenotype for the algorithm would similarly be the output image, SNR improvement, implementation time, etc. The genetic algorithms employ similar processes as used in natural evolution. The initial value is a set of random genotypes that are evaluated for a certain fitness function - the output phenotype that is being maximized. Each generation carries with a mutation to the original genotype. If the mutation is beneficial in promoting fitness, then that mutation is propagated to the next generation. The algorithm is typically run until a certain fitness metric is achieved or if the algorithm has evolved for a user-specified number of generations.

For this study, a genetic algorithm was implemented as a proof of concept to change the thresholding mask and add appropriate weights. The resultant changes in SNR can be seen in Figure 9. Every discrete step in SNR shows a successful mutation being passed onto the new generation. While the SNR gains are modest here, the fast convergence to a locally steady-state maximum in a few number of iterations is promising. In addition, these algorithms can be easily scaled to include multiple parameters that were optimized manually previously. This will increase the generations needed to achieve a significant output, but is shown to be possible here.

6. Future Improvements

Instead of a square search window for each block, a window the shape of the ROI that is being denoised could be used to look for similar images. Advances in neural networks could also be applied to optimizing parameter spaces if there is a database of denoised ROI's that it could learn from. In addition, denoising methods can be used in a routine clinical study to characterize clinical translation.

7. Acknowledgements

The 3 algorithms used here were obtained from the open source community The author of this study would like to thank the program authors for making their code free and accessible for research purposes.

BM3D:
`https://www.cs.tut.fi/~foi/GCF-BM3D/index.html`

Non-Local Means:
`https://www.mathworks.com/matlabcentral/fileexchange/52018-simple-non-local-means--nlm--filter`

Bilateral Filter:
`https://www.mathworks.com/matlabcentral/fileexchange/12191-bilateral-filtering`

8. References

1. NiellesVallespin, Snia, et al. "3D radial projection technique with ultrashort echo times for sodium MRI: clinical applications in human brain and skeletal muscle." Magnetic resonance in medicine 57.1 (2007): 74-81.

2. Gu, Tianliang, et al. "PC VIPR: a high-speed 3D phase-contrast method for flow quantification and high-resolution angiography." American journal of neuroradiology 26.4 (2005): 743-749.

3. Raya, Jos G., et al. "T2 measurement in articular cartilage: impact of the fitting method on accuracy and precision at low SNR." Magnetic resonance in medicine 63.1 (2010): 181-193.

4. Buades, Antoni, Bartomeu Coll, and Jean-Michel Morel. "A review of image denoising algorithms, with a new one." Multiscale Modeling & Simulation 4.2 (2005): 490-530.

5. Perona, Pietro, and Jitendra Malik. "Scale-space and edge detection using anisotropic diffusion." Pattern Analysis and Machine Intelligence, IEEE Transactions on 12.7 (1990): 629-639. APA

6. Nowak, Robert D. "Wavelet-based Rician noise removal for magnetic resonance imaging." Image Processing, IEEE Transactions on 8.10 (1999): 1408-1419.

7. Dabov, Kostadin, et al. "BM3D image denoising with shape-adaptive principal component analysis." SPARS'09-Signal Processing with Adaptive Sparse Structured Representations. 2009.

8. Tomasi, Carlo, and Roberto Manduchi. "Bilateral filtering for gray and color images." Computer Vision, 1998. Sixth International Conference on. IEEE, 1998.

9. Buades, Antoni, Bartomeu Coll, and Jean-Michel Morel. "A non-local algorithm for image denoising." Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on. Vol. 2. IEEE, 2005.

10. L. Yaroslavsky. Digital Picture Processing - An Introduction. Springer Verlag, 1985.

11. Gudbjartsson, Hkon, and Samuel Patz. "The Rician distribution of noisy MRI data." Magnetic resonance in medicine 34.6 (1995): 910-914.

12. Kellman, Peter, and Elliot R. McVeigh. "Image reconstruction in SNR units: A general method for SNR measurement." Magnetic resonance in medicine 54.6 (2005): 1439-1447.