

Adaptive Enhancement of Cardiac Magnetic Resonance (CMR) Images

Erich T.H. Leung and John K. Tsotsos
Department of Computer Science and Engineering and
Center for Vision Research (CVR), York University
Toronto, Ontario, M3J 1P3, Canada
{leung, tsotsos}@cs.yorku.ca

Abstract

This paper presents a wavelet-based framework for enhancing the coherent structures attributable to the target organ in cardiac Magnetic Resonance (MR) images. Previous approaches focus on the Rician nature of noise in magnitude MR images. Image noise is but only one of the confounding factors that obscure the anatomical structures of the target organ. This paper models the image noise in a magnitude MR image in terms of two noise classes which occur over different ranges of signal intensity. An adaptive enhancement scheme is developed to achieve simultaneous attenuation of the effects of these factors and improvement in image contrast.

1. Introduction

Diagnosis of cardiac disease relies upon an accurate assessment of the morphology and function of patients' hearts. The tasks of extracting information about the shape and motion of the target organs rely on the coherency of the structures that provide this information. A major issue at the focus of research is the signal-dependent noise of a magnitude Magnetic Resonance (MR) image, which is generally modeled by a Rician distribution [8, 11]. The bias introduced by Rician noise significantly reduces the image contrast essential for extracting the boundaries of the heart and related anatomical parts. Furthermore, this noise effectively reduces the lossless compression ratio, which is important to image transmission [26].

This research, on the other hand, observes that (1) noise is only one of the factors that may obscure the target structures, i.e., the anatomical structures of the target organ, and (2) not all data are equally important for clinical diagnosis, especially where some of these data are such obscuring factors. For instance, the thoracic cavity in the proximity of the heart appears as a low intensity region where the lungs, blood vessels and other tissues create patterns of dy-

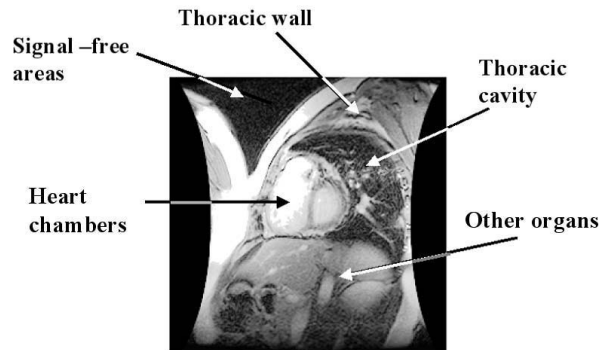


Figure 1. Cardiac magnitude MR image

namic structures, which we will refer to as the background structures hereafter. These fibrous structures lack a definite shape. Like the heart, many of them are dynamic structures in periodic motion. The interactions between these background structures and Rician noise confound the tasks of boundary detection and anatomical reconstruction both spatially and dynamically.

From this perspective, this research ventures into a more general problem, referred to hereafter as adaptive enhancement, that entails the following: (1) recovering the signals of the anatomical structures of the heart from the image noise, (2) enhancing the image contrast attributable to these signals and (3) suppressing the background structures. The challenge of the problem is signal enhancement by attenuating the effects of the confounding factors while at the same time improving image contrast.

This paper presents a wavelet-based solution to the problem of adaptive enhancement. Consider a separable wavelet orthonormal basis generated by the scaling function ϕ and wavelets ψ_κ , where κ indicates the orientations. An image I can be expressed as a superimposition of a J -scale approximation and a succession of details of spatial oscillations at

different scales:

$$I = \sum_k s_k \phi_{J,k} + \sum_\eta w_\eta \psi_\eta \quad (1)$$

$$= \mathcal{P}_v I + \mathcal{R}_w I \quad (2)$$

where $\eta \in \{j\} \times \{k\} \times \{\kappa\}$ is an index over j (the scale), k (the spatial location), and κ (the orientation) [1, 13, 14, 24]. $s_k = \langle I, \phi_{J,k} \rangle$ are the approximation or scaling function coefficients where $\langle \cdot, \cdot \rangle$ denotes the inner product. Similarly $w_\eta = \langle I, \psi_\eta \rangle$ are the wavelet or detail coefficients. For simplification of exposition, we will drop the indices where no confusion arises. $\mathcal{P}_v I$, the projection of the image on the J -scale approximation space \mathcal{V}_J , captures those coarse-scale features visible after smoothing of the signal whereas $\mathcal{R}_w I$ is a collection of the fine-scale details detectable at successive scales of increasing resolution. The set of wavelet coefficients in the second term provides a complete representation of the rapidly oscillating content localized at different scales, orientations, and spatial locations. In the cases of clinical Magnetic Resonance Imaging (MRI), the set of approximation coefficients represents the smoothly varying regions that belong to the anatomical objects. The rapidly oscillating components related to the boundaries of anatomical objects as well as image noise and other artifacts contribute to the wavelet coefficients. The energy in the wavelet representation of the signal is concentrated in a small number of wavelet coefficients where the signal singularities occur within the support of the wavelet functions. An approximation of the “noise-free” image can be re-constructed with a selected subset of attenuated coefficients. A variety of selection/shrinkage rules has been developed to produce a good approximation in the situations where Gaussian noise arises [9, 14, 17, 23]. These solutions become inadequate when the Gaussian assumption does not hold. The purpose of this paper is to explore a solution to the problem of adaptive enhancement in a situation where noise is better described by a Rician model.

2 Image Noise in MR Images

Rician noise arises from the process of constructing magnitude MR images [16, 22], which are usually used in the clinical setting as well as Cardiac Magnetic Resonance (CMR) image analysis. A complex-valued spatial-domain representation is generated from the frequency-domain data read out from an MRI machine through an inverse Fourier transform. The magnitude image is the square root of the sum of squares of the complex pair. The real and imaginary data after the inverse transform remain Gaussian distributed with zero means and equal variance σ_η^2 . The non-linear operator in the last step of image construction, however, transforms the distribution from Gaussian to Rician.

The magnitude data¹ M are characterized by the probability density function given (noise-free) signal S and variance σ_η of noise in the complex images [15, 21]

$$f_{\mathcal{M}}(M|S, \sigma_\eta) = \frac{M}{\sigma_\eta^2} e^{-\frac{M^2+S^2}{2\sigma_\eta^2}} I_0\left(\frac{MS}{\sigma_\eta^2}\right) \mathbf{1}_{\{M \geq 0\}} \quad (3)$$

where I_0 is the modified Bessel function of the first kind of zeroth order, and $\mathbf{1}_{\mathcal{A}}$ the indicator function over the set \mathcal{A} . The distribution depends on the $\frac{S}{\sigma_\eta}$, that measures the signal-to-noise ratio (SNR), and the signal intensity [11, 21]. When $S = 0$, Eq. (3) is reduced [6, 10] to a Rayleigh distribution: $f_{M|S, \sigma_\eta} = \frac{M}{\sigma_\eta^2} e^{-\frac{M^2}{2\sigma_\eta^2}} \mathbf{1}_{\{M > 0\}}$ with

$$E(M) = \sqrt{\frac{\pi}{2}} \sigma_\eta \quad (4)$$

$$\text{var}(M) = \left(2 - \frac{\pi}{2}\right) \sigma_\eta^2 \quad (5)$$

It can be shown using a truncated expansion of I_0 and some algebraic manipulation that as the signal approaches infinity, Eq. (3) tends to a Gaussian distribution with a mean equal to S and variance σ_η^2 .

The data of a squared magnitude image [7, 16] follow a non-central $\chi_2^2\left(\frac{S^2}{\sigma_\eta^2}\right)$ distribution with two degrees of freedom and non-centrality parameter $\frac{S^2}{\sigma_\eta^2}$. The mean and variance are given by:

$$E(M^2) = S^2 + 2\sigma_\eta^2 \quad (6)$$

$$\text{var}(M^2) = 4\sigma_\eta^4 + 4(S\sigma_\eta)^2 \quad (7)$$

The variance is approximately $4\sigma_\eta^4$ over the low intensity range and then rises rapidly with the signal intensity.

2.1 Related Approaches to Rician Noise in Clinical MRI

In the general context of clinical MRI, a number of filtering frameworks have been proposed for noise reduction in the Rician situations. A wavelet-domain filtering approach is proposed to shrink the wavelet coefficients toward zero with a set of filter weights that approximately minimize the mean squared error (MSE) of the coefficients [7, 16]. These filters rely on the assumptions that (1) the noise in the high SNR images is signal independent and (2) the wavelet coefficients are unbiased. Obviously, these conditions fail to hold for the low-intensity regions. For the low SNR (below 10dB) MRI, the framework of noise filtering is applied to squared magnitude images with two additional steps: (1) provide a MSE estimator of σ_η^2 using the discrete wavelet squared transform of the squared image and (2) remove the bias from the approximation coefficients of the

¹This section uses M rather than I for image data to avoid confusion.

image [7, 16]. This second algorithm, however, fails to outperform the first one [16] for any MRI with the SNR > 15 dB, a condition that is common in clinical MRI application. The technique of removing bias from squared images was adapted in some later algorithms, such as [27]. Over the low intensity regions, however, this filtering framework fails to remove the Rician noise satisfactorily. Alternative procedures are proposed to identify the background noise for removal [19, 26]. Other techniques are also proposed including wavelet packets [25], Bayesian frameworks with joint signal detection and estimation [18, 19] and spatial domain filtering [22].

3 Adaptive Enhancement of MRI

Rician noise in the magnitude MR images can be modeled in terms of two noise classes which occur over different ranges of signal intensity. Noise in the high SNR regions can be approximated by the asymptotic properties of an infinity-intensity signal – hereafter referred to as the approximation model – i.e., it is approximately signal independent noise with mean zero and variance σ_η^2 . Variants of the standard techniques being informed by the stochastic properties of Gaussian noise are adequate for signal recovery for this class of noise. Over the low SNR regions, the approximation model fails to provide a reliable approximation of the data. Noise remains active over these regions. We observe that the variance of χ^2 noise as well as that of the low-intensity structures over the low SNR regions in a squared magnitude image is relatively small. In the wavelet domain, these features will be mostly captured in the coarse scale approximation. This allows an adaptive enhancement scheme of coefficient attenuation to achieve the combined effect of noise reduction, removal of the background structures and enhancement of image contrast in two steps: namely the **smoothing (S)** step and the **enhancement (E)** step.

The **smoothing (S)** step seeks to reduce the noise with the asymptotic properties of Rician noise. Given the noise statistics of a typical clinical MR image, we observe that the image noise of a magnitude image over the high signal intensity regions, which cover most of the target organ, can be described by the asymptotic properties reasonably well. This approximation model justifies the application of many existing denoising techniques for the purpose of this step. In a nutshell, an approximation of the signal can be generated by a set of coefficients attenuated according to some shrinkage rule to remove the components due to noise [9, 14, 17, 23]. Obviously, the shrinkage rule is critical. It is the set of performance requirements of the application that dictates the choice of the wavelet shrinkage rule. These rules are further discussed in the next section. In spite of approximation, the noise remains Rician and its

stochastic properties cannot be ignored. Indeed, it plays a critical role in this step. Most of the shrinkage rules rely on the estimation of σ_η from the data in their threshold selection. In a typical clinical MR image, Fig. 1 being an example, there are signal-free regions of air outside the body. Given this subset of data, Eq. (5) gives a Rician-based estimator as shown in Eq. (8):

$$\sigma_\eta^2 = \frac{2}{4 - \pi} \text{var}(I_r), \quad (8)$$

where I_r are the data from the signal-free regions. Alternatively, from Eq. (6) we get the following estimator:

$$\sigma_\eta^2 = 0.5E(I_r^s), \quad (9)$$

where I_r^s are the data from the signal-free regions of the squared magnitude image. Both estimators provide highly reliable estimates of σ_η .

The **enhancement (E)** step attenuates the approximation coefficients of the squared magnitude image constructed from the output of the smoothing (S) step. Medical images are usually characterized by a number of piece-wise smoothly varying regions separated by sharp discontinuities. In multiresolution analysis, the coarse approximation $\mathcal{P}_v I$ represents the smooth part of the image whereas the discontinuities are captured by the wavelet coefficients. Rapid shifts in intensity across the regions contribute to image contrast which is essential for feature detection and structure finding. Contrast enhancement can therefore be achieved by amplifying these intensity gradients. One way to accomplish this is to reduce the magnitude of the smooth parts while preserving the energy in the oscillating content captured in the detail spaces. To suppress the background structures as well as the remaining noise in the low-signal intensity proximity of the heart, the enhancement scheme is necessarily adaptive to signal intensity. The adaptive factor which controls the amount of magnitude to be reduced should be a continuous function in signal intensity in order to avoid creating spurious gradients.

Given a noisy image I , the squared magnitude image can be expressed compactly as the following:

$$I^s = \sum_k s_k^{I^s} \phi_{j,k} + \sum_\eta w_\eta^{I^s} \psi_\eta = \mathcal{P}_v I^s + \mathcal{R}_w I^s \quad (10)$$

where s^{I^s} and w^{I^s} are the approximation coefficients and the wavelet coefficients of the squared image respectively. Consider the following operator on image I with threshold parameter \mathcal{T} :

$$\Phi I = \begin{cases} \sqrt{\frac{I - \mathcal{T}}{I + \mathcal{T}}} (I + \mathcal{T}) & \text{if } \mathcal{T} > 0 \\ I & \text{if } \mathcal{T} = 0. \end{cases} \quad (11)$$

The square root term is the adaptive factor. The operator returns the original image, $\Phi I = I$, given threshold $\mathcal{T} = 0$.

Given a positive threshold, i.e., $\mathcal{T} > 0$, the operator reduces the signal intensity of I . The numerator of the adaptive factor is reduced to zero when $I = \mathcal{T}$ and to a negative value over the regions where $I < \mathcal{T}$. Due to the non-linearity of the square root operator, the adaptive factor rises rapidly for signal intensities that exceed the threshold and tends to 1 as $I \gg \mathcal{T}$.

Observe that

$$(\Phi I)^2 = (I)^2 - \mathcal{T}^2 \quad \text{for } \mathcal{T} \geq 0. \quad (12)$$

where $(\Phi I)^2$ is the squared image of ΦI . Raising the data to a higher power amplifies image oscillations; consequently, image features in a squared image are projected onto a subspace at a scale at least as high as the corresponding scale in the decomposition of the magnitude image. The wavelet representation of the squared magnitude image of ΦI preserves the image features in the detail spaces in the sense that those features in $\mathcal{R}_w I$ stay in the detail spaces in the decomposition of the squared image. To preserve and enhance the ‘‘micro-features’’ of the image from the smoothing (S) step, we apply the signal adaptive operator on only the coarse scale approximation of the image without altering the detail representation of the squared image. The enhanced image is therefore

$$I_\Phi^s = \mathcal{P}_v I_\Phi^2 + \mathcal{R}_w I^s \quad (13)$$

The first term $\mathcal{P}_v I_\Phi^2$ is the projection of the square magnitude image of ΦI onto the approximation space. Combining Eq. (10), Eq. (12) and the fact [2] that $\sum \phi_k = 2^J$ for a J -scale decomposition yields the approximation coefficients as follows:

$$s_k^{\Phi I^s} = \left\langle (\Phi I)^2, \phi_k \right\rangle \quad (14)$$

$$= \sum_k I^2 \phi_k - \sum_k \mathcal{T}^2 \phi_k. \quad (15)$$

$$= s_k^{I^s} - 2^J \mathcal{T}^2 \quad (16)$$

where $s_k^{I^s}$ are the approximation coefficients of the squared image. More conveniently, the attenuated coefficients can be written as

$$\widetilde{s}^{I^s} = s^{I^s} - \tau \quad (17)$$

where

$$\tau = 2^J \mathcal{T}^2, \quad (18)$$

Let I_{se}^s be the enhanced squared image computed by the inverse wavelet transform with the attenuated coefficients $\{\widetilde{s}^{I^s}\}$. The non-positive part of the image represents the image features being suppressed.

The enhanced image depends on the threshold used in the adaptive enhancement rule. As mentioned, the scheme

adaptively erodes the signal intensity of the coarse scale approximation over the low-intensity regions. With a threshold set to a level which discriminates these structures and the boundaries of the heart, the adaptive scheme suppresses most of these structures in the proximity of the heart. This critical range is a characteristic of MRI scanning. It does not change significantly over a temporal sequence.

The remaining noise after the smoothing step contributes to some small fluctuations in the approximation and wavelet coefficients of the squared magnitude image. A study [26] of the distribution of the scaling function coefficients of squared magnitude images suggests a threshold be set to $(2^{J+1} + 6)\sigma_\eta^2$ to discriminate between coefficients with the support of the scaling functions in the signal regions and those in the Rayleigh (low-intensity) regions. To suppress the spurious fluctuations in the coefficients over the low signal regions, we set a threshold, which should be proportional to the variance of the noise, at or above this level. Consider a threshold set to $\zeta\sigma_\eta$. Where $\zeta = \sqrt{2}$, $\tau_n = 2^{J+1}\sigma_\eta^2$ is equivalent to the threshold set to remove the bias from the approximation coefficients [7, 16]. For most practical choices of J , setting ζ to two is adequate to discriminate between the approximation coefficients with a support inside the signal-free and those with a support over the signal regions of the image. A higher ζ raises the thresholding level to offset larger fluctuations in the coefficients. Our experimental findings indicate that the background noise can be more thoroughly attenuated with a ζ set to three. In general, the result is not sensitive to minor variations in the range between two and five.

4 Experimental Results

We use a four-way comparison to test the performance of the enhancement scheme with both synthetic and real MR images. Four algorithms are implemented:

Algorithm	Smoothing Step	Enhancement Step
SE Group		
SE-U	Universal threshold	Yes
SE-MSE	MSE filter	Yes
S Group		
S-U	Universal threshold	No
S-MSE	MSE filter	No

These algorithms fall into two groups: SE-group and S-group. The two algorithms, namely, SE-U and SE-MSE, in the SE-groups implement both steps of adaptive enhancement method. This is the main contribution of this paper. In contrast, the S-U and the S-MSE algorithms in the S group implement only the smoothing step but without the

enhancement step. Comparing the performance of these two groups of algorithms gives us a clear idea about the contribution of the adaptive scheme in the enhancement (E) step. Within each group, the two algorithms differ in the technique used in the smoothing (S) step. The U algorithms – S-U and SE-U – apply the classical soft-thresholding rule with the universal threshold selection scheme [3, 5, 9, 13]: For a wavelet coefficient w , the thresholding estimator \tilde{w} is defined as

$$\tilde{w} = \text{sgn}(w)(|w| - \lambda_u)\mathbf{1}_{\{|w| > \lambda_u\}}, \quad (19)$$

where $\text{sgn}(\cdot)$ denotes the signum function, with the universal threshold

$$\lambda_u = \tilde{\sigma}_\eta \sqrt{2 \ln N} \quad (20)$$

where N denotes the data size. The MSE algorithms – S-MSE and SE-MSE – apply the MSE-filtering rule [7, 16] given by $\tilde{w} = \alpha_* w$ with the filter weight $\alpha_* = \alpha \mathbf{1}_{\{\alpha \geq 0\}}$ where

$$\alpha = \frac{w^2 - \beta \tilde{\sigma}_\eta^2}{w^2} \quad (21)$$

and $\tilde{\sigma}_\eta$ is the estimate of σ_η . In our experiments, β is set to two as suggested by [16].

The major differences between the U-algorithms and the MSE-algorithms lie not only in the shrinkage rule, but also in their estimator of σ_η . The MSE algorithms estimate the noise variance σ_η^2 using the Rician-based estimator specified in Eq. (8) or Eq. (9) from the data in signal-free regions while the U-algorithms employ a Gaussian-based estimator [4, 14, 17]

$$\tilde{\sigma}_\eta = \frac{MAD(w)}{0.6745} \quad (22)$$

where $MAD(w)$ is the median of absolute deviation of $\{w\}$, the set of wavelet coefficients at the highest scale of decomposition. This latter estimator assumes that the coefficients are generated from a data set with Gaussian noise. The difference in performance between the U- and the MSE-group of algorithms demonstrates the performance impact of completely ignoring the Rician properties of the noise even in the case of high SNR images.

The performance of the algorithms is evaluated with respect to (1) contrast enhancement and (2) the fidelity (goodness) of approximation. The contrast improvement ratio (CIR) quantifies contrast enhancement achieved by an algorithm. The image contrast is measured in terms of the difference between the average intensity of the foreground object and that of the background over a selected region of interest. Given a pair of images: the input image I_{in} and the output image I_{out} , the contrast improvement ratio is defined as

$$CIR = \frac{C_{out} - C_{in}}{C_{out}} \times 100\%, \quad \text{where} \quad (23)$$

$$C_i = \frac{m_{i,F} - m_{i,B}}{m_{i,F} + m_{i,B}} \quad i \in \{\text{in}, \text{out}\} \quad (24)$$

and $m_{.,F}$ is the average intensity of the foreground object and $m_{.,B}$ the average intensity of the background. The goodness of approximation is measured in terms of two metrics: the error measurement with respect to the l_1 -norm [20]

$$\|\mathcal{I}_{in} - \mathcal{I}_{out}\|_{l_1} = \frac{1}{N} \sum |\mathcal{I}_{in} - \mathcal{I}_{out}|, \quad (25)$$

where N is the data size, and the SNR with respect to the l_2 -norm [20]

$$SNR = -10 \log(\xi_{l_2}) \quad (26)$$

where

$$\xi_{l_2} = \frac{\sum |\mathcal{I}_{in} - \mathcal{I}_{out}|^2}{\sum |\mathcal{I}_{in}|^2} \quad (27)$$

In order to isolate the error of approximation from the enhancement improvement, the output of the SE-group is adjusted for the intensity shifts that contribute to the enhancement measurement.

Each synthetic image used in the experiments is a magnitude image constructed from a pair of identical images corrupted by different realizations of Gaussian noise with zero mean and identical variance selected over a range compatible to a typical clinical image where the noise variance usually falls between three and ten. The noise-free image, a piece-wise step function in the form of concentric rings with different signal intensities, is designed to simulate the essential parts of a CMR image, i.e., the region inside the heart chamber, the heart wall, the thoracic cavity, the thoracic wall and the volume of air outside the human body. The image in Fig. 2a is an example of a noisy input with the ROI selected for measuring the contrast improvement achieved by the algorithms. Fig. 2b-2d display the performance of the algorithms with respect to the performance metrics; each bar represents the range of measurement obtained from fifty repeated trials. Significant differences are observed between the S-group and the SE-group of algorithms. Both SE algorithms provide significant improvement in contrast over those without the enhancement step. The negative CIR's for the S-group indicate that the coefficient attenuation techniques reduce image noise at a cost of image contrast. Similarly, there is an obvious watershed in terms of the absolute error between these groups of algorithms with SE algorithms producing much lower absolute errors. With respect to the l_2 -norm, the SE-MSE algorithm out-performs all the other algorithms whereas the output from the SE-U algorithm achieves the lowest SNR. Within each group, the MSE-algorithm achieves a better performance with respect to all performance metrics than the U-algorithm does. Unlike SE-MSE, SE-U does not incorporate the stochastic properties of Rician noise. First, the smoothing step of the algorithm uses the median of absolute deviation (MAD) estimator of σ_η for threshold selection, resulting in a higher discrepancy between the smoothed image

and the true signal. These errors in approximation are further amplified by the enhancement processes in the second step. Similar patterns hold for the input images corrupted by noise generated with different values of σ_{η}^2 , as shown in Fig. 3.

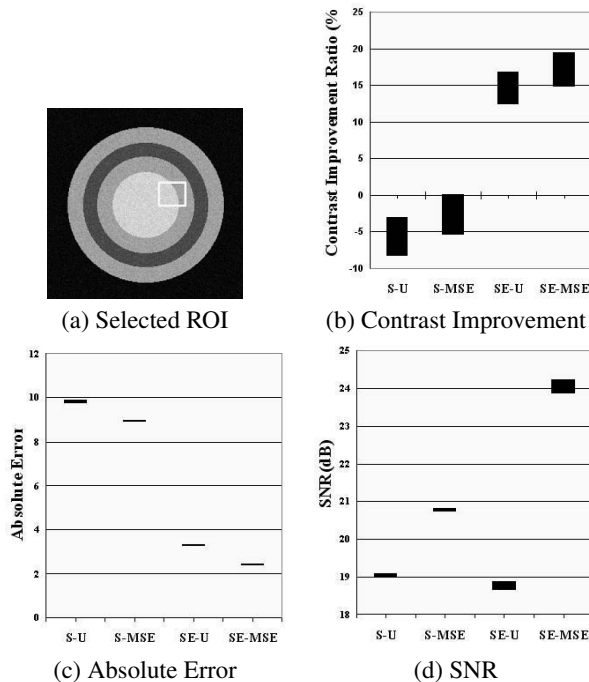


Figure 2. Performance of the algorithms on synthetic images. ($\sigma_{\eta} = 10$)

The performance of these algorithms is also evaluated on a variety of real medical images – cardiac, brain and vascular – with different wavelet bases. Similar results are found in all cases. Fig. 4 shows the performance statistics in contrast improvement achieved in a selected test case in (a) with the input image and the ROI selected for measurement in (b). The performance between the two groups of algorithms is qualitatively different. As in the case of synthetic images, only algorithms in the SE-group achieve significant enhancement. Fig. 5 shows the performance of the algorithms in reducing the confounding effect of image noise and low-intensity structures. Depicted in the column on the left are the input and output images (contrast adjusted) and the column on the right the maps of zero crossings superimposed on the input image. The purpose of these alternative representations of image data is to discern those image structures that are visible to an image analysis tool such as an edge detector. Three low-intensity regions therefore should draw attention: these are (1) the thoracic cav-

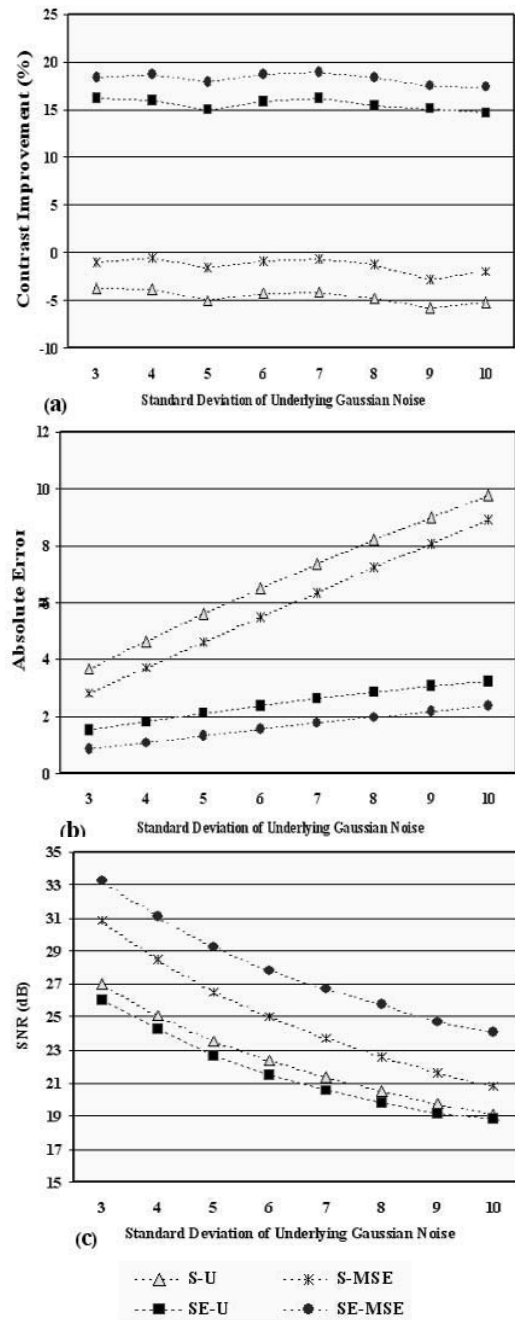


Figure 3. Performance of the algorithms on synthetic images with noise with different σ_{η} with respect to (a) contrast improvement, (b) absolute Error and (c) SNR.

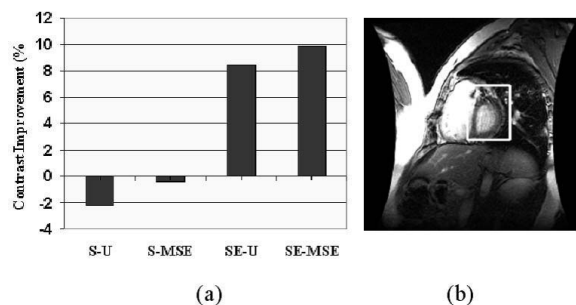


Figure 4. A selected CMR image: contrast improvement.

ity with the low-intensity structures behind the heart and (2) the signal-free areas above the cavity wall and (3) the low-intensity boundaries of the heart and its chambers. As shown in the second and the third row, S-algorithms using only the smoothing step are inadequate for handling the lower intensity regions. Shown in the bottom two rows is the output of the SE-algorithms which, by virtue of its enhancement step, (1) remove most of the image noise, including that in the low-intensity regions both in the background and inside the body, (2) preserve the edge information related to the boundaries of the heart and its anatomical parts (as shown in the map of zero-crossings) and (3) attenuate most of the background structures in the thoracic cavity. The remaining edges found in the cavity belong to the major arteries and veins.

5 Discussion and Conclusion

The method described in this paper is developed in the context of CMR image analysis. Indeed, it has a wider scope of application to other classes of medical images such as brain and Magnetic Resonance Angiography (MRA). In other sets of experiments, the performance of the SE-group is also found to be superior to spatial domain approaches such as anisotropic diffusion filtering; for details, see [12]. The four-way comparison used in our performance evaluation sheds light upon the contribution of different components of the method to the solution of the problem of adaptive enhancement. The adaptive enhancement scheme suppresses most of the noise and the background structures. Features related to the anatomical structures of the heart are preserved, and image contrast is significantly improved. This capability to preserve and enhance important image structures is demonstrated by the output of the SE-algorithm with both synthetic and MR images. The stochastic properties of the image noise – Rician noise in magnitude MR images and those of the corresponding χ^2 noise in

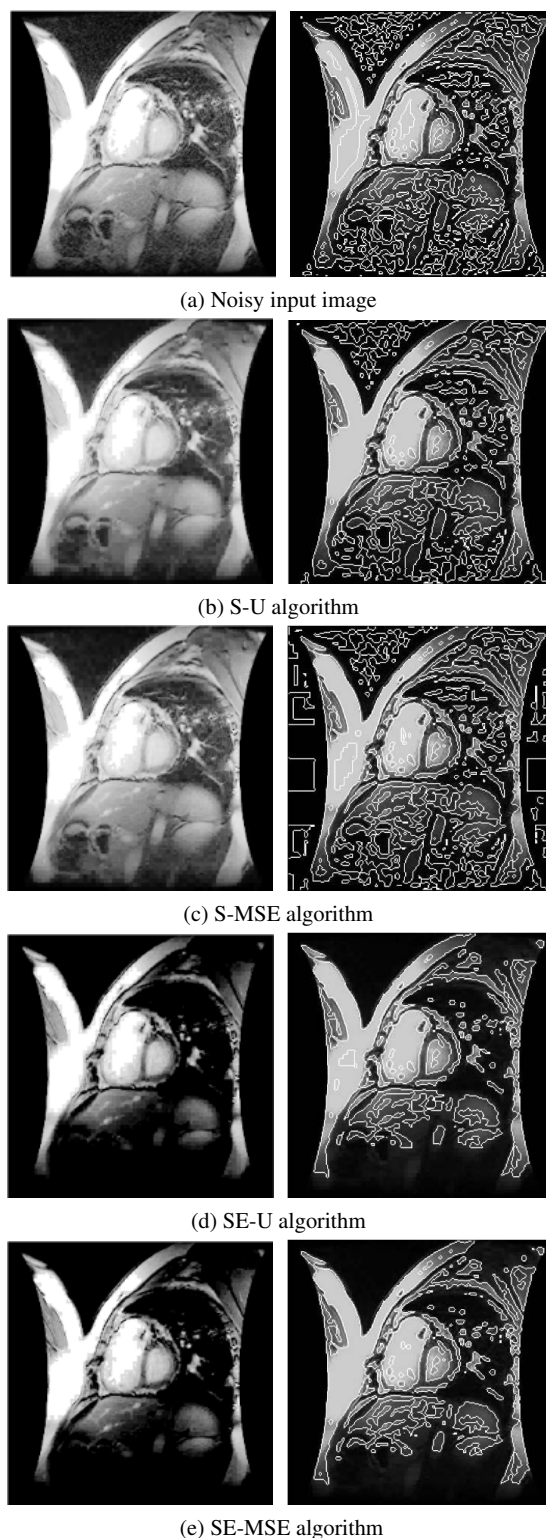


Figure 5. A selected cardiac MR image. Input/output images on the left and maps of zero crossings on the right.

squared magnitude images – play a critical role in both steps of the method. The U-algorithms, with both their shrinkage rule and the estimator of σ_η based on a strong assumption of Gaussian noise, are outperformed by their counterparts which incorporate the Rician properties of noise, especially in the case where the signal intensity is low. This paper integrates image enhancement by way of noise removal into a broader framework of adaptive enhancement. Although our experiments employ two very simple shrinkage techniques in the smoothing step, this framework can be adapted to the unique context of application, allowing different approaches to signal recovery. The method is linear in the image data. Since a typical cardiac MR image sequence consists of hundreds of images, this low complexity of the method allows fast processing of the sequence. With a complexity being compatible to that of Gaussian-based techniques, this method provides a more complete solution to the problem of adaptive enhancement of cardiac MRI and similar medical applications.

Acknowledgements

The authors thank Drs. Paul Babyn and Shi-Joon Yoo of the Hospital for Sick Children in Toronto, Canada, for images and assistance. E. T. H. Leung holds an Ontario Graduate Scholarship (OGS). J. K. Tsotsos holds the Canada Research Chair in Computational Vision.

References

- [1] P. S. Addison. *The illustrated wavelet transform handbook: introductory theory and applications in science, engineering, medicine and finance*. Institute of Physics Publishing, Bristol, 2002.
- [2] C. S. Burrus, R. A. Gopinath, and H. Guo. *Introduction to wavelets and wavelet transforms: a primer*. Prentice Hall, Upper Saddle River, N.J., 1998.
- [3] D. L. Donoho. De-noising by soft-thresholding. *IEEE Trans. Information Theory*, 41(3):613–627, 1995.
- [4] D. L. Donoho and I. M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455, 1994.
- [5] D. L. Donoho, I. M. Johnstone, G. Kerkycharian, and D. Picard. Wavelet shrinkage: Asymptopia? *J. R. Statist. Soc. B*, 57(2):301–337, 1995.
- [6] M. Evans, N. Hastings, and B. Peacock. *Statistical distributions*. J. Wiley, New York, 2nd. edition, 1993.
- [7] R. L. Gregg and R. D. Nowak. Noise removal methods for high resolution MRI. In *IEEE Nuclear Science Symposium*, volume 2, pages 1117 – 1121, 1997.
- [8] H. Gudbjartsson and S. Patz. The Rician distribution of noisy MRI data. *Magn Reson Med*, 34:910–914, 1995.
- [9] M. Jansen. *Noise reduction by wavelet thresholding*. Lecture notes in statistics volume 161. Springer, New York, 2001.
- [10] S. M. Kay. *Fundamentals of statistical signal processing*, volume II. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1998.
- [11] S. J. Kisner, T. M. Talavage, and J. Ulmer. Testing a model for MR imager noise. In *EMBS/BMES Conference*, pages 23–26, Houston, TX, 2002.
- [12] E. T. H. Leung. Adaptive enhancement of cardiac magnetic resonance (MR) images. Master’s thesis, Department of Computer Science and Engineering, York University, Toronto, Ont., 2005.
- [13] S. G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *PAMI*, 11:674–693, 1989.
- [14] S. G. Mallat. *A wavelet tour of signal processing*. Academic Press, San Diego, CA, 2nd edition, 1999.
- [15] J. Minkoff. *Signal processing fundamentals and applications for communications and sensing systems*. Artech House, Boston, MA, 2002.
- [16] R. D. Nowak. Wavelet-based Rician noise removal for magnetic resonance imaging. *IEEE Trans. Image Processing*, 8(10):1408 – 1419, Oct. 1999.
- [17] R. T. Ogden. *Essential wavelets for statistical applications and data analysis*. Birkhuser, Boston, MA, 1997.
- [18] A. Pizurica. *Image Denoising Using Wavelets and Spatial Context Modeling*. Ph.D thesis, Ghent University, Belgium, Jun. 2002.
- [19] A. Pizurica, W. Philips, I. Lemahieu, and M. Acheroy. A versatile wavelet domain noise filtration technique for medical imaging. *IEEE Trans. Med. Imag.*, 22(3):323–331, Mar. 2003.
- [20] W. K. Pratt. *Digital image processing : PIKS inside*. Wiley, New York, 3rd edition, 2001.
- [21] J. Sijbers. *Signal and noise estimation from Magnetic Resonance Images*. PhD thesis, University of Antwerp, Antwerp, Belgium, 1998.
- [22] J. Sijbers, A. J. den Dekker, E. Raman, and D. V. Dyck. Parameter estimation from magnitude MR images. *Int J Imaging Syst Technol*, 10(2):109–114, 1999.
- [23] B. Vidakovic. *Statistical modeling by wavelets*. J. Wiley, New York, 1999.
- [24] D. F. Walnut. *An introduction to wavelet analysis*. Birkhäuser, Boston, MA, 2002.
- [25] J. C. Wood and K. M. Johnson. Wavelet packet denoising of magnetic resonance images: importance of Rician noise at low SNR. *Magn Reson Med*, 41(3):631–5, Mar. 1999.
- [26] Z. Q. Wu, J. A. Ware, and J. Jiang. Wavelet-based Rayleigh background removal in MRI. *Electronics Letters*, 39:603 – 604, Apr. 2003.
- [27] J. Xue, A. Pizurica, W. Philips, E. Kerre, R. V. D. Walle, and I. Lemahieu. An integrated method of adaptive enhancement for unsupervised segmentation of MRI brain images. *Pattern Recognition Letters*, 24(15):2549–2560, Nov. 2003.