Interaction between noise suppression and inhomogeneity correction in MRI

Albert Montillo[*a], Jayaram Udupa[a], Leon Axel[b], Dimitris Metaxas[c]

[a]Univ. of Pennsylvania, Philadelphia, PA; [b]New York University, NY, NY;
[c]Rutgers University, New Brunswick, NJ

ABSTRACT

While cardiovascular disease is the leading cause of death in most developed countries, SPAMM-MRI can reduce morbidity by facilitating patient diagnosis. An image analysis method with a high degree of automation is essential for clinical adoption of SPAMM-MRI. The degree of this automation is dependent on the amount of thermal noise and surface coil-induced intensity inhomogeneity that can be removed from the images.

An ideal noise suppression algorithm removes thermal noise yet retains or enhances the strength of the edges of salient structures. In this paper, we quantitatively compare and rank several noise suppression algorithms in images from both normal and diseased subjects using measures of the residual noise and edge strength and the statistical significance levels and confidence intervals of these measures.

We also investigate the interrelationship between inhomogeneity correction and noise suppression algorithms and compare the effect of the ordering of these algorithms. The variance of thermal noise does not tend to change with position; however, inhomogeneity correction increases noise variance in deep thoracic regions. We quantify the degree to which an inhomogeneity estimate can improve noise suppression and whether noise suppression can facilitate the identification of homogeneous tissue regions, and thereby, assist in inhomogeneity correction.

Keywords: MRI, thermal noise suppression, surface coil, inhomogeneity correction, SPAMM

1. INTRODUCTION

Cardiovascular disease is the leading cause of death for both men and women in most developed countries. Well over 2,500 Americans die of the disease every day: more than the next seven causes of death combined [7]. However, SPAMM-MRI can provide data that can reduce morbidity by improving patient diagnosis, treatment and intervention assessment. Manual extraction of heart motion from the SPAMM-MRI is the bottleneck of this methodology, requiring roughly 5 hours per subject. A myocardial deformation analysis method with a higher degree of automation is essential for clinical adoption of SPAMM-MRI [8]. The degree of this automation is dependent on the amount of thermal noise and surface coil-induced intensity inhomogeneity that can be removed from the images.

MRI is typically corrupted by thermal noise from the sample and receiver coils. SPAMM MRI is acquired with surface coils which can cause the same tissue to have different intensities depending on location. Previous research has focused on either the extraction of field inhomogeneity [1,6,9,10] or thermal noise [2,5] in isolation. The purpose of this paper is to determine optimal techniques for the removal of both surface coil intensity inhomogeneity and thermal noise in 4D cardiac MRI and to evaluate the nature of the interaction between these two techniques. This research is similar in spirit to that of Madabhushi [11] who has studied a different relationship, namely, the interaction between field inhomogeneity correction and MR intensity standardization.

2. METHODOLOGY

We denote an acquired MR volume image by a pair $V = (\mathcal{V}, \mathbf{g})$ where $\mathcal{V}$ is a finite rectangular array of voxels and $\mathbf{g}$ is a function that assigns to every voxel c in $\mathcal{V}$ an intensity value $\mathbf{g}(c)$ from a set of integers. We model the combined effect of the surface coil and thermal resistance as $\mathbf{g}(x, y, z) = \mathbf{f}(x, y, z)\beta(x, y, z) + \mathbf{n}_{therm}(x, y, z)$ where $\mathbf{g}$ is the

---

[*] montillo@seas.upenn.edu; phone 215 898 1976; http://www.cis.upenn.edu; 200 S. 33[rd] St, Moore Bldg, Rm 556 Philadelphia, PA 19104-6389

observed image, $\mathbf{f}$ is the true image, and β is the multiplicative gain factor from the surface coil. We model the thermal noise, $\mathbf{n}_{therm}$, as an additive, signal independent, white, zero-mean bi-variate stationary Gaussian random field ($\mathbf{n}_{therm} \sim N(0, \sigma_n^2)$) which is uncorrelated from point to point. In order to reconstruct $\mathbf{f}$ from $\mathbf{g}$, it is necessary to develop an estimate of the surface coil inhomogeneity $\hat{\beta}$, and the thermal noise $\hat{\mathbf{n}}_{therm}$. If this could be done perfectly then one could "subtract" $\hat{\mathbf{n}}_{therm}$ and then divide the result by $\hat{\beta}$ to reconstruct $\mathbf{f}$. The estimation of these terms is an ill-posed problem, and so a perfect solution is infeasible.

We have shown previously [12] that the scale-based method [1] can be used to derive $\hat{\beta}$ from $\mathbf{g}$ and correct for much of the intensity inhomogeneity in SPAMM-MRI. This method estimates $\mathbf{f}$ by dividing $\mathbf{g}$ by $\hat{\beta}: \hat{\mathbf{f}}(x,y,z) = \mathbf{f}(x,y,z)\beta(x,y,z)/\hat{\beta}(x,y,z) + \mathbf{n}_{therm}(x,y,z)/\hat{\beta}(x,y,z)$. While this corrects much of the intensity inhomogeneity, it complicates the noise suppression task by making the second term non-stationary. In our previous work [12], we have shown that the adaptive wiener filter [3] can be used to remove some of the thermal noise. In this paper we employ systematic experimentation to determine how to best suppress both thermal noise and background intensity variation. We conduct three experiments. In the first experiment, we compare several more recently developed noise suppression methods by measuring how well each method reduces noise while retaining the strength of the edges of interest. In addition, we combine these with the wiener filter method to explore the effects of their integration. In the second experiment, we investigate whether performing noise suppression first can enable an improved estimate of $\hat{\beta}$. The hypothesis is that noise suppression can facilitate the identification of homogeneous tissue regions and thereby assist in inhomogeneity correction, and the goal of this experiment is to determine whether this is true. In our third experiment, we perform inhomogeneity correction first and then suppress noise. We investigate whether the noise suppression results are improved by using $\hat{\beta}$ to scale the estimated variance of the noise $\sigma_n^2(x,y,z)$ over the image.

2.1 First experiment: comparison of nonlinear noise suppression methods
For the first experiment, we compare the following five noise suppression algorithms:
  (1) Scale-based anisotropic diffusion [5], hereafter denoted SBD filter.
  (2) Anisotropic diffusion using Tukey's biweight influence function [4], hereafter denoted TAD filter.
  (3) Adaptive Wiener filtration [3], hereafter denoted AW filter.
  (4) AW filter followed by TAD.
  (5) TAD followed by AW filter.
We denote the set of all MRI protocols by P and the set of all body regions by B. We let V_{PB} be the set of all images that can be generated from body region B using protocols P.

Let S` denote a given set of volume images. In our case, this consists of 10 images from both normal and diseased subjects. Before applying a noise suppression method to a given input image in S`, we apply a sequence of operations to prepare the volume. First we apply intensity correction, to suppress most of the background intensity variation. Next we apply a method [13] called intensity standardization to reduce the intersubject variation in tissue intensities. For the SPAMM images in this paper, we fill the tag lines introduced in SPAMM with a grayscale morphological closing operation with a linear structuring element. Let the resulting set of volume images be denoted by S_1. To each image in S_1 we apply each noise suppression algorithm.

Each noise suppression algorithm has its own set of parameters. To initialize them we have, wherever possible, used the recommended settings from the authors of each noise suppression method. For a parameter with a range of recommended values, we have selected the parameter that makes the filter results have the same residual noise. We begin by parameterizing the AW filter. The window size is set to the size which yielded the best results in our previous experiments [12] for the task of segmentation. Next we parameterize the TAD filter. As described in [4], for a given volume, I, the parameter σ is set according to $\sigma = \sqrt{5}\left(1.4826\, median_I\left[\left\|\nabla I - median_I\left(\|\nabla I\|\right)\right\|\right]\right)$. The number of iterations is set to 100 to achieve the same approximate level of smoothing as the AW filter. Next we parameterize the SBD filter. As described in [5], for a given volume, I, we compute ∇I and find the mean, μ, and standard deviation, σ, of the lower 95% of the values in the gradient image. Then we set the SBD filter's homogeneity parameter to $\mu + 3\sigma$ as recommended in [5]. The number of iterations is set to 100 to achieve the same residual noise as the AW filter and TAD filters. For the combination filters, we let the window size be 5x5 to achieve slightly less than half the

noise suppression compared to the AW filter used in isolation and follow a similar process to initialize the remaining parameters shown in Figure 1.

SBD filter	TAD filter	AW then TAD filter	TAD then AW filter	AW filter
homogeneity see text	σ see text	σ see text	σ see text	window size 8x8
number of iterations 100	number of iterations 100	number of iterations 50	number of iterations 50	
		window size 5x5	window size 5x5	

Figure 1: Noise suppression parameters

We characterize the performance of the noise suppression algorithms using two measurements. For each volume, V_j, an expert anatomist has delineated the myocardium (Figure 2, solid line). We compute the normalized standard deviation (NSD) of the myocardium region just inside the expert's boundary (dashed line). To measure the residual edge strength after noise suppression, we compute the magnitude of intensity gradient perpendicular to the expert drawn contour at evenly spaced intervals along the expert's contour as shown in the figure. We search along the perpendicular direction to the contour along line segments roughly 4 pixels on either side of the experts contour and compute the maximum of the magnitude of the intensity gradient along these line segments. Finally we record the edge strength for this image as the trimmed mean of these maxima (excluding the highest 5% and lowest 5 %) on the line segments for a given contour. We compute these measures for all 5 filters for the image data in the set S_1.
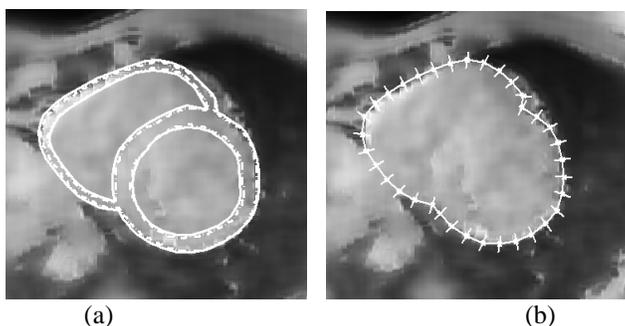

Figure 2: (a) expert's myocardium boundary (b) edge strength is measured along perpendicular profiles to expert's boundary

2.2 Second experiment: the effect of the order of noise suppression and inhomogeneity correction

2.2.1 Real noise and field inhomogeneity

In this experiment, we investigate the effect that the ordering of noise suppression and inhomogeneity correction operations has on the resulting residual intensity inhomogeneity, residual noise and edge strength. Before applying the two sequences of operations to each input image in S', we fill the tag lines as described earlier, and we denote the resulting set by S_2. To this set we apply each operation sequence. The first sequence is noise suppression followed by inhomogeneity correction, resulting in the set of images S_{2fc}. In the second sequence, we reverse the ordering yielding the set of image S_{2cf}.

2.2.2 Synthetic noise and field inhomogeneity

In this experiment, we investigate whether the same results found using real noise and inhomogeneity apply when we vary both the power of the noise and the strength of the field inhomogeneity. Before adding controlled amounts of synthetic noise and inhomogeneity, we apply a sequence of operations to prepare the images. We apply inhomogeneity correction, standardization and fill the tags with a closing and then smooth the images with a TAD filter. We denote the resulting set images by S_3. Next we modify these volumes by applying low, medium, and high strength inhomogeneity. Then we apply low, medium, and high noise power. In all we apply 9 different combinations of field strength and synthetic noise. We denote the resulting images by S_{3xy} where x denotes the strength of the applied inhomogeneity and y denotes the power of the uniform noise added. Having prepared the volumes, we then apply the above two sequences to

each image in S_{3xy} and compare the myocardium NSD and epicardium edge strength for the two resulting sets S_{3xyfc} and S_{3xycf}.

To simulate field inhomogeneity, we use the Biot-Savart law which expresses mathematically the relation between the current, I, flowing in the direction of the infinitesimal wire elements, $d\vec{L}$, to the magnetic field they generate. The current flowing in the wire along direction $\hat{L}$ over the distance $\left|d\vec{L}\right|$ creates a magnetic vector field, which at the point p, that is distance r away from dL along vector $\hat{r}$ is given by: $B(r) = \dfrac{\mu_0 I}{4\pi} \oint_{z \in wire} \dfrac{d\vec{L} \times \hat{\mathbf{r}}}{r^2} dz = \dfrac{\mu_0 I}{4\pi} \oint_{z \in wire} \dfrac{\left|d\vec{L}\right| \sin(\theta)\hat{n}}{r^2} dz$, where $\hat{r}$ is the unit vector in the radial direction from the point p in space to the wire element $d\vec{L}$. We construct two ellipses representing two surface coils and position them in the coordinate system of each input volume image, V, and solve the Biot-Savart equation to compute the magnetic field at each voxel (x,y,z) from every segment of the two tessellated ellipses. We retain the magnitude of the component of $B(x,y,z)$ perpendicular to the main magnetic field B_0 (which we assume to be in the volume's z direction) and denote this component $B_{coils}(x,y,z)$. Through the reciprocity theorem, the sensitivity of the coil to a magnetic field at a point (x,y,z) is proportional to $B_{coils}(x,y,z)$. We then inject this field inhomogeneity onto the volume by element wise multiplication: For an image $V_3 = (\mathcal{V}, \mathbf{g})$ in S_3, the resulting image $V_{3x} = (\mathcal{V}, \mathbf{g}_{3x})$ is given by, for any $(x,y,z) \in \mathcal{V}$, $g_{3x}(x,y,z) = g_3(x,y,z)\gamma B_{coils}(x,y,z)$, where γ is the scalar field strength gain. To compute low, medium and high gains, we use 0.75, 1.00, 1.25 times a nominal gain γ_* that reconstructs the original inhomogeneity found in the uncorrected volumes. Figure 3 (left) shows the field generated for an 11 slice volume (images from left to right, top to bottom correspond to the generated field for the slices from base to apex. The application of this field for one slice is shown on the right. The set resulting from this operation on S_3 is denoted S_{3x}.
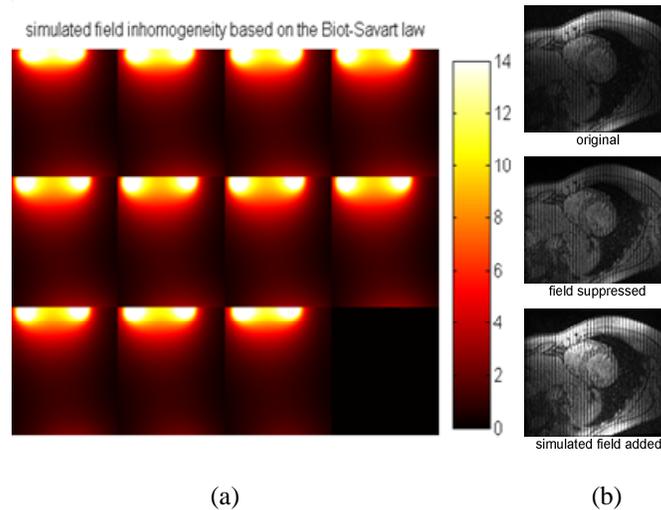

Figure 3: (a) The simulated synthetic field inhomogeneity showing all slices. (b) the effect of removing the inhomogeneity from the scanned images and adding a large strength simulated field.

To simulate noise, we compute the mean μ and standard deviation σ of the NSD in the myocardium in an image $V_{3x} \in S_{3x}$. We let the low, medium and high noise powers be $\sigma_n = \mu + \alpha\sigma$, where $\alpha \in \{-1, 0, 1\}$ for the different noise levels. We construct a signal independent, white, zero-mean, stationary Gaussian random field, $\mathbf{n}_{therm} \sim N\left(0, \sigma_n^2\right)$ and add it to each V_{3x} yielding the set S_{3xy} where y is the low, medium or high noise power used.

2.3 Third experiment: Evaluating the effect of using a field inhomogeneity estimate to regulate noise suppression

In this experiment, we investigate whether a field inhomogeneity estimate from the correction algorithm can be used to provide noise suppression with a more useful, spatially varying estimate of noise power. We prepare the images by filling the tag lines and then perform inhomogeneity correction on the images in S', yielding the input image set denoted S_4. Next we apply noise suppression methods which do and do not use the inhomogeneity estimate $\hat{\beta}$ which is saved during inhomogeneity correction. We have selected the TAD filter and the AW filters for testing. We denote the noise suppressed image sets as S_{4TADi} and S_{4AWi} when these noise suppression filters use the inhomogeneity estimate, and use the notation S_{4TAD} and S_{4AW} when we use the filters which do not use the inhomogeneity estimate. We then compute the myocardium residual noise and epicardium edge strength for 10 subjects in these volume sets.

To inject spatially changing noise variance into the TAD filter, we begin with the robust statistical formulation of anisotropic diffusion described by Black, *et al* in [4]. An estimate of the robust scale of the image, σ_e, is computed to determine how large the magnitude of the image gradient can be before it is considered an outlier, hence not part of noise, and a place where the diffusion should decrease or stop. In their formulation, the influence of outliers in the magnitude of the intensity gradient is captured by an influence function, $\psi(d,\sigma)$ and the anisotropic diffusion is expressed as $\dfrac{\partial I(x,y,z,t)}{\partial t} = div\left[\psi\left(d(x,y,z),\sigma\right)\dfrac{\nabla I}{\|\nabla I\|}\right]$, where I denotes a given image. The point where the influence of outliers in the image gradient first begins to decrease occurs when the derivative of the influence function is zero. For the Tukey's biweight influence function

$$\psi(d,\sigma) = \begin{cases} d\left[1-\left(\dfrac{d}{\sigma}\right)^2\right]^2 & |d| \leq \sigma \\ 0 & otherwise \end{cases},$$

this occurs when $\sigma = \sqrt{5}\sigma_e$, where σ_e is the estimate of the scale in the image. Black *et al*, proposed a single scalar estimate of the scale of the image: $\sigma_e = 1.4826\, median_I\left[\left\|\nabla I - median_I\left(\|\nabla I\|\right)\right\|\right]$, where the constant in this equation is derived from the median absolute deviation of a zero-mean normal distribution with unit variance which is 1/1.4826. We inject knowledge of spatially variant noise power into the estimate of σ_e. During correction, we have divided the image by $\beta(x,y,z)$. This affects both the noise and signal alike. In particular, $\sigma_{n\_new}^2(x,y,z) = \dfrac{\sigma_n^2(x,y,z)}{\beta^2(x,y,z)}$, thus where $\beta(x,y,z) > 1$ the noise power is decreased, while where $\beta(x,y,z) < 1$ the noise power is increased. Since σ_e roughly approximates the square root of the noise power in the image, σ_n, our new spatially variant estimate for σ_e is:

$\sigma_{e\_new}(x,y,z) = \dfrac{\sigma_e}{\beta(x,y,z)}$ Hence the original Tukey's biweight influence function becomes:

$$\psi(d(x,y,z),\sigma_{new}) = \begin{cases} d(x,y,z)\left[1-\left(\dfrac{d(x,y,z)}{\sigma_{new}(x,y,z)}\right)^2\right]^2 & |d(x,y,z)| \leq \sigma_{new}(x,y,z) \\ 0 & otherwise \end{cases},$$

where $\sigma_{new}(x,y,z) = \sqrt{5}\sigma_{e\_new}(x,y,z) = \sqrt{5}\dfrac{\sigma_e}{\beta(x,y,z)}$. The robust statistical formulation of anisotropic diffusion equation becomes:

$$\frac{\partial I(x,y,z,t)}{\partial t} = div\left[\psi\left(d(x,y,z),\sigma_{new}(x,y,z)\right)\frac{\nabla I}{\|\nabla I\|}\right].$$

To inject spatially changing noise variance into the AW filter, we begin with the original formulation which uses spatially invariant noise power: $g(x,y,z) = \mu + \dfrac{\sigma^2 - v^2}{\sigma^2}\left(I(x,y,z) - \mu\right)$, where μ and σ^2 are the local mean and

variance around each pixel: $\mu(x,y,z) = \dfrac{1}{MNP} \displaystyle\sum_{i,j,k \in R} I(i,j,k)$ and $\sigma^2(x,y,z) = \left(\dfrac{1}{MNP} \displaystyle\sum_{i,j,k \in R} I^2(i,j,k) \right) - \mu^2$, and where

R is the local M by N by P neighborhood around each pixel. The noise power, when spatially invariant, is estimated by the mean of the local estimated variances: $v^2 = \dfrac{1}{RST} \displaystyle\sum_{x,y,z \in R_{img}} \sigma^2(x,y,z)$, where R_{img} is the set of x,y,z coordinates of the

R by S by T image, I. The fraction of the local variance in the image which is not explained by the local noise is considered to be the fraction of the local deviation from the local mean intensity which can be trusted. It is this fraction, $\dfrac{\sigma^2 - v^2}{\sigma^2}$, which is added back to the average signal, μ. The noise power is made spatially variant by using the

estimated field inhomogeneity: $v_{new}^{\ 2}(x,y,z) = g \dfrac{v^2}{\beta^2(x,y,z)}$, where g is an overall gain constant which can be used to

scale the influence of the field inhomogeneity estimate.

3. RESULTS

3.1 Comparison of nonlinear noise suppression methods

3.1.1 Qualitative results:

We find that all noise suppression methods can be parameterized to yield low amounts of residual noise. In Figure 4 (a), we can see that the blotchiness in the filter input image is largely suppressed in all of the filtered images. This result is not surprising since we have parameterized the filters to yield about the same residual noise as the AW filter. However, we can see that the edge strength varies significantly depending on the filter used, particularly along the boundaries of the myocardium. The SBD filter appears to be preserving edges the best, followed by TAD. In (a), the SBD filter appears to provide better preservation of small structures, such as the trabeculae.

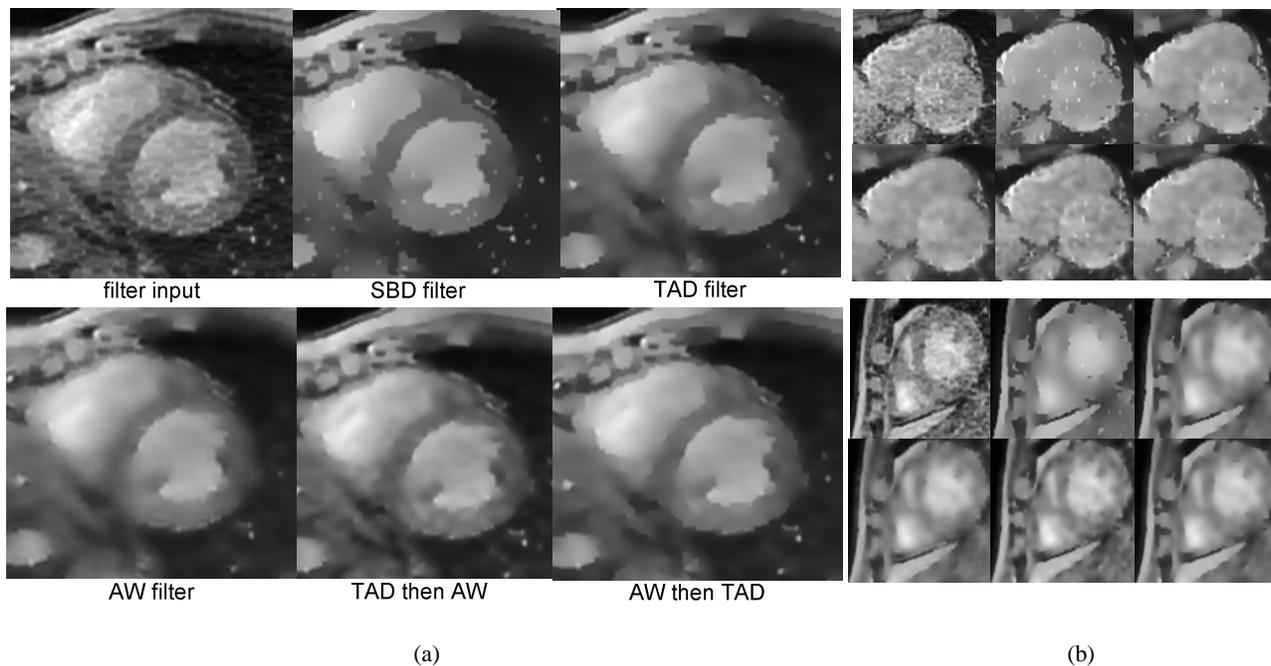(a)   (b)

Figure 4: (a) Filter results for slice near end systole and midway between base and apex of heart. (b) Additional filter results with the same layout as in (a).

All of the noise suppression methods, except the AW filter, occasionally allow some structure from the morphological closing operation to show through as white spots (Figure 4 (b,top)). In Figure 4 (b,bottom) we can see that the

moderator band, a strip of muscle dividing the right ventricle near the apex is not well preserved by any of these methods (See the results of experiment 3 for an improvement).

3.1.2 Quantitative results:

When we compare the myocardium NSD that results from these filters we find that we have parameterized them well so as to achieve about the same residual noise. In the first row in the table in Figure 5, the mean myocardium NSD is the first entry and the standard deviation in the NSD is the second entry. The second row shows the mean reduction in the myocardium NSD as a percentage of the original input myocardium NSD. The reduction of the myocardium NSD varies by only a few percentage points depending on the method. The second entry in the second row shows the 95% confidence interval for the reduction in NSD. The third row shows the mean (first entry) and standard deviation (second entry) of difference in the NSD: input NSD minus filtered NSD.

	Input	SBD	TAD	AW then TAD	TAD then AW	AW
residual noise	0.258	0.204	0.210	0.205	0.213	0.196
	0.076	0.072	0.074	0.073	0.074	0.070
% reduction in residual noise	NA	20.9%	18.7%	20.8%	17.5%	24.3%
		20.2%-21.5%	18.3%-19.1%	20.4%-21.2%	17.2%-17.8%	23.8%-24.8%
difference from input	NA	0.054	0.048	0.054	0.045	0.063
		0.031	0.020	0.019	0.015	0.022

Figure 5: Residual noise results for each noise suppression method.

In the graph in Figure 6 (left plot), each data point plotted represents the difference in the NSD in the noise suppression output for one image compared to the same image without noise suppression. That is for the x^{th} image, the value of the following expression is computed:

$y(x)$= (myocardium NSD from the same unfiltered x^{th} image - myocardium NSD in the x^{th} image after filter j)
where j = filter type. The plot indicates that the methods are reducing the myocardium NSD by about the same amount, since the data points for all methods overlap at around y=0.05.
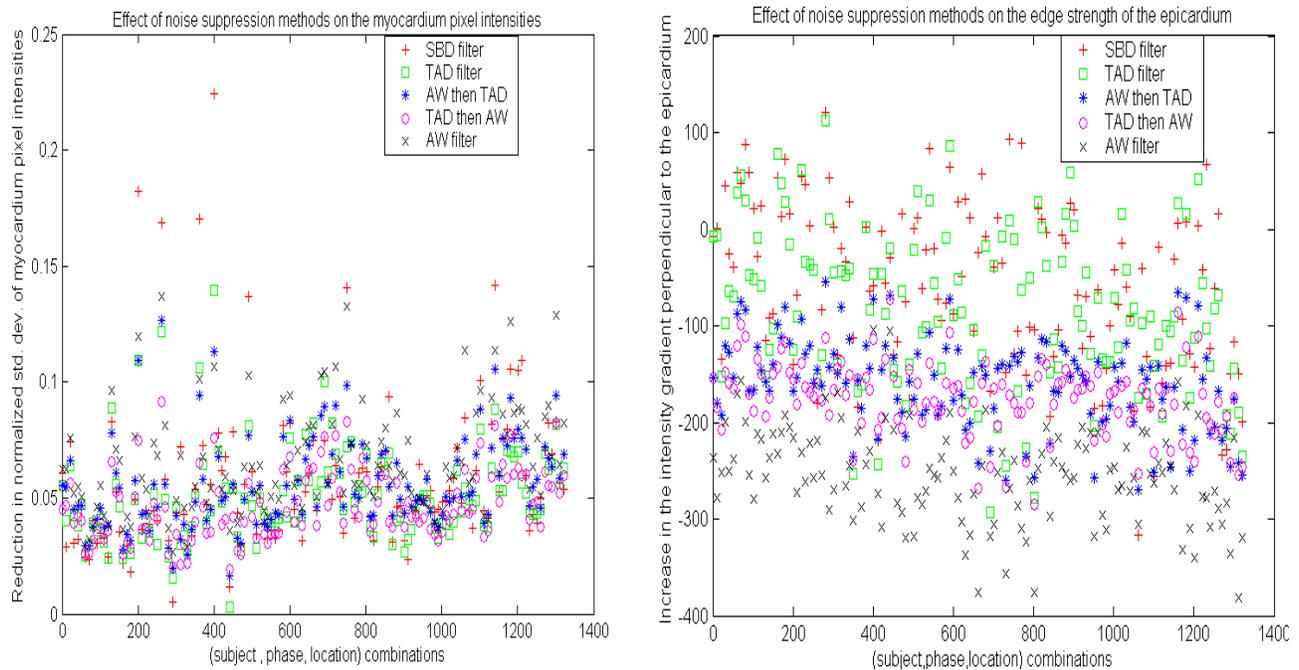

Figure 6: Left plot: The similarity in the reduction in myocardium NSD by the noise suppression methods extends across subjects, time and throughout the volume, as shown by the overlap of the data points. This plot is a subsampling of every 10th data point (for readability) from over 1300 images taken from 10 subjects recorded throughout systole. Right plot: The stratification in the

increase/preservation of myocardium edge strength by the noise suppression methods extends across subjects, time and throughout the volume, yielding the appearance of a distinct band for each method. This plot is also a subsampling of every 10th data point.

At the same level of noise reduction there are significant differences in the strength of the epicardium edge. The table in Figure 7 shows, in the first row, the mean edge strength (first entry) and edge strength standard deviation (second entry) for the input images and the filtered images. There are statistically significant differences between all pairs of methods for the mean edge strength (p<0.0001) using pair wise t-tests. The second row shows the mean increase in the edge strength as a percentage of the original input edge strength and the second entry in the row shows the 95% confidence interval. The third row shows the mean and standard deviation (second entry) of difference in the edge strength: input edge strength minus filtered edge strength.

	Input	Output				
		SBD	TAD	AW then TAD	TAD then AW	AW
edge strength	610	553	523	453	435	358
	148	199	180	151	146	134
% increase in edge strength	NA	-9.4%	-14.4%	-25.7%	-28.6%	-41.1%
		-10.1% - -8.6%	-15.1% - -13.6%	-26.2% - -25.3%	-28.9% - -28.2%	-41.6% - -40.7%
difference from input	NA	-57	-87	-157	-175	-252
		82	84	48	40	52

Figure 7: Edge strength results for each noise suppression method.

The methods are not equivalent in their ability to preserve the edges of interest. In the graph in Figure 6, (right plot) each data point plotted represents the difference in the edge strength in the noise suppression for one image compared to the same image without noise suppression. That is for the x^{th} image, the value of the following expression is computed:

   y(x)= (Edge strength in the x^{th} image after filter j   –   edge strength from the same unfiltered x^{th} image)

where j = filter type. This plot corroborates the ranking borne out in the table of mean edge strengths in Figure 7. In the plot, the data points from the methods form bands. From top to bottom (greatest increase/preservation in edge strength to least) we find the SBD filter, TAD filter, AW then TAD, TAD then AW and at the bottom, the AW filter which preserves edge strength the least.

          We have compared each pair of filters. When comparing the edge strength from the SBD filter to that from TAD filter, we find that the mean SBD filter edge strength is 30 gray levels stronger than TAD in a 12bit image. The standard deviation is 66.0. Using a paired t-test we find that the difference between the means is statistically significant: p<0.001 when we analyze >1300 images from 10 subjects. On the average, SBD filtering yields 5.0% stronger edges than TAD with a 95% confidence interval of (4.4%-5.6%).

3.2 Evaluating the effect of the ordering of operations

3.2.1 Real noise and field inhomogeneity

On images with real noise and field inhomogeneity from the MR scanner, we find significantly better correction is visible in the images when correction precedes noise suppression. In Figure 8 (a) is an image from the set S_2 before correction and noise suppression. In (b) we have applied TAD filtering followed by inhomogeneity correction yielding image from the set S_{2fc}. In (c) we have reversed the order of the last two operations yielding an image from the set S_{2cf}. There are two noticeable differences: 1) there is much greater inhomogeneity correction in (c) and 2) there is better edge strength preservation in the deep thoracic region in (c). Correction boosts the contrast between the weak edges in this region. These higher contrast edges are smoothed less upon application of the TAD filter. When noise suppression is applied first, the weak edges are smoothed and are not recovered by inhomogeneity correction.
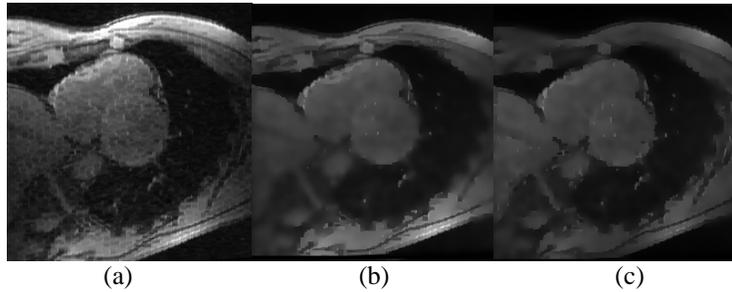
Figure 8: (a) Input image with background intensity variation. (b) Image when noise suppression precedes correction. (c)Image when correction precedes noise suppression.

3.2.2 Synthetic noise and field inhomogeneity

In order to determine whether noise suppression can be made to improve correction, we selected 2 subjects, one normal and one diseased and 1 volume from each subject from at early systole, mid systole, and at end systole. We then added 9 different combinations of synthetic field strengths (x) and noise power (y) to the prepared volumes yielding input volume set S_{3xy}. We then began comparing the performance of the two sequences of operations – inhomogeneity correction followed by noise filtering and vice versa – yielding sets S_{3xycf} and S_{3xyfc}.

It became immediately apparent that S_{3xycf} yields much better inhomogeneity correction than S_{3xyfc}. To understand this effect, we investigated the inhomogeneity correction process. The first row in Figure 9 illustrates the salient steps in the correction process. To correct the input image (first column), the scale in the image is computed where the scale at a voxel is the radius of the largest sphere centered at the voxel for which the intensities of the voxels in the sphere satisfy an intensity homogeneity constraint. The voxels having the largest scale are extracted and from these and the largest connected object is selected (third column). Next, the mean μ and standard deviation σ of the intensities of these voxels are computed. All voxels in the input image whose intensities fall into the range $[\mu - \nu\sigma, \mu + \omega\sigma]$ are used as points (fourth column) at which the field inhomogeneity is sampled and a 2nd degree polynomial is fit (fifth column) to the intensities at these points. The performance of the algorithm on a given input image is determined by the scale image that is computed. The scale image is determined by the homogeneity criteria used by the correction algorithm and the amount of noise in the image. The homogeneity criteria is largely dependent on the homogeneity value h (see below), which describes the expected standard deviation intensities in the same tissue.
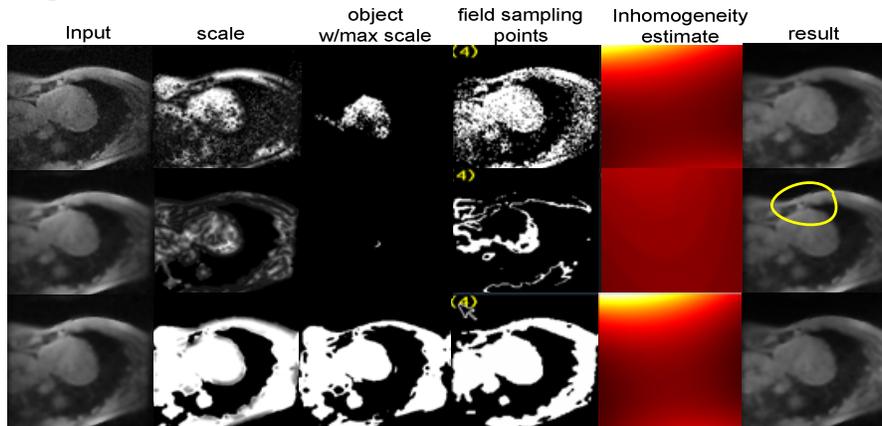

Figure 9: Intermediate results from the correction algorithm for three tests. First row: an image from S_{3xycf} with homogeneity estimated from the image in S_{3xy}. Second row: an image from S_{3xyfc} with homogeneity estimated from the image in S_{3xyf}. Third row: an image from S_{3xyfc} with homogeneity estimated from the image in S_{3xy}.

To compute the homogeneity value, the algorithm assumes that 15% of the voxels in the volume are occupied by true tissue edges. Homogeneity is computed as $h = \mu_l + 3\sigma_l$ where μ_l and σ_l are the mean and standard deviation of the lower 85% of the histogram of the intensity gradient image. In Figure 9 (first row), the sequence correction followed by

filtering is applied and the correction algorithm calculates the tissue homogeneity measure from the image in S_{3xy}. In the second row, the reverse sequence is applied and correction algorithm calculates tissue homogeneity from the image in S_{3xyf}. A much higher homogeneity is computed for the first sequence ($h = 1274$ vs $h = 145$) because the thermal noise has not been suppressed. Consequently the object with maximum scale (max scale is 12 in both sequences) is much smaller when the more stringent homogeneity ($h = 145$) is used as shown in the third column first and second rows. As a result, the standard deviation of the intensities of the pixels with max scale is much smaller (e.g., 56 vs 494 gray levels) while the mean remains roughly the same (e.g., 2051 vs 1993). Since the standard deviation is much smaller in the second sequence, a much smaller set of pixels in the image falls into the range [$\mu - \nu\sigma, \mu + \omega\sigma$], and thus, a much smaller set of inhomogeneity sampling points is used. In the intensity correction of cardiac MR, there is not one tissue which adequately samples the field inhomogeneity. To adequately sample the background intensity variation, the intensities of several tissues and structures are typically needed including: the blood in the ventricular cavities, the myocardium, the fat between the myocardium and chest wall and the chest wall. As seen in the fourth column, second row, the field sampling points do not contain much of the chest wall, and particularly those areas of the image which are brightest due to the surface coils, despite the fact that $\omega = 5$ in the sampling range. The sampling points used in the first sequence cover a greater extent of these structures and yields better correction. In SPAMM-MRI, the dominant slowly varying component of the intensities throughout the image is the field induced inhomogeneity and not the contrast between tissues. As a result, the correction based on a field estimated by sampling from multiple structures tends to be a more accurate estimate than that achieved from estimation from any one structure.

In order to investigate whether the second sequence would benefit from using the same homogeneity as the first sequence, we created a modified correction algorithm for the second sequence that uses the same homogeneity value created in the first sequence. As shown in the third row of the figure, the homogeneity is computed from images in S_{3xy} for filtering followed by correction and the object with maximum scale covers a much larger portion of the salient structures. Consequently, a much larger portion of the image is used to sample the inhomogeneity. This yields correction which is as good as correction followed by filtering. To explore the merits of this modified correction approach, we have measured the residual inhomogeneity in the myocardium (using the myocardium NSD) and the epicardium edge strength after applying the two sequences to the 54 input volumes in S_{3xy} while using the same (volume dependent) homogeneity for both sequences. In Figure 10, the first entry in each row is for inhomogeneity correction followed by filtering and the second entry is for the reverse sequence. For the myocardium, the differences in NSD are not statistically significant for any tested combination of noise and field: p>0.30 for all cases. On the other hand, the first sequence yields much stronger edges: the differences in the edge strength between the two orderings are statistically significant p<0.01 for all cases except 2 as noted by the p values.

Residual inhomogeneity	Low noise power	Medium noise power	High noise power
Low strength field inhomogeneity	0.171 (0.048)	0.171 (0.047)	0.170 (0.046)
	0.172 (0.047)	0.170 (0.046)	0.169 (0.045)
Medium strength field inhomogeneity	0.173 (0.049)	0.172 (0.049)	0.172 (0.048)
	0.173 (0.048)	0.172 (0.047)	0.171 (0.046)
High strength field inhomogeneity	0.174 (0.051)	0.173 (0.050)	0.171 (0.049)
	0.174 (0.048)	0.173 (0.047)	0.172 (0.047)
Edge strength			
Low strength field inhomogeneity	113.3 (38.47)	108.3 (34.26)	105.5 (32.96)
	106.7 (39.58)	102.2 (33.09)	100.4 (30.89)
	p<0.0001	p<0.0001	p<0.0001
Medium strength field inhomogeneity	137.7 (55.64)	129.8 (46.00)	125.0 (42.22)
	133.3 (66.13)	123.0 (50.03)	117.9 (42.0)
	p=0.20	p=0.002	p<0.0001
High strength field inhomogeneity	189.3 (118.0)	156.5 (67.28)	147.6 (56.20)
	161.7 (91.80)	152.13 (79.31)	142.47 (62.20)
	p<0.0001	p=0.27	p=0.012

Figure 10: Residual homogeneity and edge strength comparison between inhomogeneity correction followed by filtering (first entry) and the reverse sequence (second entry)

3.3 Evaluating the effect of using a field inhomogeneity estimate to regulate noise suppression

We formed volume sets S_{4TADi} and S_{4AWi} by injecting the inhomogeneity estimate into the TAD filter and the AW filter, respectively, to spatially regulate the noise suppression. We then compared these volumes to those generated using the same filters that do not utilize the inhomogeneity estimate: S_{4TAD} and S_{4AW}. Qualitatively, the filters that use the estimate yield images that appear to be significantly sharper and more detailed. While the noise which had been amplified by inhomogeneity correction is clearly reduced in the appropriate places, much more of the true structure in the images is preserved in the regions closer to the surface coil where the SNR is higher. In Figure 11(a), the second and third images in the second row illustrate the distinct advantage of employing the estimate of the field inhomogeneity to regulate noise suppression. Fine structures, such the papillary muscles, are much better preserved when the estimate is used. Two other examples are also shown. In the first row of (b), the papillary muscles of this diseased subject with right ventricular hypertrophy are more readily visible. In the second row, the moderator band is much better preserved (compare with the result in Figure 4 (b,bottom)).
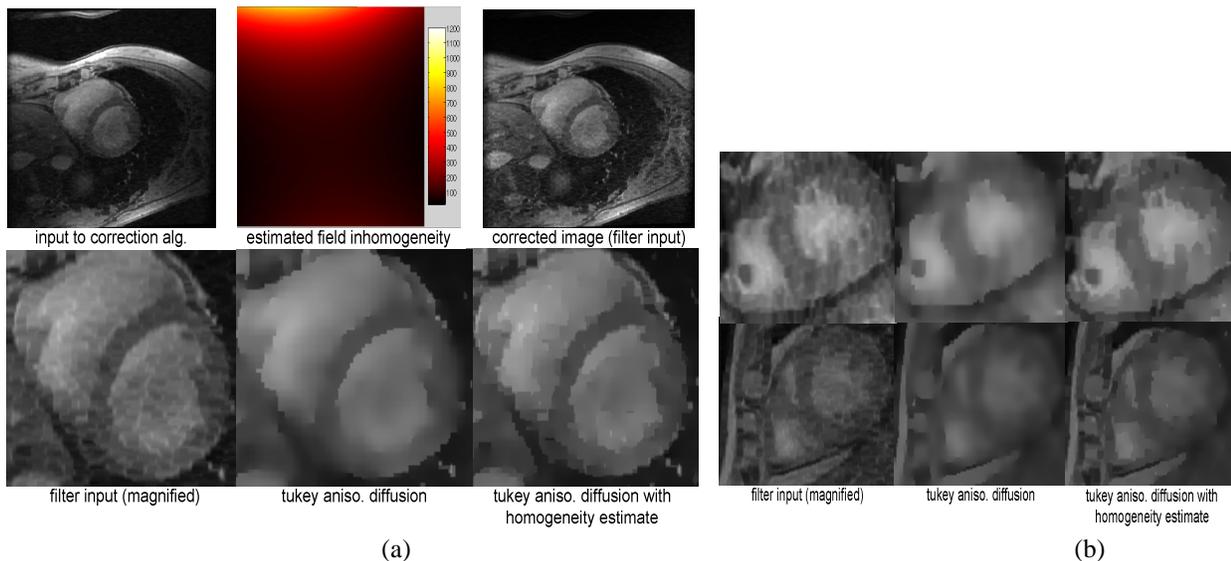

| | (a) | | (b) |

Figure 11: (a)First row: a field inhomogeneity estimate is used to correct the image. Second row: The field inhomogeneity estimate is also used to regulate noise suppression in the third image. (b)Other examples of the effect of using a field inhomogeneity estimate to regulate noise suppression.

We computed the myocardium NSD and epicardium edge strength for 10 subjects including normal and diseased. In Figure 12, we can see that there is slightly more residual noise in the myocardium when using the inhomogeneity estimate during noise suppression. On the other hand there is significantly stronger epicardium edge strength when the inhomogeneity estimate is used during noise suppression.

	TAD	TAD using inhomogeneity estimate	AW filter	AW filter using inhomogeneity estimate
Residual noise	0.178	0.209	0.172	0.191
	0.058	0.059	0.057	0.060
edge strength	178.5	233.9	128.0	178.14
	81.84	68.05	56.94	60.74

Figure 12: Comparison of the residual noise and edge strength (first entry: mean, second entry: standard deviation) .

4. CONCLUSIONS

In conclusion all the methods can be parameterized to yield the same residual noise. However using these parameter settings, they yield significantly different edge strengths. In particular, the SBD filter yields the strongest edges, followed by TAD and then the remaining combination filters with the weakest edges from the AW filter. The mean differences in edge strengths are statistically significant, $p < 0.001$, and sizable enough to be meaningful for many segmentation and analysis tasks. For the greatest reduction in background intensity variation, inhomogeneity correction

should precede noise suppression. For the best noise suppression, the inhomogeneity estimate from the correction method should be used to spatially regulate noise suppression. In future research, we plan to combine scale-based filtering with the inhomogeneity estimate for further improvements to noise suppression in MRI.

ACKNOWLEDGMENTS

REFERENCES

[1] Y. Zhuge, J. Udupa, J. Liu, P. Saha, T. Iwanaga, A Scale-Based Method for Correcting Background Intensity Variation in Acquired Images, SPIE Proceedings, vol.4684, pp.1103-1111, 2002.

[2] J. Sijbers, Signal and noise estimation from Magnetic Resonance Images, Ph.D. Thesis, Physics, University of Antwerp, 1998

[3] J. Lim, Two-Dimensional Signal and Image Processing, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1990, pp.536-540

[4] M. Black, G. Sapiro, et. al., Robust Anisotropic Diffusion, IEEE Trans. Img. Proc.,Vol 7, No 3, March 1998

[5] P.K. Saha, J.K.Udupa, Scale-based image filtering preserving boundary sharpness and fine structure, IEEE Trans. Med Im, vol.20, pp.1140-1155, 2001.

[6] S. Moyher, Surface Coil MR Imaging of the Human Brain with an Analytic Reception Profile Correction, JMRI 1995, vol. 5, pp. 139-144

[7] A.F. Frangi, Three-dimensional Model-based Analysis of Vascular and Cardiac Images, Ph.D. Thesis, Utrecht University, The Netherlands, 2001

[8] A. Amini, J. Prince eds, Measurement of Cardiac Deformations from MRI: Physical and Mathematical Models, Kluwer Academic Publishers, 2001

[9] J. Sled, et al, "A Nonparametric Method for automatic correction of Intensity Nonuniformity in MRI Data", IEEE TMI, Feb 1998.

[10] Leon Axel, Jay Costantini, and John Listerud. Intensity correction in surface-coil MR imaging. American Journal of Radiology, 148:418-420, February 1987

[11] A. Madabhushi, J. Udupa, Evaluating the effect of Intensity Standardization and Inhomogeneity Correction on MR Images, 28th IEEE North-Eastern Conference on Bioengineering, 2002, Philadelphia, PA

[12] A. Montillo, D. Metaxas, L. Axel, "Automated Segmentation of the Left and Right Ventricles in 4D Cardiac SPAMM Images", Accepted for Oral presentation, MICCAI Sept. 25, 2002

[13] L. Nyul, J. Udupa, "On Standardizing the MR Image Intensity Scale", Magnetic Resonance in Medicine, vol 42, 1999