

Complex Thresholding Methods for Eliminating Voxels That Contain Predominantly Noise in Magnetic Resonance Images

Daniel B. Rowe, Jing Jiang, and E. Mark Haacke

INTRODUCTION

There have been a number of excellent studies that have discussed the basic properties of noise in magnetic resonance (MR) imaging [1–12]. Removing noise in magnetic resonance images is important for improving the visualization of phase images and for the quantification of, for example, parametric maps such as the T_2 map [13–19]. Once noise is removed, it is also possible to revisit the data and extract boundary information, for instance. The simplest and most effective means to date to remove noise is to use a simple threshold on the magnitude images, which is to set voxels whose intensity fall below a certain threshold to zero. However, this approach has its limitations and can lead to the loss of signal information in the image and an incomplete removal of noise. In this chapter, we show that it is possible to achieve “better” results using information from both magnitude and phase images by applying a connectivity constraint [20, 21] or using a local neighborhood of voxels [22]. The motivation behind this work stems from susceptibility weighted imaging (SWI) [23] where it has been found that the tissue contrast was enhanced and the noise evidently reduced throughout

the image [24] using both magnitude and phase data. This work builds upon the feature of SWI to create a rapid means to remove as much noise from the image with as little effect on the object as possible.

We are interested in the noise behavior for the magnitude and phase images, specifically in complex-valued MR images. We assume that the original real and imaginary channels generate independent noise that is Gaussian distributed with mean zero and standard deviation σ .

With the Gaussian distributed noise assumption, the real and imaginary channels are $y_R = A \cos(\theta) + \eta_R$ and $y_I = A \sin(\theta) + \eta_I$, where y_R and y_I are the observed measurements for the real and imaginary parts, η_R and η_I are $N(0, \sigma^2)$ error terms for the real and imaginary parts, and A and θ are the noise-free (population) magnitude and phase. The joint probability distribution of a voxel's bivariate real and imaginary observation (y_R, y_I) is

$$p(y_R, y_I) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(y_R - A \cos \theta)^2}{2\sigma^2}\right] \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(y_I - A \sin \theta)^2}{2\sigma^2}\right] \quad (31.1)$$

A change of variable can be performed from rectangular coordinates in equation (31.1) to polar coordinates, and the joint probability distribution of a voxel's bivariate observed magnitude and phase (M, ϕ) is

$$p(M, \phi) = \frac{A}{2\pi\sigma^2} \exp\left\{-\frac{1}{2\sigma^2} [M^2 + A^2 - 2AM \cos(\phi - \theta)]\right\} \quad (31.2)$$

The (marginal) distribution of the magnitude $p_M(M)$ can be found by integrating equation (31.2) with respect to ϕ to obtain the Rician distribution [4–6]:

$$p_M(M) = \frac{M}{\sigma^2} e^{-(M^2 + A^2)/2\sigma^2} I_0\left(\frac{A \cdot M}{\sigma^2}\right) \quad (31.3)$$

where I_0 is the modified zeroth-order Bessel function of the first kind [1]. When $A = 0$, which corresponds to areas where there is only noise and no signal, the Rician distribution in equation (31.3) collapses to the Rayleigh distribution [2, 3]:

$$p_M(M) = \frac{M}{\sigma^2} e^{-M^2/2\sigma^2} \quad (31.4)$$

The mean of the Rayleigh distribution in equation (31.4) is $\sigma\sqrt{\pi/2}$ and the variance is $\sigma^2(2 - \pi/2)$.

For a low signal to noise ratio (SNR), that is, $\text{SNR} \leq 1$ where $\text{SNR} = A/\sigma$, the Rician distribution is far from being Gaussian. On the other hand, it starts to approximate a Gaussian distribution for $\text{SNR} \geq 3$:

$$p_M(M) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(M - \sqrt{A^2 + \sigma^2})^2/2\sigma^2} \quad (31.5)$$

with a mean of $\sqrt{A^2 + \sigma^2}$ and variance of variance σ^2 .

The phase (marginal) distribution $p(\varphi)$ can be found by integrating equation (31.2) with respect to M to obtain the phase marginal (pham) distribution [5, 10]

$$p(\varphi) = \frac{1}{2\pi} \exp\left\{-\frac{A^2}{2\sigma^2}\right\} \left[1 + \frac{A\sqrt{2\pi}}{\sigma} \cos(\varphi-\theta) \exp\left\{\frac{A^2 \cos(\varphi-\theta)}{2\sigma^2}\right\} \Phi\left(\frac{A^2 \cos(\varphi-\theta)}{\sigma}\right)\right] \quad (31.6)$$

where Φ is the standard normal cumulative distribution function. This pham distribution in equation (31.6) was described in detail by Rowe in the context of fMRI [25].

When there is signal present, the phase noise distribution in the image can be considered as a Gaussian distribution with mean $\theta = 0$ and standard deviation σ_{phase} . When $\text{SNR} \gg 1$, the standard deviation is $\sigma_{\text{phase}} = \sigma/A$, and the distribution of the phase is

$$p_{\text{object}}(\phi) = \frac{1}{\sqrt{2\pi}(\sigma/A)^2} \exp\left(-\frac{(\phi-\theta)^2}{2(\sigma/A)^2}\right) \quad (31.7)$$

As is obvious in equation (31.7), the standard deviation of the phase is [13]

$$\sigma_{\text{phase}} = \frac{1}{\text{SNR}} \quad (31.8)$$

where the units for σ_{phase} are in radians. In the phase image, when there is only noise, $A = 0$, and the distribution of the phase is the uniform distribution

$$p_{\text{noise}}(\phi) = \begin{cases} \frac{1}{2\pi}, & \text{if } -\pi < \phi < \pi \\ 0, & \text{otherwise} \end{cases} \quad (31.9)$$

The mean of the phase in equation (31.9) is zero and the variance is $\pi^2/3$.

Our aim is to exploit the noise presence in both phase and magnitude images so as to provide a more powerful thresholding technique. The complex threshold method (CTM) and the magnitude and phase threshold (MAPHT) method for removing noise voxels in complex-valued images are described in detail in the following.

THE CTM METHOD

The CTM consists of two steps. First, the magnitude image and the phase image are thresholded by $m\sigma_0$ and $n\sigma_{\text{phase}}$, where m and n are real and positive numbers, respectively. Then, through voxel connectivity, both false positive (the error of leaving noise voxels) and false negative (the error of eliminating signal voxels) rates are further minimized. A false negative error is called a Type I error and a false positive error is called a Type II error. More specifically, we define these errors as

$$\text{Probability of Type I error} = \frac{\text{total number of signal pixels removed}}{\text{total number of signal pixels}} \quad (31.10)$$

$$\text{Probability of Type II error} = \frac{\text{total number of noise pixels left}}{\text{total number of noise pixels}} \quad (31.11)$$

The processing algorithm is outlined graphically in Figure 31.1 and in detail below.

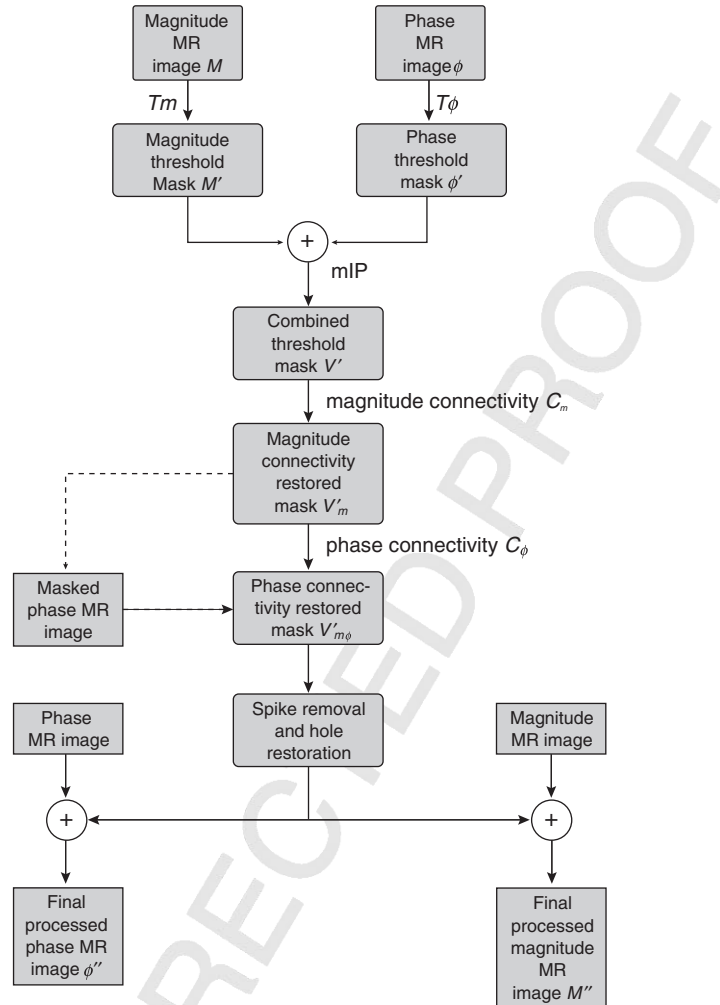


FIGURE 31.1 Flowchart of the CTM. Thresholds are applied to magnitude and phase images, and a mIP noise-removing mask is generated. A connectivity algorithm is run on the magnitude image and the noise mask is then corrected by restoring recovered non-noise object voxels. Connectivity is run again, but this time on the phase image to create a final noise mask that is then used to filter noise voxels from both magnitude and phase images.

Magnitude Threshold Step

Thresholding is applied to the magnitude image and a binary noise-removing mask image M' is created. This operation can be represented as

$$T_m : M'(x, y) = \begin{cases} 0, & \text{if } M(x, y) < m\sigma_0 \\ 1, & \text{if } M(x, y) \geq m\sigma_0 \end{cases} \quad (31.12)$$

where M is the magnitude MR image, m is the magnitude threshold, σ_0 is the standard deviation of noise as estimated from the image, and T_m is the threshold operator.

Phase Threshold Step

The useful information in the phase images is exploited by a phase thresholding method. For the phase image ϕ , a binary noise-removing mask ϕ' is created by admitting only phase values between $(-n\sigma_{\text{phase}})$ and $(n\sigma_{\text{phase}})$. Here, σ_{phase} is the sample standard deviation of noise in the phase image estimated from the SNR in the magnitude image (see equation (31.8)). Let $\phi(x, y)$ be the phase image. Then the mask ϕ' is determined as

$$T_{\phi} : \phi'(x, y) = \begin{cases} 0, & \text{if } |\phi(x, y)| > n\sigma_{\text{phase}} \\ 1, & \text{otherwise} \end{cases} \quad (31.13)$$

where T_{ϕ} is the phase threshold operator.

Combined Magnitude and Phase Thresholds

By combining magnitude and phase thresholds, it is possible to eliminate more noise voxels than when either method is used independently. This is accomplished by taking the minimum intensity projection (mIP) $\nu'(x, y)$ of the magnitude mask $M'(x, y)$ and phase mask $\phi'(x, y)$ as follows:

$$\nu'(x, y) = \begin{cases} 0, & \text{if either } M'(x, y) = 0 \text{ or } \phi'(x, y) = 0 \\ 1, & \text{otherwise} \end{cases} \quad (31.14)$$

The effect of thresholding with the distributions of both the magnitude and the phase is shown in Figure 31.2. As seen in Figure 31.2, these thresholds still remove information from the object (Type I error) and fail to remove some noise voxels (Type II error).

A short remark about Figure 31.2 is in order. When only the phase threshold is applied, Type I error can result for any value of the signal where the phase exceeds the threshold value. If both thresholds are applied, the dominant source of Type I error is from those true points to the left of the threshold being thrown out. Similarly, for Type II error, the phase-only threshold can still allow points to the left of the magnitude threshold to contribute to the noise being counted as part of the object. If both thresholds are applied, the dominant source of Type II error is from those noise points to the right of the magnitude threshold.

CONNECTIVITY

To reduce Type I error, we propose adding a local connectivity algorithm. Voxel connectivity describes a relation between the voxel under investigation and the surrounding neighborhood of voxels. Let p be a voxel with the coordinates (x, y) , then its 8-neighborhood $N_8(p)$ is defined as all those voxels that are immediate neighbors of p . The connectivity algorithm is applied only to those points that were discarded by the thresholding procedure. For the magnitude image, it is applied under the following conditional guideline: If the number of voxels connected to p in $N_8(p)$ of the magnitude image exceeding the magnitude threshold $m\sigma_0$ is greater than or equal to some integer number τ_M , then do not discard p . Similarly, for the phase image, if the number of voxels connected to p in $N_8(p)$ of the phase image exceeding the phase threshold $n\sigma_{\text{phase}}$ is greater than or equal to

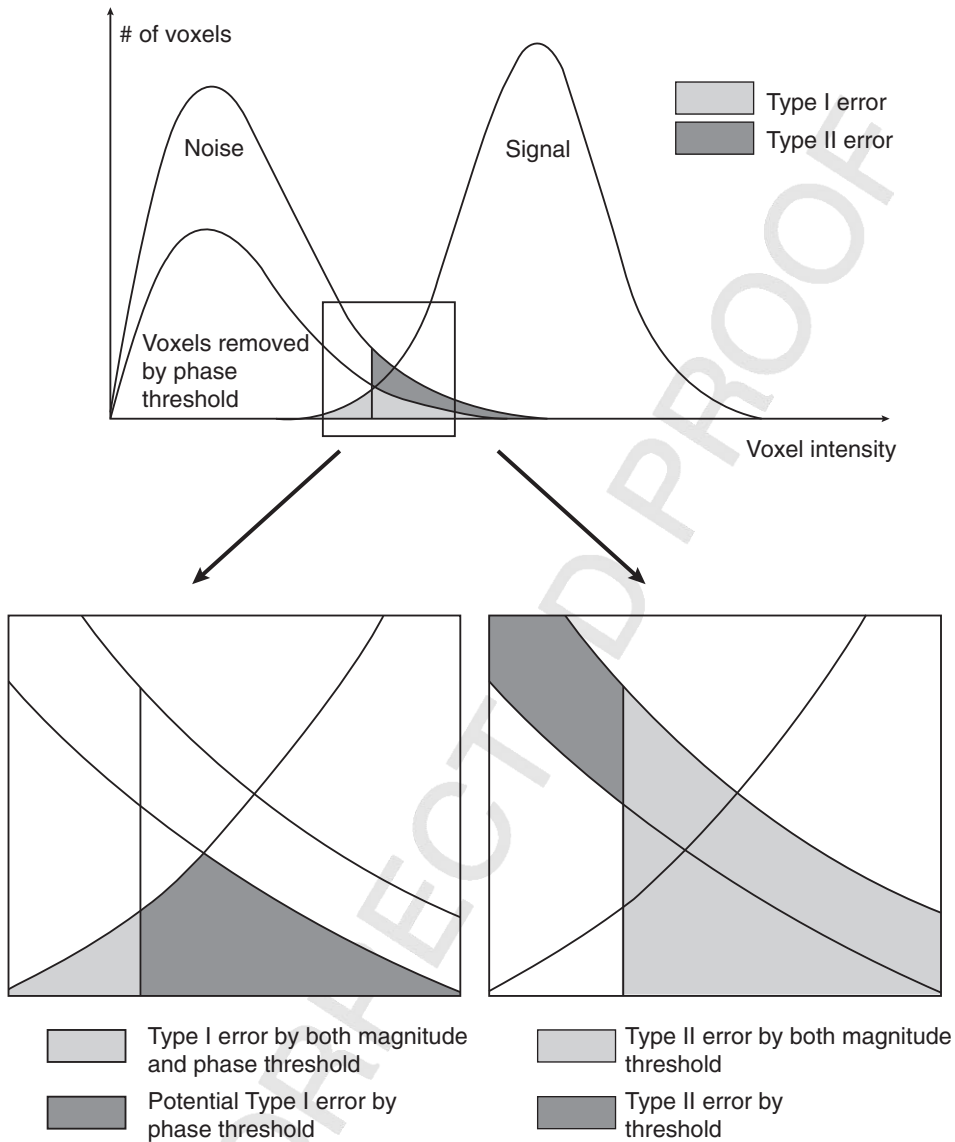



FIGURE 31.2 CTM bimodal curve showing the Rayleigh distribution for the noise (left distribution) and Rician distribution for the signal (right distribution). The bold cutoff line shows the magnitude threshold, which removes all the noise to the left of the threshold. The region under the inner curve in the noise mode represents the voxels removed by the phase threshold. In the left insert, the region shown in dark gray represents Type I error introduced by the phase threshold, while the light gray region represents the conventional Type I error from the magnitude threshold. In the right insert, the light gray region represents Type II error from the magnitude threshold while the dark gray region represents Type II error from the phase threshold.

some integer number τ_B , then do not discard p . We refer to these two connectivity operators as C_m and C_ϕ , respectively. Then the combined thresholded mask $\nu'(x, y)$ is modified to $\nu'_m(x, y) = C_m(\nu'(x, y))$ and this in turn is modified according to $\nu'_{m\phi}(x, y) = C_\phi(\nu'_m(x, y))$.

Applying CTM Connectivity for Spike Removal and Hole Restoration

As a final step, a simple spike removal and hole restoration algorithm is applied to reduce Type I and Type II errors. Since most of the noise has already been removed with the connectivity technique, the remaining points that constitute Type II error are predominantly single noisy points, whereas Type I error constitutes a few single voxels that are lost along the edges of the object. The spike removal and hole restoration algorithm works on these single voxels to remove or restore them. The algorithm works as follows: every voxel's neighborhood in the noise-reduced image is examined for connected voxels. If the magnitude connectivity shows that there is no neighboring voxels that are connected to the one under consideration, then it most likely is noise and should be removed. If the surrounding voxels are all signal, the examined voxel should be regarded as signal as well.

The MAPHT Method  local statistic based on the measured signal at the voxel of interest and its neighbors is defined. This statistic is a function of one or more random variables that does not depend on any unknown parameter and it has a distribution that approximates an F distribution [22]. Monte Carlo simulations are performed to determine critical values for this distribution, which is used to perform a statistical hypothesis test by thresholding [29, 30] to identify voxels containing predominantly noise.

With a sample of n independently and identically distributed measurements (magnitude M_i and phase ϕ_i , $i = 1, \dots, n$) from equation (31.1), the joint distribution of the measurements viewed as a function of the magnitude, phase, and error variance is the likelihood function:

$$L(A, \theta, \sigma^2) = (2\pi\sigma^2)^{-n} \left[\prod_{i=1}^n M_i \right] \exp \left\{ -\frac{1}{2\sigma^2} \sum_{i=1}^n [M_i^2 + A^2 - 2AM_i \cos(\phi_i - \theta)] \right\} \quad (31.15)$$

To identify voxels that contain predominantly noise, MAPHT performs a hypothesis test to determine whether the magnitude and the phase are statistically different from zero. A formal statistic can be derived from equation (31.15) and a joint statistical hypothesis test can be performed on the population magnitude and phase parameters.

A hypothesis test is performed to separate voxels that contain predominantly noise from those that contain predominantly signal. In hypothesis testing, there are four possible outcomes as depicted in Table 31.1. Note that the MAPHT Type 1 is similar to the CTM Type II (equation (31.10)) and that the MAPHT Type 2 is similar to the CTM Type I (equation (31.11)). Correct decisions are made in the top right and bottom left cells of Table 31.1. In the top left shaded cell, a Type 1 error is made in which the null hypothesis is rejected when it is true. The probability of a Type 1 error is the false positive rate denoted by α . In the bottom right shaded cell, a Type 2 error is made in which the null hypothesis is not rejected when it is false. The probability of a Type 2

TABLE 31.1 Four Outcomes from a Hypothesis Test

	H_0 True	H_0 False
Reject H_0	Type 1 error (α)	Correct decision ($1 - \beta$)
Do not reject H_0	Correct decision ($1 - \alpha$)	Type 2 error (β)

error is the false negative rate denoted by β . These two error rates will be examined in more detail later.

Voxels will be thresholded by testing the null hypotheses $H_0: A = 0, \theta = 0$ versus the alternative hypothesis $H_1: A > 0, \theta \neq 0$. Maximum likelihood estimators (MLEs) of the magnitude, phase, and noise variance under the two hypotheses can be computed from the likelihood function [15] by simply setting the partial derivatives of the likelihood function with respect to the sought-after parameters to zero and solving the resulting equations. Under the constrained null hypothesis $H_0: A = 0, \theta = 0$, the MLEs for the magnitude, phase, and noise variance are

$$\tilde{A} = 0, \tilde{\theta} = 0, \tilde{\sigma}^2 = \frac{1}{2n} \sum_{i=1}^n (y_{Ri}^2 + y_{Ii}^2) \quad (31.16)$$

while under the unconstrained alternative hypothesis $H_1: A > 0, \theta \neq 0$, the MLEs are

$$\begin{aligned} \tilde{A} &= [(\bar{y}_R)^2 + (\bar{y}_I)^2]^{1/2}, \quad \tilde{\theta} = \tan^{-1} \left[\frac{\sum_{i=1}^n y_{Ii}}{\sum_{i=1}^n y_{Ri}} \right], \\ \tilde{\sigma}^2 &= \frac{1}{2n} \sum_{i=1}^n (y_{Ri}^2 + y_{Ii}^2) - \frac{1}{2} [(\bar{y}_R)^2 + (\bar{y}_I)^2] \end{aligned} \quad (31.17)$$

where \bar{y}_R is the mean of the real channel measurements and \bar{y}_I is the mean of the imaginary channel measurements. These estimates in equations (31.16) and (31.17) are then inserted back into the likelihood function in equation (31.15) and the ratio of null hypothesis likelihood over alternative hypothesis likelihood is taken to form a likelihood ratio statistic [27, 28]:

$$\lambda = L(\tilde{A}, \tilde{\theta}, \tilde{\sigma}^2) / L(\hat{A}, \hat{\theta}, \hat{\sigma}^2) \quad (31.18)$$

It is easily verifiable that the likelihood ratio statistic can be simplified to get

$$\lambda = \left(\frac{\hat{\sigma}^2}{\tilde{\sigma}^2} \right)^n \quad (31.19)$$

Finally, a statistic is formed by letting

$$F = n \left(1 - \lambda^{1/n} \right) \quad (31.20)$$

According to equations (31.16) and (31.17), $0 < \hat{\sigma}^2 \leq \tilde{\sigma}^2$. Thus, $\lambda^{1/n} = \hat{\sigma}^2 / \tilde{\sigma}^2 \in [0, 1]$, which implies that $F \in [0, n]$. For voxels containing predominantly noise, $\hat{\sigma}^2$ approaches $\tilde{\sigma}^2$ and F approaches zero. On the other hand, for voxels containing predominantly signal, under the i.i.d. assumption of the measurements, it can be easily shown that $\hat{\sigma}^2$ approaches zero and thus F approaches n . The hypothesis test on F statistic is then to determine a critical value (threshold) F_α for a given significance level α (Type 1 error rate). Obviously, given the same complex signal, as the error variance σ^2 increases, $\tilde{\sigma}^2$ increases and as a result F increases and vice versa. In other words, the critical value for the F statistic should be a monotone increasing function of the noise variance. This suggests bigger critical value for an F statistic when the noise variance is large, as we would expect. However, as described at

the beginning of this chapter, the noise variance is a constant throughout an MR image, which means that a single critical value can be applied to the whole image.

Upon substitution of equations (31.16), (31.17), and (31.19) into (31.20), we get [22]

$$F = \left(\frac{n[(\bar{y}_R)^2 + (\bar{y}_I)^2]/\sigma^2}{2} \right) / \left(\frac{[\sum_{i=1}^n y_{Ri}^2 + \sum_{i=1}^n y_{Ii}^2]/\sigma^2}{2n} \right) \quad (31.21)$$

It can be shown that $n[(\bar{y}_R)^2 + (\bar{y}_I)^2]/\sigma^2$ and $[\sum_{i=1}^n y_{Ri}^2 + \sum_{i=1}^n y_{Ii}^2]/\sigma^2$ are χ^2 distributed with degrees of freedom 2 and $2n$, respectively. By dividing these two χ^2 distributed terms by their degrees of freedom and taking the ratio, the result is generally a statistic that under the null hypothesis has an F distribution with 2 and $2n$ degrees of freedom. However, this is not true in this case. These two χ^2 statistics need be statistically independent in order for this to be true. It can be shown that the correlation between these two statistics is $1/\sqrt{n}$. This correlation approaches zero for large n and the F statistic becomes F distributed with 2 numerator and $2n$ denominator degrees of freedom. In nonfunctional MR imaging applications, there may be a very small number of repeated images if any. For the case where only one single image is available, we take the voxel and its nearest neighbors as the sample set of independent measurements. Still, n is small, and this asymptotic result does not hold. Thus, critical values from the F distribution do not apply to this F statistic.

However, critical values for small n can be achieved by way of Monte Carlo simulation. For a given level of significance (Type 1 error rate α), we reject H_0 (do not threshold voxel) if the test statistic F is larger than the critical value $F_{\alpha}(2, 2n)$ and do not reject (threshold voxel) if F is smaller than the critical value $F_{\alpha}(2, 2n)$. When thresholding the statistic image, an adjustment for multiple comparisons such as with a Bonferroni corrected threshold can be performed [29, 30].

To examine the theoretical Type 1 error and Type 2 error relationship, 20 million independent simulated data values for $n = 9$ were created under the null hypothesis $A = 0$ and $\theta = 0$ and under the alternative hypothesis with $A = (1, 3, 4, 5, 7.5, 10)$ and $\theta = 0$. Normally distributed independent noise variates were generated for the real and imaginary parts with a mean of zero and a variance of $\sigma^2 = 1$. Figure 31.3 shows a histogram for the million data sets when $A = 0$ in light gray (unshaded) and $A = 1$ in dark gray colors (shaded).




FIGURE 31.3 MAPHT histograms when H_0 is true ($A = 0$) and when H_1 is true ($A = 1$). (a) magnitude, (b) phase, and (c) F statistic. The vertical line in (c) is at $F_{0.05}(2, 18) = 2.81$, which is the critical value for $\alpha = 0.05$. Note that there is not a clear separation between the two populations in either the magnitude or phase, but there is better separation with the F statistics.

Figure 31.3a presents the distribution of magnitudes, (b) presents the distribution of phases, and (c) presents the distribution of F statistics. The black vertical line in Figure 31.3c is at $F_{0.05} = 2.81$, which is the critical value when $\alpha = 0.05$. The false positive rate α is the intersection area that is to the right of F_α . The light gray and intersection areas less than F_α are the true positive rate $1 - \alpha$. The dark gray and intersection areas to the right of F_α are the true negative rate $1 - \beta$. The false negative rate β is the intersection and the dark colored areas that are to the left of F_α . It can be seen in Figure 31.3 that as the null hypothesis false positive rate α decreases, the alternative hypothesis false negative rate β increases. Also, we note that the Monte Carlo results validated that the distribution of F values never exceeds n regardless of A .

For each value F_α for the vertical line with corresponding false positive rate areas α to the right of it, we can determine the false negative rate β . A plot of α on the horizontal axis and $1 - \beta$ on the vertical axis is called a receiver operating characteristic (ROC) curve [8]. This curve can be made for each combination of A . In Figure 31.4 are the ROC curves for $A = (3, 4, 5, 7.5, 10)$. Note that as A increases for a given false positive rate α , the true positive rate ($1 - \beta$) increases. As A increases, the alternative hypothesis distribution (dark gray colored) in Figure 31.3c becomes more peaked and shifts to the right. As the alternative hypothesis distribution shifts to the right, the lower tail has less area below a given Type 1 error rate




FIGURE 31.4 MAPHT ROC curves from Monte Carlo simulation for $A = (0, 3, 4, 5, 7.5, 10)$. Note that as A increases, for a given false positive rate α , the true positive rate $1 - \beta$ increases.




FIGURE 31.5 MAPHT histogram of F statistic for 5×10^7 data sets under null hypothesis for critical values. Note that larger F values are less likely.

(false positive rate α), and thus the Type 2 error rate (false negative rate β) decreases and the true positive rate ($1 - \beta$) increases.

To determine Monte Carlo theoretical critical values for the F statistic when $n=9$, 5×10^7 data values were generated. For each data value, the F statistic was computed. A histogram of these F statistics for $n=9$ is presented in Figure 31.5. These F statistics were ordered and percentiles determined. Selected critical values are presented in Table 31.2 for $n=9$. Intermediate critical values can be reliably interpolated or extrapolated with $F_{\max} = n$ when $\alpha = 0$.

In conclusion, when the MAPHT method is applied to a complex MR image, an F statistic is computed for all the voxels. A critical value is estimated for the whole image using the results of the Monte Carlo simulations and the voxels from the signal free area (noise measurements). The F statistic map is then thresholded by the critical value. A binary mask is produced in which the number 1 (unity) corresponds to an F statistic greater than the critical value and the number 0 (zero) corresponds to an F statistic less than the critical value.

TABLE 31.2 MCMC F Statistic Significance Level and Critical Values for $n=9$

α	$F_{\alpha}(2, 18)$
0.05	2.8102
0.01	3.9377
0.001	5.1991
0.0001	6.1512
0.00001	6.8678
0.000001	7.3911
0.05/256/256	7.5627
0.05/352/512	7.5869
0.05/384/512	7.7051
0.05/512/512	7.7575

This binary mask is used to threshold the original magnitude and phase images. The magnitude and phase of thresholded voxels are set to zero but for display the thresholded voxel phase values are set to $-\pi$.

SIMULATED DATA

Simulated images were created to test the algorithm under controlled conditions for a circle with a radius of 128 voxels embedded in a field of view of 512 voxels. Using a Monte Carlo approach, the SNR in the circle was set to 3:1, 5:1, and 10:1 and the algorithm was tested in each case. A ROC [12] was produced for each SNR value in each of the steps defined in Figure 31.1. More specifically, the real and imaginary channels were created for a given A (of 3, 5, or 10) as $A \cos(\theta) + \eta_R$ and $A \sin(\theta) + \eta_I$, where η_R and η_I are $N(0, \sigma^2)$ with $\sigma = 1$ and $\theta = 0$. Magnitude and phase data were then generated from this complex-valued data set.

HUMAN DATA

To test the CTM under low SNR when phase is expected to have a zero mean, spin echo data were collected with a thin slice and high resolution at 1.5 T on a Siemens Sonata. The imaging parameters were field of view (FOV) = $(192 \times 256 \text{ mm}^2)$, matrix size = 384×512 , in-plane resolution = $0.5 \times 0.5 \text{ mm}^2$, slice thickness = 2 mm, TR/TE = 300 ms/15 ms, and flip angle (FA) = 90° . To evaluate data with a bulk of the phase with zero mean but some structures within phase different from zero, such as the veins, SWI [23] data of brain and leg were collected. The SWI brain volume was acquired at 3 T on a Siemens Verio. Imaging parameters were FOV = $196 \times 256 \text{ mm}^2$, matrix size = 384×512 , in-plane resolution = $0.5 \times 0.5 \text{ mm}^2$, slice thickness = 2 mm, TR/TE = 29 ms/20 ms, and FA = 15° . The SWI leg data were collected at 1.5 T on a Siemens Sonata. The imaging parameters were FOV = $150 \times 200 \text{ mm}^2$, matrix size = 384×512 , in-plane resolution = $0.4 \times 0.4 \text{ mm}^2$, slice thickness = 1 mm, TR/TE = 21 ms/10.2 ms, and FA = 15° .

Simulated Data for CTM

The ROC curves for magnitude and phase thresholds (both separately and combined) for the simulated circle are shown in Figure 31.6 without connectivity for an SNR of 3:1, and in Figure 31.7 with connectivity for SNRs of 3:1, 5:1, and 10:1. In Figure 31.7, we can see that both errors remain rather large for either the magnitude or the phase method with moderate improvement when both are combined. In order to keep Type II error less than 0.1, the best choice for m and n would be an m of 1.5 or 2 and an n of 2–2.5. Adding magnitude connectivity dramatically reduces the Type I error (figure not shown). Minimum error is achieved for a magnitude connectivity τ_M of 3. In order to keep Type II error less than 0.004, the best choice for m and n would now be an m of 3 or 4 and an n of roughly 2–4. In this case, all the Type I error will now be less than 0.003. Finally, adding phase connectivity reduces Type I error even further to less than 0.0001 for a phase connectivity τ_P of 2 or 3 and a magnitude connectivity τ_M of 2 or 3. In summary, the set of (τ_M, τ_P, m, n) that will work best for an SNR of 3:1 could range from the minimum of the above choices (2, 2, 3, 2) to the maximum of roughly (3, 3, 4, 4), with a Type I error of no more than 0.0001. Generally, the lower the n is, the lower the Type II error will be.




FIGURE 31.6 The CTM ROC results combining both magnitude (circles) and phase (triangles) threshold operations (SNR 3:1). There is a clear reduction of Type I and Type II errors in the combined operation (diamonds).




FIGURE 31.7 CTM Type I error versus Type II error for magnitude ($\tau_M = 2-4$) and phase ($\tau_P = 2-4$) connectivity along with magnitude threshold ($m = 0.5-4.5$) and phase threshold ($n = 1.0-5.5$) values for SNR of 3:1, 5:1, and 10:1. Choosing a Type II error less than 0.003 provides a broad range of possible (τ_M , τ_P , m , n) values with very small Type I error.

For higher SNR, τ_M and τ_P can range from 2 to 4, m from 1.5 to 4, and n from 2 to 4. Under these circumstances, Type I error will remain less than 0.0004 and Type II error will remain less than 0.006. For example, a (τ_M, τ_P, m, n) of $(3, 3, 3, 3)$ fits in this domain for an SNR of 5:1 or higher. Running the spike removal and hole restoration once yields Type I and Type II errors of 1.944×10^{-5} and 0.0035, respectively, and running it twice yields errors of 0 and 0.005, respectively. For the higher SNR cases, a magnitude and phase connectivity of 4 performs the best. An example of the full processing as applied to the simulated circle for an SNR of 3:1 is shown in Figure 31.8. The final Type I error is 0.000486 (25 voxels thrown out) and the final Type II error is 0.002112 (445 voxels not thrown out), in good agreement with the above predictions. Finally, the time to fully process one complex-valued image is just under 3 s at a processing speed of 3.06 GHz.




FIGURE 31.8 CTM results for simulated data: (a) original magnitude image, (b) original phase image, (c) thresholded magnitude image with $(\tau_M, \tau_P, m, n) = (3, 3, 2, 2)$, and (d) thresholded phase image with $(\tau_M, \tau_P, m, n) = (3, 3, 2, 2)$, showing no noise remaining outside the image. In this case, Type I error is 0.000486 (25 points are removed, mostly around the edges) and Type II error is 0.002112 (445 noise voxels remain). The SNR in the magnitude image is 3:1.

Human Data for CTM

Estimation of background noise was done by selecting a region of interest outside the brain and using the voxel intensity values to obtain σ_0 as explained in the “Methods” section. We used the approach of taking the mean of the noise (signal) outside the object as being 1.25 standard deviations of that on the inside [12]. First, we tested a set of 1.5 T spin echo images with an SNR of 3:1 (Figure 31.9). The original magnitude image (Figure 31.9a) is very noisy and consequently shows little contrast. Figure 31.9b shows the usual magnitude-only threshold with $m = 2$, and thus a great deal of noise still remains in the image. Figure 31.9c uses the full CTM processing with $(\tau_M, \tau_P, m, n) = (3, 3, 2, 4)$. Much of the noise is now suppressed. Figure 31.9d shows a heavily filtered and averaged image showing what the background contrast of this T_1 -weighted spin echo scan would look like if there were enough SNR. (This was accomplished by using a $4 \times 4 \times 3$ average filter). Figure 31.9e is the original phase image and, as expected for a spin echo scan, the phase is flat (except in the vicinity of vessels that may not be fully flow compensated). The final masked phase image is shown in Figure 31.9f. It is now much easier to adjust window level settings and the image is beginning to look more like a conventional MR image.




FIGURE 31.9 CTM results for spin echo brain data. (a) Original magnitude image with a resolution of $0.5 \times 0.5 \times 2 \text{ mm}^3$ and SNR of 3:1. (b) Magnitude result after CTM filter with $(\tau_M, \tau_P, m, n) = (0, 0, 2, 0)$, that is, magnitude threshold of $m = 2$ and no connectivity. (c) Magnitude result after CTM filter with $(\tau_M, \tau_P, m, n) = (3, 3, 2, 4)$. (d) Magnitude image after $4 \times 4 \times 3$ average filter on original magnitude, resolution decreased to $2 \times 2 \times 6 \text{ mm}^3$. (e) The original phase image and (f) phase result after CTM filter with $(\tau_M, \tau_P, m, n) = (3, 3, 2, 4)$.




FIGURE 31.10 CTM results for SWI brain data with an SNR of roughly $>30:1$. (A1) Original magnitude image of midbrain, (B1) original phase image of midbrain, (C1) magnitude result after CTM filter with parameters $(\tau_M, \tau_P, m, n) = (3, 0, 5, 5)$, that is, thresholding with $m = n = 5$, in both magnitude and phase, $p = 5$ in magnitude connectivity and no connectivity, of A1 and B1, (D1) phase result after CTM filter with parameters with $(\tau_M, \tau_P, m, n) = (3, 0, 5, 5)$. (A2) Original magnitude image of thalamostriate area; (B2) original phase image of thalamostriate area; (C2) magnitude result after CTM filter with parameters $(\tau_M, \tau_P, m, n) = (3, 0, 5, 5)$ of A2, B2; (D2) phase result after CTM filter with parameters $(\tau_M, \tau_P, m, n) = (3, 0, 5, 5)$ of A2, B2. (A3) Original magnitude image of motor cortex area; (B3) original phase image of motor cortex area; (C3) magnitude result after CTM filter with parameters $(\tau_M, \tau_P, m, n) = (3, 0, 5, 5)$ of A3, B3; (D3) phase result after CTM filter with parameters $(\tau_M, \tau_P, m, n) = (3, 0, 5, 5)$ of A3, B3.




FIGURE 31.10 (Continued).

Another example in the human brain is shown from an SWI data set at 3 T (Figure 31.10). Three slices are chosen, representing the midbrain (Figure 31.10A1 and B1), thalamostriate area (Figure 31.10A2 and B2), and motor cortex area (Figure 31.10A3 and B3). The SNR is very good, above 30:1, in the original magnitude image (Figure 31.10A1–A3). Thus, the CTM filtered magnitude result does not have much change with the original input (Figure 31.10C1–C3). The SWI filtered phase image, however, is fairly uniform except for the phase variations caused by air/tissue interfaces at the top (Figure 31.10B1 and B2). Phase aliasing or wrapping artifacts can be seen clearly (Figure 31.10B1). First, a threshold value of $m = n = 5$ on magnitude and phase is applied to keep most of the data inside the




FIGURE 31.10 (Continued).

brain. To delete most of the spike or islands of noise, a magnitude connectivity of $\tau_M = 3$ is used. No phase connectivity is run in this example since the spike removal part was run. In the resulting processed phase image, the noise outside brain has been discarded successfully without sacrificing signal loss from within the brain (Figure 31.10D3). The phase wrapping effects cause some signal lost at the air/tissue interface (Figure 31.10D1 and D2). The CTM parameters used here were (3, 0, 5, 5).

Figure 31.11 shows the final example in the human leg SWI data, in which the method works in a 1.5 T dataset using (3, 3, 2, 2). The original unthresholded magnitude image is in




FIGURE 31.11 CTM results for the SWI leg data. (a) Original magnitude image, (b) original phase image, (c) thresholded magnitude image $(\tau_M, \tau_P, m, n) = (3, 3, 2, 2)$, and (d) thresholded phase image $(\tau_M, \tau_P, m, n) = (3, 3, 2, 2)$. This thresholded phase image takes on more characteristics of the conventional magnitude image, except that it shows the veins clearly, whereas no veins are seen in the conventional image. The thresholded phase image is no longer hampered by the noise points, and so it is easier to adjust the window and the level and to avoid being distracted by the presence of phase noise cluttering the image. The SNR in the magnitude image is 5:1.

Figure 31.11a and original unthresholded phase image is in Figure 31.11b. A CTM threshold of $(\tau_M, \tau_P, m, n) = (3, 3, 2, 2)$ is applied and in Figure 31.11c is the resulting thresholded magnitude image and in Figure 31.11d is the resulting thresholded phase image. The phase-thresholded image in Figure 31.11d not only shows a significant removal of most noise but also shows those areas where the veins have higher phase than the set threshold. This observation can be useful for processing the SWI data.

Although Type I errors remain along the edges of the brain and the leg in the presented examples due to partial volume signal effects and result from the phase variations there, for display purposes, the overall feature of removing the noise from the phase image is quite robust.

Simulated Data for MAPHT

The ROC curves characterizing the theoretical Type 1 and Type 2 error rates for the MAPHT method were presented in Figure 31.4. To examine the empirical performance of the MAPHT method in comparison to the CTM, new simulated data was generated according to the same process as the simulated data that the CTM method was applied to. The simulated image has a matrix size of 512×512 , and thus a Bonferroni corrected threshold of $\alpha = 0.05/512/512$ was applied with $n = 9$. Figure 31.12a shows the original magnitude image with noise outside the object (simulated brain). Note that in Figure 31.12b there is high noise in the phase image outside the object and in some internal areas, while the magnitude noise in Figure 31.12a is relatively low. Figure 31.12c shows the computed F statistic map, and the anatomical structure of the object can be seen. In Figure 31.12d, the histograms of the image magnitudes (top), the image phases (middle), and the F statistic values (bottom) are shown. Note that in the histogram for F statistics, there appear to be two populations, but there is no such presence in histograms of either the magnitude or the phase. The population on the left for smaller F values




FIGURE 31.12 MAPHT results for simulated data. (a) Original magnitude image, (b) original phase image, (c) unthresholded F statistic map, (d) histograms, magnitude data (top), phase data (middle), and F statistics (bottom), (e) $F = 7.7575$ Bonferroni thresholded magnitude image, (f) $F = 7.7575$ Bonferroni thresholded phase image, (g) $F = 5.75$ thresholded magnitude image, and (h) $F = 5.75$ thresholded phase image.

(similar to the ones in Figures 31.3 and 31.5) corresponds to voxels that contain predominantly noise and the other on the right for larger F values (similar to the one in Figure 31.3) corresponds to voxels that contain tissue signal plus noise. The Bonferroni threshold $F = 7.7575$ is depicted in Figure 31.12d (bottom) as a dash-dot vertical line. In Figure 31.12e and f are the thresholded simulated magnitude and phase images. Note that in Figure 31.12f there is not only an elimination of the voxels that contain predominantly noise but also an elimination of many voxels that contain signal plus noise. As seen in Figure 31.12d (bottom), many voxels that contain tissue signal plus noise appear to be below the Bonferroni threshold indicated by the dash-dot line and thus eliminated. This is due to low SNR of the object and the lack of a clear separation of two populations by magnitude in Figure 31.12d (top). The threshold can be lowered to a value such as $F = 5.5$ denoted by a solid vertical line as in Figure 31.12d (bottom) at the expense of a very small number of false positives (voxels that predominantly contain noise not being thresholded). Figure 31.12g and h show the lower thresholded magnitude and phase images. More anatomical detail (true object) is present in the phase image of Figure 31.12h and less tissue plus noise voxels are eliminated.

Human Data for MAPHT

To examine the empirical performance of the MAPHT method compared to the CTM method, it was applied to the same human brain spin echo data. The human brain spin echo image has a matrix size of 384×512 , and thus a Bonferroni corrected threshold of $\alpha = 0.05/512/384$ was applied with $n = 9$. Figure 31.13b shows that there is high noise in the phase image outside the brain and in some internal areas while the magnitude noise in Figure 31.13a is relatively low. Some anatomical structure can be seen in the F statistic

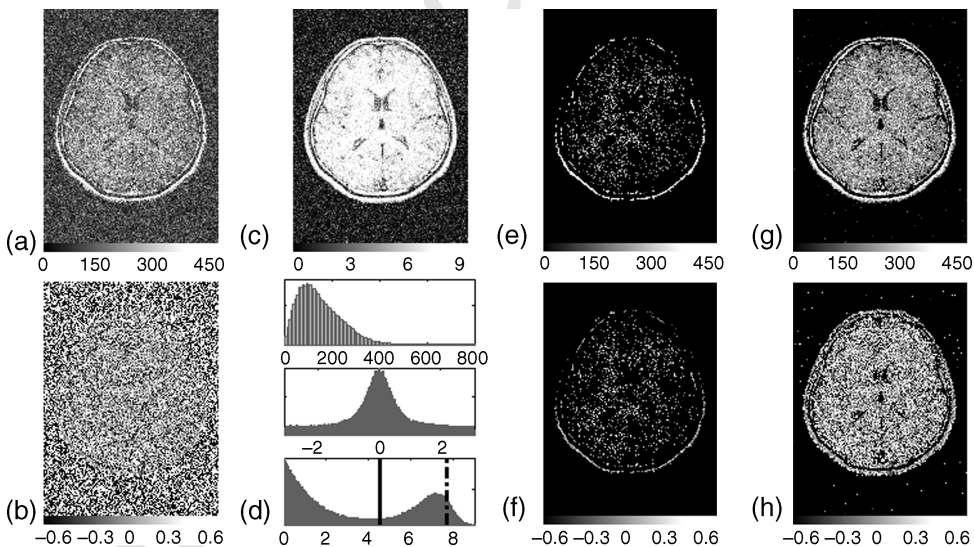


FIGURE 31.13 MAPHT results for the spin echo brain data. (a) Original magnitude image, (b) original phase image, (c) unthresholded F statistic map, (d) histograms, magnitude data (top), phase data (middle), and F statistics (bottom), (e) $F = 7.7051$ Bonferroni thresholded magnitude image, (f) $F = 7.7051$ Bonferroni thresholded phase image, (g) $F = 4.5$ thresholded magnitude image, and (h) $F = 4.5$ thresholded phase image.

map (Figure 31.13c). Figure 31.13d shows the histograms for magnitudes, phases, and F statistics. Note again that there appear to be two populations in the histogram of F statistics, but there is no such presence in histograms of either the magnitude or the phase. The population on the left for smaller F values corresponds to voxels that contain predominantly noise and the other for larger F values corresponds to voxels that contain tissue signal plus noise. The Bonferroni threshold $F=7.7051$ is depicted in Figure 31.13d (bottom) as a dash-dot vertical line. Figure 31.13e and f show the thresholded SWI magnitude and phase images. Note in Figure 31.13f that there is not only an elimination of the voxels that contain predominantly noise but also an elimination of many voxels that contain signal plus noise, which renders the thresholded image useless. As seen in Figure 31.13d (bottom), many voxels that contain tissue signal plus noise appear to be below the Bonferroni threshold indicated by the dash-dot line and thus eliminated. This is due to low SNR of the brain and the lack of a clear separation of two populations by magnitude in Figure 31.13d (top). The threshold can be lowered to a value such as $F=4.5$ denoted by a solid vertical line as in Figure 31.13d (bottom) at the expense of a small number of false positives (voxels that predominantly contain noise not being thresholded). Figure 31.13g and h shows the lower thresholded magnitude and phase images, respectively. More anatomical detail is present in the phase image of Figure 31.13h and less tissue plus noise voxels are eliminated.

The MAPHT method is also applied to a human brain SWI data. The human brain SWI image has a matrix size of 352×512 , and thus a Bonferroni corrected threshold of $\alpha = 0.05/352/512$ was applied with $n=9$. The results are shown in Figure 31.14. As seen in Figure 31.14d (bottom), the Bonferroni threshold identified by a vertical dash-dot line can

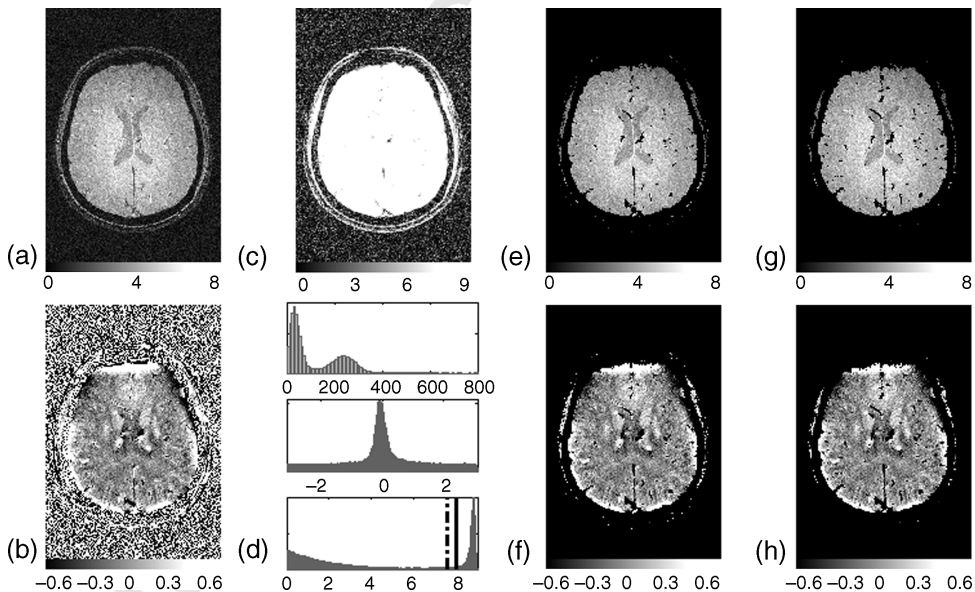


FIGURE 31.14 MAPHT results for the SWI brain data. (a) Original magnitude image, (b) original phase image, (c) unthresholded F statistic map, (d) histograms, magnitude data (top), phase data (middle), and F statistics (bottom), (e) $F = 7.5869$ Bonferroni thresholded magnitude image, (f) $F = 7.5869$ Bonferroni thresholded phase image, (g) $F = 8$ thresholded magnitude image, and (h) $F = 8$ thresholded phase image.




FIGURE 31.15 MAPHT processing on SWI leg image. (a) Original magnitude image, (b) original phase image, (c) unthresholded F statistic map, (d) histograms, magnitude data (top), phase data (middle), and F statistics (bottom), (e) $F = 7.7051$ Bonferroni thresholded magnitude image, (f) $F = 7.7051$ Bonferroni thresholded phase image, (g) $F = 5.5$ thresholded magnitude image, and (h) $F = 5.5$ thresholded phase image.

be raised to the solid line in order to obtain better tissue contrast with little elimination of voxels that contain tissue signal plus noise. This is due to high SNR of the object. Figure 31.14g and h shows the higher thresholded magnitude and phase images. More anatomical detail (tissue contrast) is present in the phase image of Figure 31.14h.

Finally, the MAPHT method is applied to the same human leg SWI data. The human leg SWI image has a matrix size of 384×512 , and thus a Bonferroni corrected threshold of $\alpha = 0.05/384/512$ was applied with $n = 9$. The results are shown in Figure 31.15. The Bonferroni threshold $F = 7.7051$ is depicted in Figure 31.15d as a dash-dot vertical line. In Figure 31.15e is the thresholded SWI magnitude and in Figure 31.15f is the thresholded phase image. Note in Figure 31.15f that there is an elimination of the voxels that contain predominantly noise, but there is a possible elimination of a small number of voxels that contain signal plus noise. As seen in Figure 31.15d (bottom), a small number of voxels that contain signal plus noise appear to be below the Bonferroni threshold indicated by dash-dot line and thus eliminated. This is due to some low SNR voxels in the object. The threshold can be lowered to a value such as $F = 5.5$ denoted by a solid vertical line as in Figure 31.15d (bottom). The results of lower threshold value are depicted in Figure 31.15g and h.

CONCLUDING REMARKS

Removing noise involves finding the delicate balance between removing unwanted signal components and components that are part of the object of interest. For low SNR, a simple

magnitude threshold will remove a considerable amount of noise, but at the expense of removing signal as well.

Complex Threshold Method

The goal of using phase and connectivity is to increase the probability of retaining as much information about the object as possible (smallest Type I error) and removing as much noise as possible (smallest Type II error). The results of this study show that it is possible to reduce Type I and Type II errors to almost zero even for very noisy data with an SNR of 3:1. The remnant areas that are hard to remove represent first some voxels near the edges that are kept by applying the connectivity algorithm, and second, clusters of noise points that exceed the diameter of three neighboring points. The choice of connectivity = 3 appears to be optimal and makes sense geometrically. If there are three connected points, then objects thrown out inside a rectangle-like object will be reinstated while circle-like objects may not. This keeps some noise points inside the object but will tend to enlarge the boundaries of the circle. A connectivity of 4 tends to throw out more of these noise points, but at the expense of throwing out signal points when the SNR is too low. However, if one extends this on a second pass to a connectivity of 4, then edges will remain essentially untouched. Other methods [15, 18] recognize the need for edge preservation as well.

Although our focus has been on the low SNR cases (because these cases are perfect examples of where the simple magnitude threshold methods fail), the higher SNR data can be further optimized as well using CTM. From a practical point of view, ideally there exists a fixed set of values for connectivity and thresholds that would give a robust result. For an SNR of 3:1, the best choice of connectivity and threshold values (τ_M, τ_P, m, n) ranges from (3, 3, 2, 2) to (3, 3, 2, 4), as shown in Figures 31.8 and 31.9. For an SNR of 5:1 or higher, the best choice for (τ_M, τ_P, m, n) ranges from (2, 2, 3, 2) to (3, 3, 4, 4), as shown in Figure 31.11. Although many images in MRI have high SNR, with the recent push to higher resolution in MR angiography, susceptibility weighted imaging, and anatomical imaging (especially at high fields where the rf response is nonuniform), the SNR values will drop considerably (perhaps approaching 3:1), making the CTM filter more useful.

For an SNR of 5:1 in the SWI leg data, (3, 3, 2, 2) suppresses the phase of veins more efficiently (Figure 31.11d), while (3, 3, 2, 4) keeps many more voxels within the veins and at the edge of air/tissue interfaces while still removing the background noise (image not shown). The choice of a phase threshold of 2 or 4 depends on whether one wants to simply remove phase noise higher than 4 standard deviations or to also remove signal from veins, for instance, as in SWI. In such a case, one may wish to push the phase threshold down to 2. In the SWI brain example (SNR > 30:1), (3, 3, 2, 4) works fine and one can afford to increase the magnitude threshold to 5 or even higher and use (3, 0, 5, 5). The choice of $m = 5$ removes the outer boundary of the skull because of its lower signal but maintains the signal inside the brain. The higher connectivity of 4 tends not to restore many voxels because the connectivity is too stringent a constraint, especially for the very low SNR of 3:1. When the SNR is higher than 10:1, (τ_M, τ_P, m, n) has a large range to choose from. As demonstrated in Figure 31.10, the phase image can be well separated from the noise even if there are general baseline shifts.

As a practical point, the condition for equation (31.7), in which the phase has a mean of zero, is valid only for a spin echo sequence with a perfectly centered echo. However, for an asymmetric placement of the π pulse relative to the echo time or for a gradient echo

sequence (if the phase is high pass filtered), the mean phase will again tend to be zero [23]. There are other methods where the phase need not be zero for this approach to be useful. This could include removing major veins in SWI (as shown in Figures 31.11 and 31.14) and removing vessels in flow quantification techniques that use phase. Finally, this approach can also be used to set the voxels determined as noise to have a high signal rather than zero, such as a maximum value available to the system. Then, when a series of SWI data is evaluated using a mIP method, the noise in one slice no longer causes the removal of regions having signal in other slices. This is particularly valuable near the top of the brain where the head narrows and eventually disappears. In this example, even if an image is completely noisy, the algorithm will prevent the failure of the mIP operation.

The use of phase information can be important not only for gradient echo methods such as SWI but also for spin echo methods if used as additional information for removing noise points. The CTM along with a connectivity criterion has been shown to give excellent results for an SNR as low as 3:1. The most prominent advantage of the CTM method is its easy implementation and reasonable computational time. The use of the phase information has proven successful in removing noise that would otherwise not be recognized if only noise in the magnitude image had been considered. Practically, using this approach makes it possible not only to improve the magnitude image but also to make phase images appear more like magnitude images. The noise in the phase images leads to wild swings in the phase values, causing a visually unappealing appearance in the phase images and difficulty in adjustments of window level settings. Using the complex threshold technique to suppress the noise, especially in the phase images used in susceptibility weighted imaging, makes phase data more easily viewable. The CTM method is robust with a limited set of connectivity parameters serving for a low SNR of 3:1 and another set for a broader range of SNRs of 5:1 and higher.

Magnitude and Phase Thresholding Method

The MAPHT method was successfully applied to both simulated and human MRI data. It was shown to perform well in separating voxels that contain tissue signal plus noise from voxels that contain predominantly noise. The MAPHT thresholding procedure is based on a likelihood ratio test and its theoretical statistical properties were shown through Monte Carlo simulation in terms of both false positives and false negatives. The MAPHT method for thresholding complex-valued magnetic resonance images was successfully applied to a simulated data set and three human data sets and shown to produce increased tissue contrast by eliminating false positives. It was seen that when the SNR is low, voxels cannot be reliably separated by only the magnitude data because the histogram of predominantly noise voxels overlaps with the histogram of tissue plus noise voxels. However, when the SNR was low, the MAPHT method produced an F statistic that could reliably separate the predominantly noise voxels from the tissue plus signal voxels. It was found that a Bonferroni threshold that is corrected for multiple comparisons may be too conservative. As an alternative to a Bonferroni threshold, an FDR threshold can be used as it is less conservative [29, 30].

The CTM approach to the elimination of voxels that contain predominantly noise utilized not only complex thresholding but also local connectivity to enhance suppression of noise or prevent the incorrect assignment of signal to noise in order to reduce Type I errors. The MAPHT method does not use connectivity but a local neighborhood of voxel values and looks at local variance on a voxel by voxel basis. The CTM approach may suffer when the

phase itself deviates from zero or any set zero point. This can lead to the suppression of areas where the phase is offset from flow effects or susceptibility effects (the later being acceptable in some cases when the goal is to also suppress veins). However, the phase has a rapid variation near the air/tissue interfaces, and in these areas the CTM might fail and that part of the brain will be suppressed. It is not so with the MAPHT method that uses the local variance and keeps the local value as the phase offset (i.e., there is not a global offset of zero in phase). For this reason, this new approach is more robust to variations in phase caused by unwanted field inhomogeneity effects. This method can also easily be adapted for 3D images.

In conclusion, two magnitude and phase thresholding procedures were described and successfully applied to both simulated and human images. Both the CTM and MAPHT methods were shown to produce increased tissue contrast by eliminating false positives. The joint use of magnitude and phase images improves the removal of noise voxels in magnetic resonance images. This can be useful in automated visualization of phase images without the highly distractive noise voxels or in susceptibility weighted imaging when taking the minimum intensity projections of variably sized regions.

REFERENCES

1. Rice SO. Mathematical analysis of random noise. *Bell Syst. Tech. J.* 1944;23:283.
2. Edelstein WA, Glover GH, Hardy CJ, Redington RW. The intrinsic signal-to-noise ratio in NMR imaging. *Magn. Reson. Med.* 1986;3:604–618.
3. Henkelman RM. Measurement of signal intensities in the presence of noise in MR images. *Med. Phys.* 1985;12:232–233.
4. Bernstein MA, Thomasson DM, Perman WH. Improved detectability in low signal-to-noise ratio magnetic resonance images by means of phase corrected real reconstruction. *Med. Phys.* 1989;16:813–817.
5. Gudbjartsson H, Patz S. The Rician distribution of noisy MRI data. *Magn. Reson. Med.* 1995;34:910–914.
6. Andersen AH. On the Rician distribution of noisy MRI data. *Magn. Reson. Med.* 1996;36:331–333.
7. Sijbers J, Den Dekker AJ, Van Audekerke J, Verhoye M, Van Dyck D. Estimation of noise in magnitude MR images. *Magn. Reson. Imaging* 1998;16:87–90.
8. Sijbers J, Poot D, den Dekker AJ, Pintjens W. Automatic estimation of the noise variance from the histogram of a magnetic resonance image. *Phys. Med. Biol.* 2007;52:1335–1348.
9. Chang LC, Rohde G, Pierpaoli C. An automatic method for estimating noise-induced signal variance in magnitude-reconstructed magnetic resonance images. *SPIE Med. Imaging: Image Process.* 2005;5747:1136–1142.
10. Rowe DB Logan BR. A complex way to compute fMRI activation. *Neuroimage* 2004;23:1078–1092.
11. Constable RT, Henkelman RM. Contrast, resolution and detectability in MR imaging. *J. Comput. Assist. Tomogr.* 1991;15:297–303.
12. Hendrick RE, Haacke EM. Basic physics of MR contrast agents and maximization of image contrast. *J. Magn. Reson. Imaging* 1993;3:137–148.
13. Haacke EM, Brown RW, Thompson MR, Venkatesan R. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*, Wiley, New York, 1999.

14. Madore B, Henkelman RM. A new way of averaging with applications to MRI. *Med. Phys.* 1996;23:109–113.
15. Macovski A. Noise in MRI. *Magn. Reson. Imaging* 1996;38:494–497.
16. Nowak R.D. Wavelet based Rician noise removal for magnetic resonance imaging. *IEEE Trans. Image Process.* 1999;8:1408–1419.
17. Sijbers J, den Dekker AJ, Van Der Linden A, Verhoye M, Van Dyck D. Adaptive anisotropic noise filtering for magnitude MRI data. *Magn. Reson. Imaging* 1999;17:1533–1539.
18. Lysaker M, Lundervold A, Tai XC. Noise removal using fourth-order partial differential equation with applications to medical magnetic resonance images in space and time. *IEEE Trans. Image Process.* 2003;12:1579–1589.
19. Chen B, Hsu EW. Noise removal in magnetic resonance diffusion tensor imaging. *Magn. Reson. Med.* 2005;54:393–407.
20. Cline HE, Dumoulin CL, Lorensen WE, Souza SP, Adams WJ. Volume rendering and connectivity algorithms for MR angiography. *Magn. Reson. Med.* 1991;18:384–394.
21. Lin W, Haacke EM, Maaryk TJ, Smith AS. Automated local maximum-intensity projection with three-dimensional vessel tracking. *J. Magn. Reson. Imaging* 1992;2:519–526.
22. Rowe DB, Haacke EM. MAGnitude and PHase Thresholding (MAPHT) of noisy complex-valued magnetic resonance images. *Magn. Reson. Imaging* 2009;27(9):1271–1280.
23. Reichenbach JR, Venkatesan R, Schillinger DJ, Kido DK, Haacke EM. Small vessels in the human brain: MR venography with deoxyhemoglobin as an intrinsic contrast agent. *Radiology* 1997;204:272–277.
24. Haacke EM, Xu Y, Cheng YCN, Reichenbach J. Susceptibility weighted imaging (SWI). *Magn. Reson. Med.* 2004;52:612–618.
25. Rowe D.B., Meller C.P., Hoffmann R.G. Characterizing phase-only fMRI data with an angular regression model. *J. Neurosci. Methods* 2007;161:331–341.
26. Pandian D, Ciulla C, Haacke EM, Jiang J, Ayaz, M. A complex threshold method for identifying voxels that contain predominantly noise in magnetic resonance images. *J. Magn. Reson. Imaging* 2008;28:727–735.
27. Hogg RV, Craig AT. *Introduction to Mathematical Statistics*, 4th edn, Macmillan, New York, 1978.
28. Bain LM, Engelhardt M. *Introduction to Probability and Mathematical Statistics*, 2nd edn, PSW-Kent, Boston, MA, 1992.
29. Logan BR, Rowe DB. An evaluation of thresholding techniques in fMRI analysis. *Neuroimage* 2004;22:95–108.
30. Logan BR, Geliaczova MP, Rowe DB. An evaluation of spatial thresholding techniques in fMRI analysis. *Hum. Brain Mapp.* 2008;29:1379–1389.

UNCORRECTED PROOF