

Continuous MRI Coil Quality Control Using Clinical Imaging Data

by

Trevor McKeown

Graduate Program in Medical Physics  
Duke University

Date: \_\_\_\_\_

Approved:

\_\_\_\_\_  
Ehsan Samei

\_\_\_\_\_  
Scott Robertson

\_\_\_\_\_  
Christopher Lascola

Thesis submitted in partial fulfillment of  
the requirements for the degree of  
Master of Science in the Graduate Program in  
Medical Physics in the Graduate School  
of Duke University

2022

ABSTRACT

Continuous MRI Coil Quality Control Using Clinical Imaging Data

by

Trevor McKeown

Graduate Program in Medical Physics  
Duke University

Date: \_\_\_\_\_

Approved:

\_\_\_\_\_  
Ehsan Samei

\_\_\_\_\_  
Scott Robertson

\_\_\_\_\_  
Christopher Lascola

An abstract of a thesis submitted in partial  
fulfillment of the requirements for the degree of  
Master of Science in the Graduate Program in  
Medical Physics in the Graduate School of  
Duke University

2022

Copyright by  
Trevor McKeown  
2022

# **Abstract**

## **Purpose**

RF coil technology has evolved significantly over recent years. The number of individual coil elements has steadily increased to achieve higher image quality and reduce acquisition times. However, the advantages associated with additional coil elements come with a great deal of added complexity and potential for failures. Over time coils degrade from near continuous use and occasional mishandling, causing elements to become faulty or even die. Thanks to the robustness in modern MR image reconstruction faulty coil elements may not immediately present themselves in clinical images. Once a sufficient number of elements have failed or a sufficient degree of degradation has occurred, diagnosis becomes impaired. In order to avoid such occurrences, RF coils undergo regular physics testing. This testing is time consuming and requires dedicated scanner time. Consequently, each RF coil is typically only tested once per year. This project serves to create an automated quality control process that tracks the performance of individual coil elements across a fleet of MRI scanners on a daily basis using raw clinical MRI data.

## **Methods**

Raw MR imaging datasets were automatically collected from all localizer scans and transferred to a network drive using the Yarra Raw Data Storage (RDS) client. Each

of these raw localizer scans were automatically processed by a separate computer. This processing computer reconstructed images from the raw k-space data, segmented regions of signal from noise, and calculated the signal to noise ratio (SNR) of each element. The calculated SNR was corrected for differences in acquisition parameters such that it could be trended across patients, scanners, and over time. SNR thresholds were derived from vendor-supplied phantom tests to provide pass/fail criteria based on the clinical MRI data.

### **Results**

MR images were successfully reconstructed across a variety of acquisition protocols using various acceleration techniques. Images were accurately segmented across a wide range of anatomy without any physicist or technician input. The SNR values were consistent over time, except in the case of failure. Clinical coil SNR thresholds identified three cases of element failure. These failures were confirmed by performing vendor-recommended phantom testing.

### **Conclusion**

This work demonstrates that faulty coil elements can be detected directly from raw clinical imaging data rather than requiring dedicated coil testing. Catching these failures early allows for replacement parts to be ordered and prevents patients from needing to be rescheduled. This automated process also requires no additional MR data, which saves valuable staff and scanner time.

# Contents

Abstract .....	iv
List of Tables .....	viii
List of Figures .....	ix
1. Introduction .....	1
1.1 Radio frequency coils .....	1
1.2 Clinical impact of faulty RF coils .....	3
1.2.1 Clinical scans most affected .....	3
1.3 Advantages to using clinical data for quality control .....	4
1.4 Aim of this work .....	5
2. Data workflow .....	6
3. Image reconstruction .....	7
3.1 MRI data .....	7
3.2 Reconstruction methods .....	8
3.2.1 Reconstructing images acquired with multiple coils .....	9
3.2.1.1 Partial Fourier .....	10
3.2.1.2 Parallel imaging .....	11
4 Image segmentation .....	13
4.1 Segmenting sum of squares image .....	13
4.2 Segmenting individual coil sensitivity images .....	14
4.3 Segmenting individual coil images .....	16

4.4 Mask smoothing .....	18
4.5 Validation of segmentation across various anatomical regions .....	19
5. Quantifying coil performance .....	20
5.1 SNR calculation.....	20
5.2 Correcting SNR for acquisition parameters.....	21
5.3 Defining action thresholds .....	22
6. Verification of coil failures.....	23
7. Limitations .....	24
7.1 Limitations in data workflow .....	24
7.2 Limitations in thresholds.....	24
8. Future work .....	25
9. Conclusion .....	26
References .....	27

## List of Tables

Table 1: Distribution of data collection .....	7
--	---



## List of Figures

Figure 1: Anatomy specific coils in clinical use .....	2
Figure 2: Workflow .....	6
Figure 3: Methods of data acquisition and resulting images:.....	9
Figure 4: Coil image, sum-of-squares image, coil sensitivity image.....	10
Figure 5: Sum of Squares image segmentation .....	14
Figure 6: Sensitivity and standard deviation filtered sensitivity images.....	15
Figure 7: Segmentation based on the standard deviation filtered sensitivity image of a single coil element.....	16
Figure 8: Combination of sum-of-squares, standard deviation filtered sensitivity image, and individual coil channel segmentation .....	17
Figure 9: Final segmentations before and after smoothing.....	18
Figure 10: Different anatomical regions segmented during this project.....	19
Figure 11: Noise distributions in real, imaginary, and magnitude images .....	20
Figure 12: SNR corrections applied to phantom data.....	21
Figure 13: SNR of individual coil elements over time showing dead elements before and after replacement.....	22
Figure 14: SNR correlation between clinical and phantom scans, error bars showing quartile ranges .....	23

# 1. Introduction

## 1.1 *Radio frequency coils*

In MRI, radio frequency (RF) coils perform two primary functions: 1) transmitting RF energy into the patient to excite the protons in water atoms and 2) receiving RF energy that is emitted from the patient such that the signal can be digitized and recorded. Typically, RF energy is transmitted into the patient using a large body coil that is integrated into the bore of the magnet. This body coil is designed to have a homogenous transmit field so that tissue is uniformly excited. However, the body coil is large in diameter, which makes it fairly insensitive to signals being imaged in the center of the bore. Placing a coil closer to the patient is an effective way to boost image SNR. To that end, RF coils have been designed in many shapes and sizes to conform to various anatomy. A typical MRI scanner has on the order of 10 separate RF coils for various types of imaging scenarios.

Modern receive RF coils commonly employ multiple coil elements. Individual coils with 8-64 coil elements are routinely used in clinical practice. Research into increasingly higher number of coil elements continues, with active research in up to 256 coil elements [Hendriks 2018]. One advantage to having multiple small coil elements instead of a single large coil element is that the smaller coil elements provide higher SNR. Another advantage is that image acquisition can be accelerated through techniques

such as parallel imaging (SENSE [Pruessmann 1999], GRAPPA [Griswold 2002]), compressed sensing imaging [Donoho 2006], and Partial Fourier imaging [Margosian 1986].



**Figure 1: Anatomy specific coils in clinical use. Clockwise from top left: transmit/receive head coil, receive only head coil, foot/ankle coil, and flexible body matrix coil**

## **1.2 Clinical impact of faulty RF coils**

One consequence of having so many RF coils is that technologists frequently must swap coils between patients. Even with extra care by technologists, this process occasionally results in coils being bent, twisted, bumped and even dropped. Over time this can cause issues in electronics, wiring, and plug connections. As the damage accumulates, eventually the coil's performance becomes affected. Such degradation can be immediate or result in a slow loss in SNR over time. In the latter case, physicians and technologists may grow accustomed to the image quality as it degrades, masking the root issue. However, eventually, a perceptible loss of image quality will prompt an investigation in the coil. Once a faulty coil is reported, the coil must undergo testing, then either repair or replacement, which can take around a week. Since these coils can cost tens of thousands of dollars, there is not always a spare coil available, which can disrupt the clinical workflow.

### **1.2.1 Clinical scans most affected**

The SNR in MRI images is spatially dependent upon the coil performance and patient positioning. Faulty coils will reduce the SNR in regions corresponding the location of the coil. These effects are most severe in scans that utilize suppression of signal or highly accelerated sequences. It is critical in these protocols to ensure that coil elements are functioning properly to allow for accurate diagnoses of patients.

### ***1.3 Advantages to using clinical data for quality control***

Currently, the ACR requires coils to be tested on a yearly basis for appropriate SNR, signal uniformity, and degree of ghosting using dedicated phantoms. This is far too infrequent to catch faulty coil elements within a reasonable timeframe. However, more frequent phantom testing would impose a large time burden on the staff and would require dedicated scanner time for each additional scan. Typically, physics tests are scheduled in advance or performed at non-peak times to minimally affect patient care. At Duke, coil elements are evaluated individually for each RF coil. Collectively, it takes between 2 and 3 hours to test all the RF coils on a single MRI system. As a result of this infrequent testing, most faulty coils are first reported by physicians or technologists noticing a change in image quality.

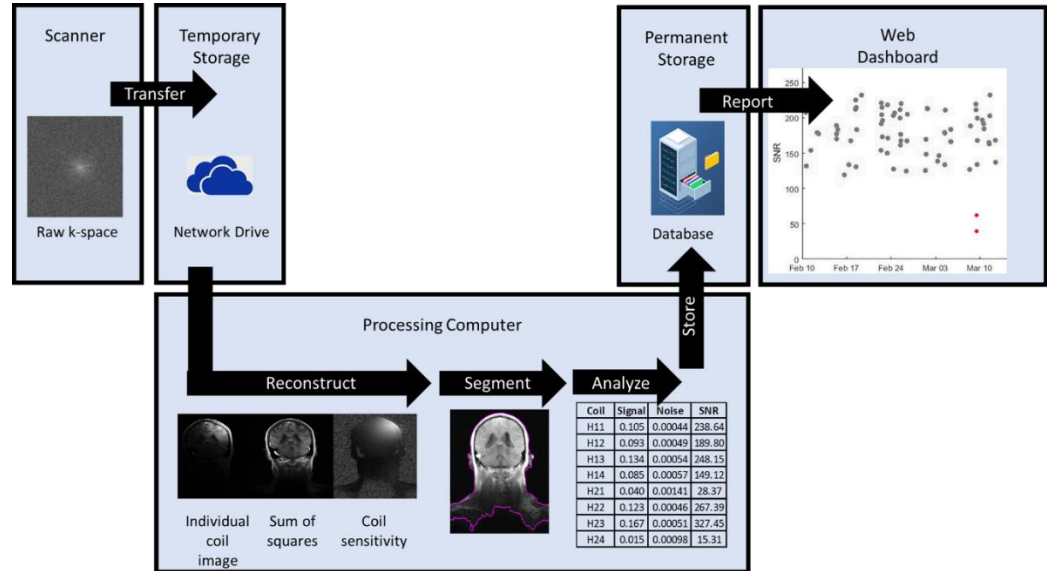
There are several benefits to using clinical data for quality control purposes. First is the continuous availability of data. Multiple RF coils are used every day, providing a testing frequency that naturally matches the clinical volume of cases. Often a given coil is used multiple times in a single day, which provides more frequent measurements than could be practically obtained through phantom testing. Additionally, this data does not require any additional testing. Finally, it provides an opportunity for clinical image quality review.

The primary disadvantage to using clinical data for evaluating RF coil performance is the inherent messiness of clinical data. In clinical data, there exists additional sources of variability in SNR measurements, such as patient size, coil positioning, and patient compliance. This variability would be significantly reduced in phantom testing, which allows for tighter tolerances to be made in vendor supplied testing. Such testing uses phantoms that provide uniform signal that is stable year to year. Typically, vendors request performing these phantom tests before any service tickets are opened. Though individual measurements of SNR will be more variable when obtained from clinical rather than phantom data, the higher frequency of clinical measurements allows them to be averaged over time, which inherently reduces variability.

#### ***1.4 Aim of this work***

This work sought to analyze the localizer scans that are taken at the beginning of each MRI exam to assess and trend individual coil element performance. In this way, no additional QC scans or phantom setups are required for technologists or physicist.

## 2. Data workflow



**Figure 2: Workflow**

Figure 2 illustrates the overall workflow of the automated QC program. The raw MRI data from each localizer scan are collected and transferred to a network drive for temporary storage using the Yarra Raw-Data Service (Yarra Client version 0.6, Center for Advanced Imaging Innovations and Research). This data transfer is scheduled to automatically take place after the MRI scanner first boots up in the morning and before the clinical workflow starts up. The transfer typically takes less than 10 minutes. Each individual localizer scan contains roughly 0.5 Gb of imaging data, and a typical day will transfer roughly 15 Gb of data per scanner. Given the large size of data, only localizer scans are currently being transferred, though in theory any clinical scan could be analyzed.

The localizer scans are then accessed via the network drive from another computer that automatically processes the data. This processing is comprised of three steps: 1) image reconstruction, 2) signal and noise segmentation, and 3) SNR measurement. These steps are described in the following sections. Once each dataset is processed, the results are stored in a database which can be viewed by a web-based dashboard. The raw imaging data are then deleted from the network to free up storage space. Table 1 summarizes the cumulative amount of raw clinical datasets that have been processed to date.

**Table 1: Distribution of data collection**

Scanner	Unique Coils Assessed	Total Exams Assessed
Siemens 1.5T Aera	3	38
Siemens 1.5T Aera	7	895
Siemens 3T Skyra	10	606
Siemens 3T Skyra	7	156
Siemens 3T Skyra	6	66

### **3. Image reconstruction**

#### **3.1 MRI data**


In MR, spatial data is not collected directly as it is in several other imaging modalities. Instead, spatial frequencies are measured, then and inverse Fourier



Transform is used to transform the spatial frequencies into their spatial domain representation, the reconstructed image. Acquiring spatial frequency data takes time, so much effort goes into optimizing the quantity of spatial frequency data required to produce an image. The order and pattern that these spatial frequencies are acquired therefore serve as important design constraints to each MR imaging protocol. Our code is designed to be flexible to various acquisition strategies. It is important to note that the goal of this project was to be able to robustly handle the various clinical acquisition strategies, rather than use the most advanced reconstruction methods. Speed is also an important component of this project, due to the large number of scans performed each day, our code needs to be able to keep up with the influx of clinical data.

### ***3.2 Reconstruction methods***

The following sections will describe common acquisition methods and their corresponding reconstruction techniques. For a single element coil, the simplest, but most time-consuming acquisition is to fully acquire k-space frequencies up to the Nyquist limit in each dimension. This is illustrated on the upper left side of Figure 3 on the next page. The image-domain representation of the coil channel can then be reconstructed by the Fourier transform. An example image is shown on the bottom left of Figure 3. The fully sampled acquisition represents a useful baseline to compare various acceleration strategies.



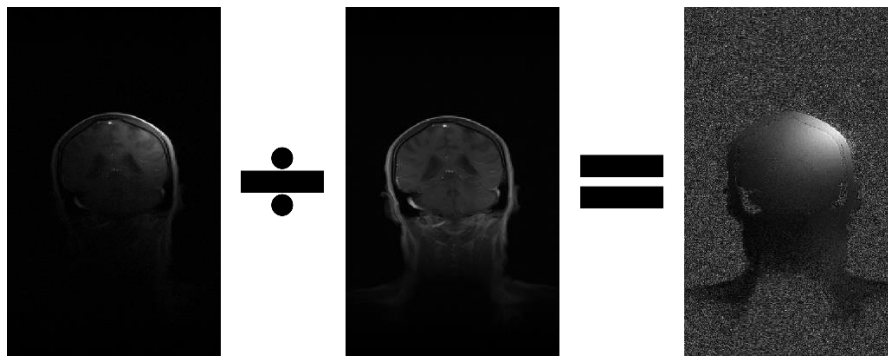
**Figure 3: Methods of data acquisition and resulting images: full k-space collection (left) partial Fourier collection (middle) parallel imaging (right)**

### 3.2.1 Reconstructing images acquired with multiple coils

Once the k-space data from each coil has been reconstructed, the set of element images need to be merged into a single combined image. This is necessary because each individual coil image is shaded, and it would not be practical for clinicians to have to scroll through the data from each coil. The shading occurs because coils are most sensitive to anatomy nearest the coil, and the sensitivity falls off with increasing distance. Mathematically, we can describe the shading as a multiplicative factor that scales the combined image into the shaded coil image. The simplest way to combine the individual coil channels is to perform a sum-of-squares reconstruction. The equation for this method is listed below.

$$S(r) = \sqrt{\sum_{i=1}^{N_c} S_i^*(r)S_i(r)}$$

As is illustrated in the center image of Figure 4 below, the sum-of-squares reconstruction has more uniform intensities in the body than the individual coil image shown on the left. Dividing each coil image by the sum-of-squares provides an estimate of the coil sensitivity, shown by the rightmost image in Figure 4. The intensities seen in the coil sensitivity images are similarly shaded to those in the original coil image, however the intensity variations due to anatomy are eliminated. Smoothing, masking, and additional processing can be applied to the coil sensitivity to make a more accurate approximation to the true coil sensitivity.



**Figure 4: Coil image, sum-of-squares image, coil sensitivity image**

### 3.2.1.1 Partial Fourier

In Partial Fourier acquisition, illustrated on the middle of Figure 3, only an asymmetric portion of k-space is acquired, leaving a percentage of k-space uncollected. This uncollected space can be simply left as zeros, a form of reconstruction called zero-padding. Zero-padding introduces a small amount of blurring in the phase encode direction, but still accurately reflects the measured data. Other algorithms exist that estimate the frequencies of uncollected data. Typically, these algorithms first estimate

the phase of the image from the symmetric central region of collected data. This symmetric region only consists of lower spatial frequencies, but since phase typically varies slowly across the image, the low spatial frequencies provide a decent estimate of image phase. After approximating image phase, the phase is removed from the imaging data, which then allows the reconstruction to exploit a mathematical relationship, Hermitian symmetry, that exists in k-space where there is no phase in the image. Hermitian symmetry allows the missing frequencies to be estimated using the measured frequencies at positions that are mirrored on the opposite sides of k-space. Once these frequencies are estimated, the phase can be reincluded to provide the final reconstruction. In our initial implementation, we are only using zero-padding for simplicity.

### **3.2.1.2 Parallel imaging**

Parallel imaging uses a similar concept as Partial Fourier where in order to speed up acquisition, certain spatial frequencies are not collected. However instead of leaving one edge of k-space uncollected, phase encoding lines are acquired in a pattern leaving other phase encoding lines empty. Depending on the level of acceleration, every other line may be collected or only every fourth line of k-space is collected. The consequence of regularly sampling k-space phase encodes is image aliasing, which manifests as repeated overlapping images that occur in a predictable pattern.

The SENSE algorithm uses the predictable pattern of aliasing as well as the spatially varying sensitivity profiles to reconstruct an unaliased image. The sensitivity images are estimated from a fully-sampled subset of k-space. This fully sampled region lies in the low spatial frequencies in the center of k-space, thus the images are lower resolution. Coil sensitivities tend to be slowly varying, so the sensitivities are typically well-represented by this subset of frequencies. We take the Fourier transform of this data to get the low resolution unaliased images, calculate a sum of squares image, and from this calculate sensitivity profiles. The SENSE algorithm then takes the sensitivity profiles and aliased images and solves for the unaliased image using a least squares method.

## 4 Image segmentation

Once all images are reconstructed, regions of signal and noise must be segmented such that the SNR can be calculated for each coil channel. Segmenting MR images is challenging due to the high variability of signal intensity that is present in the body and also the shading imposed by the natural coil sensitivities. To create a robust segmentation, our code uses four steps to generate a final segmentation.

### ***4.1 Segmenting sum of squares image***

To generate a rough initial segmentation, Otsu's method is applied to the sum-of-squares image to identify a single threshold that identifies regions of significant signal. Otsu's method calculates a threshold intensity that separates the intensities into two different groups with the lowest intra-class variance. Once this threshold is calculated, all pixels above this value are classified as regions of signal as illustrated in the figure below. This segmentation well-identifies regions of high intensity, however in regions where there is shading such as the shoulder and lower neck in Figure 5, the intensities are below the intensity threshold and are missed.





Figure 5: Sum of Squares image segmentation

## ***4.2 Segmenting individual coil sensitivity images***

To segment regions of signal that have inherently low intensity, we rely on the variability of coil sensitivity images. As shown below on the left side of Figure 6, coil sensitivity images have little noise within the body, but are extremely noisy outside of the body. The extreme noise in the air regions originates from the calculation of the sensitivity images, which are derived by dividing the coil image by the sum-of-squares image. In the air, the intensity of both the coil image and sum-of-squares image are near zero, making their division an ill-defined operation that causes very noisy results. Our segmentation exploits the highly varying signal in the air by using a standard deviation

filter that calculates the standard deviation of local intensities throughout the image. The standard deviation filtered sensitivity image is shown on the right side of Figure 6.



**Figure 6: Sensitivity and standard deviation filtered sensitivity images**

The variability of the sensitivity image is significantly lower within the body than in the surrounding air, with one exception. At the edge of the body where the coil is most sensitive there is an abrupt transition from the bright coil sensitivity to the lower signal air region. This sharp edge causes the local standard deviation to be high. However, in other coil elements, the location of the sharp edge tends to differ according to the position of the other coils. To segment regions of signal from noise, we apply Otsu's threshold to the standard deviation filtered sensitivity image of each coil as



illustrated in Figure 7. Then we combine the segmentation masks across all coils with a binary union operator.




**Figure 7: Segmentation based on the standard deviation filtered sensitivity image of a single coil element**

### ***4.3 Segmenting individual coil images***

Since the segmentation of the standard deviation filtered sensitivity images misses regions nearest to the coil, we added an additional thresholding step to segment individual coil images. This augments the segmentation in regions where the coil images are brightest. Using the combined segmentation of the sum-of-squares and standard deviation filtered sensitivity images, we identify all noise-like pixels that are not within the mask. The noise-like pixels are separated into real and imaginary components, then


the mean and standard deviation are calculated to characterize the noise in each coil channel. Finally, an intensity threshold is defined as the mean of noise pixels plus five standard deviations in the magnitude image. The value of five standard deviation above the mean considers that the real and imaginary channels combine in quadrature. With two gaussian distributions adding in quadrature, well over 99 percent of noise pixels would be expected to have a magnitude that falls below this threshold. In this way, our threshold is very insensitive to picking up noise pixels, but is still able to identify regions of signal, even if they are low intensity. The combination of the sum-of-squares, standard deviation filtered sensitivity image, and individual coil channel segmentation is shown in Figure 8.



**Figure 8: Combination of sum-of-squares, standard deviation filtered sensitivity image, and individual coil channel segmentation**

#### **4.4 Mask smoothing**


As a final processing step, we eliminate noise pixels and rough edges by performing a binary erosion followed by a binary dilation. Figure 9 below compares the effects of erosion and dilation. The resulting segmentation closely follows the body contours, even in shaded regions.



**Figure 9: Final segmentations before and after smoothing**

## **4.5 Validation of segmentation across various anatomical regions**

Our segmentation method has proven to be remarkable robust across a variety of RF coils and different patient anatomy. Figure 10 below illustrates several anatomical regions that have been segmented by this code. The magenta outline represents the edges of the segmentation. Even hypointense regions are accurately segmented within the body, which is an advantage to using the joint intensity and coil sensitivity information while segmenting.



**Figure 10: Different anatomical regions segmented during this project listed from left to right. Top: hand, head, ankle. Bottom: breast, knee, shoulder**

## 5. Quantifying coil performance

### 5.1 SNR calculation

The SNR calculations made in this project are determined by dividing the highest local mean value by the standard deviation of the combined real and imaginary values within the noise region. The noise is calculated from the real and imaginary components because they will both have a Gaussian distribution centered around a mean of zero. We chose to not use the noise from the magnitude images because the magnitude operation causes all pixel intensities to be non-negative. This skews the noise distribution rightward into a more complicated Rician distribution that does not have a zero mean [Gudbjartsson 1995]. Figure 11 illustrates the differences in noise distributions between real/imaginary and magnitude data.

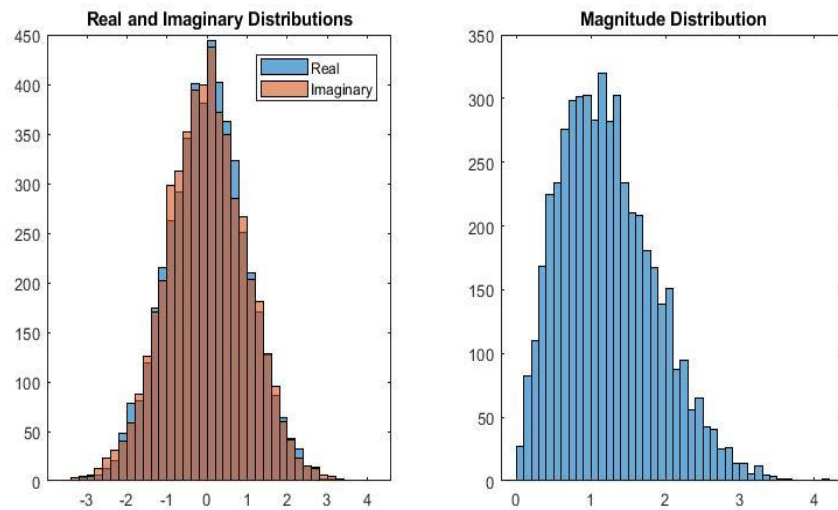


Figure 11: Noise distributions in real, imaginary, and magnitude images

## 5.2 Correcting SNR for acquisition parameters

The SNR calculated from an image is representative not only of the intrinsic SNR of the coil and imaging hardware, but also a function of the acquisition parameters. To minimize discrepancies between different protocols, the SNR was adjusted for the voxel volume and total digitization time. Through phantom studies, we have found that there exists a linear relationship between SNR and the voxel volume. There is also a linear relationship between SNR and the square root of the time that the scanner is collecting data which will be referred to as  $T_{\text{adc}}$ . In this work we adjusted all SNR values to reference acquisition parameters (1 mm<sup>3</sup> voxel volume, matrix size of 256 x 256, and a receiver bandwidth of 25 kHz, giving a  $T_{\text{adc}}$  value of 1.31 s). Figure 12 shows phantom measurements done to verify the corrections.

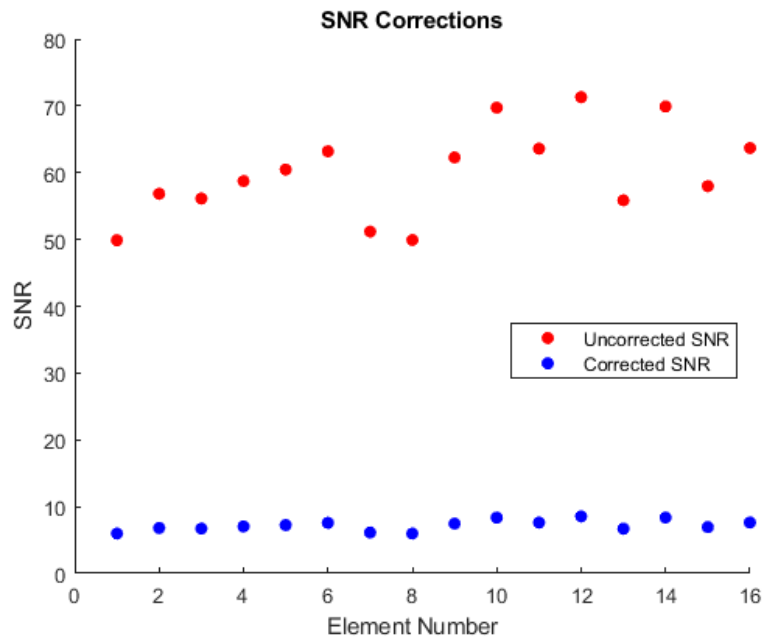


Figure 12: SNR corrections applied to phantom data

### 5.3 Defining action thresholds

Analyzing the trends of SNR values for each coil element, threshold values were empirically determined in order to make sure that the system was sensitive to not only complete coil failure but also to coil degradation over time. These thresholds were then assigned to each individual coil element for the purpose of notifying physics when the SNR values fell below the threshold that were not only completely dead but also that had begun to drop in performance.




Figure 13: SNR of individual coil elements over time showing dead elements before and after replacement.

## 6. Verification of coil failures

This project identified three coils with low SNR elements. In order to test the accuracy of this project in identifying the faulty elements, the vendor specified coil element scan was performed. This vendor specific test serves as the gold standard in evaluating coil performance and in determining when a replacement is justified. In all three instances the coils that were identified as failures using our automated QC were also identified as failures on the vendor tests. The following correlation plot shows that our measurements of SNR correlate well to the SNR calculated in the vendors phantom-based test.

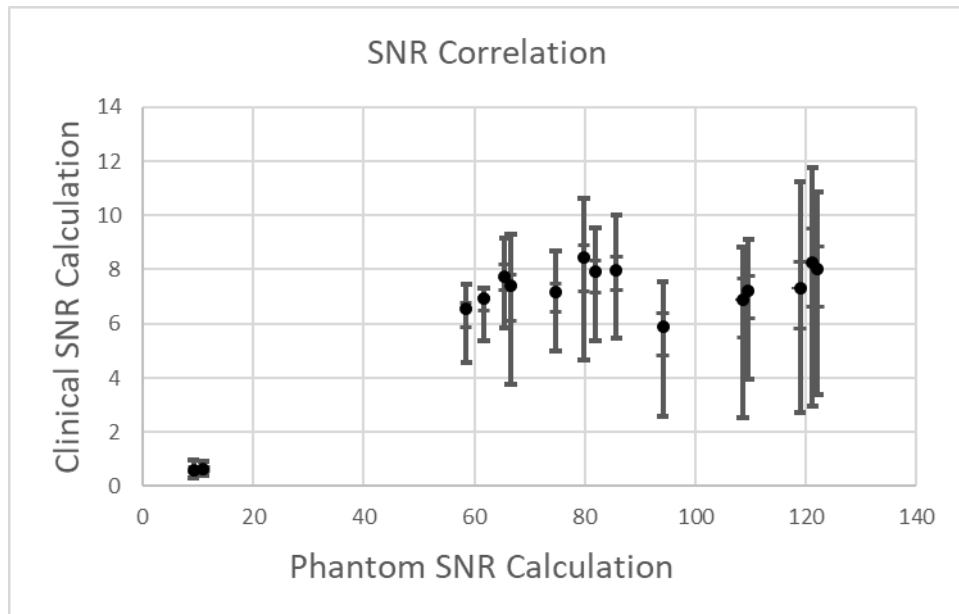


Figure 14: SNR correlation between clinical and phantom scans, error bars showing quartile ranges



## **7. Limitations**

### ***7.1 Limitations in data workflow***

Currently the data transfer and reconstruction only support Siemens MRI scanners. A clear future step would be to extend the code to support other vendors. Unfortunately, Yarra only supports Siemens MRI data, so additional effort would be needed to transfer raw data from other scanners. Also, the raw data formats are proprietary to each vendor, so there would be a learning curve to reconstruct data from other manufacturers.

### ***7.2 Limitations in thresholds***

Another limitation of this work is the empirically defined thresholds currently being used. One potential solution to this would be use a plot similar to the one shown in Figure 14 that compare the SNR from phantom scans to clinical scans. Once the relationship between the two is determined the thresholds that are set by the vendor in the phantom scans could be converted into thresholds that are relevant to clinical imaging scans. Additionally, it would be better to establish action thresholds during acceptance testing when the coils are functioning at peak performance.

## 8. Future work

The robustness of this work to handle a wide range of anatomy and image acquisition methods allows for many unique applications. While currently our tool is used to analyze coil performance on one system across time, there is the potential to use this tool as a way of comparing SNR of comparable coils across the fleet of MRI machines in use at Duke. This would be particularly useful during acceptance, as it could provide a minimum standard for acceptance.

Another avenue of future work that this work enables is to look at other features of image quality. Presently, we only calculate the maximum SNR, but this tool could additionally be used as a tool to evaluate patient and coil positioning or identify common image artifacts like noise spikes or RF zippers.

## **9. Conclusion**

This quality control program demonstrates that we can accurately measure coil performance from clinical data. This provides performance metrics at a much higher frequency and does not require any additional scanning. This tool has become an essential tool for quickly identifying problems and saving valuable staff and scanner time. To date, the quality control program has detected three coil failures before they were reported by physicians or technologists and the failures were validated using standard phantom-based methods. This tool could be extended into other metrics of clinical image quality, which aligns well with the goals of Medical Physics 3.0.

## References

- Bushberg, J.T., Seibert, J.A., Leidholdt, E.M. and Boone, J.M. (2012) *The Essential Physics of Medical Imaging*. 3rd Edition, Lippincott Williams and Wilkins, Philadelphia.
- Constantinides CD, Atalar E, McVeigh ER. Signal-to-noise measurements from magnitude images in NMR phased-arrays. *Magnetic Resonance in Medicine*. 1997;38:852-857.
- Donoho, D.L. "Compressed Sensing." *IEEE Transactions on Information Theory* 52, no. 4 (April 2006): 1289–1306. <https://doi.org/10.1109/TIT.2006.871582>.
- Firbank MJ, Coulthard A, Harrison RM, Williams ED. A comparison of two methods for measuring the signal-to-noise ratio on MR images. *Physics in Medicine and Biology*. 1999;44:N261-N264.
- Gilbert G. Measurement of signal-to-noise ratios in sum-of-squares MR images. *Journal of Magnetic Resonance Imaging*. 2007;26(6):1678-1678. doi:10.1002/jmri.21171
- Griswold, Mark A., Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. "Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA)." *Magnetic Resonance in Medicine* 47, no. 6 (2002): 1202–10. <https://doi.org/10.1002/mrm.10171>.
- Gudbjartsson, Hákon, and Samuel Patz. "The Rician Distribution of Noisy MRI Data." *Magnetic Resonance in Medicine* 34, no. 6 (December 1995): 910–14.
- Hendriks AD, Luijten PR, Klomp DWJ, Petridou N. Potential acceleration performance of a 256-channel whole-brain receive array at 7 T. *Magn Reson Med*. 2019;81(3):1659-1670. doi:10.1002/mrm.27519
- Henkelman RM. Measurement of signal intensities in the presence of noise in MR images. *Med Phys*. 1985 Mar-Apr;12(2):232-3. doi: 10.1118/1.595711. PMID: 4000083.
- Kaufman L, Kramer DM, Crooks LE, Ortendahl DA. Measuring signal-to-noise ratios in MR imaging. *Radiology*. 1989;173:265-267.
- Lattanzi R, Grant AK, Polimeni JR, et al. Performance evaluation of a 32-element head array with respect to the ultimate intrinsic SNR. *NMR in Biomedicine*. 2010;23(2):142-151. doi:10.1002/nbm.1435

Margosian, Paul, and Franz Schmitt. "Faster MR Imaging Methods." In *Medical Image Processing*, 0593:6–13. SPIE, 1986. <https://doi.org/10.1117/12.952139>.

Montin E, Lattanzi R. Seeking a Widely Adoptable Practical Standard to Estimate Signal-to-Noise Ratio in Magnetic Resonance Imaging for Multiple-Coil Reconstructions. *Journal of Magnetic Resonance Imaging*. 2021;54(6):1952-1964. doi:10.1002/jmri.27816

Price R, Allison J, Clarke G, et al. *ACR Magnetic Resonance Imaging Quality Control Manual 2015*. :126.

Pruessmann, Klaas P., Markus Weiger, Markus B. Scheidegger, and Peter Boesiger. "SENSE: Sensitivity Encoding for Fast MRI." *Magnetic Resonance in Medicine* 42, no. 5 (1999): 952–62.

Sijbers J, den Dekker AJ, van Audekerke J, Verhoye M, van Dyck D. Estimation of the noise in magnitude MR images. *Magnetic Resonance Imaging*. 1998;16:87-90.

Yanasak NE, Kelly MJ. MR Imaging Artifacts and Parallel Imaging Techniques with Calibration Scanning: A New Twist on Old Problems. *RadioGraphics*. 2014;34(2):532-548. doi:10.1148/rg.342135051