Phantom Test Guidance for Use of the Small MRI Phantom for the

MRI Accreditation Program

Contents

Small Phantom Test Guidance for the ACR MRI Accreditation

0.0 INTRODUCTION

0.1 Overview and Purpose

This document provides information about the phantom tests that are part of the American College of Radiology Magnetic Resonance Imaging Accreditation Program.

This document will be useful to facilities wishing to determine whether they will pass the phantom tests prior to data submission, and to facilities wishing to utilize their ACR small MRI phantom for quality control and system performance testing. This document will also enable facilities that have failed the phantom tests to understand the significance of the failure, to take steps to correct it, and to determine whether their corrective actions have been successful.

We begin with an introduction in which we briefly describe the phantom and the image data acquired for the tests, introduce terminology used in referring to the images, and list the tests that constitute the phantom assessment portion of the accreditation process. Then we discuss each test in turn, describing how it is done, giving an acceptance criterion, naming common causes of failure, and offering advice on corrective actions that might be taken.

The acceptance criteria provided here are indicative of a minimum level of performance one can reasonably expect from a well functioning specialty MRI system. However, as minimum levels of performance, these criteria are not to be construed as indicators of typical or normal levels of performance.

Information about the phantom, phantom placement and image acquisition are given here only insofar as they are necessary to the discussion. For a more detailed treatment of these topics see the ACR document called Accreditation Program Testing Instructions for use of Small MR Phantom which is available on the acraccreditation.org website.

0.2 The Phantom

The ACR small MRI phantom is a short, hollow acrylic cylinder of acrylic plastic closed at both ends. The inside length is 100 mm; the inside diameter is 100 mm. It is filled with a solution of nickel chloride and sodium chloride: 10 mM $NiCl_2$ and 0.45% by weight aqueous NaCl. The separate vial is filled with 20 mM $NiCl_2$ but no aqueous NaCl.

Inside the phantom are several structures designed to facilitate a variety of tests of scanner performance. These structures will be described as the tests in which they are used are discussed below.

0.3 The Required Images

The phantom portion of the MRI accreditation program requires the acquisition of a sagittal localizer and four axial series of images. The same set of seven slice locations within the phantom is acquired in each of the four axial series. These images are acquired using the scanner's standard knee coil. The scan parameters for the localizer and the first two axial series of images are fully prescribed by the ACR in the scanning instructions. Therefore, we refer to them as the **ACR sequences** or **ACR images**.

The third and fourth series of axial images are based on your site's own protocols, and are referred to as the **site sequences** or **site images**. To discuss the image data it is convenient to introduce names for the different sets of images and numbering for the slice locations within the phantom.

The localizer is a sagittal 20 mm thick single-slice spin-echo acquisition through the center of the phantom, and is referred to simply as the **localizer**.

The first axial series is a spin-echo acquisition with ACR specified scan parameters that are typical of T1 weighted acquisitions. This series is called the **ACR T1** series.

The second axial series is a spin-echo acquisition with ACR specified scan parameters that are typical of T2 weighted acquisitions. The set of second-echo images from this acquisition is called the **ACR T2** series.

The third and fourth axial series are based on the scan parameters the site normally uses in its clinical protocols for axial knee T1 and T2 weighting respectively. These series are called the **site T1** and **site T2** series.

Each of the axial series has seven required slice locations. The locations are numbered starting at the inferior end of the phantom; so, slice location 1 is at the end of the phantom with the slice thickness and high contrast resolution inserts. Different scanners may number images differently. Regardless of how the scanner numbers the images, we always refer to them by their series name and slice location number. For example, ACR T2 slice 7 is the image of the ACR prescribed T2 weighted acquisition at slice location 7.

For all four axial series the required slice thickness is 5 mm and the slice gap is 3 mm. The set of seven slices spans a distance of 48 mm from the center of the first slice to the center of the last slice. (Some scanners will not allow prescription of 5 mm slices with 3 mm gaps; the *MR Accreditation Program Testing Instructions* explain what to do in those cases.)

Figure 1a shows a sagittal localizer with the various internal structures labeled. Figure 1b shows the same localizer image with the seven axial slice locations cross-referenced on it. Slice 1 is prescribed to be centered on the vertex of the angle formed by the 45° crossed wedges at the indicated end of the phantom and slice 7 aligns with the most superior low contrast resolution disc.

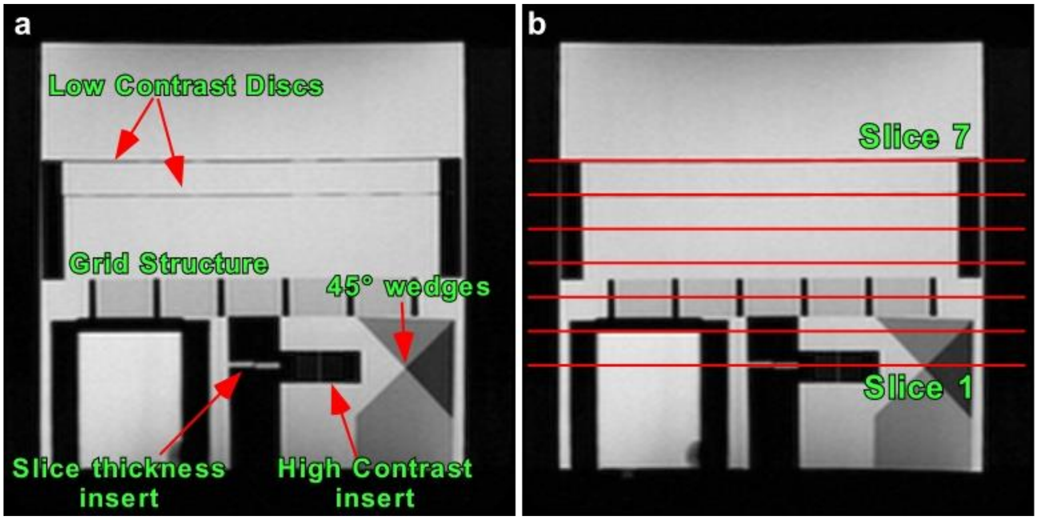

Figure 1. Sagittal localizer showing the phantom's internal structures and the seven required axial slice locations.

The image data must be submitted to the ACR in digital form (DICOM-formatted image files written onto CD) or electronically uploaded. *(Please see the MR Accreditation Program Testing Instructions and the Instructions for Electronic Upload for detailed information on submitting images for accreditation evaluation.)*

0.4 The Image Analysis

For quantitative analysis the digital data are displayed on a computer image workstation capable of the basic image manipulation functions found on all scanner consoles: window and level adjustments, magnification (zoom), mean and standard deviation measurements in a region-of-interest (ROI), and length measurements. In most cases the most convenient place for a facility to make these measurements on their own data will be the scanner console or an associated image review station. However if the images are to be submitted on CD in uncompressed DICOM format or directly on-line, it is recommended to review the images in the form in which they will be submitted in order to replicate the way the data will be analyzed by the ACR reviewer. There are seven quantitative tests made using measurements on the digital data. They are:

1. Geometric accuracy
2. High contrast spatial resolution
3. Slice thickness accuracy
4. Slice position accuracy
5. Image intensity uniformity
6. Percent signal ghosting
7. Low contrast object detectability

Each of these tests will be described in turn in the sections below. The sections are numbered to correspond with the numbering of the tests in this list.

It won't be discussed further here, but the reader should be aware that in addition to these quantitative tests the ACR phantom reviewers examine the images for artifacts. Artifacts that could have an adverse effect on diagnostic accuracy, and artifacts suggestive of system problems that could affect diagnostic accuracy may result in a failure for accreditation even though the system passes the quantitative tests.

1.0 GEOMETRIC ACCURACY

1.1 What It Is

The geometric accuracy test assesses the accuracy with which the image represents known dimensions in the imaged subject. This is also sometimes called the geometric error test. It consists of making length measurements on the images, between readily identified locations in the phantom, and comparing the results with the known values for those lengths.

A failure means that dimensions in the images differ from the true dimensions substantially more than is usual for a properly functioning scanner.

1.2 What Measurements Are Made

Seven measurements of known lengths within the phantom are made using the display station's on-screen length measurement tool. The display window and level settings can affect the length measurements, so it is important to set them properly. For that purpose a separate ancillary procedure for adjusting the display window and level settings is provided following the main procedure.

Measurements for this test are made according to the following procedure:

1. Display the localizer. Adjust the display window and level as described below.
2. Measure the end-to-end length of the phantom as it appears in the localizer. This should be measured along a line near the middle of the phantom (Figure 2).
3. Display slice 1 of the ACR T1 series. Adjust the display window and level as described below.
4. Measure the diameter of the phantom in two directions: top-to-bottom and left-to-right (Figure 3).
5. Display slice 3 of the ACR T1 series. Adjust the display window and level as described below.
6. Measure the diameter of the phantom in four directions: top-to bottom, left-to-right, and both diagonals (Figure 4).

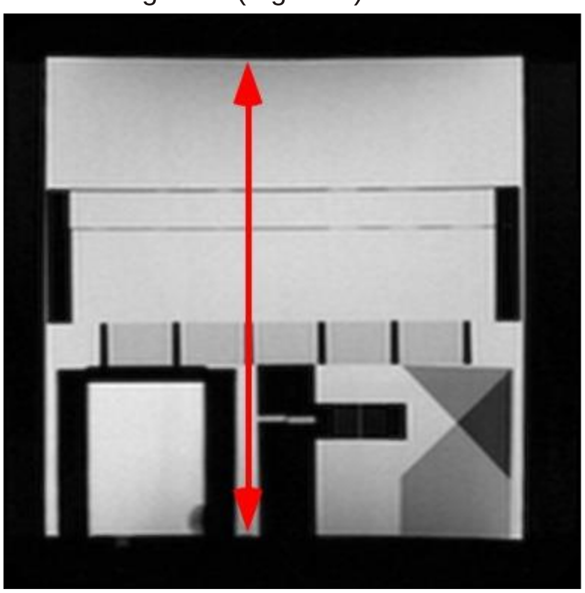

Figure 2. Sagittal localizer with end-to-end measurement illustrated (arrow).

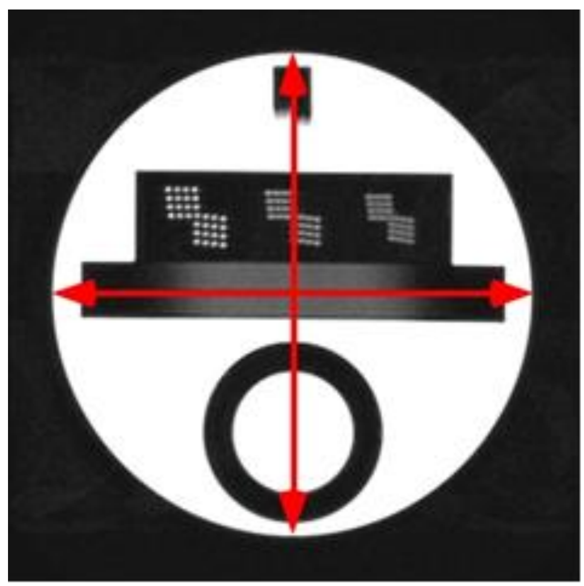

Figure 3. Slice 1 with diameter measurements illustrated (arrows)

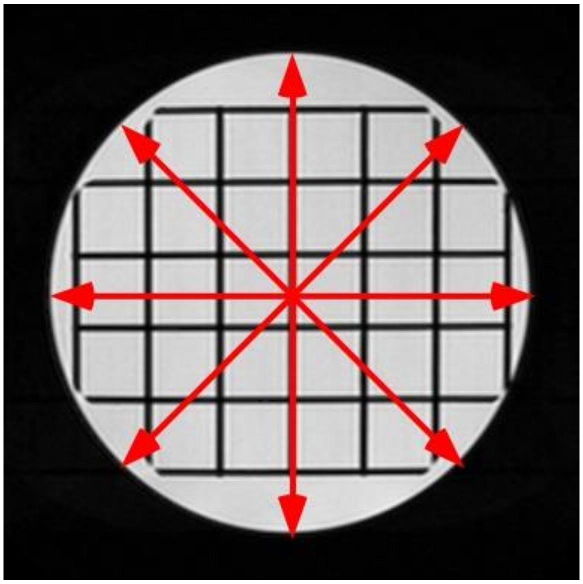

Figure 4. Slice with diameter measurements illustrated (arrows)

Figure 5. Sagittal localizer shown with display window set to zero and level adjusted for appropriate measurement of mean signal level.

The display window and level settings can affect the apparent location of the edges of the phantom and thus cause errors in the length measurements. To avoid this, the measurements should be made with the display window set as narrow as possible and the level set to a value equal to half the mean signal value of the water only regions of the image.

On most scanners the following procedure can be used for setting the display window and level:

1. Adjust the window to its narrowest setting, which is zero or one on most scanners.
2. Observe the regions of the phantom that have only water, i.e., the regions that are not subject to partial-volume effects with some of the internal structures of the phantom. These regions have the highest signal. Lower the display level until the signal in these water-only regions is all white.
3. Raise the display level until about half of the total area of the water-only regions has turned dark. This is illustrated in Figure 5 for the localizer. The level is now set to a numerical value approximating the mean signal of the water-only regions; note this value. (This is really estimating the median signal, but that will be sufficiently close to the mean for our purpose.)
4. Lower the level setting to half of the mean signal value found in step 3. Increase the window width setting to equal the mean signal value.

1.3 How the Measurements Are Analyzed

The length measurements are compared with the known values of the distances in the phantom.

The inside end-to-end length of the phantom is 100 mm.

The inside diameter of the phantom is 100 mm.

1.4 Recommended Action Criteria

All measured lengths must be within ± 2 mm of their true values.

1.5 Causes of Failure and Corrective Actions

The most common cause of failure of this test is one or more miscalibrated gradients. A miscalibrated gradient causes its associated dimension (x, y or z) in the images to appear longer or shorter than reality. It will also cause slice position errors. It is normal for gradient calibration to drift over time and to require recalibration by the service engineer.

> **NOTE**: Gradient amplifiers need time to warm up and stabilize when they are turned on. Some sites power off their scanner hardware, including gradient amplifiers, overnight. Those sites should make sure their hardware has been on at least an hour before acquiring images of the phantom.

Some MR vendors allow the user to select whether to apply gradient distortion correction. For these systems be sure that distortion correction is turned on as this is a common cause of geometric accuracy failures for these systems.

Another possible cause of failure is use of too low an acquisition bandwidth. It is common practice on some scanners to reduce the acquisition bandwidth in order to increase signal-to-noise ratio. This can be pushed to the point that the normal inhomogeneities in B_0 manifest themselves as spatial distortions in the image. On most scanners the default bandwidth for T1 acquisitions is set high enough to avoid this problem.

If the geometric accuracy test fails, and the ACR T1 series was acquired at low bandwidth, try acquiring that series again at a higher bandwidth to see if the problem is eliminated.

Although uncommon, it is possible that abnormally high B_0 inhomogeneities could cause significant dimensional errors in the phantom images. Such B_0 inhomogeneities could be caused by improper adjustment of the gradient offsets, improper adjustment of passive and/or active magnet shims, or a ferromagnetic object such as a paper clip or large hair clip lodged in the magnet bore. Regardless of the cause, the service engineer can then measure the magnetic field homogeneity, and any inhomogeneity large enough to cause failure of the geometric accuracy test should be immediately apparent.

2.0 HIGH CONTRAST SPATIAL RESOLUTION

2.1 What It Is

The high contrast spatial resolution test assesses the scanner's ability to resolve small objects when the contrast-to-noise ratio is sufficiently high that it does not play a role in limiting that ability. This is sometimes called limiting high contrast spatial resolution.

A failure of this test means that for a given field of view and acquisition matrix size the scanner is not resolving small details as well as normal for a properly functioning scanner.

2.2 What Measurements Are Made

For this test, one visually assesses the distinguishability of closely spaced small objects. These objects consist of water-filled holes drilled in a small block of plastic called the **resolution insert,** which appears in slice 1. Before describing how to make the visual assessment of resolution using the images of the insert, it is first necessary to describe the insert in some detail.

The resolution insert

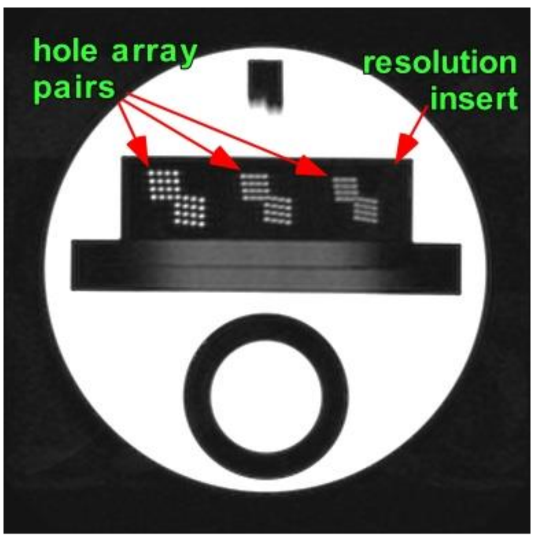

Figure 6. Slice 1 with resolution insert and hole array patterns indicated.

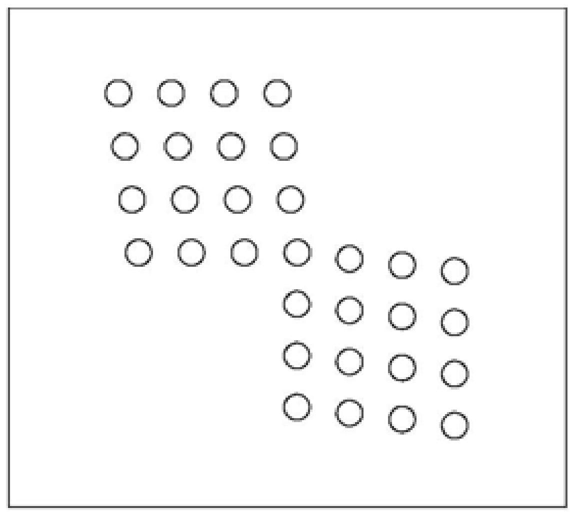

Figure 7: Illustration of one of the pairs of hole arrays in the resolution insert

Figure 6 shows an image of slice 1 with the resolution insert identified. There are three pairs of not-quite-square arrays of holes in the insert. One pair of hole arrays is illustrated in Figure 7.

Note that it consists of an upper left (UL) hole array and a lower right (LR) hole array. Here right and left refer to the viewer's right and left. The UL and LR arrays share one hole in common at the corner where they meet. The UL array is used to assess resolution in the right-left direction, and the LR array is used to assess resolution in the top-bottom direction (anterior-posterior if this were a knee).

The UL array comprises four rows of four holes each. The center-to-center hole separation within a row is twice the hole diameter. The center-to-center row separation is also twice the hole diameter. Each row is staggered slightly to the right of the one above, which is why the array is not quite square. The staggering ensures that the holes in at least one row will align exactly with the display matrix so that each hole in that

row will be centered within a pixel. Holes that do not align with the display matrix will experience partial volume affects and appear to be blurred, irregularly shaped spots of signal.
The LR array comprises four columns of four holes each. The center-to-center hole separation within each column and the center-to-center spacing between columns are twice the hole diameter. Each column is staggered slightly downward from the one to its left. As with the rows, staggering of columns ensures that the holes in at least one column will align exactly with the display matrix and not experience partial volume effects.

The hole diameter differs between the array pairs: for the left pair it is 0.9 mm, for the center pair it is 0.8 mm, and for the right pair it is 0.7 mm. Thus, using this insert, one can determine whether resolution has been achieved for each of these three hole sizes.

For this test, resolution in slice 1 of each of the two ACR axial series is evaluated. The following procedure is repeated for each of those series:

1. Display the slice 1 image.
2. Magnify the image by a factor of between 2 and 4, keeping the resolution insert visible in the display. This is illustrated in Figure 8.
3. Begin with the left most pair of hole arrays, which is the pair with the largest hole size, 0.9 mm.
4. Look at the rows of holes in the **UL** array, and adjust the display window and level to best show the holes as distinct from one another.
5. If all four holes in any single **row** are distinguishable from one another, score the image as resolved right-to-left at this particular hole size.

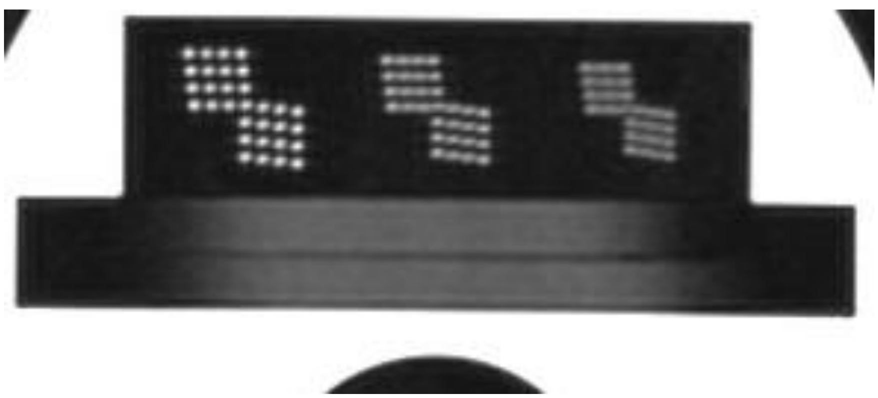

Figure 8. Magnified portion of slice 1 displayed appropriately for visually assessing high contrast resolution.

To be “distinguishable” or resolved, it is not required that image intensity drop to zero between the holes. To be distinguishable a single window and level setting can be found such that all four holes in at least one row are recognizable as points of brighter signal intensity than the spaces between them. Figure 9a shows the typical appearance of well resolved holes.

When the hole size is comparable to the resolution in the image, there is a tendency for groups of two or more holes in a row to blur together and appear as a single irregularly shaped spot of signal. In this case the holes in that row or column are considered unresolved. An example of this is shown in row 1 of the UL array of figure 9b.

Sometimes one or more holes which are distinguishable from neighbors in their own row blur together with their neighbors in adjacent rows. This is acceptable and does not affect the scoring for the row. An example of this is shown in the second row of the UL array of figure 9b, where the holes at each end of the row blur with their neighbors in adjacent rows.

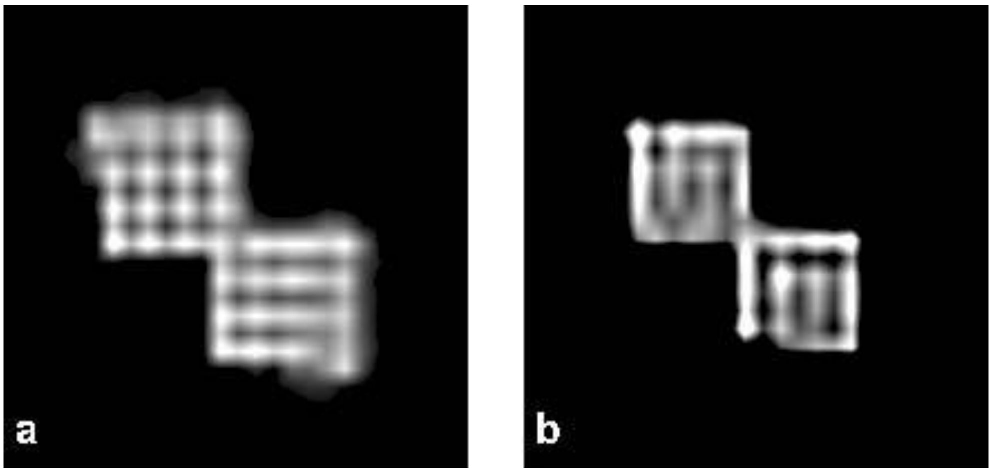

Figure 9

(a) Typical appearance of well resolved holes. Rows 2 through 4 of the UL array are resolved, and columns 1 through 3 of the LR array are resolved. (Rows and columns are numbered starting from the upper left corner of each array).

(b) Example of barely resolved holes and unresolved columns. Row 2 of the UL array is resolved because all four holes are discernable from each other, even though the holes at either end of the row blur together with their neighbors in the row below. So, the horizontal direction would be scored as resolved at this whole size. None of the columns of the LR array show more than three discernable spots within the column, so the vertical direction is not resolved at this hole size.

6. Look at the holes in the **LR** array and adjust the display window and level to best show the holes as distinct from one another.
7. If all four holes in any single **column** are seen to be distinguishable from one another, score the image as resolved top-to-bottom at that particular hole size.
8. The remarks made in step 6 about distinguishing holes within rows apply here to holes within columns.
9. Move on to the pair of arrays with the next smaller hole size, and evaluate as in steps 4 through 8. Continue until the smallest resolvable hole sizes have been found for the right-to-left and top-to-bottom directions.
10. Make a note of the smallest hole size resolved in each direction; that is the measured resolution for that direction.

2.3 How the Measurements Are Analyzed

There is no analysis. One simply notes the measured resolution in both directions for both of the axial ACR series.

2.4 Recommended Action Criteria

The field of view and matrix size for the axial ACR series are chosen to yield a resolution of 0.8 mm in both directions. The measured resolution of both axial ACR series must be 0.8 mm or better. If the resolution score for either of the ACR series is more than 0.8 mm, then evaluate the site series, If both site series can resolve 0.8 mm then the scanner passes this test.

A scanner must pass on both the ACR T1 and T2 series, or on both the site T1 and T2 series. A scanner cannot pass on just one of ACR series and one site series.

2.5 Causes of Failure and Corrective Actions

Excessive image filtering can cause failure. Many types of filtering that are used to make the images appear less noisy also smooth the image, which blurs small structures. A site that has failed the high contrast resolution test should check that any user selectable image filtering is either turned off, or at least set to the low end of the available filter settings.

Poor eddy current compensation can cause failure. The scanner's service engineer should check and adjust the eddy current compensation if this problem is suspected.

Excessive image ghosting can cause failure. The presence of excessive ghosting will be obvious elsewhere in the image if it is sufficient to cause failure of the high contrast resolution test. Ghosting is a very non-specific symptom of a hardware problem. In general, it is caused by instability of the measured signal from pulse cycle to pulse cycle, which can have its origin in the receiver, transmitter, or gradient subsystems. Motion of the phantom can also cause ghosting. Make sure the phantom is stable in the knee coil and not free to move or vibrate. Having ruled out phantom motion, it will usually be necessary to ask the service engineer to track down and correct the cause of ghosting.

Geometric errors from gradient miscalibration, B_0 inhomogeneity and too-low acquisition bandwidth can cause failure of this test. However, it is unusual for the geometric errors to be large enough to do that. In such cases the image will usually be obviously misshapen, e.g., the circular phantom may appear elliptical or egg-shaped. If the scanner passes the geometric accuracy test, it is very unlikely that geometric error could be the cause of failure of the high contrast spatial resolution test. On the other hand, if the scanner fails the geometric accuracy test by a large margin, then the failure of this test and the geometric accuracy test may have a common cause. Refer to section 1.5 of the geometric accuracy test for discussion of causes and corrective actions for geometric error.

3.0 SLICE THICKNESS ACCURACY

3.1 What It Is

The slice thickness accuracy test assesses the accuracy with which a slice of specified thickness is achieved. The prescribed slice thickness is compared with the measured slice thickness.

A failure of this test means that the scanner is producing slices of substantially different thickness from that being prescribed. This problem will generally not occur in isolation since the scanner deficiencies that can cause it will also cause other image problems. Therefore, the implications of a failure are not just that the slices are too thick or thin, but can extend to things such as incorrect image contrast and low signal-to-noise ratio.

3.2 What Measurements Are Made

For this test the lengths of two signal ramps in slice 1 are measured. This is done for both ACR series. The ramps appear in a structure called the **slice thickness insert**. Figure 10 shows an image of slice 1 with the slice thickness insert and signal ramps identified. The two ramps are crossed: one has a negative slope and the other a positive slope with respect to the plane of slice 1. They are produced by cutting 1 mm wide slots in a block of plastic. The slots are open to the interior of the phantom and are filled with the same solution that fills the bulk of the phantom.

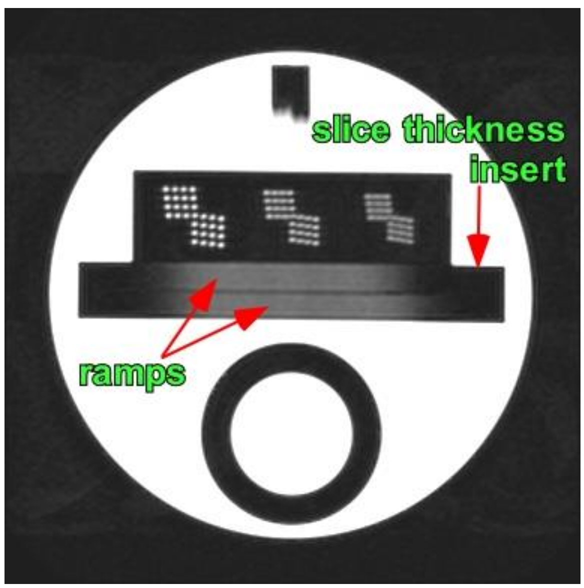

Figure 10. Slice 1 with the slice thickness insert and thickness ramps indicated.

The signal ramps have a slope of 10 to 1 with respect to the plane of slice 1, which is an angle of about 5.71^{o}. Therefore, the signal ramps will appear in the image of slice 1 with a length that is 10 times the thickness of the slice. If the phantom is tilted in the right-left direction, one ramp will appear longer than the other. Having crossed ramps allows for correction of the error introduced by right-left tilt, and the slice thickness formula provided in the next section takes that into account.

If Slice 1 in either ACR axial series is not correctly located within 2mm of the center of the paired wedges seen in the localizer the slice thickness signal ramps (see below) may be shifted beyond the insert. Under these conditions the signal ramps extend to or beyond the insert, rendering measurement impossible. This will result in failure of this test. The Slice Position Accuracy test described in section 4 below can help to verify proper location of each ACR series and, thus, avoid such problems.

For each ACR series, the length of the signal ramps in slice 1 is measured according to the following procedure:

1. Display slice 1, and magnify the image by a factor of 2 to 4, keeping the slice thickness insert fully visible on the screen.
2. Adjust the display level so that the signal ramps are well visualized. The ramp signal is much lower than that of surrounding water, so usually it will be necessary to lower the display level substantially and narrow the window. Except for the ROIs, which have not yet been placed, Figure 11 shows what the image should look like at this point.
3. Place a rectangular ROI at the middle of each signal ramp as shown in Figure 11. Note the mean signal values for each of these two ROIs, then average those two values together. The result is a number approximating the mean signal in the middle of the ramps. An elliptical ROI may be used if a rectangular one is unavailable. (When making these measurements be careful to fully cover the widths of the ramp with the ROIs in the top-bottom direction, but not to allow the ROIs to stray outside the ramps into adjacent high or low signal regions. If there is a large difference — more than 20% — between the signal values obtained for the two ROIs, it is often due to one or both of the ROIs including regions outside the ramps.)

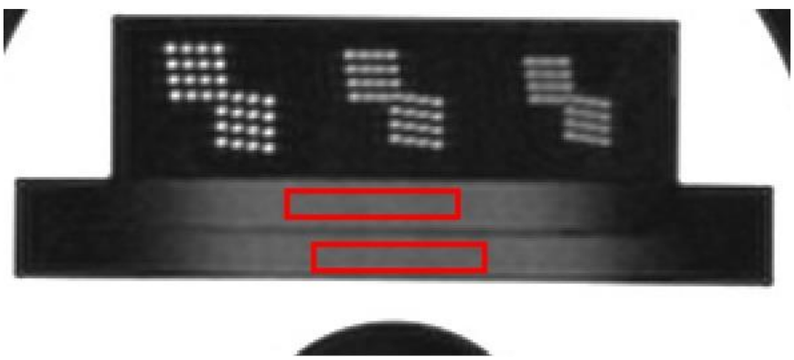

Figure 11. Magnified region of slice 1 showing slice thickness signal ramps with ROIs placed for measuring average signal in the ramps.

4. Lower the display level to half of the average ramp signal calculated in step 3. Leave the display window set to its minimum.
5. Use the on-screen length measurement tool of the display station to measure the lengths of the top and bottom ramps. This is illustrated in Figure 12. Record these lengths. They are the only measurements required for this test.

Figure 12. Magnified region of slice 1 showing signal ramps with display window at zero and the level at half the average signal level of the ramps. The length measurements for the ramps are shown on the image (lines).

Often there are horizontal striations in the signal intensity of the ramps that cause the ends to appear scalloped or ragged. The striations are a manifestation of truncation (Gibbs) artifact, and are normal. In this case one must estimate the average locations of the ends of the ramps in order to measure the ramp lengths. Figure 12 is an example of this problem and how the measurements should be made. Estimating the ends of the ramps introduces a source of error, but a millimeter of error in the ramp length measurement corresponds to only a tenth of a millimeter error in the slice thickness, so the errors introduced are small in effect.

3.3 How the Measurements Are Analyzed

The slice thickness is calculated using the following formula:

$$\text{slice thickness} = 0.2 \times (\text{top x bottom})/(\text{top} + \text{bottom})$$

where top and bottom are the measured lengths of the top and bottom signal ramps. For example, if the top signal ramp were 59.5 mm long and the bottom ramp were 47.2 mm long, then the calculated slice thickness would be

$$\text{slice thickness} = 0.2 \times (59.5 \times 47.2)/(59.5 + 47.2) = 5.26 \text{ mm.}$$

3.4 Recommended Action Criteria

The measured slice thickness of both axial ACR series should be 5.0 mm ± 0.7 mm. Errors greater than 1.0 mm will fail. If the thickness error for either ACR series fails, then evaluate both site series. If slice thickness for both site series is 5.0 mm +/- 1.0 mm, then the scanner passes this test.

3.5 Causes of Failure and Corrective Actions

Radiofrequency (RF) amplifier non-linearity can cause distorted RF pulse shapes and failure of this test. On many scanners the service engineer must empirically calibrate the RF power amplifier for linearity. If this calibration were lost or done incorrectly, it could possibly cause failure of this test.

Distorted RF pulse shapes can also arise from malfunctions anywhere in the high power RF portion of the transmitter, i.e., in the RF power amplifier, the cables and RF switches that convey power from the amplifier to the transmitter coil, or in the transmitter coil itself.

A bad gradient calibration or poor gradient switching performance can also cause failure of this test.

All of these possible causes for failure require corrective action by the service engineer.

4.0 SLICE POSITION ACCURACY

4.1 What It Is

The slice position accuracy test assesses the accuracy with which slices can be prescribed at specific locations utilizing the localizer image for positional reference.

A failure of this test means that the actual locations of acquired slices differ from their prescribed locations by substantially more than is normal for a well functioning scanner.

4.2 What Measurements Are Made

For this test the differences between the prescribed and actual position of slice 1 are measured. These measurements are made for the ACR T1 and T2 series.

Recall from the Introduction that slice 1 in each ACR series is prescribed to align with the center of the crossed 45° wedges at the starting end of the phantom (Figure 1). On slice 1 the crossed wedges appear as a pair of adjacent, dark, vertical bars at the top (anterior side) of the phantom. Figure 13 shows slice 1 with the vertical bars of the crossed wedges indicated.

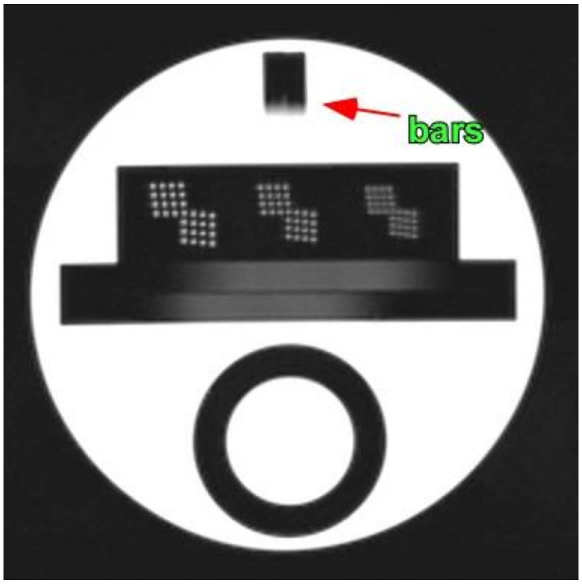

Figure 13. Slice 1 with the pair of vertical bars from the 45° crossed wedges. On this image the length difference between the right and the left bars is small and typical of well positioned slices.

If slice 1 is exactly aligned with the center of the crossed wedges, then the wedges will appear as dark bars of equal length on the image. By design of the wedges, if the slice is displaced superiorly with respect to the center, the bar on the observer's right (anatomical left) will be longer (Figure 14a). If the slice is displaced inferiorly with respect to the center, the bar on the left will be longer (Figure 14b).

Measurements are made for slice 1 of the ACR T1 and ACR T2 series. Use the following procedure for each image:

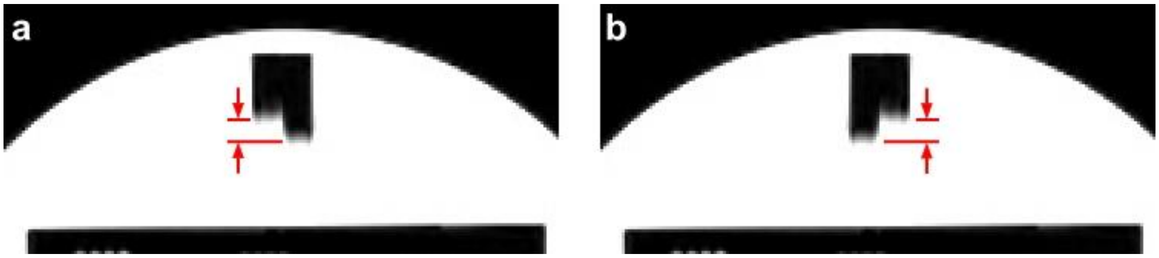

Figure 14

Images of slice 1 illustrating measurement of slice position error. The arrows indicate the bar length difference measurement that is to be made.

(a) The bar on the right is longer, indicating that the slice is mispositioned superiorly; this bar is assigned a positive value.

(b) The bar on the left is longer indicating the slice is mispositioned inferiorly; this bar is assigned a negative value.

1. Display the slice. Magnify the image by a factor of 2 to 4, keeping the vertical bars of the crossed wedges within the displayed portion of the magnified image.
2. Adjust the display window so that the ends of the vertical bars are well defined — not fuzzy. This will mean using a fairly narrow display window. The display level setting is not critical, but should be set to a level roughly half that of the signal in the bright, all-water portions of the phantom.
3. Use the on-screen length measurement tool to measure the difference in length between the left and right bars. The length to measure is indicated by the arrows in Figure 14.

If the left bar is longer, then assign a minus sign to the length. For example, if the bar length difference is 5.0 mm and the left bar is longer, then record the measurement as –5.0 mm.

In total, there are two length measurements, one for the ACR T1 series and one for the ACR T2 series.

4.3 How the Measurements Are Analyzed

This test requires no analysis of the measurements. The action criteria are specified in terms of limits on the bar length difference measurements. However, it is worth pointing out that because the crossed wedges have 45^o slopes, the bar length difference is twice the actual slice displacement error. For example a bar length difference of –5.0 mm implies the slice is displaced inferiorly by 2.5 mm from the center of the crossed wedges.

4.4 Recommended Action Criteria

The magnitude of each bar length difference must be less than or equal to 5 mm. As explained in section

3.2 above and in section 7.5, a bar length difference of more than 4 mm for slice 1 may adversely affect the slice thickness measurements and the low contrast object detectability score. So, although 5 mm is acceptable for this test, it is advisable to keep the bar length difference to 4 mm or less.

4.5 Causes of Failure and Corrective Actions

An error by the scanner operator in the prescription of slice location can cause a failure. This is probably the most common cause of failure of this test. This type of error will be evident when the axial images are cross-referenced on the localizer: slice 1 will not be aligned with the crossed wedge center on the localizer image. Other causes of failure do not show up as slice position error on the cross-referenced images. It is important to prescribe the slices as carefully as possible since errors introduced here can add to other sources of error and push an acceptable level of performance to an unacceptable level.

A particularly bad gradient calibration or poor B_0 homogeneity can cause failure of this test. In this case the problem also will usually be apparent as a failure or near failure of the geometric accuracy test.

Sometimes a failure of this test is an unfortunate combination of two or three of the problems just mentioned — inaccurate slice prescription, error in the table positioning mechanism, and poor gradient calibration or B_0 homogeneity — with none of the problems in itself being sufficiently bad to cause a failure on its own. Therefore, if no one thing seems to be responsible for causing a failure of this test, try having the service engineer shim B_0, recalibrate the gradients, and check the table positioning mechanism for excessive play. Then acquire a new image data set prescribing the slices as carefully as possible.

5.0 IMAGE INTENSITY UNIFORMITY

5.1 What It Is

The image intensity uniformity test measures the uniformity of the image intensity over a large water-only region of the phantom lying near the middle of the imaged volume and thus near the middle of the knee coil.

Knee coils for clinical use have fairly uniform spatial sensitivity near the middle of the coil when loaded as typical for a human knee. Failure of this test means that the scanner has significantly greater variation in image intensity than is normal for a properly functioning system. Lack of image intensity uniformity indicates a deficiency in the scanner, often a defective knee coil or problem in the radio-frequency subsystems.

5.2 What Measurements Are Made

For this test the high and low signal levels within a large, physically uniform, water-only region of the phantom are measured. This is done for the ACR T1 and T2 series.

For each of the two series, the measurements are made according to the following procedure:

1. Display slice location 4.
2. Place a large, circular region-of-interest (ROI) on the image as shown in Figure 15. This ROI must have an area of between 54 cm^2 and 56 cm^2 (5,400 to 5,600 mm^2). This ROI defines the boundary of the region in which the image uniformity is measured.

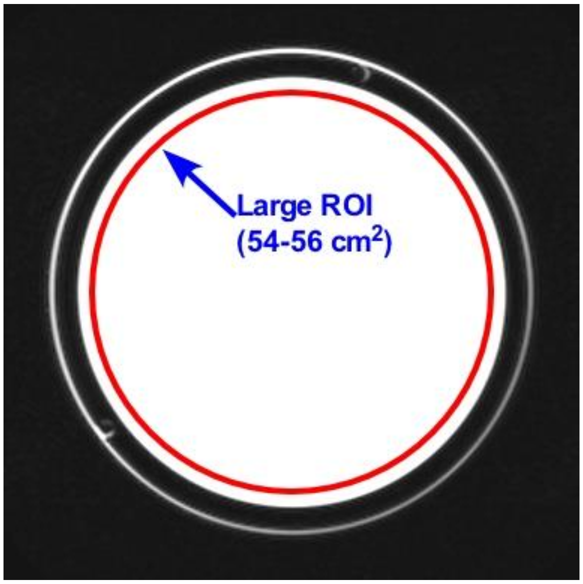

Figure 15. Image of slice 5 illustrating size and placement of the large 55 cm^2 ROI that defines the boundary inside which image uniformity measurements are made.

3. Although the mean pixel intensity inside this ROI is not needed for the uniformity test, it is used in the percent signal ghosting test (section 6.0), so it should be noted. Set the display window to its minimum, and lower the level until the entire area inside the large ROI is white. The goal now is to raise the level slowly until a small roughly 1 $cm^{2,}$ region of dark pixels develops inside the ROI. This is the region of lowest signal in the large ROI.

   Sometimes more than one region of dark pixels will appear. In that case, focus attention on the largest dark region.

   It can happen that rather than having a well-defined dark region, one ends up with one or more wide, poorly defined dark areas or areas of mixed black and white pixels. In that case one makes a visual estimate of the location of the darkest 1 cm^2 portion of the largest dark area.

4. Place a 1 cm^2 circular ROI on the low-signal region identified in step 3. Figure 16a shows what a typical image looks like at this point.

   Record the mean pixel value for this 1 cm^2 ROI. This is the measured low signal value.

   If there is uncertainty about where to place the ROI because there is no single obviously darkest location, one can try several locations and select the one having the lowest mean pixel value.

5. Now raise the level until all but a small, roughly 1 cm^2 region of white pixels remains inside the large ROI. This is the region of highest signal.

   Sometimes more than one region of white pixels will remain. In that case, focus attention on the largest white region.

   It can happen that rather than having a well defined white region, one ends up with one or more diffuse areas of mixed black and white pixels. In that case one makes a best estimate of the location of the brightest 1 cm^2 portion of the largest bright area.

6. Place a 1 cm^2 circular ROI on the high-signal region identified in step 5. Figure 16b shows what a typical image looks like at this point.

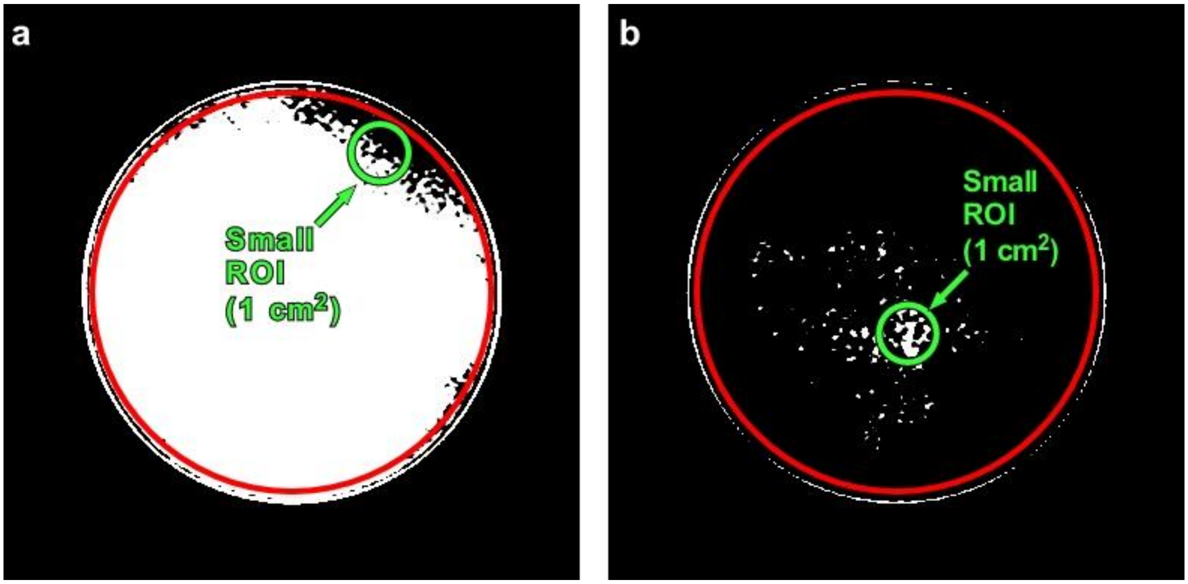

Figure 16. Image of slice 5 showing windowing of the image and the placement of the small 1 cm^2 ROIs for image uniformity measurements

(a) Example of windowing and ROI placement for measurement of the low signal value. In this case the proper ROI placement is not entirely clear. The guidance provided in step 3 above has been followed. Visual estimation was used to place the ROI on the largest 1cm^2 dark area inside the large ROI.

(b) Example of windowing and ROI placement for the highest signal value measurement. Following the guidance in step 5 above, the ROI has been placed at what is visually estimated to be the brightest 1cm^2 inside the large ROI.

.

Record the average pixel value for this 1 cm^2 ROI. This is the measured high signal value.

If there is uncertainty about where to place the ROI because there is no single obviously brightest location, one can try several locations and select the one having the highest mean pixel value.

Some display workstations have ROI tools that report the maximum and minimum pixel values within an ROI. It is tempting to use these as the high and low signal values. However, we advise against it. Due to the presence of noise in the image, using the maximum and minimum pixel values introduces systematic over- estimation of the high signal and under-estimation of the low signal. This systematic error can be significant, and biases the test towards failure.

5.3 How the Measurements Are Analyzed

The measured high and low signal values for each of the ACR series are combined to produce a value called **percent integral uniformity** (PIU). Use the following formula to calculate PIU:

$$PIU = 100 \times (1 - \{ (high - low)/(high + low) \}).$$

In this formula high is the measured high signal value and low the measured low signal value.

5.4 Recommended Action Criteria

PIU must be greater than or equal to 85%.

5.5 Causes of Failure and Corrective Actions

When scanning the phantom it is important to center it in the knee coil. If the phantom is closer to one side of the knee coil than another, one can expect uneven image intensities and potentially a failure of this test. This problem occurs most often with poor centering in the anterior-posterior (AP) direction.

It may be necessary to change some cushioning to get the phantom centered AP. In retrospect, poor centering may be evident in the filmed images. Another indicator that centering may be a problem is the appearance of bright spots in the image where the phantom is too close to the coil's conducting elements. If the scanner seems to be working well, and making ghost-free images with the usual amount of signal-to-noise ratio, poor phantom centering may be the cause of failure.

Image ghosting can cause image intensity variations and hence failure of this test. Ghosting sufficient to cause failure of this test will be readily apparent in the image, and likely will cause failure of another test such as percent signal ghosting (section 6.0). Ghosting is a very nonspecific symptom of a hardware problem. In general, it is caused by instability of the measured signal from pulse cycle to pulse cycle, which can have its origin in the receiver, transmitter, or gradient subsystems. Motion of the phantom can also cause ghosting. Make sure the phantom is stable in the coil and not free to move or vibrate. Having ruled out phantom motion, it will usually be necessary to ask the service engineer to identify and correct the cause of ghosting.

Degraded image intensity uniformity can result from failure of components in the knee coil. In these cases the images usually become noticeably lower in signal-to-noise ratio, i.e. they usually appear grainier. A service engineer is required to diagnose and correct these problems.

6.0 PERCENT SIGNAL GHOSTING

6.1 What It Is

The percent signal ghosting test assesses the level of ghosting in the images. Ghosting is an artifact in which a faint copy (ghost) of the imaged object appears superimposed on the image, displaced from its true location. If there are many low-level ghosts they may not be recognizable as copies of the object but simply appear as a smear of signal emanating in the phase encode direction from the brighter regions of the true image. Ghosting is a consequence of signal instability between pulse cycle repetitions. For this test the ghost signal level is measured and reported as a percentage of the signal level in the true (primary) image.

Although ghosts are most noticeable in the background areas of an image where there should be no signal, they can also overlay the main portions of the image as well, altering the true image intensities. A failure of this test means that there is signal ghosting at a level significantly higher than that observed in a properly functioning scanner.

6.2 What Measurements Are Made

For this test measurements are made on slice 4 of the ACR T1 series. Using the system's ROI tool, five intensity measurements are made: the average intensity in the primary image of the phantom, and the average intensity in the background at four locations outside of the phantom. The ROIs are placed as shown in Figure 17.

The phase encode shadow of an object in an image is the area of the image that is swept out by translating the object along the phase encode direction. Ghosts of an object can only fall in its phase encode shadow. Since the background ROIs are placed along the four edges of the field-of-view, two will be in the phantom's phase encode shadow and two will not. So, two of the background ROIs will sample the ghost signal and two will be free of ghost signal. It is necessary to have the two ghost-free background ROIs to serve as a control on the mean background intensity which can be affected by several factors, most notably noise.

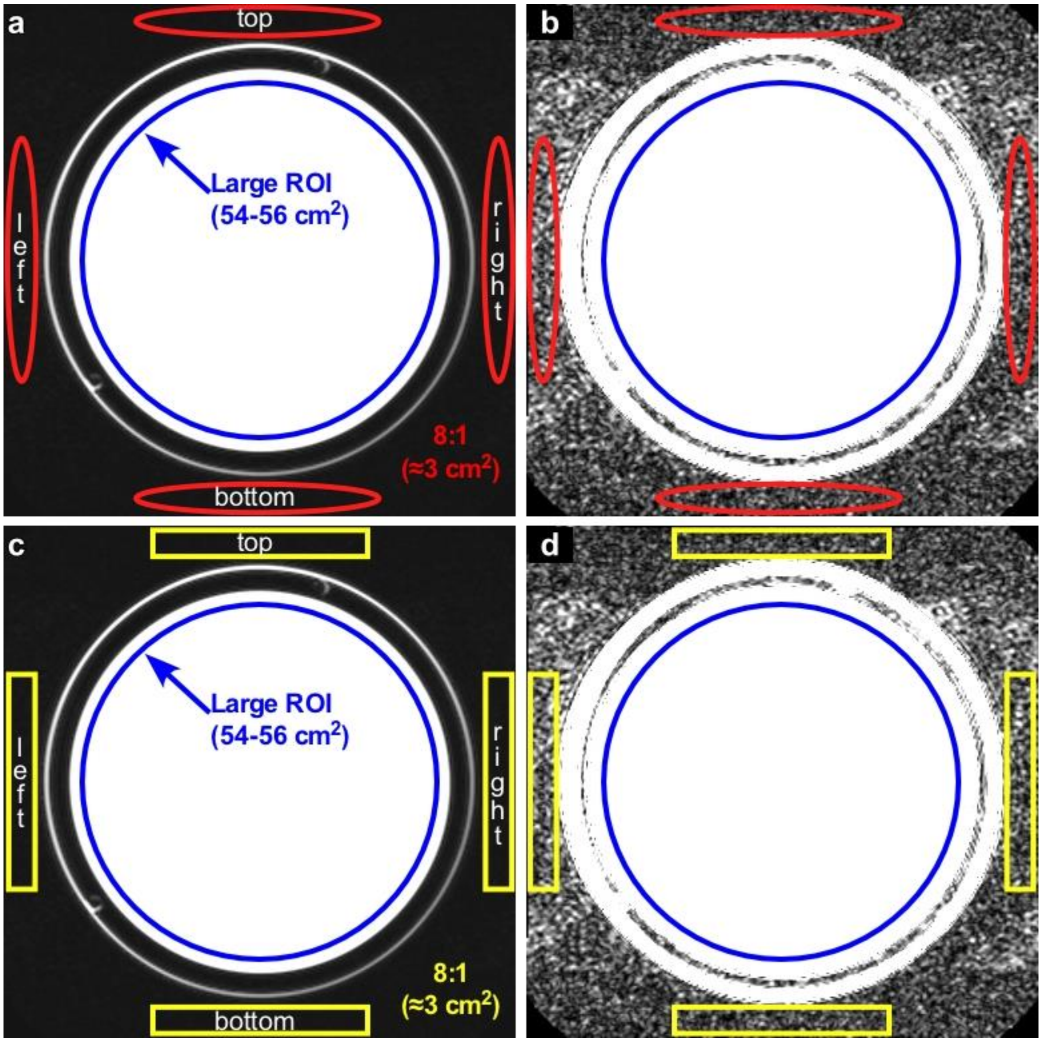

Figure 17. Image of slice 5 ROI placement for percent signal ghosting measurements.

The procedure for making these measurements is:

1. Display slice 5 of the ACR T1 series.
2. Place a large, circular ROI on the image as shown in Figure 17. This ROI must have an area of

between 54 cm^2 and 56 cm^2 (5,400 to 5,600 mm^2). The ROI should be centered on the phantom, as shown.

Record the mean pixel value for this ROI.

If the workstation cannot produce a circular ROI, a square ROI may be used. The area of the square ROI should be about 39 cm^2 (3,900 mm^2).

3. Place ROI's along, but not immediately against, the 4 edges of the field of view (FOV), as shown in Figure 17a and 17c.

   The ROIs should have a total area of approximately 3 cm^2 (300 mm^2). They can be either elliptical or rectangular and should have a length-to-width ratio of about 8:1. For an ellipse this is approximately 56 mm x 7 mm, for a rectangle this is approximately 50 mm x 6 mm. We will refer to these ROIs as they are labeled in Figure 17: top, bottom, left and right.

   Record the mean pixel value for each ROI, keeping track of which value goes with which ROI. It is important that the ROIs be placed in regions of actual noise and not against the edges of the phantom or against the edges of the FOV. Setting the window/level to visualize the background noise (as shown in 17b and 17d) helps to ensure that the ROIs are properly positioned. Additionally, if the phantom is not well centered it may be necessary to reduce the ROI width and increase the length to keep the ROI over the noise pixels. Be sure to maintain approximately 3 cm^2 area.

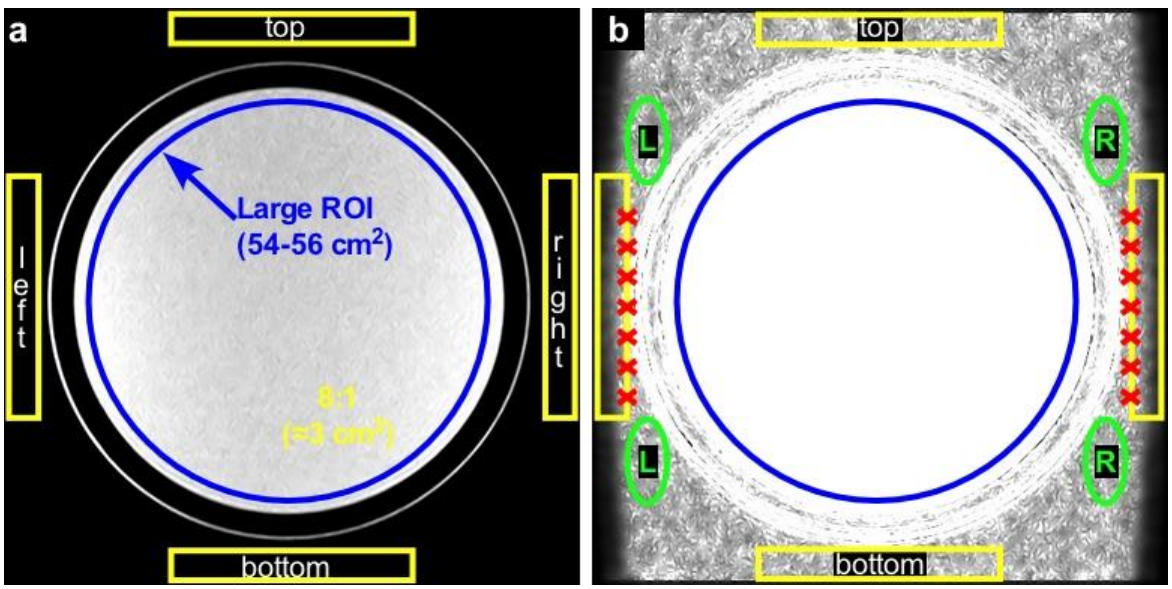

Figure 18. How to measure background noise when reconstructed FOV is too small.

On some specialty MRI systems either the acquired field-of-view (FOV) or the reconstructed FOV may be limited to less than 12cm. Under such conditions one must attempt to place the ROIs as close as possible to the proper location but without including any signals from the phantom itself or edge effects such as filtered roll off or Gibbs or truncation artifacts. In Figure 18a the four rectangular ROIs are placed as described earlier. By setting the Window/Level to highlight the background (Figure 18b) we see that the background noise on the left and right sides suffer from filter roll-off. These ROIs should not be used. Instead, we place smaller ROIs above and below the phantom arc and we average each these background values.

6.3 How the Measurements Are Analyzed

The value for the ghosting, as a fraction of the primary signal, is calculated using the following formula:

$$\text{ghosting ratio} = \left| \left((\text{top} + \text{btm}) - (\text{left} + \text{right}) \right) / \left(2 \times (\text{large ROI}) \right) \right|$$

where top, btm, left, right and large ROI are the average pixel values for the ROIs of the same names. The vertical bars enclosing the right hand side of the equation mean to take the magnitude of the enclosed value.

6.4 Recommended Action Criteria

The ghosting ratio must be less than or equal to 0.03 (3.0%).

6.5 Causes of Failure and Corrective Actions

Ghosting can be caused by motion or vibration of the phantom during the acquisition. Make sure the phantom is securely positioned in the knee coil and not free to move.

Ghosting is a very non-specific symptom of a hardware problem. In general, it is caused by instability of the measured signal from pulse cycle to pulse cycle, which can have its origin in the receiver, transmitter, or gradient subsystems. Having ruled out phantom motion, it will usually be necessary to ask the service engineer to track down and correct the cause of ghosting.

Scanners with non-digital receivers are prone to a kind of ghost called quadrature receiver imbalance ghost, or receive quadrature ghost. This type of ghost is not caused by instability of the measured signal from pulse cycle to pulse cycle, but by imbalance of the quadrature channels of the analog receiver. This ghost is usually easy to distinguish from other types of ghost because there will only be one ghost and it will be a reflection of the primary image through the origin. For example, an object in the lower left corner of the image will have it’s receive quadrature ghost in the upper right corner at an equal distance from the image center. Receive quadrature ghosting is corrected by balancing the receiver’s quadrature channels, which can be done by the service engineer.

7.0 LOW CONTRAST OBJECT DETECTABILITY

7.1 What It Is

The low contrast object detectability test assesses the extent to which objects of low contrast are discernible in the images. For this purpose the phantom has a set of low contrast objects of varying size and contrast.

The ability to detect low contrast objects is primarily determined by the contrast-to-noise ratio achieved in the image, and may be degraded by the presence of artifacts such as ghosting.

Scanners at different field strengths differ widely in their contrast-to-noise ratio performance, and scan protocols are normally adjusted to take these differences into account. Therefore, in addition to the ACR series, this low contrast object detectability test is applied to the site’s series. Most scanners can pass the test on the ACR series, but it is sufficient for a scanner to pass on the site’s series.

A failure of this test means the images produced by the scanner show significantly fewer low contrast objects than most properly functioning clinical scanners. Further, this deficiency holds even when the site's own clinical protocol is employed.

7.2 What Measurements Are Made

Measurements are made for the ACR and site's series. The low contrast objects appear on two slices: slices 6 and 7. In each slice the low contrast objects appear as rows of small disks, with the rows radiating from the center of a circle like spokes in a wheel. Each spoke is made up of three disks, and there are ten spokes in each circle. Figure 19 shows slice 7 with the circle of ten spokes visible.

All the disks on a given slice have the same level of contrast. In slice 6 the contrast value is 3.6% and in slice 7 it is 5.1%. All the disks in a given spoke have the same diameter. Starting at the 12 o'clock position and moving clockwise, the disk diameter decreases progressively from 7.0 mm at the first spoke to 1.5 mm at the tenth spoke.

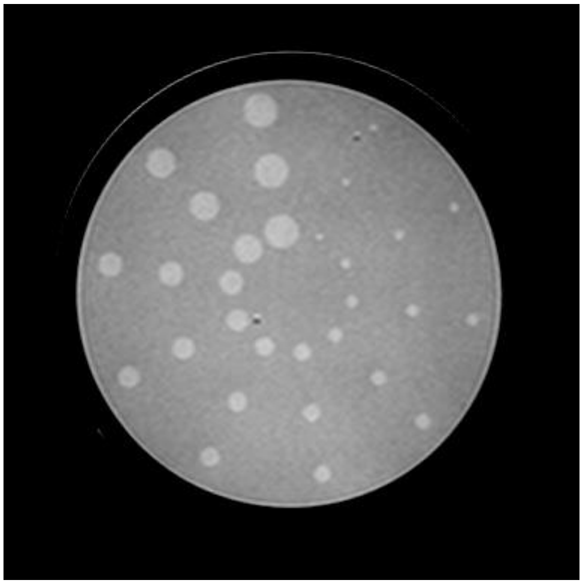

Figure 19. Slice 7 with 10 Low Contrast spokes visible.

The low contrast disks are actually holes drilled in thin sheets of plastic mounted in the phantom at the locations of the two slices. Since the contrast is derived from the displacement of solution from the slices by the plastic sheets, the contrast is independent of pulse sequence, TR, flip angle and field strength.

The measurements for this test consist of counting the number of complete spokes seen in each of the two slices. This is done for each of the four axial series.

Use the following procedure to score the number of complete spokes seen in a slice:

1. Display the slice to be scored. It helps to start with slice 7, which has the highest contrast objects.
2. Adjust the display window width and level settings for best visibility of the low-contrast objects. This will usually require a fairly narrow window width and careful adjustment of the level to best distinguish the objects from the background.

3. The task now is to count the number of complete spokes. Begin counting with the spoke having the largest diameter disks; this spoke is at 12 o'clock or slightly to the right of 12 o'clock, and is

referred to as spoke 1. Count clockwise from spoke 1 until a spoke is reached where one or more of the disks is not discernible from the background.

The number of complete spokes counted is the score for this slice. Record the score.

A spoke is complete only if all three of its disks are discernible. Count complete spokes, not individual disks.

As an example, Figure 20 shows an image of slice 6 in which less than all ten spokes are complete. The score for this image is 3 complete spokes.

Sometimes there will be a complete spoke of smaller size following a spoke that is not complete. Do not count it. Stop counting at the first incomplete spoke.

Disks on the threshold of discernibility can present a difficult scoring decision. They may appear ragged or misshapen; that is OK. The question is whether or not there is some sort of smudge or spot at the known location of the disk which is different enough from the background that one can say with a reasonable degree of confidence that there is something really there. In making this decision it can be helpful to look at areas where there are no disks in order to gauge the fluctuations in intensity from noise and artifacts that might mimic a barely discernible disk. A disk that looks no different than the brighter background noise fluctuations would not be deemed discernible.

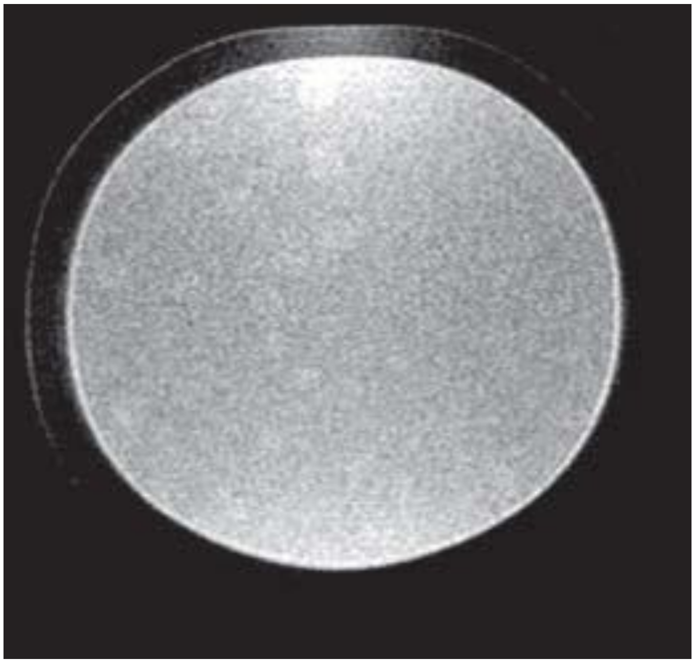

Figure 20

Some scanners greatly exceed the minimum passing score given in the action criteria section below. In most cases it isn't necessary to spend time pondering difficult decisions on barely visible disks; just score the test conservatively and revisit the scoring in the unlikely event the final score is below passing.

7.3 How the Measurements Are Analyzed

For each series, sum the number of complete spokes scored on each slice. For example, if the ACR T2 series scored 3 spokes in slice 6 and 8 spokes in slice 7, the total score for the ACR T2 series would be 3

+ 8 = 11.

7.4 Recommended Action Criteria

Both ACR series must have a total score of at least 9 spokes to pass If either ACR series has less than 9 spokes, then evaluate both site series. If the LCD score for both site series pass, then the scanner passes this test.

A scanner must pass on the ACR T1 and T2 series, or on the site's T1 and T2 series. A scanner cannot pass on just one of the ACR series and one of the site's series.

7.5 Causes of Failure and Corrective Actions

The most frequent cause of failure is incorrectly positioned slices. Slices 6 and 7 must fall close to their proper locations within the phantom in order for the thin plastic sheets that create the low contrast objects to do their job. If a slice is mispositioned by more than 2 mm there will be substantial reduction in the contrast of the low contrast objects in that slice. The easiest way to check if this is a problem is to look at slice 1. Recall from section 4.0 (Slice Position Accuracy) that the bar length difference of the crossed 45° wedges on slice 1 is twice the slice position error. So, make sure the bar length difference is less than 4.0 mm on slice 1. If it is not, reacquire the images with adjustments to the slice prescription as needed to bring the crossed wedge bar length difference in slice 1 to less than 4.0 mm.

Failure of this test can be caused by the phantom being tilted. A tilted phantom leads to parts of the slices being out of their proper location, and therefore to the same sort of problem as just described for incorrectly positioned slices. In this regard tilt (rotation) about the inferior-superior axis of the phantom is not a problem, but tilt about the other two axes can be. Tilt about the right-left axis will be readily apparent on the localizer. If the phantom doesn't look square with the edges of the field of view on the localizer, it should be repositioned before continuing to acquire data. The non-magnetic bubble level that accompanies the phantom can be used to level the phantom in the knee coil to avoid this kind of tilt. Tilt about the anterior-posterior axis is harder to see on the images. It can often be detected as right-to-left fade of structures that should show uniform partial-volume effects across an axial slice. The best way to avoid this kind of tilt is to carefully align the phantom squarely in the knee coil.

Some scanners may not allow specification of 5 mm slices with 3 mm gaps. In such cases it is not possible to acquire all of the required slice locations in a single acquisition. The *Site Scanning Instructions* included with the accreditation materials give guidance on allowable alternative acquisitions. Sites that cannot specify 5 mm slices with 3 mm gaps must be especially alert to the potential for mispositioned slices adversely affecting the low contrast object detectability test.

Ghosting artifacts can affect the ability to see low-contrast objects and cause a failure of the test. If ghosting is the cause of failure, it should be apparent on inspection of the images. Make sure the phantom is stable and can't move or vibrate during image acquisition. If ghosting is still a problem, then the service engineer should be asked to find and correct the cause. On some scanners, a small but noticeable amount of ghosting is normal on conventional T2 spin-echo acquisitions and fast spin-echo acquisitions. If in doubt whether the ghosting seen is normal, ask the service engineer to perform the manufacturer's diagnostic tests that relate to signal stability and ghosting.

If the images are free of ghosts, and the slices are positioned accurately, then a failure of this test is most likely due to inadequate signal-to-noise ratio (SNR) in the image. Ask the service engineer to check that the scanner's SNR performance is within manufacturer specifications.

8.0 ADVICE AND RECOMMENDATIONS

Make sure the service engineer has a few weeks of warning to check the system and make everything right before acquiring images for accreditation.

The scanner hardware should be powered up for at least an hour before acquiring data for accreditation. This advice is primarily aimed at sites that power off their scanners overnight.

Position the phantom in the center of the knee coil. This may not be where the patient knee normally rests. Line up the phantom with the center of the coil along the inferior-superior direction. Also make sure the phantom is centered in the coil right-left and anterior-posterior. This helps with image uniformity.

Use cushions and padding as necessary to stabilize the phantom against motion to avoid ghosting artifacts.

As explained in section 7.5, it is important to make sure that slices 6 and 7 lie as close as possible to their proper locations in the phantom, and that the phantom is closely aligned with the principal axes, not tilted or rotated.

If a small bubble is in the phantom, avoid shaking the phantom as it causes the bubble to break up into smaller bubbles which may adhere to structures within the phantom where they may interfere with scoring the test.