

Chapter 2

Partial k-Space Reconstruction

2.1 Motivation for Partial k-Space Reconstruction

(READ)

In theory, most MRI images depict the spin density as a function of position, and hence should be real valued. If this were true, then by the symmetry of the Fourier transform, only half of the spatial-frequency data will need to be collected. Since real functions have conjugate symmetry in spatial frequency space, the uncollected data could be synthesized by reflecting conjugate data across the origin. Unfortunately, there are many sources of phase errors that cause the real-valued assumption to be violated. These include variations in the resonance frequency, flow, and motion. As a result, partial k-space reconstructions always require some type of phase correction, to correct for these sources of incidental phase variation. This allows real images to be reconstructed.

An example of a gradient-recalled axial head image shown in Fig. 2.1 illustrates the problem. The magnitude reconstruction of a full k-space acquisition is shown in Fig. 2.1a and, the phase in Fig. 2.1b. The linear component of the phase has been corrected, leaving only the non-linear components. The absolute values of the real and imaginary components are shown in Fig. 2.1c and d. Clearly, a significant amount of phase correction is required before the conjugate phase symmetry can be exploited.




Figure 2.1: Axial gradient echo image acquired at 0.5T with an echo time of 13.8 ms. At this time water and fat are rephased. The linear shim terms have been corrected, leaving the non-linear components due to susceptibility shifts.

The reason for this phase is that the precession frequency varies across the head. The image phase is approximately

$$\phi(x, y) = \omega(x, y)T_E \quad (2.1)$$

where $\omega(x, y)$ is the local resonant frequency in rad/s, and T_E is the echo time. This can be due to inhomogeneity of the main magnet. These changes




Figure 2.2: Partial k-space acquisition for reducing scan time by reducing the number of phase encodes required.

in frequency vary slowly with spatial position, and can in theory be calibrated out of the system with proper shimming. More fundamental, and more problematic, is the variation in frequency due to the magnetic susceptibility difference between tissue and air. At air-tissue interfaces it is common to see local frequency shifts of several parts-per-million (ppm). These are seen in these images in the brain tissue adjacent to the ears, which is directly above the auditory canal, and around the nasal passages. An additional problem is that these shifts can occur over relatively short distances, and this governs the amount of resolution required for phase compensation, the amount of coverage in k-space that will be required, and ultimately how “partial” a partial k-space acquisition can be.

2.1.1 2DFT Applications

There are two general applications of 2DFT imaging where it is desirable to collect only a fraction of the full k-space data. The first is for reducing




Figure 2.3: Pulse sequence with a reduced echo time, resulting in a partial echo readout.

scan time by reducing the number of acquisitions that are required to construct an image of a given resolution. This is illustrated in Fig. 2.2. Slightly more than half of the complete k-space data is collected, allowing the scan time to be reduced by almost a factor of two.

The second application is for reducing echo times. Here the area of the readout dephasing gradient is reduced so that the echo comes earlier in the readout window as is shown in Fig. 2.3. This can be important for reducing flow dependent dephasing, and through plane susceptibility-induced signal loss. This case is illustrated in Fig. 2.4.

2.1.2 Applications to Other k-Space Acquisition Methods

Partial k-space acquisitions are also of interest for most other k-space acquisition methods. For EPI a partial k-space acquisition reduces the echo time, which can be quite long for a fully symmetric acquisition. Reconstruction of this type of partial k-space EPI data is very similar to the reduced phase-encode 2DFT case.

Other trajectories are more interesting. Projection acquisitions can exploit partial k-space symmetries, as can spiral acquisitions. The proper way to reconstruct these data sets is still an open re-

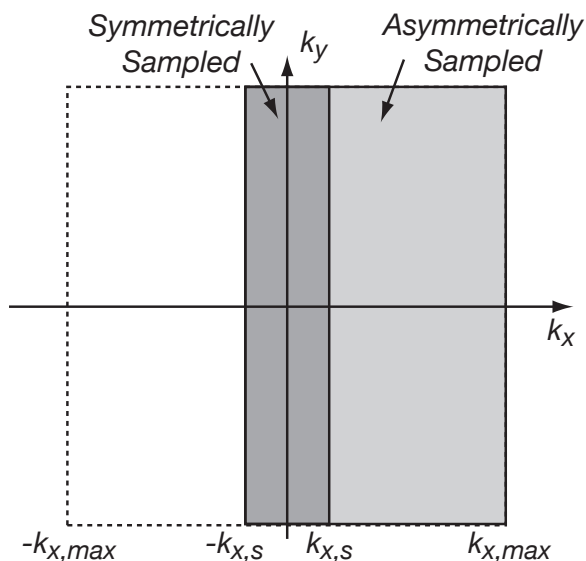


Figure 2.4: Partial k-space acquisition for reducing echo times by collecting only a fraction of the full echo.

search question. We will briefly describe the issues and possible solutions below.

2.1.3 Approaches to 2DFT Partial k-Space Reconstruction

We are going to consider two different types of approaches to partial k-space reconstruction. First are direct methods. These operate by constructing a real image in a single pass. One example is the homodyne algorithm [1]. These methods have limitations due to the interaction of phase compensation and synthesis of the conjugate data. The second type of approach addresses these limitations by using an iterative algorithm. One example of an iterative algorithm is called POCS for “projection onto convex sets” [2-4]. This operates by iteratively synthesizing the missing data that would be consistent with the data that was collected. We start with the direct algorithms, describe their limitations, and then move on to iterative algorithms.

2.2 Direct Partial k-Space Reconstruction

(READ)

2.2.1 Trivial Reconstruction by Zero Padding

The simplest way to reconstruct a partial k-space data set is to simply fill the uncollected data (phase-encodes or readout samples) with zeroes. Then, perform the 2DFT and display the magnitude. This works acceptably if the collected k-space fraction is close to 1, and works poorly as this fraction approaches 0.5. This is illustrated in Fig. 2.5 for a k-space fraction of $9/16^{th}$ s. The reconstruction of the full k-space data is shown in (a), and the reconstruction of the zero-padded partial k-space data in (b). The result is significant blurring in the phase-encode direction. Clearly this is unacceptable, and this motivates the search for other solutions.

The reason for the blurring can be identified by considering the data set to be the product of a full k-space data set multiplied by a weighting function. In this case an offset step function, where the offset corresponds to the k-space fraction. This will be denoted $W(k_y)$, and is illustrated in Fig. 2.6. The inverse Fourier transform of this function is the impulse response that produces the blurring. If we look at the real component, we see a sharp impulse at the desired resolution plus a broader component that corresponds to the width of the symmetrically acquired data. There is also a significant undesired imaginary component.

2.2.2 Phase Correction and Conjugate Synthesis

In order to correct for the blurring from the trivial reconstruction we need to fill in the missing uncollected data. From Fig. 2.1 it is clear that in gen-




Figure 2.5: Comparison of a reconstruction of a full k-space data set, and a trivial partial k-space reconstruction of the same data set where only 144 of 256 phase encodes have been used, and the remaining 112 have been replaced by zeros. Note the significant blurring in the phase-encode (left-right) direction.

eral phase correction must be applied before the k-space symmetry can be exploited to synthesize the missing data. In order to do the phase correction, we will use the narrow strip of data for which we have symmetric coverage. The phase of this low resolution image is then used to phase correct the partial k-space data. After inverse transforming the phase correction, the image reconstructed from the partial k-space data is transformed back

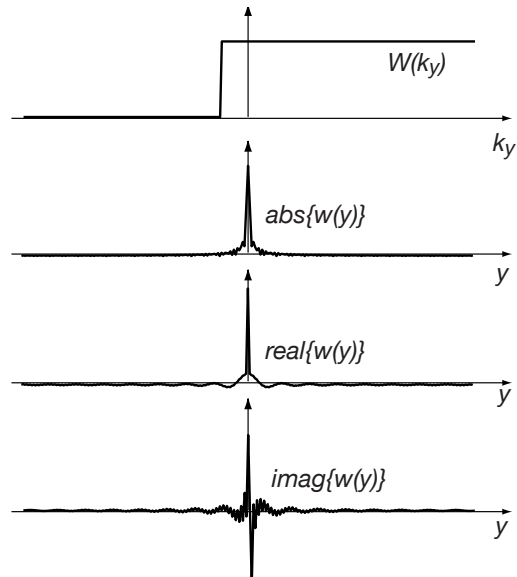


Figure 2.6: A k-space weighting function $W(k_y)$ that truncates a full k-space data set into a partial k-space data set. In this case the k-space fraction is $9/16^{th}$ s, as in Fig. 2.5. The blurring evident in Fig. 2.5 is due to the convolution of the inverse transform of $W(k_y)$, which is plotted here in absolute value, and real and imaginary components.

to the spatial frequency domain, where the data corresponding to the missing data is synthesized by conjugate symmetry,

$$M(k_x, k_y) = M^*(-k_x, -k_y). \quad (2.2)$$

This process is illustrated in Fig. 2.7. The partial k-space data is $M_{pk}(k_x, k_y)$, $M_s(k_x, k_y)$ is the narrow strip of symmetric data, and $m_{pk}(x, y)$ and $m_s(x, y)$ are the corresponding images produced by an inverse Fourier transform. The phase correction function is a unit amplitude image with a phase that is the conjugate of $m_x(x, y)$,

$$p^*(x, y) = e^{-i\angle m_s(x, y)}. \quad (2.3)$$

The problem with this approach is due to the effects of the phase compensation step near the boundary of the acquired data. The multiplication by the phase compensation function in the image domain is a convolution in the frequency

domain, and the size of this convolution function can be significant. The fact that the convolution is operating on zero data for the uncollected phase encodes produces errors near the boundary. Below we will describe a conjugate synthesis method based on k-space weighting that reduces this problem, and present some examples of reconstructed images.

Before proceeding, it is interesting to consider what sorts of features we are likely to lose if we rely on a phase correction function with limited spatial resolution. Figure 2.8 compares a full k-space reconstruction $m_f(x, y)$, left, with the real part of the phase corrected full-k-space data $Re\{m_f(x, y)p^*(x, y)\}$. The difference image is shown on the right. This shows which features will tend to be lost in a partial k-space reconstruction. One of the main effects is loss of vessel signal. This is to be expected since the vessels are too small to be resolved by the phase compensation function, and because motion through the slice select and imaging gradients produces velocity dependent phase shifts. The other areas of signal loss are near the air-tissue boundaries such as the sinuses.

The difference image has interesting noise characteristics. The background in the phase corrected image has lower noise because one component of the complex noise has been suppressed. Where the image has significant signal, the difference between the two images is very close to zero. The reason for this is illustrated in Fig. 2.9, which shows a vector diagram for the sum of two complex numbers, S the signal from some voxel, and n , the noise component from that voxel. If $|S| \gg |n|$, only the n_i , the component of the noise that is in-phase with the signal, contributes to $|S + n|$. Hence, $|S + n| \approx |S + n_i|$, and the magnitude operation has suppressed the noise component that is in quadrature with the signal. A more detailed description of this effect, as well as a discussion of the intermediate case where $|S|$ is on the same order as the magnitude of $|n|$ is given in [1].

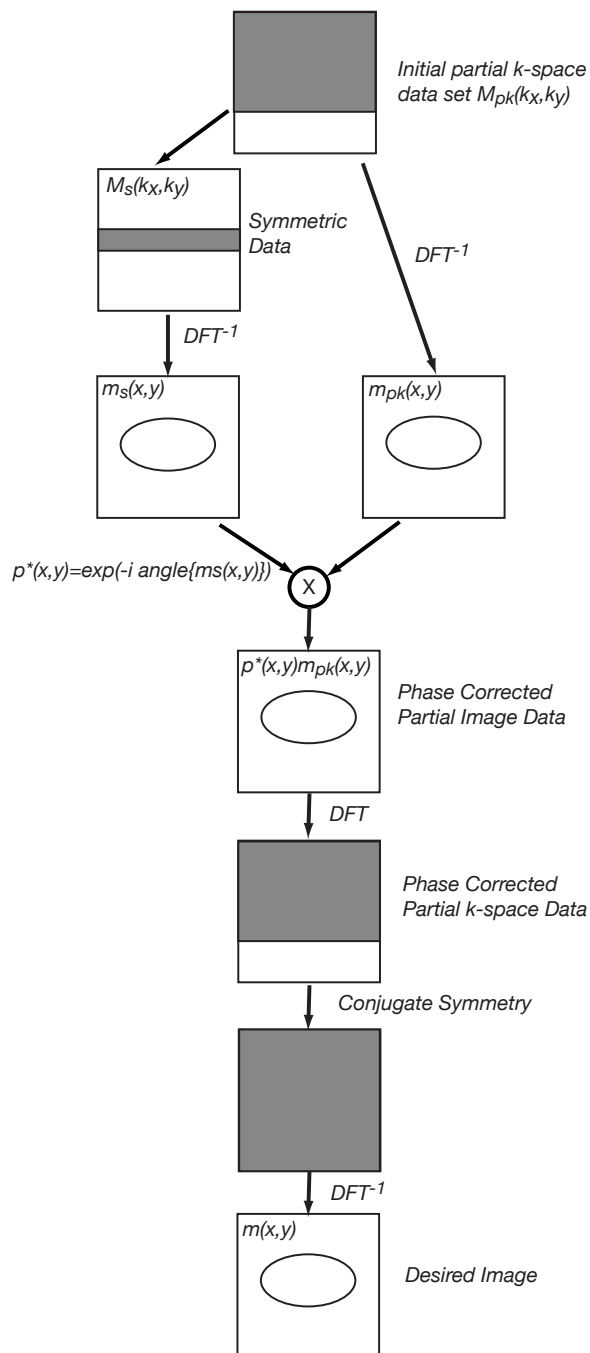


Figure 2.7: Summary of the phase correction and conjugate synthesis algorithm.




Figure 2.9: A full k-space reconstruction (left), and the real part of the phase corrected image. The phase correction function was computed using $\pm 1/16^{th}$ of the k-space data (corresponding to a $9/16^{th}$ s k-space reconstruction). The difference image shows vessels, as well as areas of rapid local change in phase, such as the air-tissue interfaces in the sinuses.




Figure 2.8: Magnitude operator suppression of the noise component that is in quadrature with the signal vector S . The magnitude of the sum $|S+n|$ is approximately $|S+n_i|$, where n_i is the component of the noise that is in-phase with S .

2.2.3 Homodyne Reconstruction

The drawback with the previous method is that, after phase correction in image space, the data must be transformed back to the frequency do-

main in order to fill in the missing data, and then an inverse transform back to the image domain is required to reconstruct the final image. Another approach, called homodyne, eliminates these last two transforms. The ways this is done is based on the symmetry properties of the Fourier transform.

The real part of an image corresponds to the conjugate symmetric component of the transform. The imaginary part corresponds to the conjugate anti-symmetric component. The weighting function that truncates the full k-space data set to produce the partial k-space data set shown in Fig. 2.6 can be decomposed into symmetric and anti-symmetric components, as shown in Fig. 2.10.

The imaginary part of the impulse response shown in Fig. 2.6, due to the antisymmetric component in spatial frequency space in Fig 2.10, will be suppressed by retaining the real part of the image. The real part of the impulse response is the inverse transform of the symmetric component in Fig. 2.10, and shows two readily identifiable elements. One is the desired impulse at the origin. The other is a much broader sinc, due to the over-weighting of the low spatial frequencies.

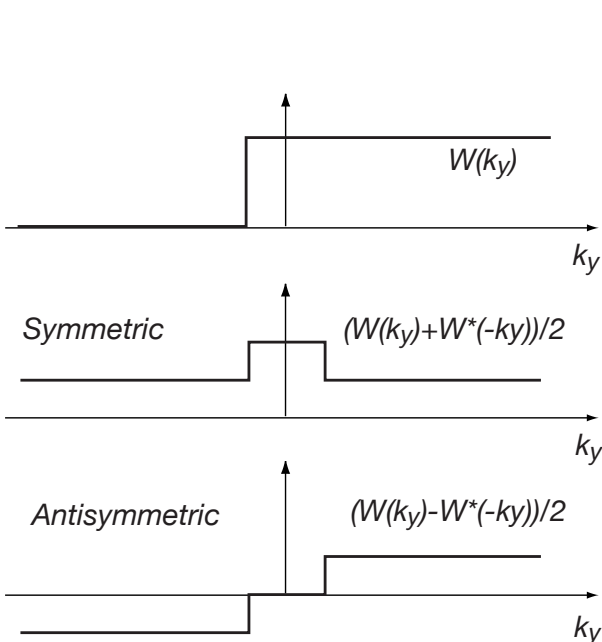


Figure 2.10: The weighting function that truncates a full k-space data set to a partial k-space data set can be decomposed into symmetric and antisymmetric components, corresponding to the real and imaginary components of the impulse response, shown in Fig. 2.6.

The key idea in the homodyne algorithm is to preweight the k-space data so that when we take the real part of the image data, it corresponds to a uniform weighting in k-space. The simplest approach is illustrated in Fig. 2.11. Here the amplitude of the high spatial frequencies have been doubled relative to the symmetrically acquired low frequency data.

The only problem with this approach is that the phase correction in image space corresponds to a convolution in k-space. The weighting of Fig. 2.11 has sharp discontinuities that produce transients from this convolution, and this can produce image artifacts. As a result, the weighting shown in Fig. 2.12 is preferred. The central k-space data is weighted linearly from zero up to 2. Again, the symmetric component is uniform. The anti-

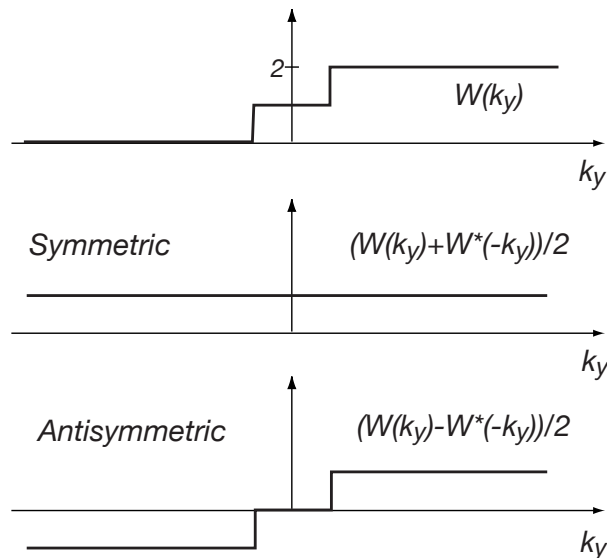


Figure 2.11: Doubling the high spatial frequencies relative to the central k-space data results in a uniform weighting for the symmetric component of the weighting.

symmetric component is also smoother, and this results in the imaginary component of the impulse response having fewer oscillations, which can also reduce image artifacts. Other, even smoother weightings can also be used.

A comparison of homodyne reconstructions of the gradient recalled echo data from Fig. 2.1 is shown in Fig. 2.13. On the left is a full k-space reconstruction. In the middle is a homodyne reconstruction of $9/16^{th}$ s of the data using a step weighting. Note the ghost of the scalp interfering with the brain. Using a ramp weighting, shown on the right, effectively eliminates this artifact. However, it produces an additional artifact above the auditory canal, where the phase is changing rapidly as a function of space.

The homodyne algorithm is summarized in Fig. 2.14. The central, symmetrically sampled data is reconstructed as a low resolution image, $m_s(x, y)$. A phase correction image $p^*(x, y)$ is computed as in Eq. 2.3. The partial k-space




Figure 2.13: Comparison of the performance of different homodyne k-space weighting functions for a 9/16ths data set. This is aggressive for phase gradients in this data set. On the left is the full k-space reconstruction. In the middle a step k-space weighting has been used. Note the distinct ghosts from the subcutaneous fat in the scalp (arrow). Using a ramp weighting, right, effectively eliminates this artifact. However, an additional artifact appears above the auditory canal, where the image phase is rapidly changing (arrow).




Figure 2.12: Using a linear ramp weighting over the central symmetrically sampled strip in k-space reduces the transients at the boundaries of the different k-space regions, and still produces a uniform k-space weighting for what will be the real component of the image.

data is preweighted in the partial k-space direction. The weighted partial k-space data is inverse Fourier transformed to produce an image $m_{pk}(x, y) * w(x, y)$. This is phase corrected by mul-

tiplying by $p^*(x, y)$. The final image is obtained by taking the real part of the result

$$m_{hd}(x, y) = \text{Re}\{p^*(x, y) (m_{pk}(x, y) * w(x, y))\} \quad (2.4)$$

Note that ideally, the weighting convolution and the image domain phase correction should occur in the other order. If the image phase is varying rapidly in space, phase correction at one pixel can result in incomplete suppression of the tails of the imaginary component from a nearby pixel. This effect is illustrated in Fig. 2.15. The signal from two voxels separated by five voxels is plotted for the case where both voxels are on resonance, and in phase, and for the case where the frequency difference between the two voxels produces $\pi/2$ phase shift. On resonance the two voxels are resolved as expected. Off resonance, the tail of the quadrature component of one voxel interferes with the in-phase component of the other voxel. After phase correction, this interference remains. Hence, for the homodyne algorithm to work well, the gradient of the image phase must be small compared to the length of the tails of the imaginary component of the impulse response.




Figure 2.14: Summary of the homodyne algorithm.

2.2.4 Conjugate Synthesis by k-Space Weighting

In the phase correction conjugate synthesis algorithm, the unknown data is generated by copying data from the conjugate symmetric location in k-space. The boundary between the acquired data and the synthesized data is a potential source of artifacts. These can be reduced by using the same k-space weighting method as in the homodyne algorithm. Examples are shown in Fig. 2.16. In this case the k-space fraction has been increase to $5/8^{th}$ s of k-space because the algorithm didn't produce acceptable results at $9/16^{th}$ s.




Figure 2.15: A limitation of the homodyne algorithm. Two voxels are separated by five voxels. If both are on resonance, suppressing the imaginary component resolves the two voxels as desired, shown on the left. If the difference frequency between the two produces a $\pi/2$ phase shift, the tails of quadrature component from one voxel interferes with the in-phase component of the other voxel. Phase correction rotates both interfering components into the real channel, producing artifacts.

2.2.5 Summary of Direct Methods

Both the homodyne algorithm, and the phase corrected conjugate synthesis approaches work well if the rate of change of image phase is limited. The problems with the homodyne approach are the result of performing phases correction after conjugate synthesis. The problems with the phase corrected conjugate synthesis approach are due to performing the conjugate synthesis after the phase correction.

2.3 Iterative Partial k-Space Reconstruction

(OPTIONAL)

The methods of the previous section perform the reconstruction in one pass. Problems arise




Figure 2.16: Comparison of the performance of different k-space weighting functions for a $5/8^{th}$ s data set, for a phase compensated conjugate synthesis reconstruction where the conjugate synthesis is done using the same k-space weighting technique as is used in the homodyne reconstruction. On the left is the full k-space reconstruction. In the middle a step k-space weighting has been used. On the right a ramp has been used. Each of these results in reasonable reconstructions. At $9/16^{th}$ s ghosting is produced with either weighting function (not shown), similar to the step weighted homodyne example of Fig. 2.13.

from the interaction between phase correction and the conjugate synthesis method, as was described above. Another approach is to estimate the missing k-space data by iteratively applying phase correction and conjugate synthesis. In the image domain, the image phase is constrained to be that of the low resolution estimate. In the frequency domain, the k-space data is constrained to match the acquired data when available. Iterating produces an estimate that approximately satisfies both sets of constraints.

There are several variations on this idea, depending on how the constraints are applied, and how the iteration is performed. We describe here a simple version of the POCS (for Projection onto Convex Sets) algorithm [3]. It is closely related to the earlier Cuppen algorithm [2].

The algorithm operates by iteratively transforming between the image domain and the spatial frequency domain. In the spatial frequency domain, the phase encodes that were actually acquired replace those of the current k-space estimate $M_i(k_x, k_y)$. This updated data set is inverse

Fourier transformed to produce the new estimated image $m_i(x, y)$. The phase of this image is forced to conform to the phase of the symmetrically acquired image $m_s(x, y)$ by computing

$$m_{i,pc} = |m_i(x, y)| p(x, y) \quad (2.5)$$

where

$$p(x, y) = e^{i\angle m_s(x, y)} \quad (2.6)$$

The corresponding Fourier data $M_{i,pc}(k_x, k_y)$ is computed by Fourier transform, and the entries corresponding to the uncollected phase are propagated forward to $M_{i+1}(k_x, k_y)$. The output image is $m_i(x, y)$ on the last iteration. The algorithm is summarized in Fig. 2.17. This is a complex image. Either the magnitude or the real part of the phase corrected image $Re\{m_i(x, y)p^*(x, y)\}$ can be used.

Typically the algorithm converges very rapidly, requiring four to five iterations before the changes from one iteration to the next are an order of magnitude below the noise floor of the MRI data.

Examples of both POCS and homodyne reconstructions of the data set of Fig. 2.1 are shown

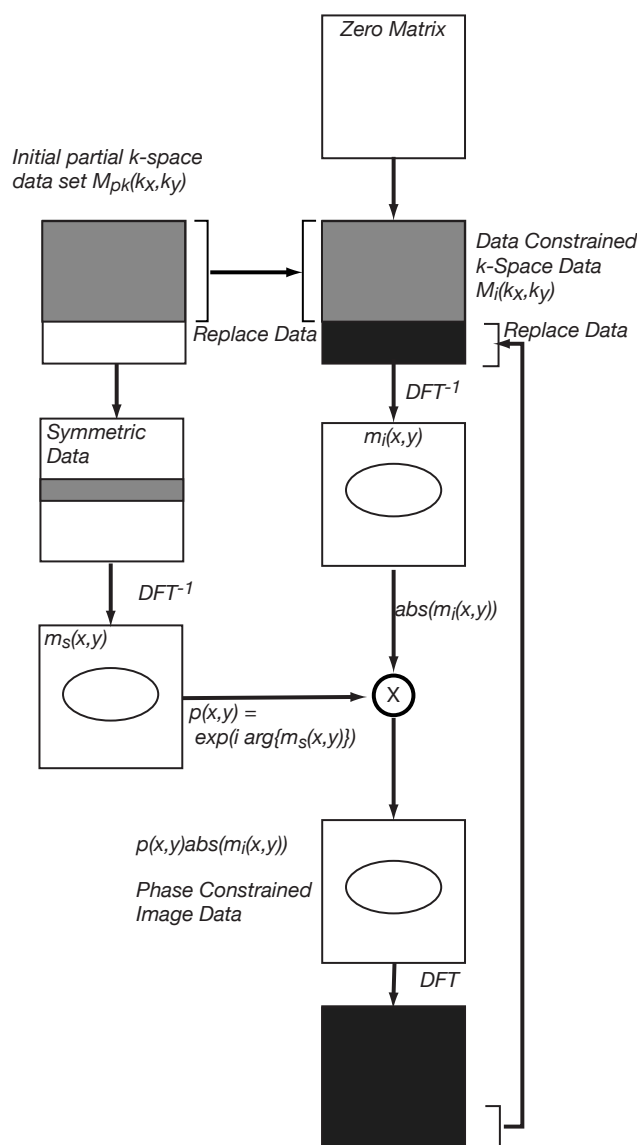


Figure 2.17: Summary of the POCS algorithm.

in Fig. 2.18, along with difference images computed with respect to the full k-space reconstruction. With either method, at $5/8^{th}$ s data sets, both methods perform well. At $9/16^{th}$ s, the POCS reconstruction has fewer artifacts than the step-windowed (shown) or ramp windowed homodyne reconstructions.

2.4 Conclusions

All of these algorithms work well for the case where the image phase variations are smooth. When the image phase changes rapidly, the homodyne algorithm produces ghosting. The POCS algorithm performs somewhat better as the k-space fraction decreases.

2.5 References

1. D.C. Noll, D.G. Nishimura, and A. Macovski, *IEEE Transactions on Medical Imag.*, **MI-10(2)**, 154, (1991).
2. J. Cuppen and A. van Est, *Magn. Reson. Imaging*, **5**, 526, (1987).
3. E.D. Lindskog, E.M. Haacke, and W. Lin, *em J. Magn Reson.*, **92**, 126 (1991).
4. Z.-P. Liang and P.C. Lauterbur *Principles of Magnetic Resonance Imaging: A Signal Processing Approach*, , IEEE Press, 2000.
5. G. McGibney, M.R. Smith, S.T. Nichols, and A. Crawley, *Magn. Reson. Med*, **30(1)**, 51, (1993).




Figure 2.18: Comparison of partial k-space reconstructions of the axial head data of Fig. 2.1. The top two are homodyne reconstructions at 9/16ths and 5/8ths k-space, and the lower two are POCS reconstructions for the same k-space fractions. Difference images are computed relative to the full-k-space reconstructions. The homodyne reconstructions at 9/16ths has clear ghosting. The POCS reconstruction at the same k-space fraction produces fewer artifacts. Either works well at 5/8ths.

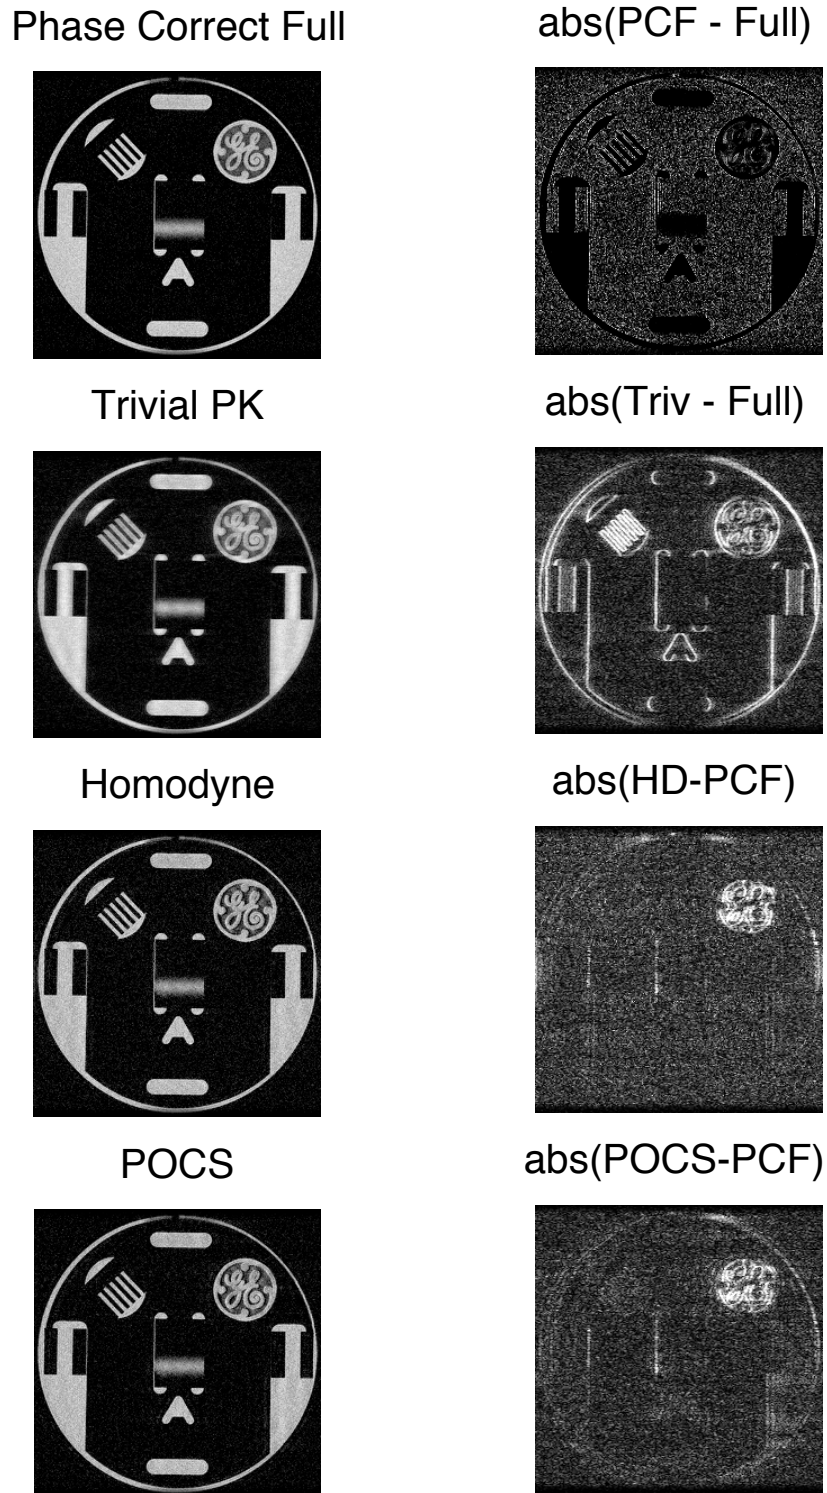


Figure 2.19: Comparison of partial k-space reconstructions for a gradient recalled echo phantom data set with $9/16^{\text{th}}$ s k-space coverage. Here the homodyne algorithm and the POCS algorithm perform similarly. Both fail in the vicinity of the “GE” logo, where the the phase compensation function $p^*(x, y)$ doesn’t have the resolution to track the rapid local changes in phase.

