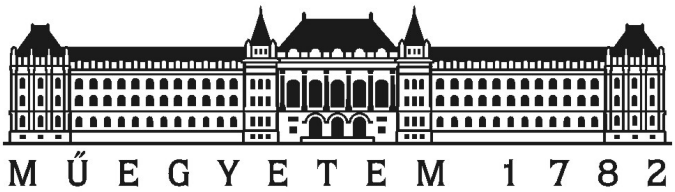

DIPLOMA THESIS

Partial Fourier reconstruction for MRI

Roberts Vivienne Astra

Supervisor: Légrády Dávid
associate professor
BME Nuclear Technical Institute
Nuclear Technical Department
Kettinger Ádám
Hungarian Academy of Sciences
Brain Imaging Centre

BME

2018

Contents

1 Absztrakt

Az MRI képrekonstrukció matematikájának alapvető eleme a Fourier transzformáció. A parciális Fourier transzformációban időigény csökkentése érdekében a k-térbeli adatsornak csak egy részét alkalmazzuk a rekonstrukcióban. Témám során ezt a homodyne módszerrel kiviteleztem, miszerint a valós fantom k-térbeli adatainak a hermitikus konjugált szimmetriáját alkalmazza. Ennek segítségével bepótolja a hiányzó adatokat, majd korrekciót végez arra nézve, hogy a gyakorlatban komplex fantomokkal dolgozunk. Ennek a matematikai kivitelezése egy felüláteresztő szűrő, majd egy aluláteresztő szűrő alkalmazásával megy végbe.

A homodyne módszert alkalmazó rekonstrukciós kódot írtam MATLAB környezetben, majd ezt a kódot optimalizáltam eleinte egy tekercses esetre végeztem el, megtaláltam a legalkalmasabb paramétereket, majd kibővítettem egy több tekercses esetre alkalmas kódra, amely egy moduláris függvénnyel végezte a rekonstrukciót, vagy az általam írt homodyne módszerrel vagy egy már megírt POCS módszerrel működő kóddal. Munkám folyamán a homodyne, POCS és zero filling módszereket hasonlítottam össze, illetve a különböző tekercs rekombinációs folyamatokat.

Azt találtam, hogy a homodyne módszer kép amplitúdójának minősége nem tér el kifejezetten a POCS módszerétől nagyobb partial Fourier faktor mellett, viszont jóval gyorsabb, mivel nem iteratív módszer, a rekombinációs folyamatoknál viszont jobb, ha a rekonstrukciót a rekombináció előtt végezzük.

2 Abstract

A fundamental element of MRI image reconstruction is the Fourier transform.[6] The partial Fourier transform, in which only part of the k-space data is used in the reconstruction, is used to reduce the amount of data that needs to be acquired in a measurement. As part of my own work, I realized this technique using the homodyne method, which applies hermitic conjugate symmetry of a real object's k-space data in order to fill in the missing data, then applies a correction to accommodate the seemingly complex real world phantoms. This mathematical implementation is carried out using a high pass filter followed by a low pass filter.

I wrote the reconstruction code using the homodyne method in a MATLAB environment[7], optimized this code for a single coil case first, found the most suitable parameters, and then extended it to a multiple coil code that had a modular reconstruction function using either my own homodyne method or previously written code with the POCS method. During my work, homodyne, POCS, and zero filling methods were compared, as well as different coil recombination processes.

I found that the quality of the image amplitude of the homodyne method does not differ significantly from that of the POCS method for higher partial Fourier factors, while working much faster, since it is not an iterative method. However recombination processes show better quality if the reconstruction is performed before recombination.

3 Theoretical background

Magnetic Resonance Imaging uses the spin property of atomic nuclei, often utilizing the abundance of hydrogen in either water and fat, common components in soft tissue. In an external field, these proton spins start a precession motion: in addition to spinning on its axis, the axis itself will also rotate around the external field at the Larmor frequency:

$$\omega_0 = \gamma B_0$$

where B_0 is the magnitude of the magnetic field and γ is the gyromagnetic factor. The gyromagnetic factor is isotope specific, however it generally is given in units of MHz/T.

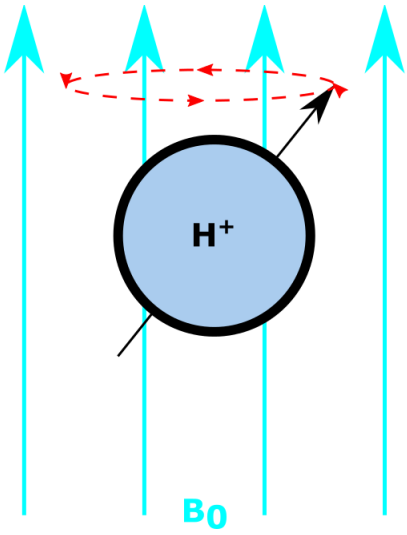

Figure 1: Precession in an external field

Using a radiofrequency pulse, we rotate the spins away from the external

field in order to rotate the spins into the transverse plane.

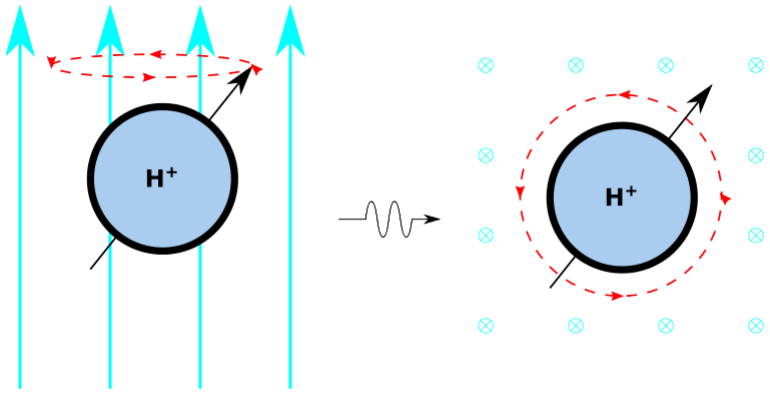

Figure 2: Effect of a 90 degree radiofrequency pulse on a proton

This precession induces a measurable voltage due to the current produced by changing magnetization. Following demodulation and assuming an ideal case (without relaxation and with perfect demodulation) the signal will be of the form:

$$S \sim \int d^3 r M_\perp(\mathbf{r}) e^{-i\phi(\mathbf{r},t)}$$

Where S is the signal, $M_\perp(\mathbf{r})$ is the transversal magnetization and $\phi(\mathbf{r},t)$ is given by the time integral of the Larmor frequency, and is thus determined by the spatial distribution of the magnetic field. For simple Nuclear Magnetic Resonance (NMR) spectroscopy, a homogeneous B_0 field is common, giving us a constant phase that we can then set to zero by definition.

3.1 MRI image reconstruction

In order to construct an image, the signal must be connected back to the spatial point that it came from. This is made possible by giving the external field a spatial dependance. This causes the Larmor frequency to varies throughout the body, making it possible to connect the frequency of the signal to its origin within the body.

This spatial variance is given by gradient coils. The gradient coils are operated as a $\mathbf{G}(t)$ function of time, giving a field of form $\mathbf{B}(\mathbf{r}, t) = \mathbf{B}_0 + \mathbf{G}(t)\mathbf{r}\mathbf{e}_z$. The spatially variant second term will also appear in the Larmor frequency and thus the signal as well, giving us frequency encoded information about the source of the signal.

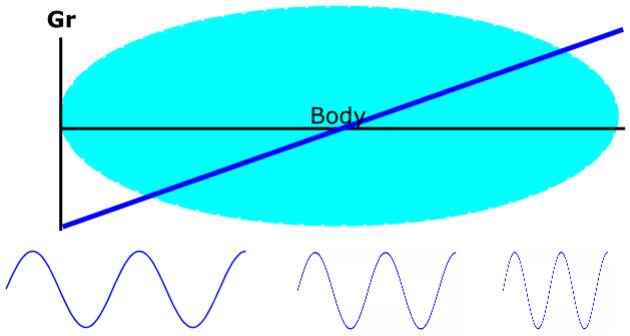

Figure 3: Effect of a gradient field within the body

In order to describe this effect mathematically, it is necessary to calculate the phase introduced earlier. In the case describe above, this phase will be given by the spatially variant term, that is:

$$\phi(\mathbf{r}, t) = \mathbf{r}\gamma \int_0^t \mathbf{G}(t')dt'$$

A useful variable to introduce is $\mathbf{k}(t)$:

$$\mathbf{k} = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(t')dt'$$

Making this substitution, our earlier formula will take the form:

$$S \sim \int d^3 r M_\perp(\mathbf{r}) e^{-i2\pi \mathbf{kr}}$$

Because of the useful nature of the $\mathbf{k}(t)$ variable, frequency data is known as k-space data, which is made up of the spectral components of spatial data (the image itself), which is proportional to $M_\perp(\mathbf{r})$. As is clear in the formula above, the signal is a simple function of the magnetization. In fact, the expression above is Fourier transform. Simply put, the following is true of the measured signal:

$$S(\mathbf{k}) = \mathcal{F}\left[M_\perp(\mathbf{r})\right]$$

In order to get back our spatial data and create an image, we have to connect each spectral component back to its spatial origin, which is similarly done by using the Fourier transform, or, more precisely, an inverse Fourier transform. Therefore, the proton density, proportional to the magnetization, is given by taking the inverse Fourier transform of k-space data. [3]

3.2 Fourier transform and partial reconstruction

The Fourier transform is a well known mathematical operation. In my work only one main property of the Fourier transform was truly important: its Hermitian symmetry. If the spatial object is real, one half of the k-space will be the conjugate of the other.

In MRI imaging could theoretically be a very useful property. After all, MRI scans measure a real body and collect k-space data directly, thus a possibility to measure only half the data and calculate the rest numerically, scan time could be reduced by almost half. However, in reality, the original object does not appear as a real object.

Figure 4: Image reconstructed with half of the k-space data

The original object will pick up phase shifts due to patient motion, resonance frequency shifts, hardware, electronics, eddy currents or other field inhomogeneity. Thus, these phase shifts create a reconstructed object that is

complex instead of purely real, causing errors in any reconstruction assuming a real object.

However, collecting less data to save time is still a goal within MRI field. Therefore, while sampling exclusively one half of the k-space is impractical, using one half to reconstruct the missing data, as well as a small section of the other half to make a reasonable estimation of the phase is still an option. This method is called Partial Fourier Reconstruction.

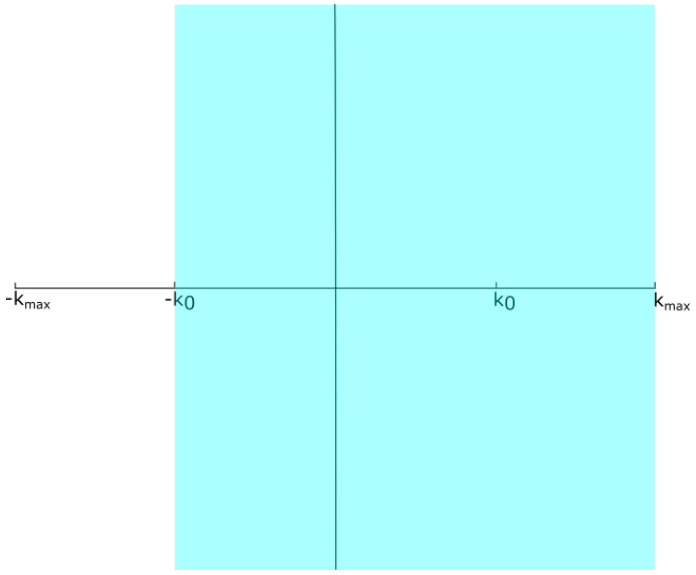

Figure 5: k-space data in partial Fourier reconstruction. Data is sampled inside the blue area.

Typically, we characterize the partial Fourier transformation with a partial Fourier factor, the fraction of the k-space that is sampled.

3.2.1 Zero filling

There are a few ways to do this. One is to simply remove the data from a section of the k-space and replace it with zeros, then perform the Fourier transform. This is called zero filling. Zero filling gives a fairly good image in the low-spatial-frequency range, but it requires large fractions of the k-space, around 0.75, to be reasonably phase accurate with good image quality.

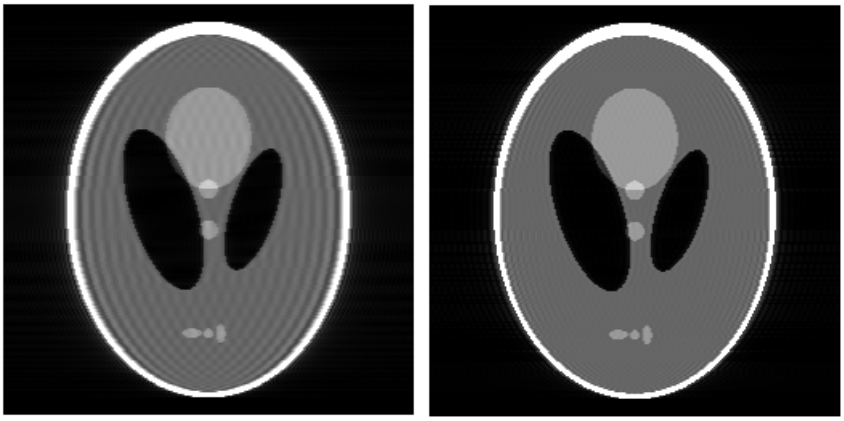

Figure 6: Zero filling with partial Fourier factor of 0.625 and 0.75 respectively

Even at 0.75, the zero filling method causes Gibbs ringing near sharp edges.

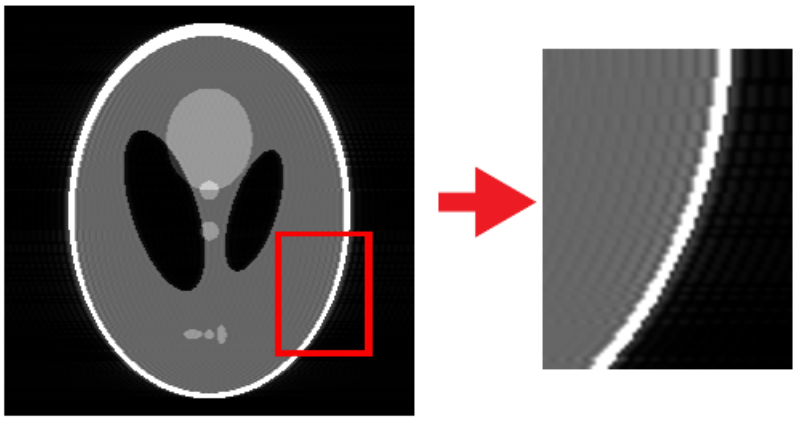

Figure 7: Gibbs ringing present in zero filling with partial Fourier factor of 0.75

The zero filling method remains one of the more practical methods for phase sensitive techniques, such as shimming. It requires very little calculation and thus computation time and at higher partial Fourier factors its image quality is comparable to those of more involved techniques. However, for high quality images with low partial Fourier factor and in events where the phase information is not necessary for following techniques, one of the simpler, more efficient techniques is that of homodyne processing.

3.2.2 Homodyne

A slightly more involved technique and the predominate technique in my own work, homodyne processing involves constructing a phase map from the symmetrically sampled k-space and adding it to the image reconstructed with the assumption of an ideal, real object.

In the direction of the partial Fourier transform, $-k_{max}$ to $-k_0$ is not sampled, $-k_0$ to k_0 is symmetrically sampled, and k_0 to k_{max} is sampled only on one side of the k-space.

In the ideal case, the image $I(x)$ is real without unwanted phase shifts, and thus the Hermitian symmetry, $S(k) = S*(-k)$, holds true for the k-space data, $S(k)$:

$$I(x) = \int_{-k_{max}}^{k_{max}} S(k)e^{i2\pi kx}dk \tag{1}$$

In this case, using the Hermitian symmetry, the Fourier Transform (1) can manipulated in the following way. First, the the integral above can be separated into two parts, one on the interval $[-k_{max}, -k_0]$, where we make the replacement allowed by Hermitian symmetry, and another on the interval $[-k_0, k_{max}]$, where we do not.

$$I(x) = \int_{-k_{max}}^{-k_0} S^*(-k)e^{i2\pi kx}dk + \int_{-k_0}^{k_{max}} S(k)e^{i2\pi kx}dk \tag{2}$$

For the first integral, the variable $k' = -k$ can be introduced. It is easy to see this will give a factor of $e^{-i2\pi kx}$ in the integral, which is the conjugate of the original factor. Thus, since the integral also contains $S^*(k')$, the resulting

integral can be interpreted as the conjugate of a regular Fourier transform of $S(k')$ (for the given interval).

$$I(x) = \left[\int_{k_0}^{k_{max}} S(k') e^{i2\pi k' x} dk' \right]^* + \int_{-k_0}^{k_{max}} S(k) e^{i2\pi kx} dk \tag{3}$$

Since there is no instance of k remaining in the first integral, there is no reason why k' cannot be denoted by k for consistency. The second integral can be further separated into two more integrals, one along the interval $[-k_0, k_0]$, the other along the interval $[k_0, k_{max}]$.

$$I(x) = \left[\int_{k_0}^{k_{max}} S(k) e^{i2\pi kx} dk \right]^* + \int_{-k_0}^{k_0} S(k) e^{i2\pi kx} dk + \int_{k_0}^{k_{max}} S(k) e^{i2\pi kx} dk \tag{4}$$

The first and third integrals can be combined, as one is merely the conjugate of the other.

$$I(x) = \int_{-k_0}^{k_0} S(k) e^{i2\pi kx} dk + 2Re \left[\int_{k_0}^{k_{max}} S(k) e^{i2\pi kx} dk \right] \tag{5}$$

Given that the initial assumption is that the image is real, and the second element is clearly real, the first integral also be real. Therefore, taking real part of both integrals gives the same effect.

$$I(x) = Re \left[\int_{-k_0}^{k_0} S(k) e^{i2\pi kx} dk + 2 \int_{k_0}^{k_{max}} S(k) e^{i2\pi kx} dk \right] \tag{6}$$

This combination of integrals can be replaced by a single integral, where the integral includes a step function.

$$I_H(x) = \int_{-k_{max}}^{k_{max}} H(k)S(k)e^{i2\pi kx}dk \tag{7}$$

This step function, H(k), is essentially a high pass filter.

$$H(k) = \begin{cases} 0 & k < -k_0 \\ 1 & -k_0 \leq k < k_0 \\ 2 & k \geq k_0 \end{cases}$$

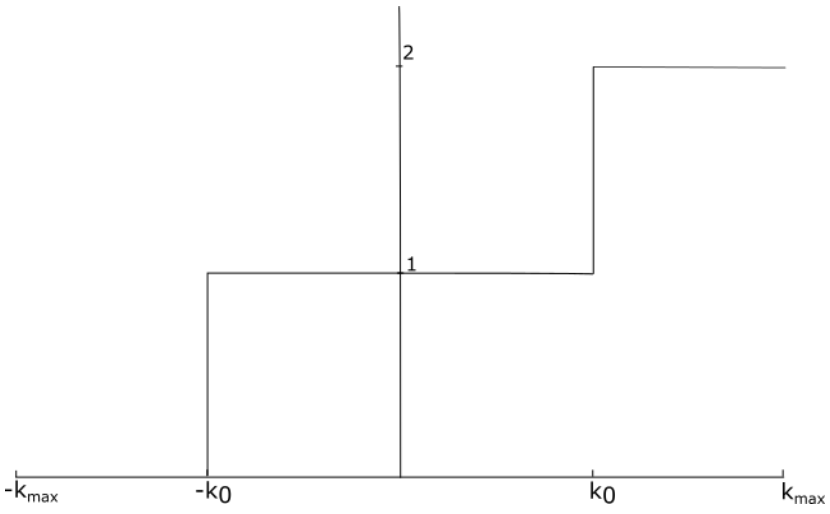

Figure 8: High pass filter, H(k)

By using this filter, data collected on the interval $[-k_0, k_{max}]$ can be reconstructed so that it is identical to the result given by a reconstruction of

a full dataset.

$$I(x) = Re[I_H(x)] \tag{8}$$

However, this is only true for a real object.

In order to take into account possible phase shifting, phase correction is used. The phase correction is calculated from symmetrically sampled k-space data by making a low-frequency image, using a low-pass filter, $L(k)$.

$$L(k) = \begin{cases} 1 & |k| \leq k_0 \\ 0 & |k| > k_0 \end{cases}$$

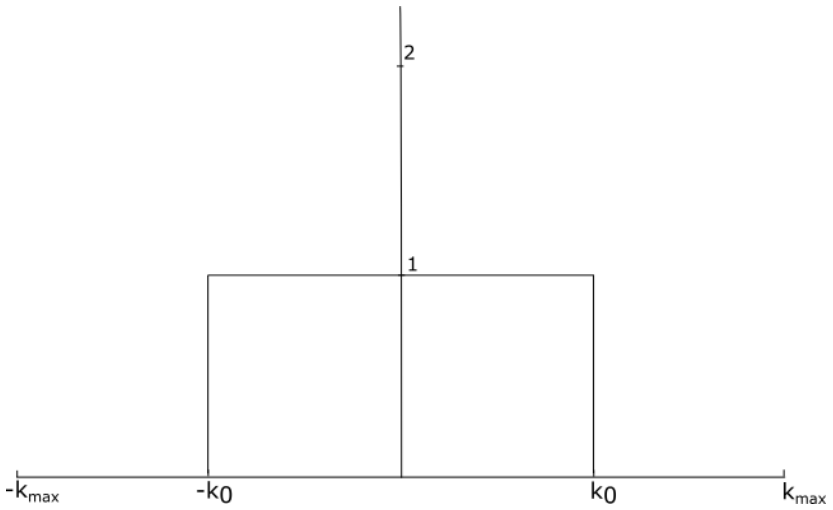

Figure 9: Low pass filter, L(k)

$$I_L(x) = \int_{-k_0}^{k_0} S(k)e^{i2\pi kx}dk = \int_{-k_{max}}^{k_{max}} L(k)S(k)e^{i2\pi kx}dk \tag{9}$$

The phase of this image is a good approximation of the phase of for the original image, reconstructed from a full data set, and thus, using the phase from this image with the earlier ideal image, the final image is calculated as follows:

$$I(x) = Re\left[I_H(x)e^{-i\phi_L(x)}\right] \tag{10}$$

where $\phi_L(x)$ is the phase of $I_L(x)$. Calculating the phase directly can cause errors due to numerical applications of the arctangent function, therefore it is better to get the phase by conjugating the low-frequency image and dividing it by its magnitude to get the phase:

$$I_H(x)e^{-i\phi_L(x)} = I_H(x)\frac{I_L * (x)}{|I_L(x)|} \tag{11}$$

The final image will be the real part of this, meaning the original phase is still lost, therefore for phase sensitive techniques the zero-filling more practical, but homodyne processing reproduces the image with sufficient quality using smaller partial Fourier factors. [4]

It is important to note that in reality, the simple step functions seen above are augmented with smoother transitions.This is because hard edges and corners, such as the ones seen in step functions, cause ringing around the edges. [8] Using cosines, it is possible to make these edges "rounder" and

thus avoid ringing, giving the following forms:

$$H(k) = \begin{cases} 0 & k \leq -k_{min} - w/2 \\ cos\left(\pi \dfrac{|k| - k_{min} + \frac{w}{2}}{2w}\right)^2 & -k_{min} - \frac{w}{2} < k < -k_{min} + \frac{w}{2} \\ 1 & -k_{min} + \frac{w}{2} \leq k \leq k_{min} - \frac{w}{2} \\ cos\left(\pi \dfrac{|k| - k_{min} - \frac{w}{2}}{2w}\right)^2 + 1 & k_{min} - \frac{w}{2} < k < k_{min} + \frac{w}{2} \\ 2 & k_{min} + \frac{w}{2} \leq k \end{cases}$$

$$L(k) = \begin{cases} 1 & k \leq k_{min} - \frac{w}{2} \\ cos\left(\pi \dfrac{|k| - k_{min} + \frac{w}{2}}{2w}\right)^2 & k_{min} - \frac{w}{2} < k < k_{min} + \frac{w}{2} \\ 0 & k_{min} + \frac{w}{2} \leq k \end{cases}$$

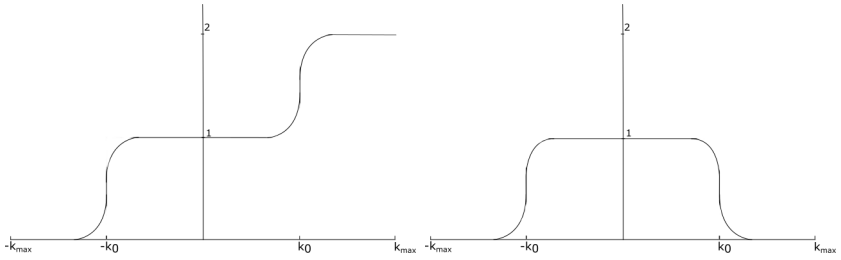

Figure 10: Final forms of filter functions with smoothing

4 Realization

During my work, I wrote a few versions of the homodyne program. The first version was a single coil program that took the following variables as input:

- The k-space data
- k_0, k_{max}
- The w parameter seen in the earlier "smoother" step functions, referred to as "smoothing factor" in further sections

The goal of this first program was to optimize the homodyne program. The most important parameter in this step was the aforementioned smoothing factor.

4.1 Smoothing factor

In order to find the value that would give optimal image quality, I wrote a program that varied the smoothing factor and produced an image for each value of this factor. I then took the image amplitude and compared it to the amplitude of the original image by taking the root mean square error. I also did a comparison of the phases of both images, but since this was not affected by the smoothing factor, I decided to focus instead on the amplitude. By plotting the root mean square error as a function of the smoothing factor, found that the graphs showed the following curve.

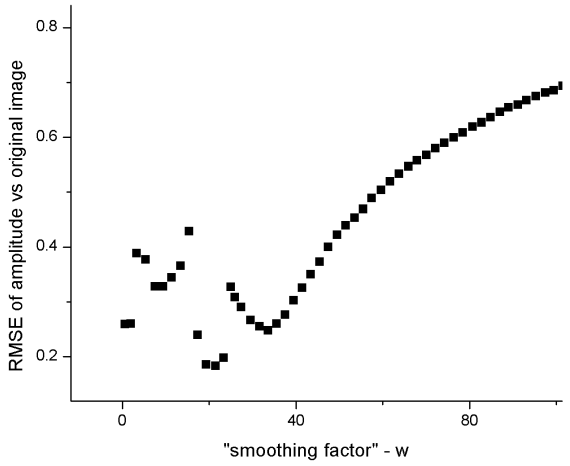

Figure 11: RMSE of images produced with varying smoothing factor, with reference to original image, a Shepp-Logan phantom of size 128 x 128 pixels

As seen on the image above, there is a pronounced dip in error values at approximately 33 pixels. I repeated this process with several images of varying size, finding that the shape of the curve was always the same as seen above, with irregular values in the first third of image size, followed by a dip and then steady increasing error. This dip was consistently between 25-35 percent of the image size.

This is also reflected by the images themselves:

Figure 12: Images produced with smoothing factor at 0.1, 0.2, 0.5 times the image size

Images with small smoothing factors show ringing, images with larger smoothing factors were blurred and low-contrast.

Thus, in the second version of the single coil homodyne reconstructive algorithm, I set this value to be a factor of 0.3 times the image size, though I kept this parameter at the top of the code where it can be easily changed in order to further optimize image quality. In order to facilitate further comparisons with other methods, I also changed the inputs that determined the sampling interval to a partial Fourier factor, where a partial Fourier factor of 0 meant no sampling, 0.5 corresponded to half the k-space being sampled, and 1 to the event of full sampling. [1]

4.2 Comparison with POCS, automation

Using this second version, I started the analysis of this method by comparing it to the projection onto convex sets (POCS) method. [5]

The POCS method treats partial Fourier reconstruction as a constrained optimization problem. Each constraint (in this case phase and data consistency, respectively) is defined as a convex set, and are then taken into account with an iterative projection algoritm. Points in the intersection of two convex sets can be found by alternating projections onto each set until the method converges. This method, while slower, provides a better reconstruction of the phase and better image quality, and thus made a good reference for comparison.

I used a POCS algorithm written by Martin Blaimer in order to do this comparison. [2]

The comparison, initially done in a single coil scenario, took into account four aspects of the images produced:

- Visual comparison of results
- Differences between amplitudes of the images, between methods and with a reference image, reconstructed from a full k-space, as well as RMSE values between them.
- Visual comparison of phases by taking the conjugate of one image and multiplying it with the other
- Noise propagation, consisting of twenty reconstructions with random noise added to each k-space, then taking the standard deviation of

values for each pixel. Later we took the averages of these values as well.

This resulted in several hundred images when run with a sufficient range of partial Fourier factors. Therefore, I automated the process of creating and saving the images and raw data, put the partial Fourier factor and averages on the image for simple comparison and picked a color scheme that gave the most information for each of these, which provided the framework for future comparisons in a multi-coil scenario.

4.3 Multi-coil reconstruction

In order to be used in parallel imaging, the reconstruction had to be usable for a multi-coil scenario. In a multi-coil scenario, each coil has its own profile, and images are created with the weighted linear combination of the images given by each individual coil. For the purposes of this work, the coil profiles were simulated by a MATLAB program. The work was done with 8 virtual coils of radius 2.56 cm.

There are two ways to do a multi-coil reconstruction:

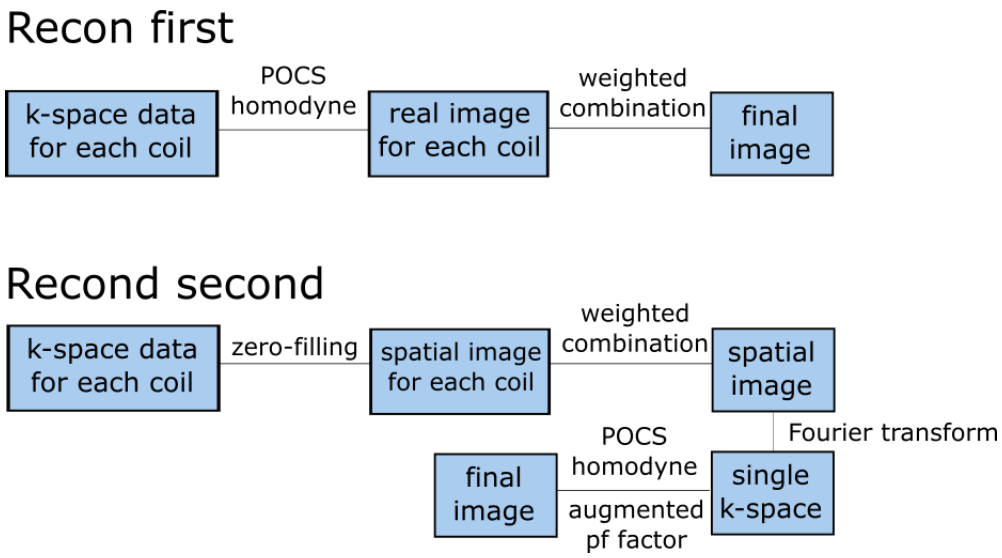

Figure 13: Recombination methods

Reconstruction first The k-space from each coil can be individually reconstructed using a homodyne or POCS reconstruction. The reconstructed images can then be combined using the coil profiles. For the homodyne reconstruction, as the homodyne method gives a real output, this is

the magnitude of the coil profiles, for the POCS method, this is the conjugate of the coil profiles.

Reconstruction second The k-space data from each coil can also be reconstructed by doing a zero filling reconstruction of each image. These spatial images are then recombined, weighted with the conjugate of the coil profiles. By doing another Fourier transform, a single k-space dataset is created.

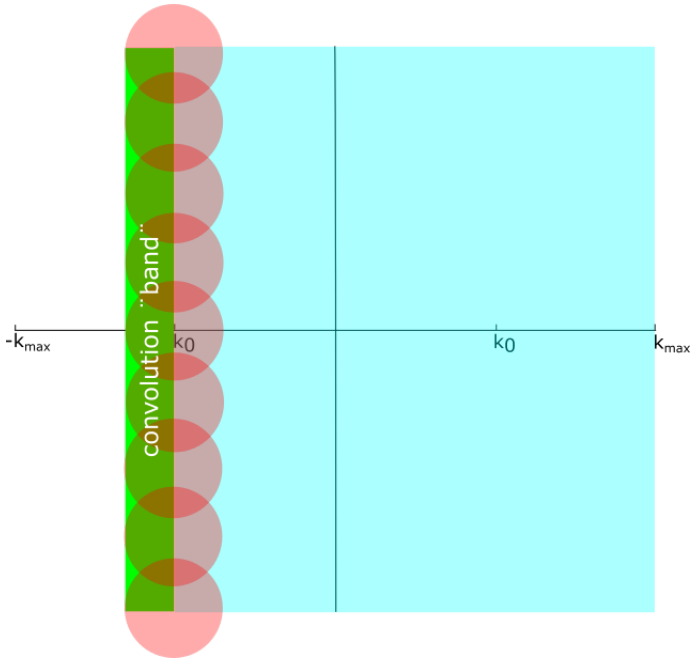

Figure 14: The effect of multiplication with the coil profiles.

It is important to note that this k-space data, while still only a partial data set, also is the product of convolution with the Fourier transform of the coil profiles (shown in red on the figure above) in the k-space, and thus has an extra "band" added below $-k_0$. Therefore, instead

of the original partial Fourier factor (corresponding to the area shown in blue), we must add half the size of the coil profiles to this partial Fourier factor (the area shown in green) and use this in the homodyne or POCS reconstruction, which gives the final image.

5 Results

Due to the nature of the homodyne process, only the real part of the image is usable. This means that the phase comparison within my work was mostly a check to see if the code was functioning as expected. This condition was met, as the homodyne had a simple, linear phase shift (the phase given to the phantom) while the POCS and zero filling functions had a zero phase, with the exception of "empty" areas were the noise dominated.

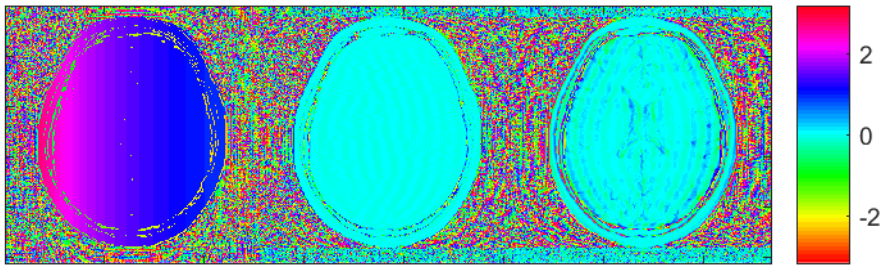

Figure 15: Phase comparison to full data set reconstruction for homodyne, POCS and zero filling methods respectively.

The rest of the results can be broken up into two categories: comparisons of the homodyne, POCS and zero filling methods done with the reconstruction before coil recombination, or done after recombination.

5.1 Comparison of POCS and homodyne with reconstruction first

The reconstruction was completed for partial Fourier factors between 9/16 and 15/16 with steps of 1/16. The RMSE values between the amplitudes of partial Fourier reconstructed images and the images reconstructed with a full data set were recorded and divided by the norm of the original image, as well as the averages of the noise propagation over the entire image, scaled to the noise added to the images (a random, complex noise created from a Gaussian distribution and of magnitude 10^{-3} over an image of size 256 by 256 pixels). These were then plotted as a function of the partial Fourier factor.

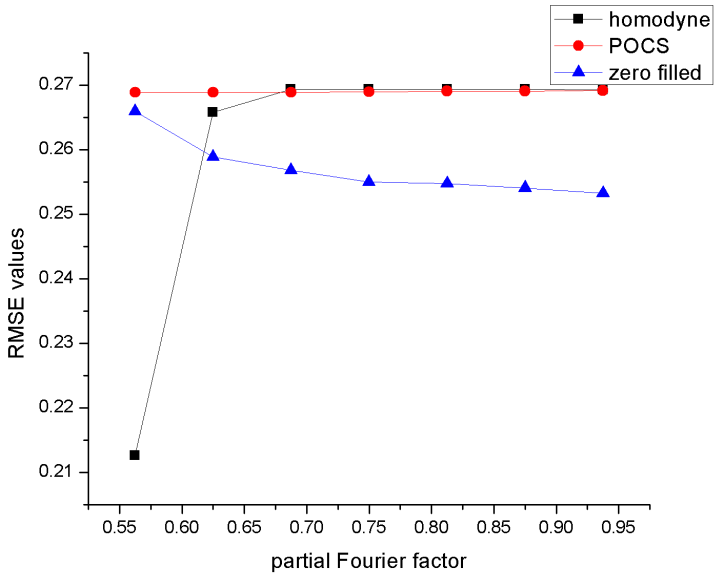

Figure 16: Relative RMSE of amplitudes for a brain slice for POCS, homodyne and zero-filling methods, reconstruction first.

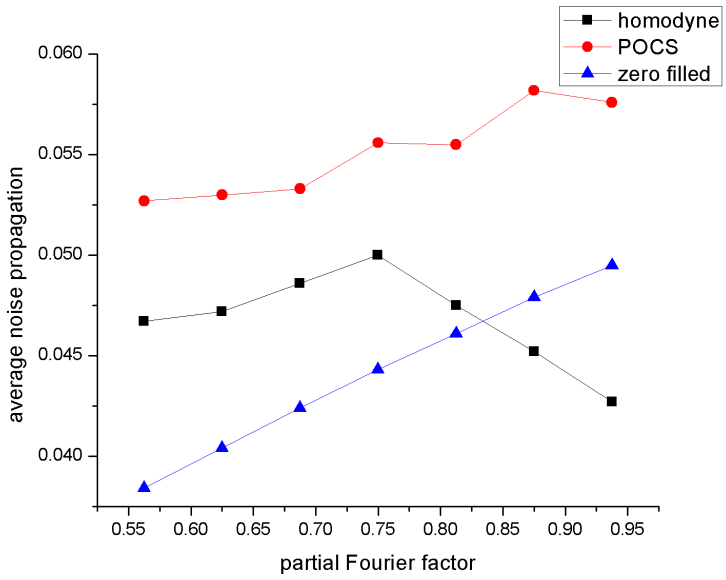

Figure 17: Relative average noise propagation for a brain slice for POCS, homodyne and zero-filling methods, reconstruction first.

For visual comparison, only partial Fourier factors between 5/8 and 7/8 are shown here, as smaller partial Fourier factors (for example, 9/16, shown below) had such large amounts of ringing they were functionally unusable, and higher partial Fourier factors gave no added information.

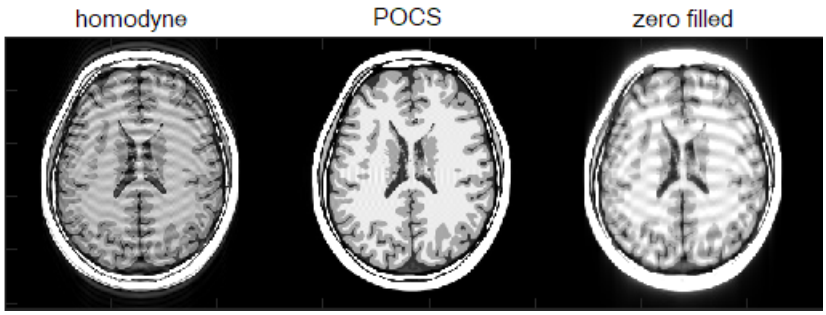

Figure 18: Images produced with small partial Fourier factor.

Below are the results of reconstruction in the region of interest. Amplitude differences were multiplied by a factor of 5 for increased visibility.

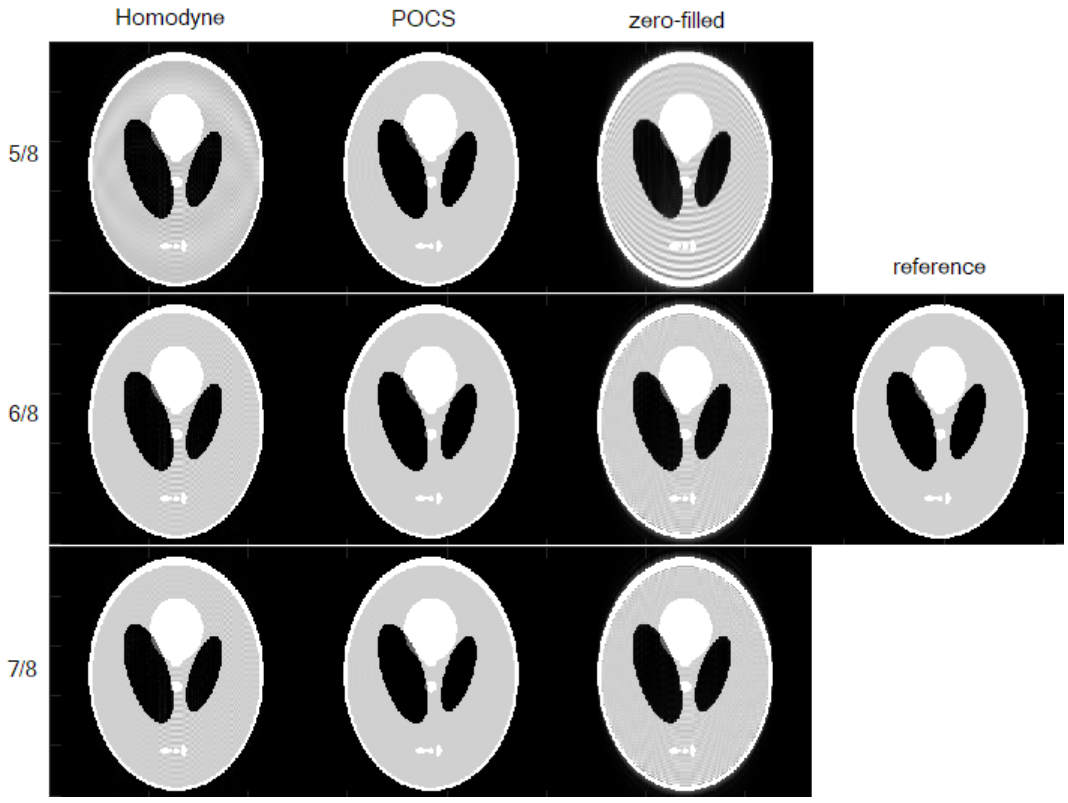

Figure 19: Shepp Logan phantom reconstructed with partial Fourier factors 5/8, 6/8 and 7/8, with a full data set as reference.

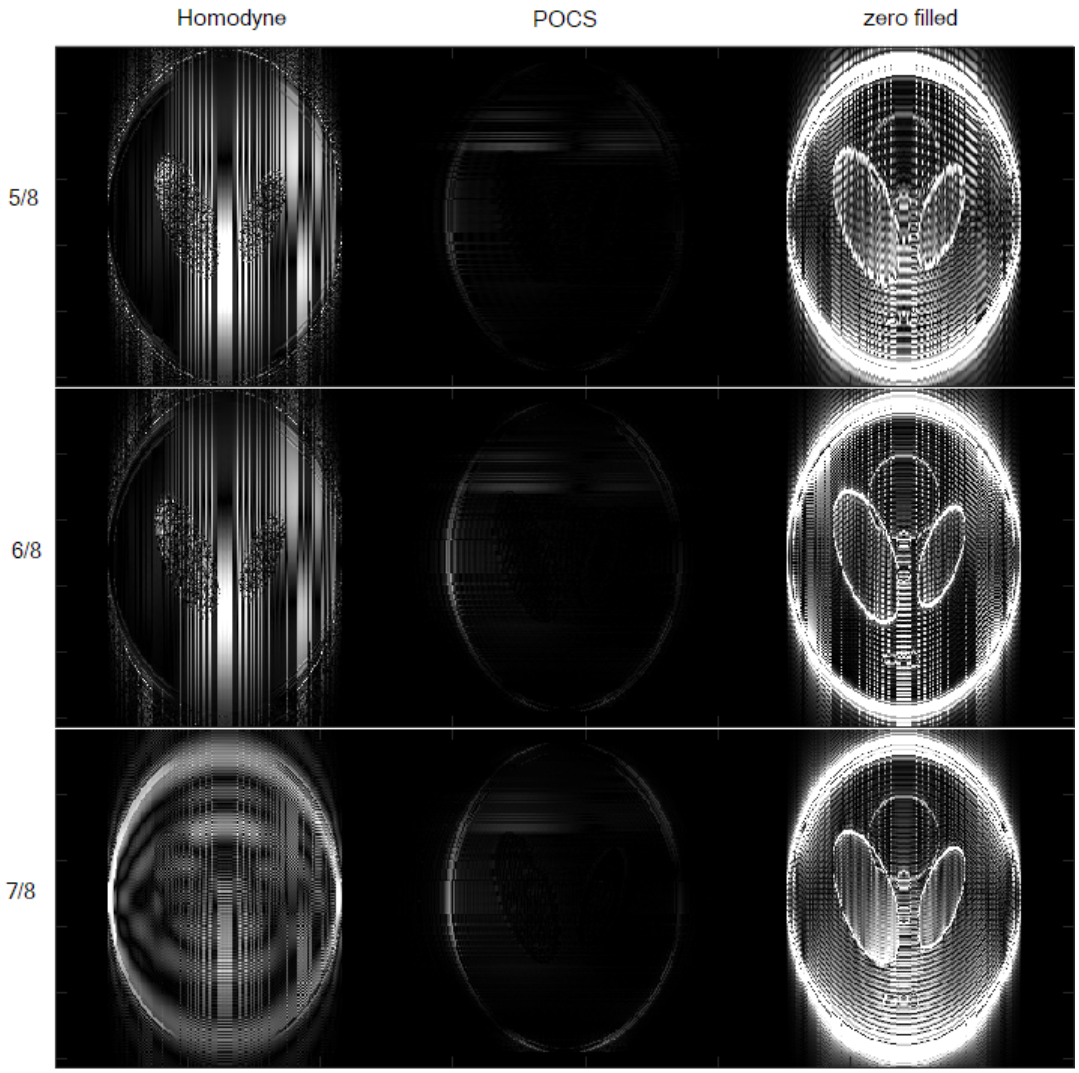

Figure 20: Amplitude differences between Shepp Logan phantom reconstructed with partial Fourier factors 5/8, 6/8 and 7/8 and reference image made with full data set.

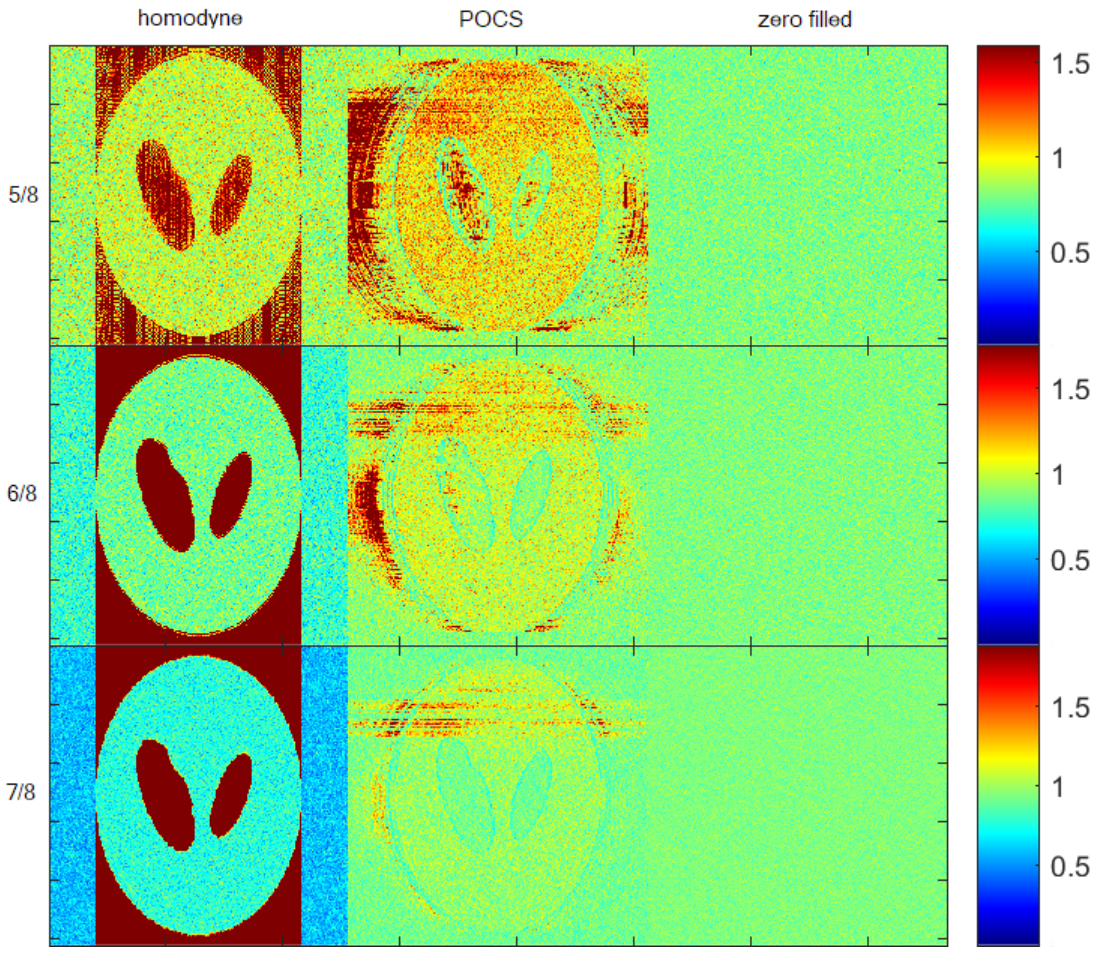

Figure 21: Noise propagation for Shepp Logan phantom reconstructed with partial Fourier factors 5/8, 6/8 and 7/8.

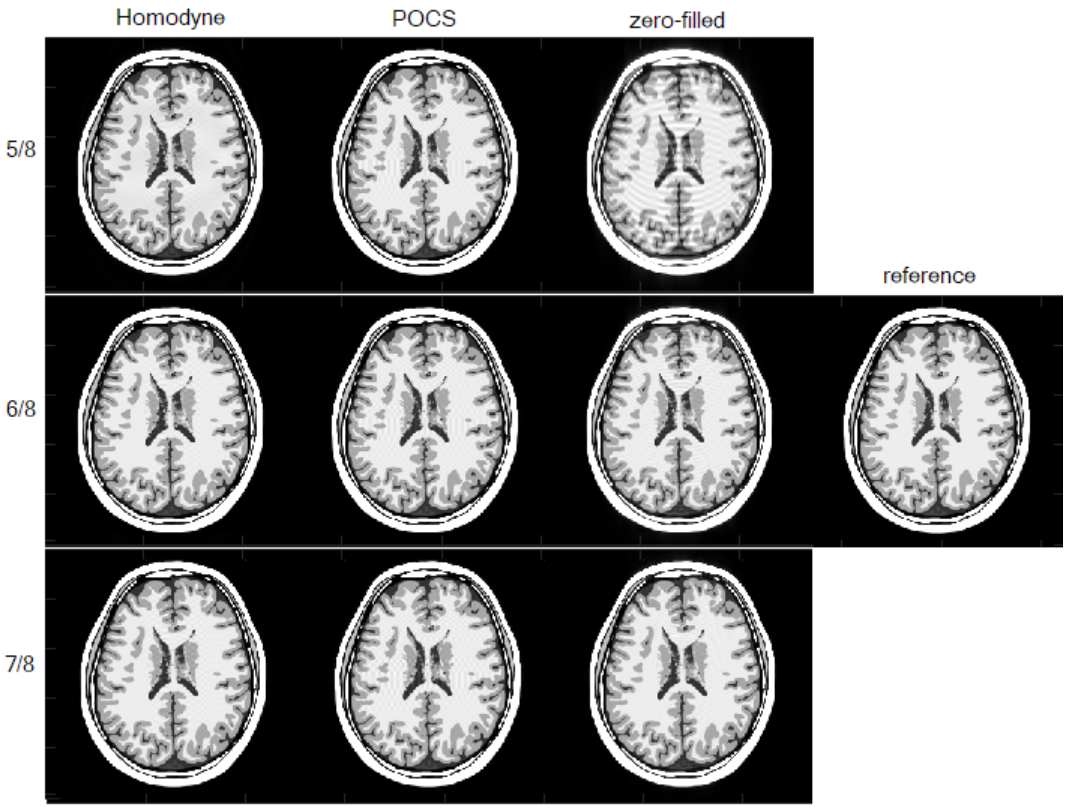

Figure 22: Brain slice reconstructed with partial Fourier factors 5/8, 6/8 and 7/8, with a full data set as reference.

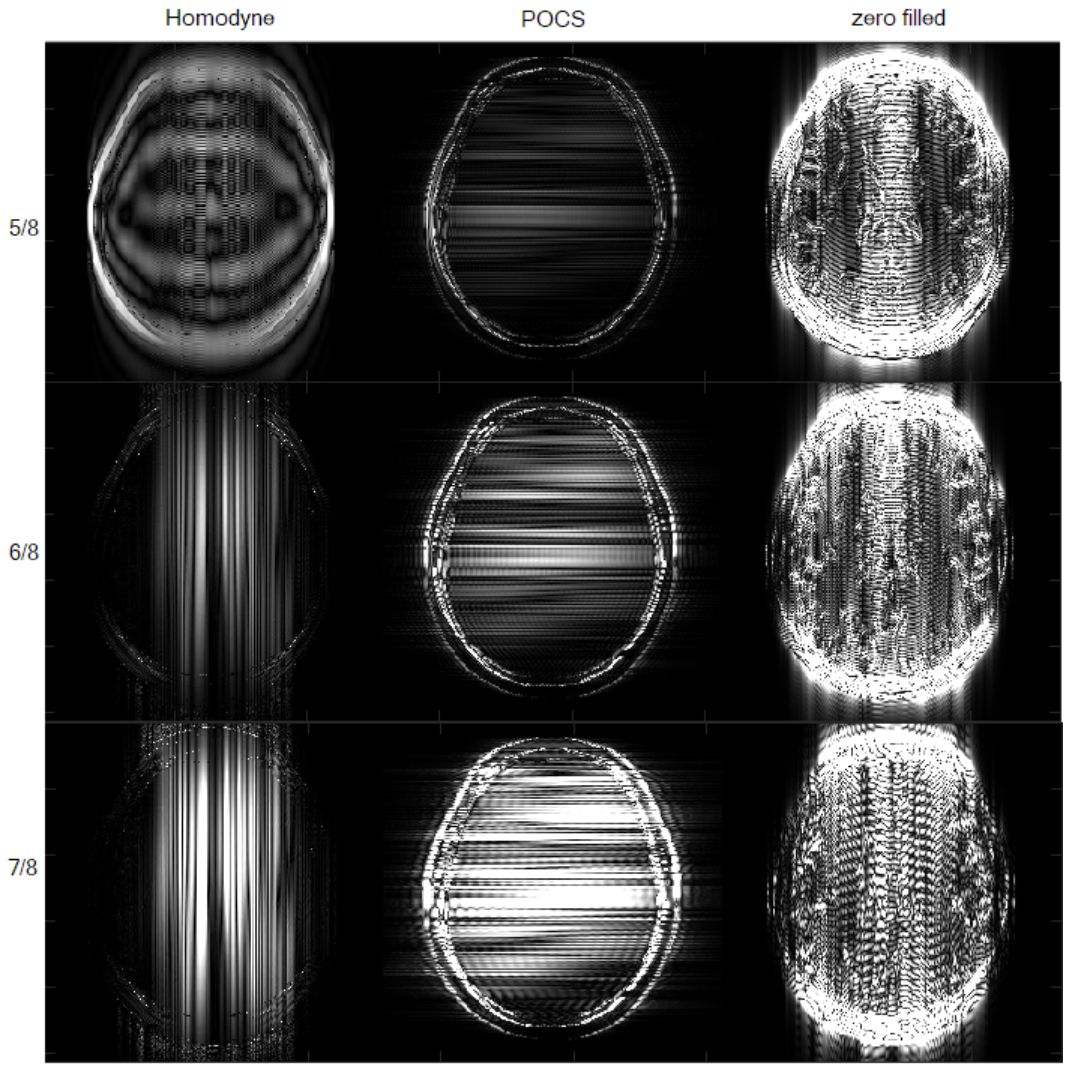

Figure 23: Amplitude differences between brain slice reconstructed with partial Fourier factors 5/8, 6/8 and 7/8 and reference image made with full data set.

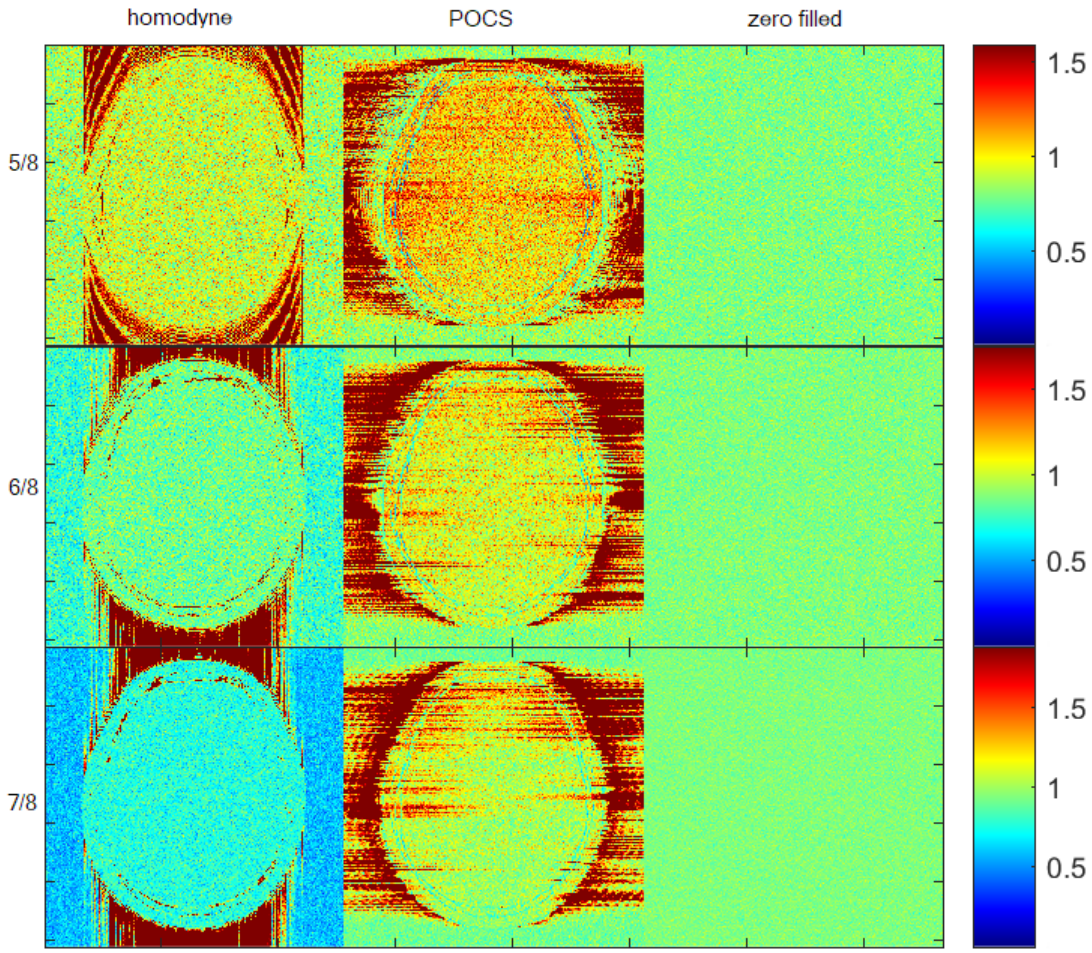

Figure 24: Noise propagation for Shepp Logan phantom reconstructed with partial Fourier factors 5/8, 6/8 and 7/8.

5.2 Comparison of POCS and homodyne with reconstruction second

The methodology in this section was identical to the one before, this time with the reconstruction method in question (homodyne, POCS, zero filling) being performed on a single k-space composed of all the coil information recombined.

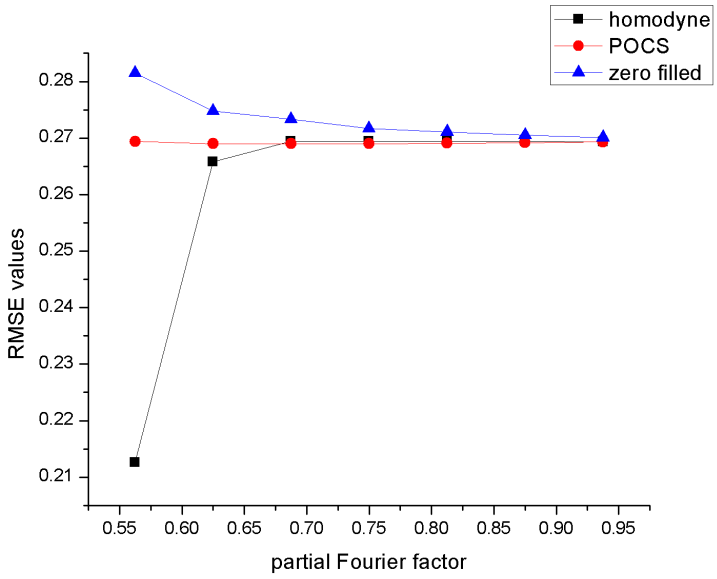

Figure 25: Relative RMSE of amplitudes for a brain slice for POCS, homodyne and zero-filling methods, reconstruction second.

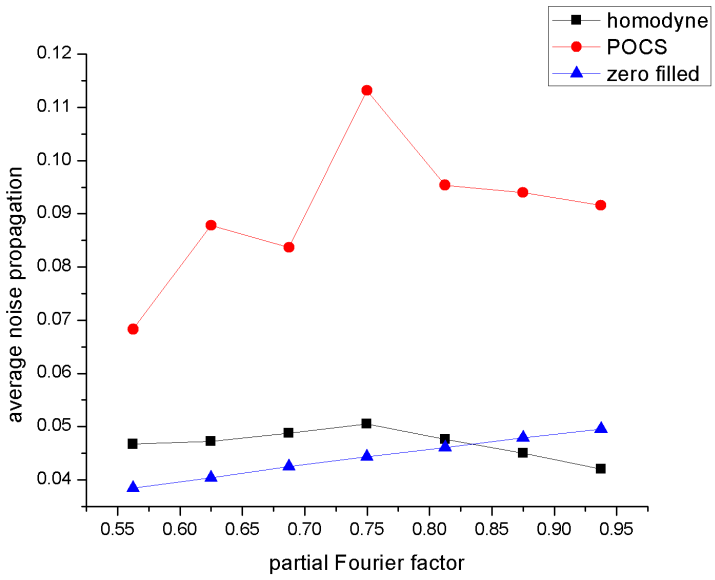

Figure 26: Relative average noise propagation for a brain slice for POCS, homodyne and zero-filling methods, reconstruction second.

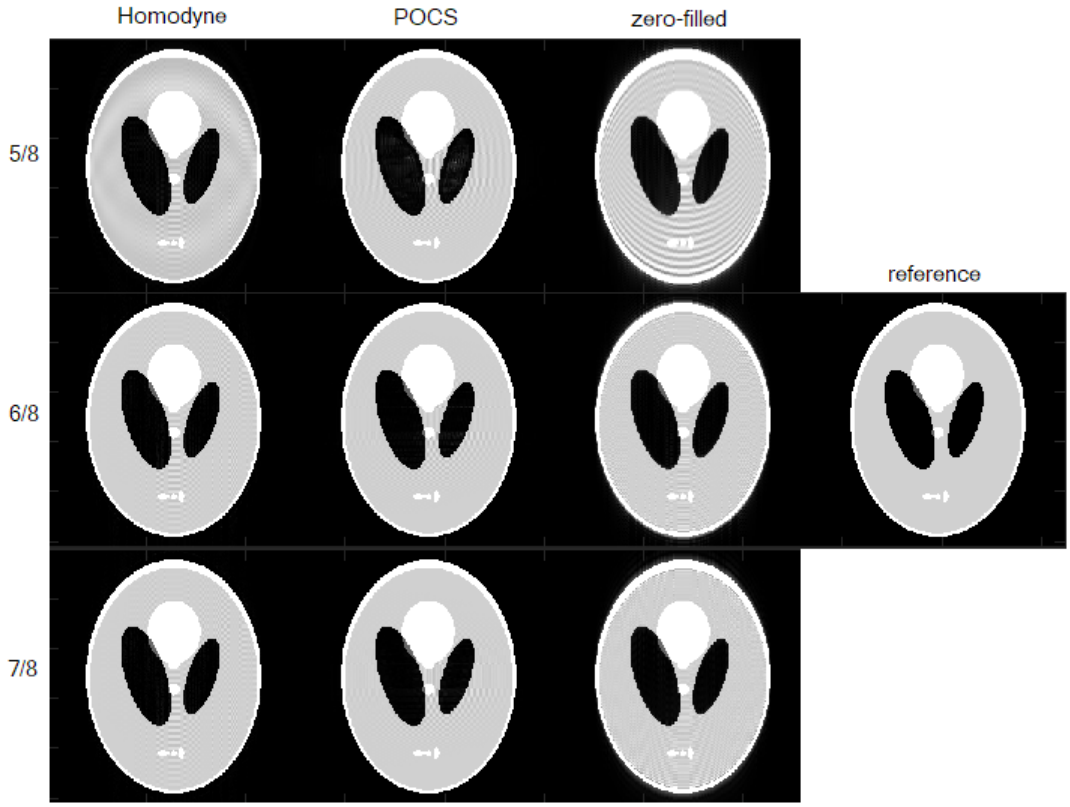

Figure 27: Shepp Logan phantom reconstructed with partial Fourier factors 5/8, 6/8 and 7/8, with a full data set as reference.

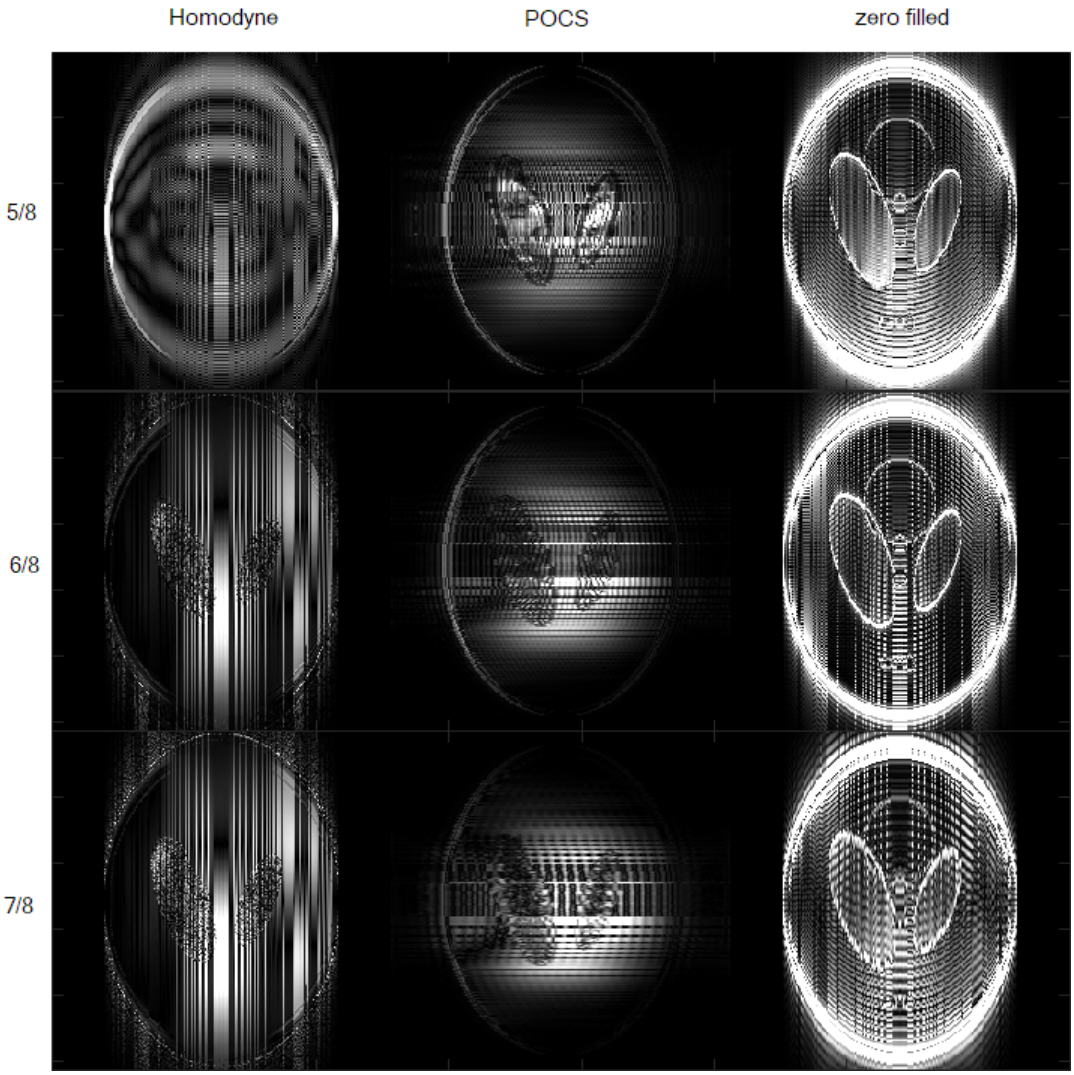

Figure 28: Amplitude differences between Shepp Logan phantom reconstructed with partial Fourier factors 5/8, 6/8 and 7/8 and reference image made with full data set.

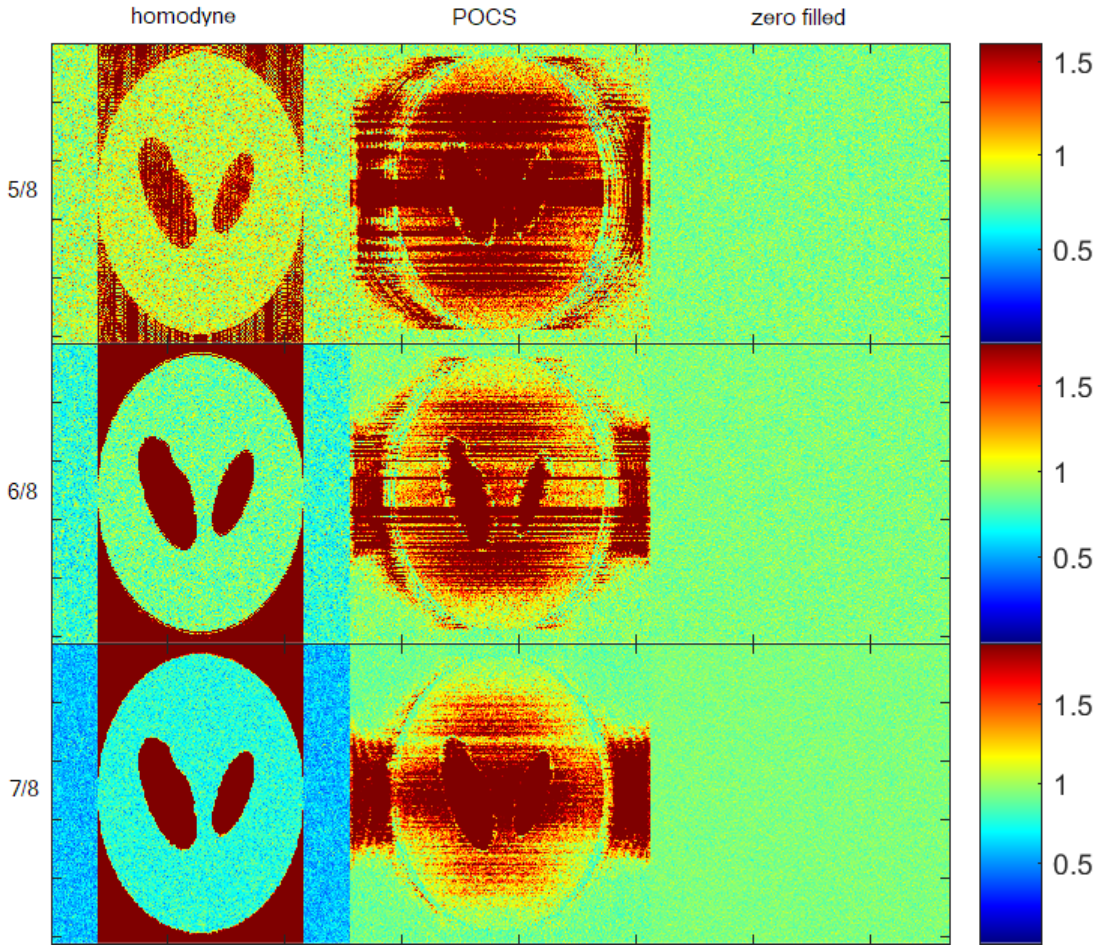

Figure 29: Noise propagation for Shepp Logan phantom reconstructed with partial Fourier factors 5/8, 6/8 and 7/8.

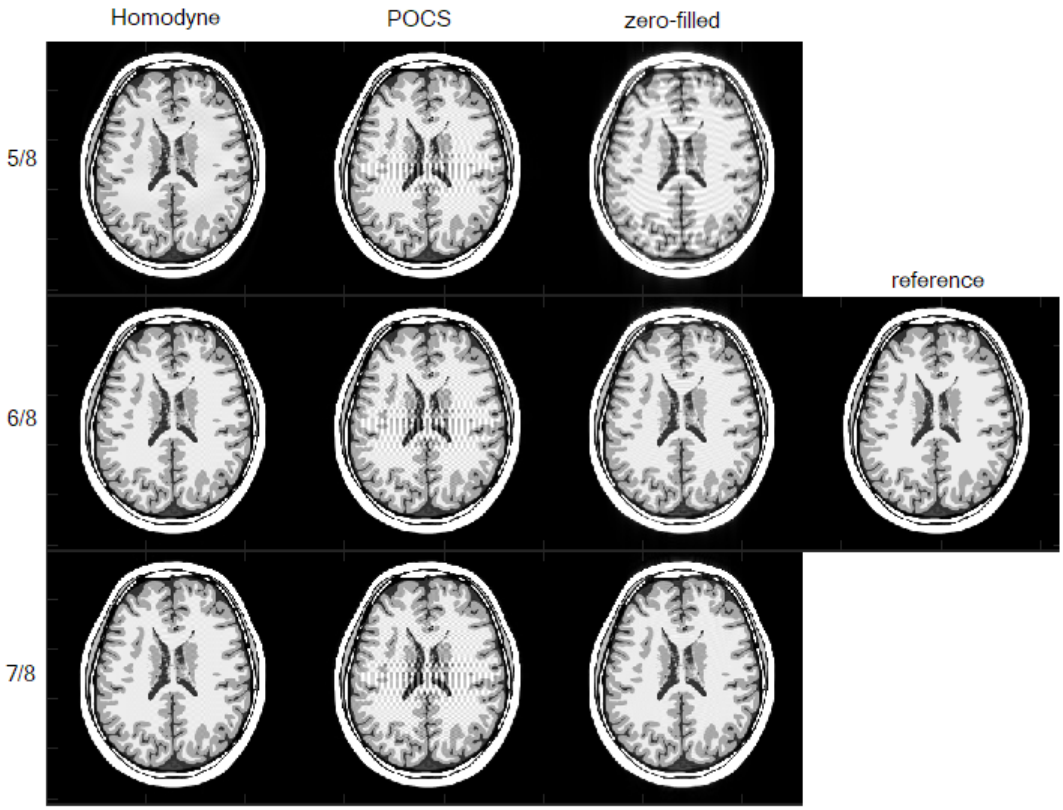

Figure 30: Brain slice reconstructed with partial Fourier factors 5/8, 6/8 and 7/8, with a full data set as reference.

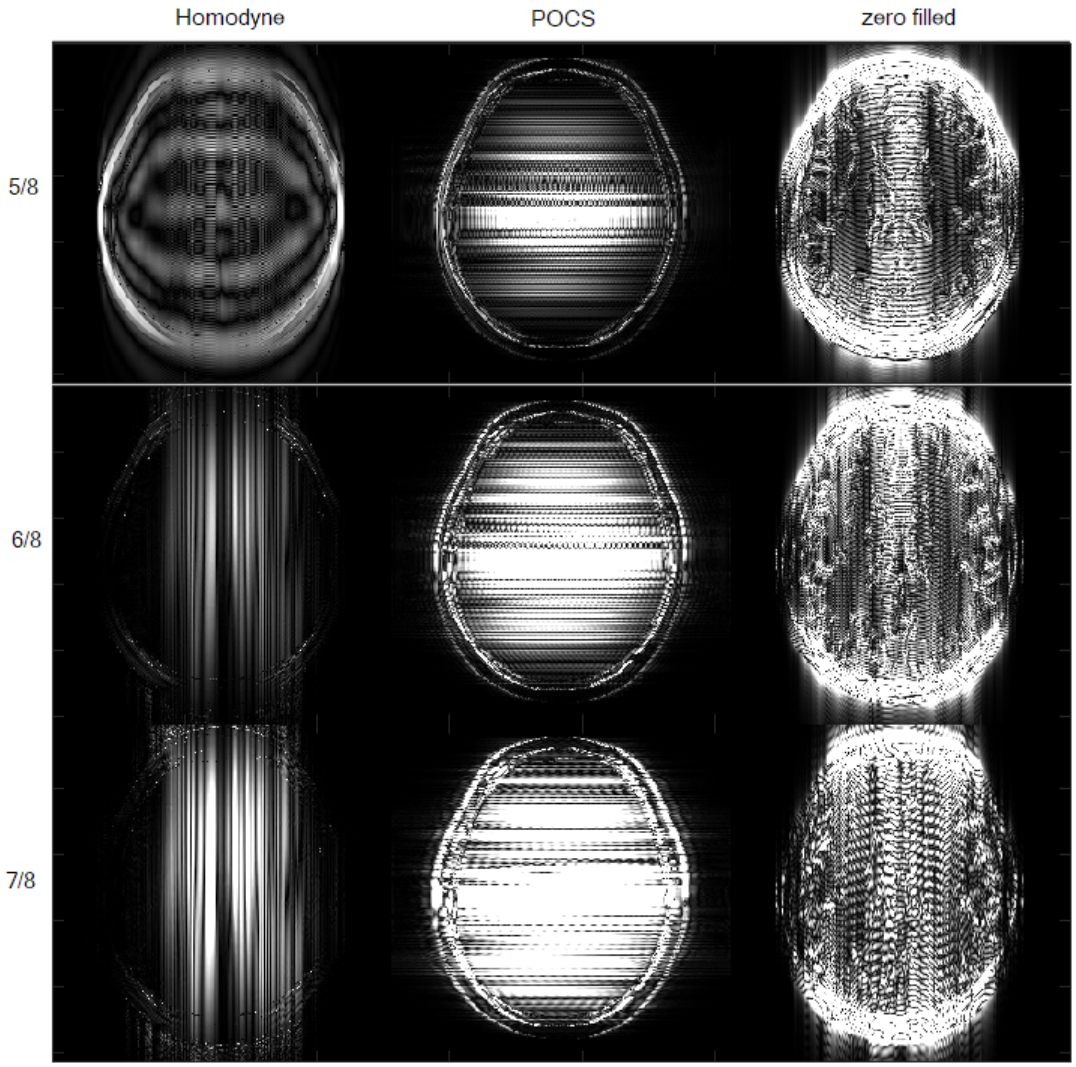

Figure 31: Amplitude differences between brain slice reconstructed with partial Fourier factors 5/8, 6/8 and 7/8 and reference image made with full data set.

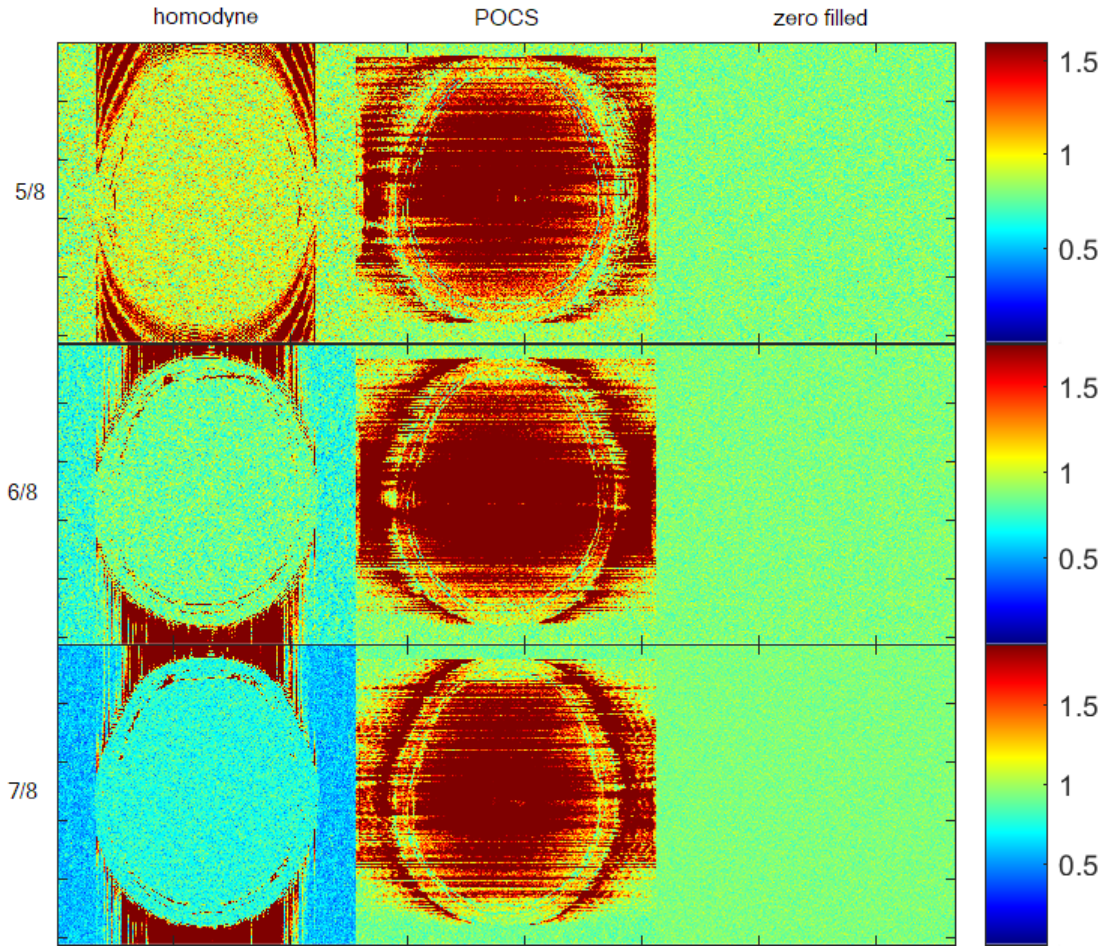

Figure 32: Noise propagation for Shepp Logan phantom reconstructed with partial Fourier factors 5/8, 6/8 and 7/8.

6 Conclusion

The homodyne method has many of the same issues as zero filling, with similar Gibbs ringing for lower partial Fourier factors. Its primary advantage over the POCS method is that it is significantly faster. This is unsurprising, as the POCS method is iterative and thus by nature a slower method.

We can see that at partial Fourier factors around 5/8, ringing is present in the homodyne methods, as can be seen on Figure 23, for example. However, the ringing is much less pronounced than that of the zero filling method. For a partial Fourier factor around 6/8 or more, this ringing is replaced by streaks similar to those seen in the POCS method, and while the homodyne method does not produce artefact free images at any partial Fourier factor, it produces much cleaner images and far less Gibbs ringing around edges than the zero filling method, and produces cleaner edges than the POCS method as well.

The RMSE values show similar behaviour in the POCS and homodyne methods. The RMSE values of the homodyne method at 9/16 are not indicative of the phantom produced, however later behaviour shows that the average error evens out at a partial Fourier factor of roughly 0.70, showing little different between the error in the homodyne and POCS methods.

The average noise propagation for the image, however, shows a greater difference between methods, with significantly larger values for the POCS method, as much as twice that of the homodyne.

Finally, the comparison between recombination orders shows that the event where reconstruction with the primary method is performed for each

coil separately, before recombination, gives a slightly better RMSE value. However, this difference, while consistent, is not significant and does not visibly reduce the amount of artefacts produced.

References

[1] M. Berstein, K.F King, and X.J. Zhou. *Handbook of MRI Pulse Sequences*. Pearson Education/Prentice Hall, 2004.

[2] M. Blaimer. Mri partial fourier reconstruction with pocs. 2012.

[3] R. Brown, E. Haacke Y. Cheng, M. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging 2nd edition*. 2014.

[4] F. Wübbeling F. Natterer. *Mathematical Methods in Image Reconstruction*. The MathWorks Inc., 2001.

[5] E.M. Haacke, E.D. Lindskogj, and W.Lin. A fast, iterative, partial-fourier technique capable of local phase recovery. *Journal of Magnetic Resonance*, 92(1):126–145, 1991.

[6] J. O. Smith III. *Mathematics of the Discrete Fourier Transform (DFT), with Audio Applications, 2nd Edition*. W3K Publishing, 2007.

[7] MATLAB. *R2016a*. The MathWorks Inc., Natick, Massachusetts, 2016.

[8] Alexander D. Poularikas. *Handbook of Formulas and Tables for Signal Processing*. CRC Press, 1998.