

Reduction of Noise and Artefacts in MRI- Scan of Metal Implant Reconstructed by Frequency Mapping

Pontus Wahlqvist

2022

Supervisor: Månsson and Magnus Cinthio

Examiner: Johannes Töger

Faculty of Engineering LTH
Master's Thesis in Biomedical Engineering
Department of Biomedical Engineering



LUND UNIVERSITY

Contents

Abstract	IV
Abbreviations	V
Acknowledgements	VI
1. Introduction	1
2. Aim	3
3. Theory	4
3.1 Magnetic resonance imaging	4
3.1.1 Magnetic fields	4
3.1.2 Gradient fields	4
3.1.3 Physics and signal	4
3.1.4 Slice selection and localization	5
3.1.5 Pulse sequences	5
3.1.6 Metal-induced artefacts	7
3.2 Modern methods of metal artefact reduction	8
3.2.1 View angle tilting	8
3.2.2 Slice encoding for metal artefact corrections	8
3.3 Fourier shift	9
4. Prior Work	11
4.1 Phantom	11
4.1.1 Frequency Map	14
4.2 Shifting the Data	16
4.2.1 The Shifting Method	16
4.2.2 Reconstructed Data	20
5. Method	22
5.1 Simulations on synthetic phantoms	22
5.1.1 Synthetic Datasets: Cylinder and Sphere	22
5.2 Prosthetic Phantom	23
5.2.1 Filtering of the frequency map	23
5.2.2 Interpolation	24
5.3 Image evaluation	26
6. Results	27
6.1 Simulations on synthetic phantoms	27

6.1.1 Cylinder 1	27
6.1.2 Cylinder 2	30
6.1.3 Sphere	32
6.2 Phantom of the hip prosthesis	34
6.2.1 Filtering of the frequency map.....	34
6.2.2 Interpolation	35
6.2.3 Combined interpolation and filtering.....	36
6.2.4 Assessment of noise levels.....	39
6.3 Final Comparisons	41
7. Discussion.....	43
7.1 Synthetic Data: Cylinder and Sphere	43
7.2 Prosthesis	44
7.2.1 Initial Filtering	44
7.2.2 Interpolation	44
7.2.3 Filtering combined with Interpolation and Analysis of Noise Levels	45
7.2.4 Study limitations	46
7.2.5 Future improvements	46
8. Conclusion	48
References.....	49
9. Appendix	51
Interpolation.....	51
Combined interpolation and Filtering	54
Final data	57

Abstract

Metal objects can cause significant artefacts in magnetic resonance imaging (MRI) scans, cause loss of signal and what is known as geometric warping. There are ways to address these problems using different MRI sequences. This project looks at a newer method of artefact correction proposed by Månsson and Morin.

The aim is to improve upon the main two faults of the method. The method makes use of a frequency map taken alongside the image dataset which describes the warping of the magnetic field caused by the metal object. This frequency map is then used to estimate how far data should be shifted back to correct for the geometric warping. The data is then shifted back using Fourier Shift. The main limitations of the method are a noticeable increase in noise present in the image post correction as well as the unexpected appearance of zebra-stripe artefacts.

The improvements are made from two parts: filtering and interpolation. These methods were tested on both synthetic datasets generated using MATLAB, as well as a previously created phantom.

By filtering the frequency map using a three-dimensional median filter prior to applying the correction method, the noise could be reduced by about a factor 2 in the final image. Additionally, interpolating both datasets prior to applying the method caused a significant reduction to the zebra-stripe artefacts.

Abbreviations

SNR – Signal-to-noise Ratio

MRI – Magnetic resonance Imaging

CT – Computed tomography

RF – Radio frequency

SEMAC – Slice encoding for metal artefact correction

VAT – View angle tilting

ROI – Region of Interest

TE – Echo Time

SE – Spine-echo

Acknowledgements

Firstly, I would like to thank both of my supervisors, Månsson and Magnus Cinthio, for their assistance and guidance throughout the various steps of the project. I would also like to thank the multitude of people who read and helped improve the report. Finally, I would also like to thank Morin who, together with Månsson, created the shifting method that was the start of this master's thesis.

1. Introduction

Hip and knee prosthetics are a powerful medical tool to help assist the large group of people who struggle with arthritis. This group of patients can encompass people of all ages, however; it primarily affects people over 50 [1]. According to the Swedish hip prosthesis register, the number of hip prostheses implanted in Sweden in 2019 alone was around 20 000 [1]. The procedure helps millions of people around the world and is a powerful tool. To make sure the implantation is as smooth as possible different imaging modalities must be used to image the prosthesis once it is in place in the body.

These image modalities include magnetic resonance imaging (MRI), and x-ray-based imaging modalities like computed tomography (CT). The main advantage of MRI is better visibility of soft tissue, and that no radiation enters the body, and it is therefore safer to do [2]. However, using an MRI scan to image prosthetics has a major drawback when the implant is made of metal [3]. When a metal object (even when it is considered MRI-conditional) is placed in the MRI scanner, it produces no detectable signal, but more importantly the local magnetic field around the metal object warps significantly, thereby affecting the resulting image [3]. The warping is three dimensional and causes distortions both within the imaged slice and also causing the imaged slice to bend capturing data that should appear in neighboring slices instead of the intended data.

There are several methods that work well to reduce the artefacts in the image, one of the most notable is called “slice encoding for metal artefact correction” or SEMAC [4], which reformats the imaging volume to obtain straight slices. However, this report will focus on further developing a potential alternative to SEMAC. This alternative first measures the changes in the warped magnetic field due to the nearby metal object and then, using these measurements, attempts to correct the image. This method has been shown to work in a previous project in 2012 [5]. However, the method is in its early development and has some noticeable issues that include a significant loss in the signal-to-noise ratio (SNR), as well as the appearance of zebra-striped artefacts. These artefacts are further visualized in Figure 1. This project will focus on finding the source of these problems as well as attempt to address them.

Assuming improvements made, this method could have several distinct advantages over the current method of artefact correction. First and most notable is the reduction in scan time. This method of artefact reduction requires far less time spent in the MRI scanner compared to the current method. Furthermore, this method also uses much fewer radio frequency (RF) pulses in the imaging process. This leads to less heating of the patient, which can be a problem in certain clinical MRI scans especially if the patient has a metallic implant [6].




Figure 1: The zebra stripe artefacts are circled in red, and the noise is present throughout most of the figure but is further highlighted in the blue square.

2. Aim

The purpose of this project is to improve upon the results of the metal-induced artefact restoration, proposed by Morin and Månsson [5]. Specifically, the goal is to reduce the noise exhibited by the results as well as reduce or remove the unexpected zebra-stripe artefacts that appeared after reconstruction. This will be done through further testing on the same dataset as well as using synthetically created datasets. The focus will be on finding the source of the artefacts as well as attempting to address them.

3. Theory

3.1 Magnetic resonance imaging

The following section will detail the basics of how an MR image is created and some of the problems that appear will be crucial to understand the purpose of the report.

3.1.1 Magnetic fields

A contemporary clinical MRI system consists of several different magnetic coils [7]. The main magnet generates a constant, stable magnetic field inside the patient bore, this field is typically denoted B_0 . In clinical MRI systems this field is typically either 1.5 T or 3 T [7].

3.1.2 Gradient fields

Aside from the main field, MRI systems make use of three different coils, known as gradient coils. These coils generate three varying magnetic fields in orthogonal directions. The magnetic fields generated by the gradient coils are superimposed upon the main field, B_0 , and thereby cause the field to vary slightly with position [7]. Simply put, a patient lying in the MRI scanner will experience slightly different magnetic fields at different positions, for example at the chest compared to at the head. Unlike B_0 , these gradient fields can be adjusted or turned off and on as circumstances dictate [7].

3.1.3 Physics and signal

The MRI signal in humans comes from the magnetic moments of hydrogen nuclei in the body. Hydrogen nuclei consist of only a single proton with a positive charge. Protons are always spinning, and when the proton spins it generates a moving charge, also known as a current [7]. Furthermore, any current will also give rise to a corresponding magnetic field which can then be measured. Due to this magnetic field the individual protons behave like a point dipole (like a simple bar magnet) with a north and a south pole separated by a distance [7]. Generally, protons will spin at with random phases and point in random directions. However, in presence of a strong magnetic field (like in an MR system) some will align parallel to the field causing a slight magnetization vector in the direction of the field [8].

When the protons are placed in the magnetic field they will also spin in a very particular way, like a spinning top, known as precession. The speed of this precession is directly correlated with the strength of the magnetic field according to the formula, $\omega_0 = \gamma * B_0$ [7]. This precession frequency is also known as the Larmor frequency (ω_0 , in MHz). B_0 describes the strength of the magnetic field and γ is a constant dependent on the atom type used to image. For MRI this atom will almost always be hydrogen and $\gamma = 42.6$ MHz [7, 9].

Initially, the net magnetic moment of the protons cannot be measured as it is aligned with B_0 . Therefore, another step needs to be taken to disturb this alignment. When a proton rotating with a certain precession frequency is hit by a radio frequency (RF) pulse with the same frequency, energy is transferred from the pulse to the proton, causing the proton to fall out of

alignment with B_0 [7]. The main effect of the energy transfer is to shift the spin distribution of the protons [8]. The measurable effect is that the net magnetic moment starts to rotate around the direction of the B_0 field noticeably. MRI is more complex than this, but this is a basic summary of where the signal comes from. A much more in-depth description of the process can be found in the article “Understanding MRI: basic MR physics for physicians”, written by Currie et al [7].

3.1.4 Slice selection and localization

The gradients, mentioned earlier, play different roles in localization of signal. The one most relevant to this report is the slice selection gradient. This gradient remains turned on while the RF pulse is sent out and is used to determine where in the body (or object being imaged) the slice will be taken. The gradient causes a linear variation in magnetic field strength from one end of the scan area to the other [9]. Since the precession frequency of the protons depends on the strength of the magnetic field, this frequency will also change along the gradient. Since the RF pulse must be matched to the precession frequency of the protons, only a small slice of the imaged object will be excited by the pulse. In effect, this means that signal collected will only be from this slice (see Figure 2).




Figure 2: A demonstration of what happens when a correctly tuned RF pulse is sent into an object. Since the RF pulse is tuned to a certain frequency only the protons with this precession frequency (red) will respond to it.

Should the RF pulse used be monochromatic, (only containing one wavelength) the slice selected would be infinitely thin. This would cause problems such as large amounts of data, and longer scan times. Therefore, to obtain a finite thickness of the slice, the RF pulse must contain a range of frequencies. Most importantly, this whole localization process rests upon the assumption that any change in the magnetic field over the scan area is due to the gradient coils and that the field B_0 is entirely constant this can be violated when metal is present and the field bends [9].

3.1.5 Pulse sequences

The gradients are turned on and off or adjusted according to the selected MRI pulse sequence. The pulse sequence is a pre-programmed set of commands for how the scan will work. It contains information regarding parameters such as, echo time (TE), repetition time (TR), flip angle and others. TE is the time at which the signal produced is measured [7]. TR is the time

from the first 90-degree pulse to the second 90-degree pulse, in other words, the time between sequences. Flip angle what decides the degree of pulses.

A commonly used pulse sequence is the spin-echo (SE) sequence. The basics of the SE sequence is that the protons in the body are first excited by a 90-degree pulse [10, 7]. Then after a time (half of the TE) they are again excited by a 180-degree pulse [10]. This pulse is also called a refocusing pulse. The reason for using the SE sequence is to correct for field or tissue inhomogeneities. Once the 90 degree pulse has hit the protons immediately start to lose their phase coherence at different speeds due to in part tissue differences. However, when the refocusing pulse is applied they start to refocus at the same speed that they defocused at. After a certain period of time they will be almost completely focused again, this is when the measurement is done. This is visualized by Figure 3.

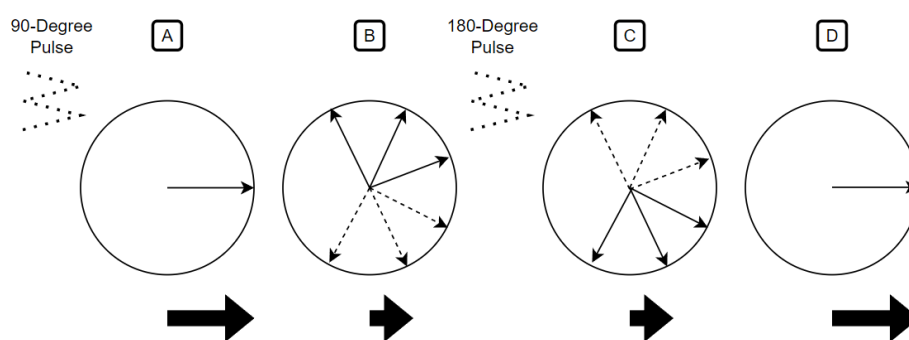


Figure 3: Demonstration of the importance of the refocusing pulse. The thick arrows represent the net magnetization vector. First when the 90-degree pulse hits, all the individual magnetization vectors are in unison as seen in [A]. Then after a short period of time they slowly start to lose coherence [B]. This leads to a decrease in the net magnetization vector. The refocusing pulse (180 degree) essentially causes the individual magnetization vectors to rotate 180 degrees effectively causing them to switch direction [C]. Then, all the individual magnetization vectors rotate at the same speed as before but in the opposite direction which causes them all to become in phase for a short time in [D]. This is where the measurement occurs or the so-called echo time.

As mentioned earlier, the pulse sequence also decides when the different gradients are turned on or off. As can be seen in Figure 4, the slice selection gradient is turned on during both the 90-degree pulse as well as the refocusing pulse. However, the frequency encoding gradient is turned on at multiple points during the schematic including during the echo. Therefore, it is also sometimes called the readout gradient [10]. Some rectangles are shaded grey in the schematic. This is because they “cancel” out one another and is called prewinding of the gradients. In the readout gradient the prewinding have the same polarity, this is because the refocusing pulse will reverse the effect of the prewinding.

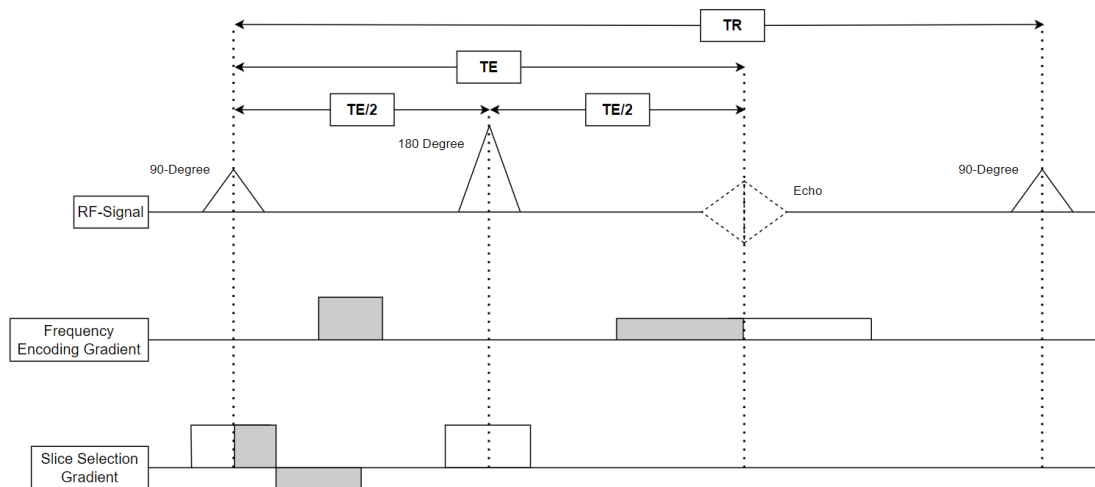


Figure 4: A slightly simplified version of a spin-echo sequence. Note: In reality, the time in between the echo and the TR is much longer than what is shown in the image.

3.1.6 Metal-induced artefacts

Placing magnetic or ferromagnetic objects such as various metals into a magnetic field can cause a host of problems. If ferromagnetic, the metallic object may experience strong forces acting upon it, it may also cause localized heating around itself. These are issues that must be considered when designing metallic implants for use in the human body. For the most part, modern implants are designed with the express intention that they can be used in an MRI by, for example, choosing appropriate materials. However, many implants are still MRI unsafe [3].

A metallic object placed in a uniform static magnetic field, like the one in an MRI system, will warp the magnetic field around it causing local variations in the otherwise uniform field [3]. The strength of the warping is decided mainly by the magnetic susceptibility of the metallic object, or more specifically, the difference in magnetic susceptibility between the object and the surroundings it is placed in [3]. There are other factors that play a part as well, such as shape of the object and its orientation compared to the magnetic field.

As mentioned earlier, the localization of signal heavily depends on the fact that the magnetic field, B_0 is uniform except for the gradient field added to it [9]. Unfortunately, as the warping of the magnetic field causes variations in the magnetic field, a constant frequency will not excite a flat slice see Figure 5. To disturb the alignment of the protons in the targeted volume the RF pulse must be the same as their precession frequency, which is dependent on the strength of the magnetic field. The RF pulse also acts upon the whole body even if only the area with the correct frequency is directly affected. Therefore, if a metallic object is causing warping the magnetic field within a volume some distance away from the target volume it can be imaged within the volume (see Figure 5). These artefacts are known as geometric distortion [3].

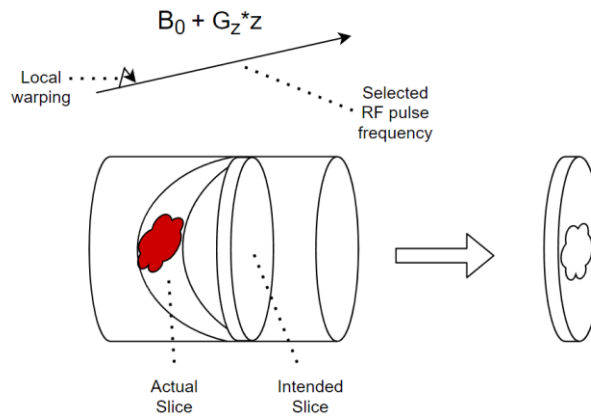


Figure 5: Geometric distortion: the red cloud which is placed some distance away from the intended slice is captured within the actual slice as, due to the magnetic field warping, it has the same precession frequency. In the final image (on the right) the cloud appears to be in the volume imaged in the slice.

3.2 Modern methods of metal artefact reduction

This section describes the basics of two methods of metal artefact reduction commonly used in MRI. The information will remain at a very basic level as a deep understanding of these methods is not relevant to the rest of this report and would require significant understanding of the intricacies of MRI. For the curious, the articles cited give a much more thorough explanation of the methods.

3.2.1 View angle tilting

View angle tilting (VAT) is a technique that is used commonly to correct for in-plane distortions. It does this by applying a gradient in the slice-selection direction simultaneously with the readout direction, with identical amplitude as the slice gradient applied during the slice-selection. This causes the actual viewing of the image to be at a slight angle, decided by the strengths of the two gradients, which can then be used to compensate for the in-plane distortion. [11].

3.2.2 Slice encoding for metal artefact corrections

To correct for through-plane metal induced artefacts, a method called slice encoding for metal artefact corrections (SEMAC) can be used. The method is conventionally used together with view angle tilting, as SEMAC corrects for the through-plane and VAT corrects for the in-plane distortions. SEMAC is by no means a poor method and is a very powerful tool for metal artefact correction. However, the scan time for SEMAC is quite long when compared to that of a simple scan using exclusively VAT [12, 4]. Furthermore, SEMAC as a method uses a lot of RF pulses to image the patient. This can be a problem as it will cause some slight heating of the patient and can lead to some discomfort.

3.3 Fourier shift

Simplified, the Fourier shift theorem states that for a function a translation in real space corresponds to a multiplication of a linear phase in the Fourier space. This is expressed mathematically as $h(x) = f(x - x_0) \Rightarrow \hat{h}(\xi) = e^{-2\pi i x_0 \xi} * \hat{f}(\xi)$ where $\hat{h}(\xi)$ and $\hat{f}(\xi)$ denote the Fourier transforms of $h(x)$ and $f(x)$ respectively. Additionally, x_0 must be a real number. For example, for a series of values like the one shown below in Figure 6 where everything is 0 except for a pulse in position 0, we can call this $f(x)$.




Figure 6: A dataset where all values are 0 apart from the pulse in position 0, where the value is 1.

The Fourier transform of Figure 6 it will show up as a straight line of only ones (see Figure 7). This we can denote $\hat{f}(\xi)$.




Figure 7: The Fourier transform of the dataset in Figure 6.

The shift is completed by a multiplication of this Fourier transform by a phase, given by the equation earlier as $2\pi x_0$ where x_0 is the distance the pulse should be translated, in this case x_0 will be 2. In Figure 8 the real part of the entire exponential function $e^{-2\pi i x_0}$ is shown.




Figure 8: The real part of the exponential function from the Fourier shift equation.

Multiplying the phase with the Fourier transform from Figure 7 and then performing an inverse Fourier transform on the result will show that the peak was shifted from the 0 position two step in the positive direction to position 2 (Figure 9, top).




Figure 9: Top: The shifted data with $x_0 = 2$. Bottom: The shifted data with $x_0 = 2.4$.

If the value for x_0 is changed to for example, -3, the peak will instead shift backwards three steps. If the shifting number x_0 is not a whole number, the result is affected dramatically. This can be seen in the bottom Figure of Figure 9 where the peak has lost some of its shape and the surrounding numbers are no longer equal to 0. This method of shifting data is used extensively in this project and is the basis for the whole methodology.

4. Prior Work

The work done prior to this thesis consisted of creating a phantom model made up of a stainless-steel hip prosthesis submerged into an agarose gel. Next, an image mapping the fluctuations of the magnetic field in the phantom was created. Finally, the program to shift the data back to its origin was created and used by Morin, Månsson and Müller [5].

Everything in this section of the report is data that was collected and processed prior to the start of this master's thesis but is essential to understanding the work that was done. This work was done primarily in two different master's theses, one written by Fredrik Wellman and the other by Morin [13, 5].

4.1 Phantom

A Phantom is an object created to mimic an actual test object. The phantom is used to gain a better understanding of something before testing it on the real object (a patient in this case).

The phantom used for this project was originally created during the master's Thesis of Fredrik Wellman and performed at the department for Medical Radiation Physics at Lund University [13].

The phantom was created by placing a Stryker Exeter stainless steel hip implant (Stryker Howmedica Osteonics, Allendale, NJ, USA) into a plastic box along with a rectilinear Perspex (plexiglass) grid as demonstrated by Figure 10. The box was then filled with an agarose gel, in part placed there for structural fixation. Importantly, it is also a medium with tissue-like properties [13].




Figure 10: A figure demonstrating how the phantom looks. The blue grid is made of plexiglass and the cylinder shows where the hip prosthesis is placed. The rest of the volume in the box is filled with agarose gel.

The full process of making the gel and the phantom is described both in text and with images in Fredrik Wellman's Thesis [13]. To further clarify, Figure 11 shows an in-progress photo of the phantom being made.



Figure 11: Image showing the prosthetic half-submerged in agarose gel along with the plexiglass grid. Taken from [14] with permission from the publisher (Elsevier). For scale, the box is approximately 20cm by 30cm.

The dataset was taken on a 1.5T-Siemens scanner and consists of two parts. One part is the normal MR image of the phantom this is shown below in Figure 12. This measurement was a normal VAT spin-echo sequence and therefore removes in-plane geometric distortions. The second part of the dataset is a map of how the magnetic field varies with position, this is called the field map or the frequency map.

Both datasets were taken with a TR/TE of 600/5 ms, matrix of 128x96 and a bandwidth of 977 Hz/px. They both consist of 32 slices with a 3mm thickness. The sequences were both multiple 2D slice sequences which were then combined into a 3D image.

This data is what one would expect from this type of measurement when SEMAC (or another correction method) is not applied. Slices 16 and 17 of Figure 12. demonstrate that the plexiglass grid surrounding the prosthetic has been warped and shifted to the surrounding slices. It can also be seen that the grid nearest to the head of the prosthesis has been shifted the furthest (see slice 16 to 20). This is likely due to more warping of the magnetic field leading to larger inhomogeneities.

Raw MRI Data




Figure 12: The initial MRI of the phantom showing all 32 slices.

4.1.1 Frequency Map

The frequency map (or field map) was created using a modified version of the spin echo sequence demonstrated earlier (see Figure 4). It was constructed from the phase difference of two spin-echo scans, one where the spin echo was centered in the acquisition window (a normal spin-echo sequence) and one where the spin echo was slightly offset.

The frequency map, also referred to as field map, is an accompanying dataset to the actual data. It shows how the magnetic field of the dataset has warped at every datapoint in the image. It is worth noting that for optimal scans during an MRI the magnetic field should be completely uniform apart from the gradients applied for localization.

The frequency map was taken so that the phantom was in the exact same position as the original dataset (Figure 12) to ensure each point in the datasets correspond completely.

In Figure 13, slices 9 to 24 (of 32) of the measured frequency map are shown.

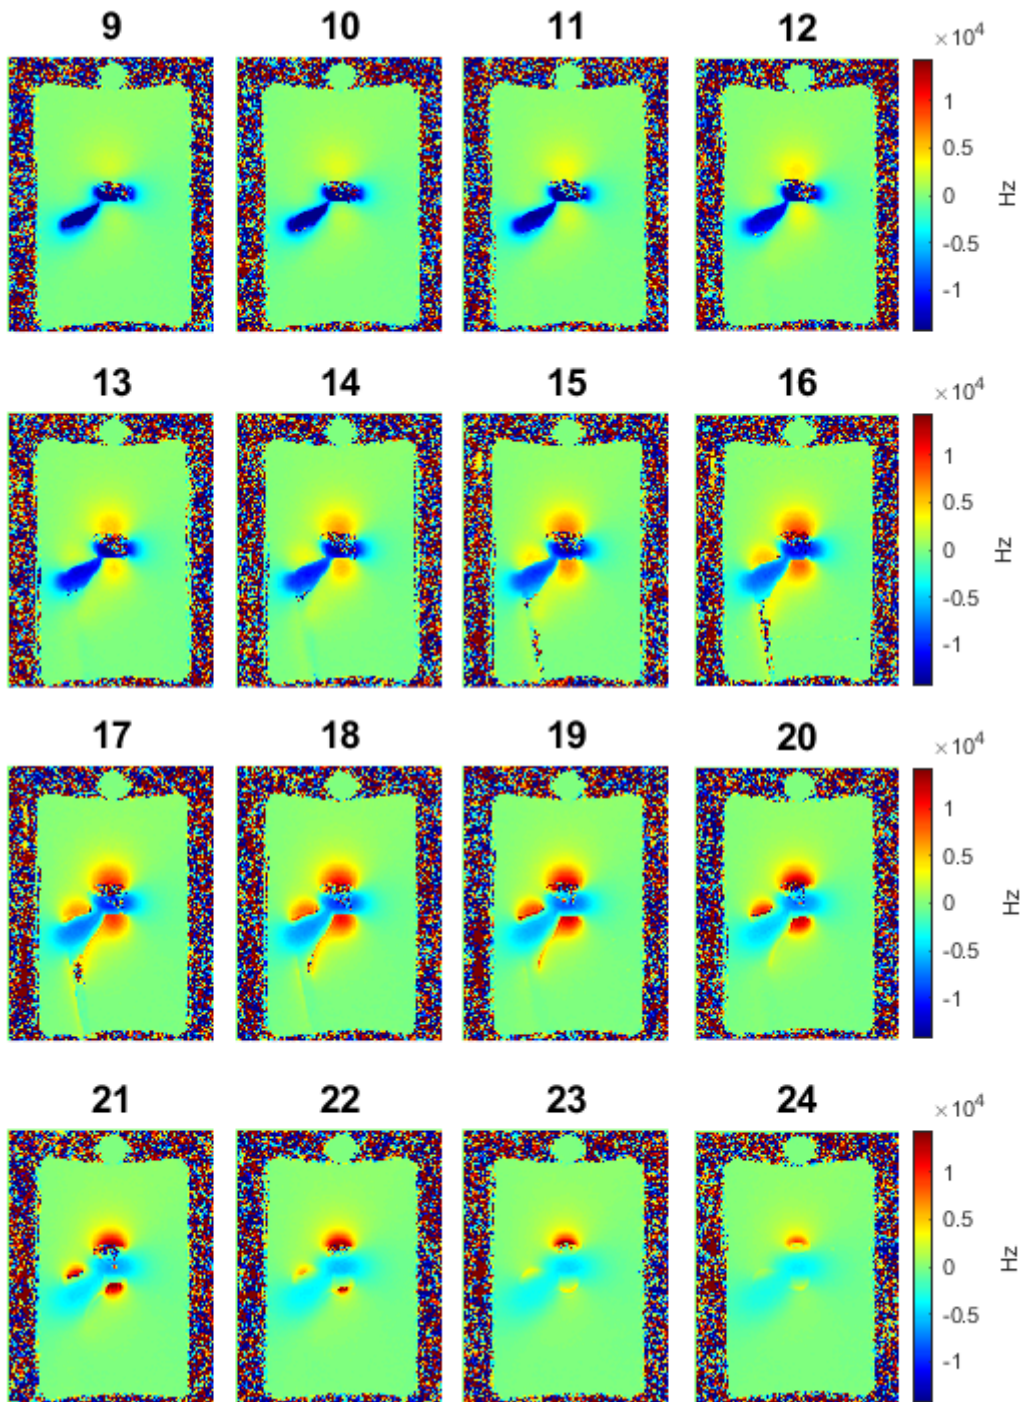


Figure 13: Slices 9-24 of the measured frequency map of the phantom.

4.2 Shifting the Data

Using the data from the frequency map, it is possible to know how far the corresponding data is shifted due to the metal-induced warping as the Shift will be proportional to the value of the frequency map. Using this information and the Fourier shift theorem described in Section 3.3 Fourier shift, data was shifted back to where it should be.

4.2.1 The Shifting Method

To run the Fourier shift method two datasets are required, firstly the raw data of the MR scan is needed, but aside from this, a dataset describing the magnetic field in each position is also required.

Assume there is an object like what is shown in Figure 14 (could be a thin box lying in water). Assume also, that the accompanying field map looks like the one shown in Figure 15. Realistically, this would not be the case; however, it is good for demonstration purposes.




Figure 14: The first example dataset to demonstrate the shifting method.




Figure 15: A frequency map for the images shown in Figure 14.

The shifting method will work through the four slices sequentially, then add them all up in the end to get the total result. As a simplification, the method checks the values of the frequency map in one slice to see what needs to be shifted. Like with the earlier example explained in the Section 3.3 Fourier shift, a positive value of 1 will shift the data 1 step to the next slice. A value of negative 1 will shift it in the opposite direction. The shifting is done pixel by pixel for every pixel on the slice. Then, the method moves on to the next slice. In the end it adds all the results together to get the full image, see Figure 16 and Figure 17.

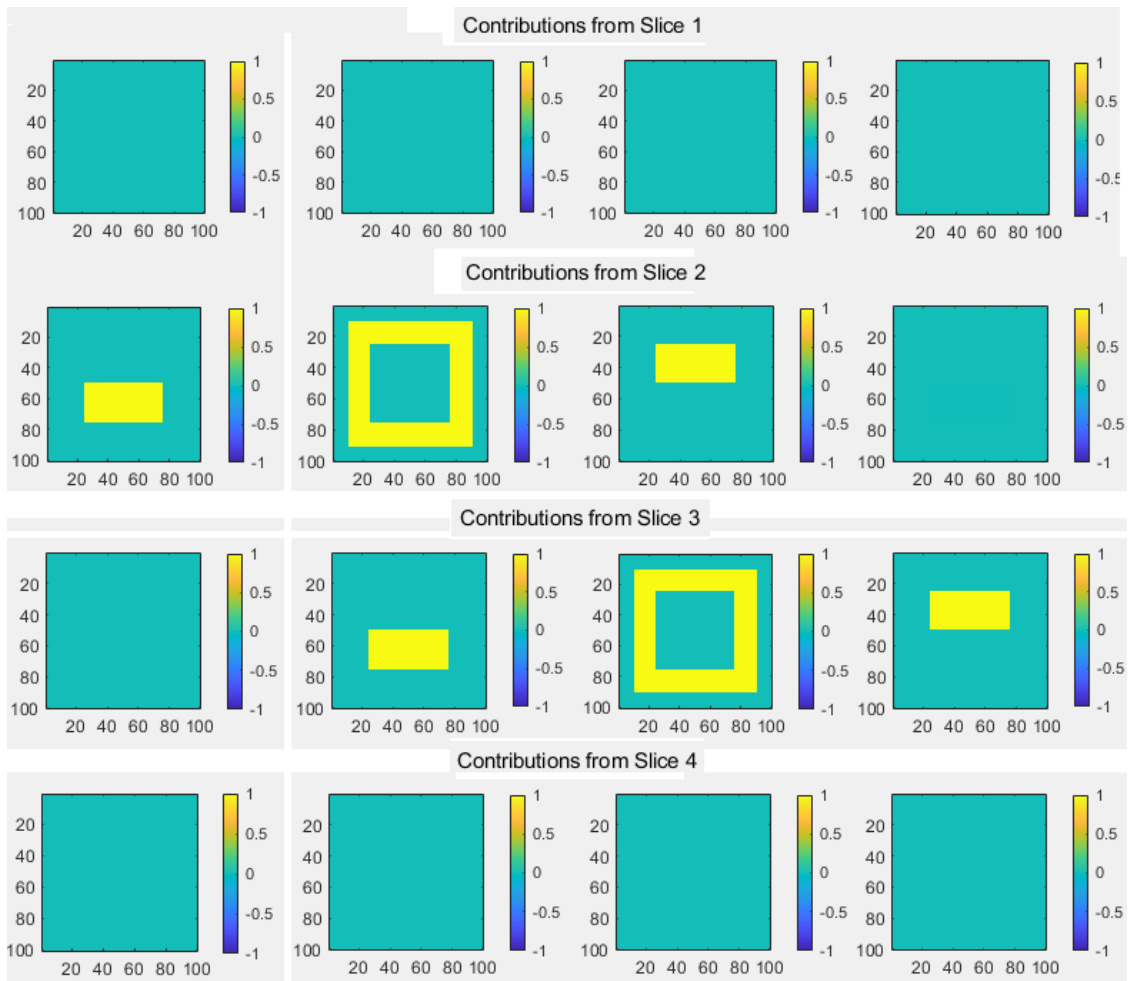


Figure 16: The individual contributions from each slice. Data in the first square of each row provides data for slice 1 in the result (in Figure 17).




Figure 17: The resulting figure when combining all the individual contributions from Figure 16.

Figure 17 is generated simply by adding the results from Figure 16. This ends up being a very successful shift in that all the data is successfully shifted without any artefacts. However, if we use a similar case, but with double the number of slices (for clarity see Figure 18) and change the values in the frequency map to something that is not a whole number, like 1.3 and -1.3, the contribution will look a bit different.




Figure 18: The initial new (second dataset) box with 8 slices.

The contribution from slices 1 to 3 and 6 to 8 will be zero like the previous case as there is no data there to shift. But the contributions from slices 4 and 5 give an unexpected result (Figure 19). Here it can be seen that instead of just shifting the data from one slice to another, the data is smoothed out over all the slices. This will be important to keep in mind when examining the results later. Another thing to note is that data wraps around from the last slice to the first slice. The result for this shift is irrelevant as it is just a demonstration of the method. However, for the curious, it is shown in Figure 20.

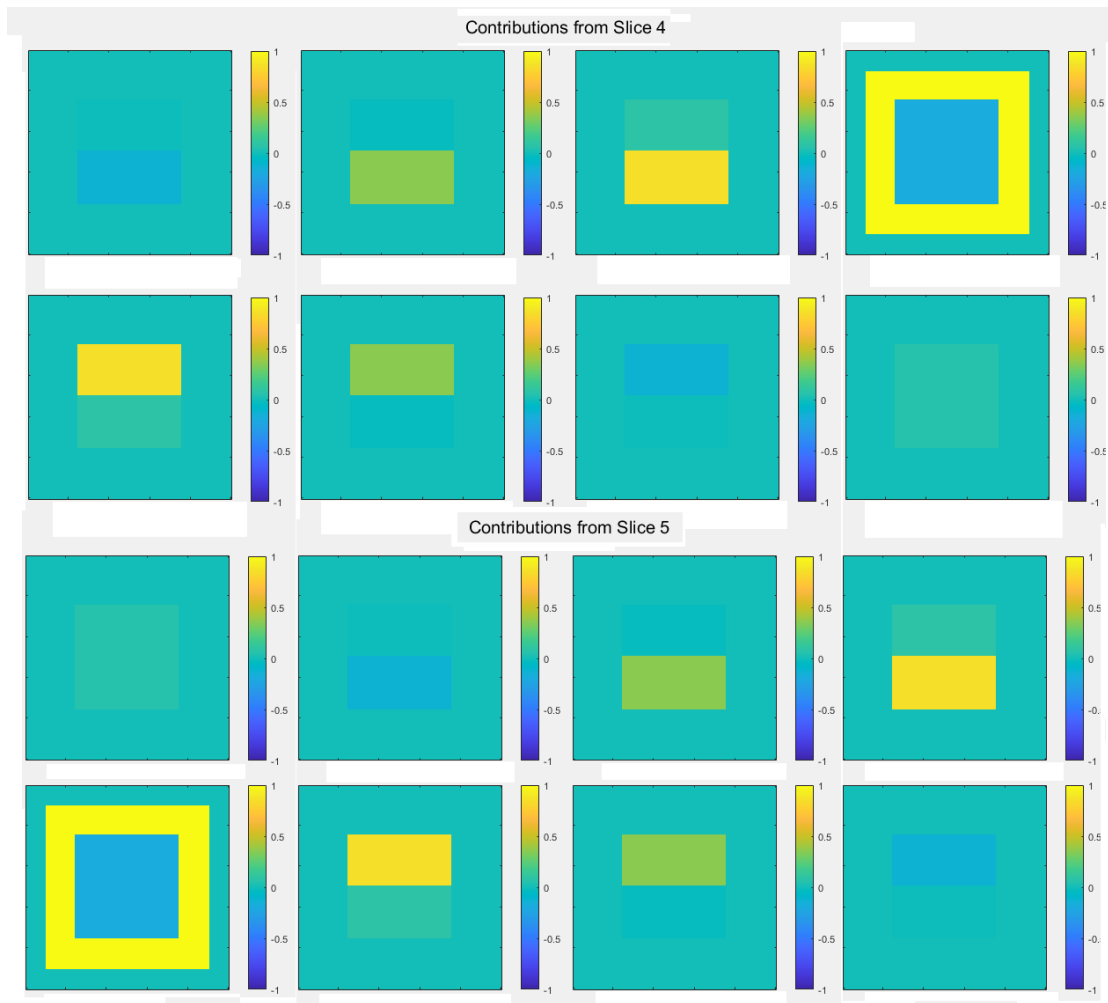


Figure 19: The contributions from each slice towards the reconstructed image (seen in Figure 20).

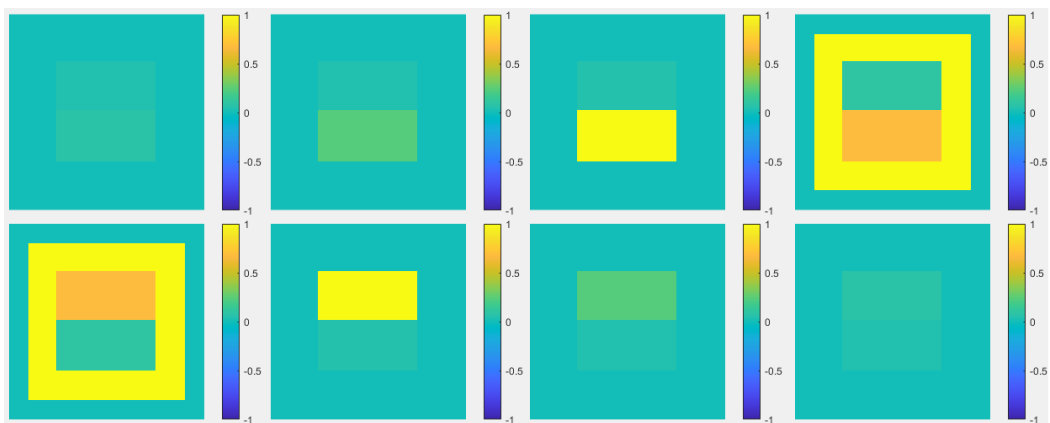


Figure 20: The box after passing through the Fourier Shift.

4.2.2 Reconstructed Data

The reconstructed data of the phantom after performing the Fourier shift program on the raw data is shown in Figure 21. This reconstructed dataset is where the prior work by Morin and Månsson finished.




Figure 21: Slices 9-24 of the phantom after the reconstruction method of Morin and Månsson has been applied.

The main advantage of this measurement is that it is clear from observing the grid that much of the data has been shifted back to its original position. However, as seen more clearly in Figure 22 the processing appears to have induced a lot of noise in the area where the agarose gel would be. Additionally, around the head of the prosthesis and a few other areas, a zebra pattern artefact appears to have formed. These zebra-patterns are more visible in other slices of Figure 21.

Raw Data vs Processed Data in Slice 16

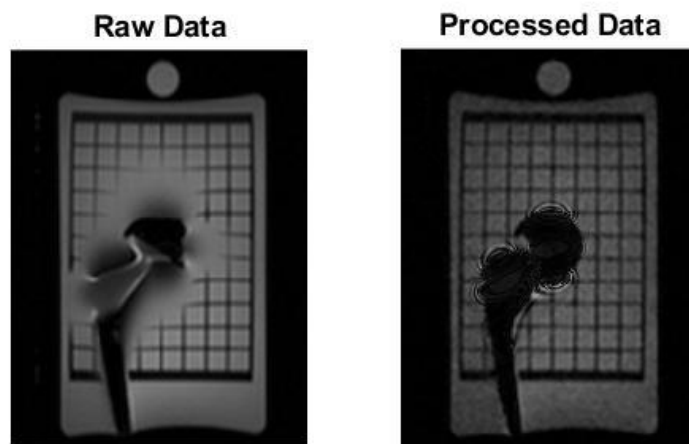


Figure 22: A comparison of slice 16 in the unprocessed and processed datasets.

5. Method

5.1 Simulations on synthetic phantoms

There were several tasks examined during the process of this project. First, synthetic datasets were created in MATLAB using simple shapes: a sphere and a cylinder. These shapes were chosen due to the relative ease to create an accompanying, analytically calculated synthetic frequency maps [15]. The goal here was to attempt to reproduce the same artefacts in a completely synthetic environment to determine whether the method itself causes the problem or something else. The testing process can be seen in Figure 23. First the original synthetic data is distorted by applying the shifting method with a frequency map where all the positive values have been changed to negative and vice versa. This is to simulate the data being distorted by geometric artefacts in an MRI scan. Next, the resulting distorted data is shifted back using the frequency map with normal values. The result is then the Restored data.

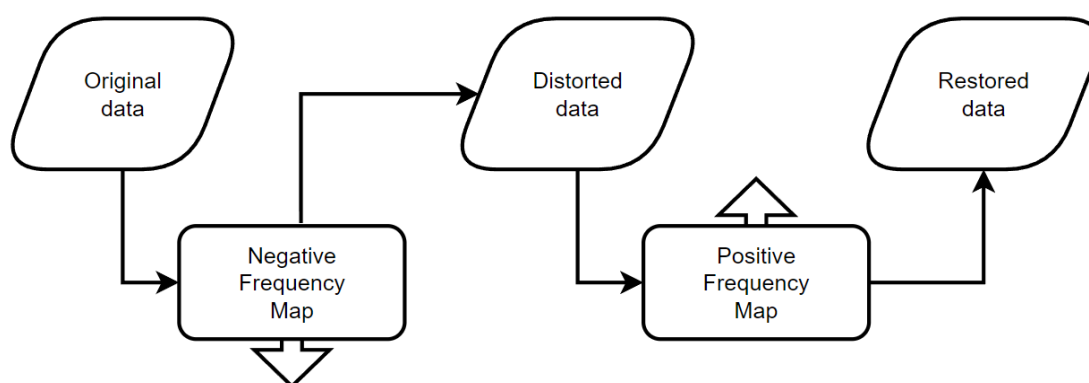


Figure 23: A flowchart detailing how the tests on the synthetic phantoms were done. The original data is first distorted using the shifting method with a negative frequency map. The distorted data is then corrected using the shifting method with a positive frequency map. This gives the restored data.

5.1.1 Synthetic Datasets: Cylinder and Sphere

Both the cylinder and the sphere datasets are simple in that they consist only of ones and zeros. The object in the middle is completely comprised of zeros to mimic the signal void from a metallic object in an MRI. In a box surrounding the objects the value is one, to mimic the agarose gel in the prosthesis phantom (see Figure 24). There is also a grid added to the middle slice, this is not shown in the figure.




Figure 24: How the synthetic dataset is designed. The cylinder is replaced by a sphere for parts of the testing. The blue area is comprised of zeros and the white is comprised of ones. Outside of the dotted box there is a smaller area of zeros to emulate air.

As they are synthetic datasets, they do not have an exact frequency map measured out like the phantom. Therefore, an analytically calculated frequency map was created using the equations for the magnetic field of a cylinder and a sphere described by Schenk et al [15].

The testing of the cylinder and sphere started with applying the Fourier shift method on them in reverse first. This is to attempt to mirror the effect the metallic objects would have in a real MRI scan. This state will be referred to as the distorted objects from here on. Following this, the same Fourier shift method was applied to the distorted objects with the values of the frequency map.

This whole process was done to examine the efficacy of the Fourier shift method on something that was not the prosthetic phantom. The hope was to establish the same problems that affected said prosthetic and thereby attempt to solve them in a more controlled environment. The process was done on two different cylinders and one spherical object. The cylinders had the same base radius, but the lengths were different. One of the cylinders extended past the scanning volume and one was contained within the volume. These tests were done to see if there would be any difference in the reconstruction of two shapes that were only slightly different. Something to note is that the analytically created frequency map is only valid for a cylinder of infinite length as the flat surfaces of a cylinder will cause the field to change. However, even if the frequency map is unrealistic, the shifting method was tried regardless.

5.2 Prosthetic Phantom

As mentioned earlier, the two main problems with the Prosthetic MR image were that after applying the Fourier shift, the image was left noticeably noisier and areas that had previously had no signal were replaced by zebra-stripe patterns. This section will describe the steps done to attempt to alleviate the issues.

The data used was in the form of 256x192x32 matrices. The raw data for the normal MRI images can be seen in Figure 12. This figure consists of all 32 images.

5.2.1 Filtering of the frequency map

Filtering was tried in several different ways. First, a simple median filter was applied to the frequency map of the prosthesis. A simple example of how a median filter works is shown in

Figure 25. This can be done in three dimensions with the same principle. For this part of the work the filter tested was a 3x3x3 median filter applied onto the frequency map.

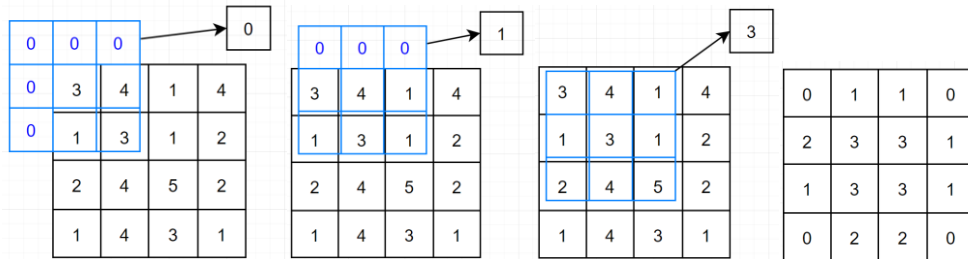


Figure 25: The 3 by 3 median filter places a box of size 3 by 3 numbers around every data point. Then takes the median of those 9 numbers. Far right image shows the completed result.

5.2.2 Interpolation

Since the size of the datasets are 256x192x32 data points, there is a risk that filtering in the third direction may affect the data more than filtering in the other two directions. Therefore, it would be ideal to increase the size of the third dimension. To do this we interpolated the dataset before applying the Fourier shift onto it. Three different factors of interpolation were tested 2x, 4x and 16x interpolation. This changes the matrix sizes to 256x192x64, 256x192x128, and 256x192x512 respectively.

For the dataset, the interpolation method used was the nearest neighbor interpolation, see Figure 26. The interpolation method for the frequency map was a cubic interpolation along the z-axis (slices).




Figure 26: Top: the interpolation method used for the dataset. Bottom: the cubic interpolation method used for the frequency map. Orange points denote the original value while blue points are interpolated.

Lastly, different strengths of median filtering (3x3x3, 3x3x9 and 3x3x17) were applied to the dataset and frequency map with 16x interpolation to combine the improvements from both methods. The entire step by step can be seen below in Figure 27. First the preprocessed data is

interpolated using the aforementioned “previous” method. Then, the frequency map is interpolated using the cubic interpolation. Next, the interpolated frequency map is filtered with a median filter. The interpolated data is then corrected using the Fourier shift method with the filtered and interpolated frequency map as the blueprint. The slices are then recombined so the matrix size returns to 256x192x32. This was done by just taking the mean of the surrounding slices. This is then the Final data.

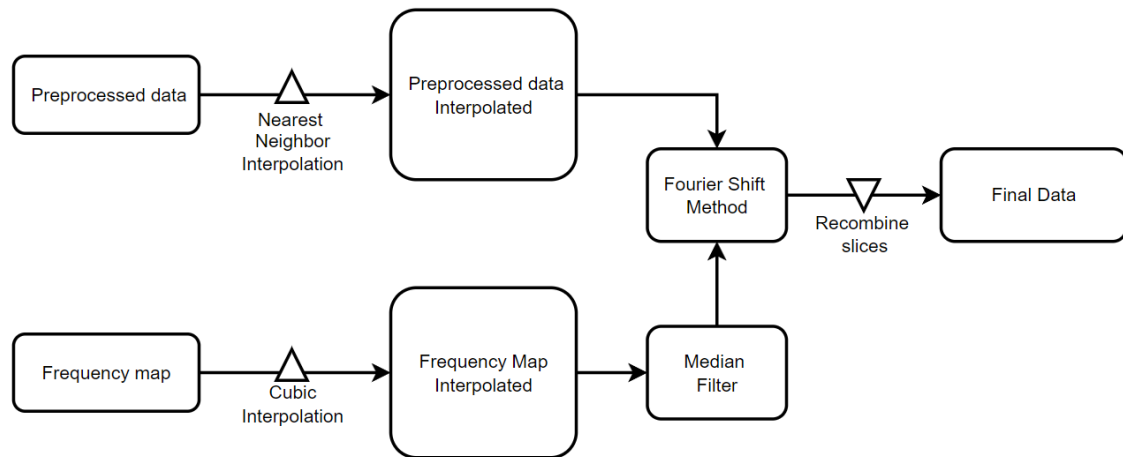


Figure 27: The series of steps taken to go from pre-processed data and the frequency map to the final processed data.

5.3 Image evaluation

To see whether these methods had an effect was rather easy when it came to the zebra stripe artefacts. This could be seen just by looking at it. However, determining which amount of filtering gave the best result proved to be rather difficult mainly due to the difficulty in finding a quantitative measurement of the signal-noise ratio (SNR).

An additional problem is that the type of noise resulting from this method is not known. The type of noise was thought to be either Gaussian or Rician noise.

Despite this, something to quantify the SNR in a figure is necessary. To make sure the zebra stripe artefacts do not affect the result a region of interest (RoI) was defined. The region that was decided on is seen below in Figure 28. For noise analysis purposes, slice 16 (where the grid is placed) should be ignored from these measurements as the grid could not be accounted for otherwise.




Figure 28: Where the RoI for the noise measurement is selected.

Three different methods were then used to estimate the noise level in the region of interest. The first method was created to find Rician noise, which is the noise type usually found in magnitude reconstructed MR images. This method and the accompanying code were developed by Jeny Rajan and taken from the MATLAB file exchange [16]. The next method uses an idea proposed by J. Immerkær [17], using the Laplacian operator it makes a quick estimation of the gaussian noise in an image. The code was taken from the MATLAB file exchange and made by Tolga Birdal [18] The last method is just a simple standard deviation divided by the mean within the region of interest.

6. Results

6.1 Simulations on synthetic phantoms

6.1.1 Cylinder 1

To demonstrate what the first cylinder looks like, Figure 29 (left) shows the 12 slides of the cylinder, (right) shows the corresponding generated frequency map according to the method proposed by Schenck [15].




Figure 29: (left) The 12 slices of the first cylinder used for testing. (right) The 12 slices of the corresponding frequency map.

After distorting the original image using the generated magnetic field, the resulting distorted image can be seen in Figure 30. As expected, the grid from slice 6 has been shifted around to all the surrounding slices. In Figure 29, the light green denotes 0, the dark red shows positive values, and blue shows negative values.

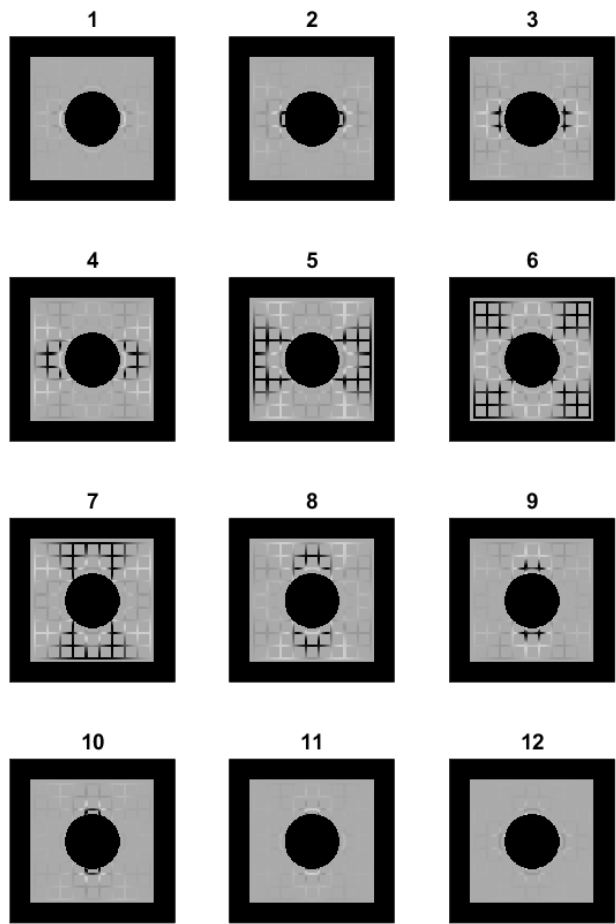


Figure 30: The distorted image using the frequency map (but inverted) from Figure 29.

Applying the frequency map (but with negative values to simulate geometric distortion from the MRI sequence) upon the distorted results shown in Figure 30 gives the images shown in Figure 31. The reconstruction appears to be very successful when comparing Figure 31 to the original image in Figure 29. A clearer comparison can be seen in Figure 32 where slice 6 of the reconstructed data is compared to the initial data.




Figure 31: The reconstructed cylinder.




Figure 32: Comparison of initial (left), distorted (middle), and reconstructed (right) slice 6.

6.1.2 Cylinder 2

The other cylinder that was tested is essentially the same as the previous one, however, the cylinder is only in slices 2 to 11. This can be seen in Figure 33 (left) and the accompanying changes to the frequency map can be seen in Figure 33 (right).

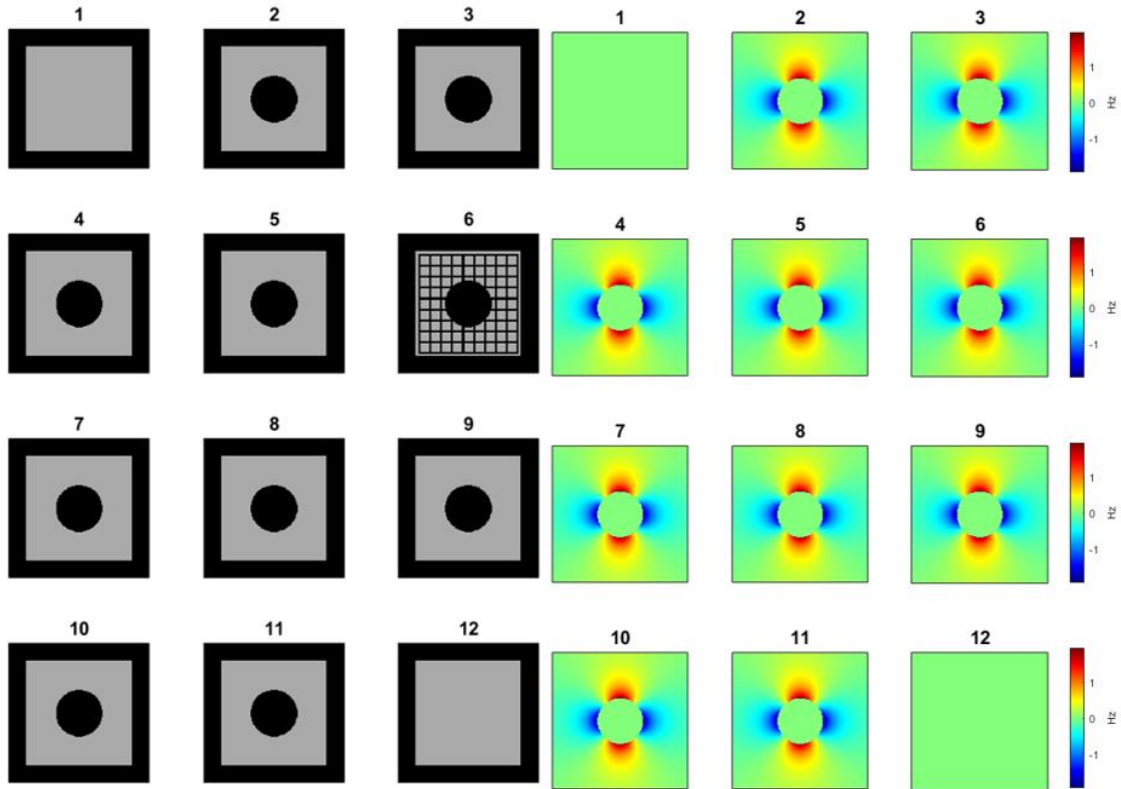


Figure 33: (left) The initial cylinder and (right) the accompanying frequency map.

The results for the distorted cylinder and reconstructed cylinder are shown in Figure 34.




Figure 34: (left) The distorted image of cylinder 2. (right) The reconstructed image of cylinder 2.

Here the reconstruction is noticeably worse than in the earlier example. However, as demonstrated by Figure 35, slice 6 seems still to be reconstructed very well. A thing to note is that the distorted images seem to contain a grid that is lighter than the background. This is a phenomenon that would not be possible in a real-life scenario.




Figure 35: A comparison between the initial, distorted, and reconstructed Slice 6 of the second cylinder.

6.1.3 Sphere

The initial sphere and its accompanying frequency map can be seen in Figure 36.

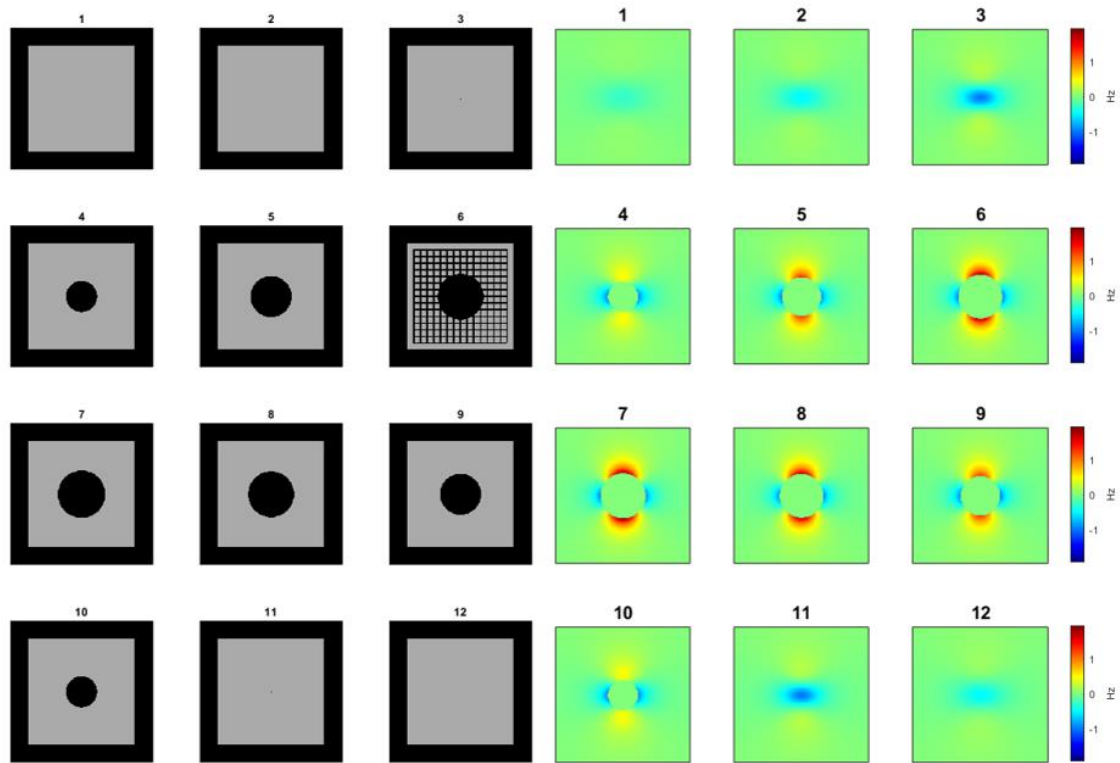


Figure 36: The sphere and the accompanying frequency map

After applying the frequency map onto the sphere, the results can be seen on the left in Figure 37. Then, the reconstructed image is shown on the right in Figure 37. Like the cylinder before, the slices show a brighter grid than the background.




Figure 37: (left) The distorted image of the sphere, (right) the reconstructed sphere.

Figure 38 Shows a comparison between the initial, distorted, and reconstructed middle slice of the sphere.




Figure 38: Comparison of the initial, distorted, and reconstructed slice 6 of the cylinder.

6.2 Phantom of the hip prosthesis

As stated, the project was started with some steps already taken. This was detailed in earlier sections; however, in Figure 39, an example is given to summarize the steps taken by Morin et al [5]. The full dataset can be seen in Figure 12 in Section 4.1 Phantom.

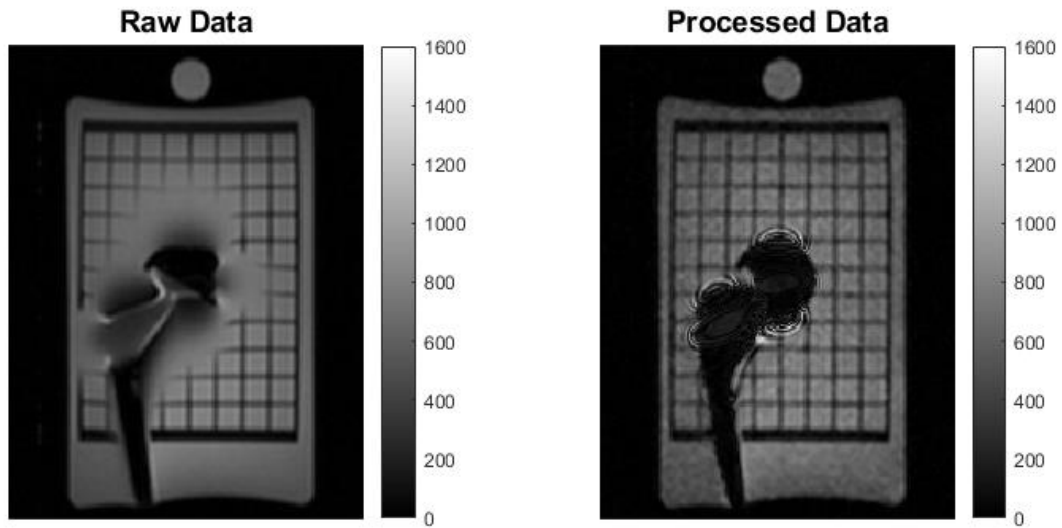


Figure 39: A comparison between slice 16 of the original dataset (left) to the processed data (right).

6.2.1 Filtering of the frequency map

The three-dimensional median filter of the frequency map provided a noticeable decrease in the noise level compared to the initial processed dataset. This can be seen through a direct comparison in Figure 40.




Figure 40: A comparison between slice 14 of the processed data with and without the median-filter.

6.2.2 Interpolation

Figure 41 and Figure 42 show the differences between the increasing interpolation factors in selected slices 18 and 21 respectively. The difference between them is not very noticeable, however, when observing the zebra stripes in Figure 41 and Figure 42, there appears to be a small improvement with each increase in interpolation factor. For more slides see Section Interpolation in the Appendix.




Figure 41: A comparison between the different interpolation factors. All three images depict slice 18 in their respective dataset.




Figure 42: A comparison between the different interpolation factors. All three images depict slice 21 in their respective dataset.

6.2.3 Combined interpolation and filtering

As the results from the dataset with interpolation factor 16 seemed to be the best, this section will only depict filtering done on said dataset. As before, more slides can be seen in the appendix in the Section Combined interpolation and Filtering.

Figure 43-Figure 45 show direct comparisons of slices 15, 16, and 18 for each of the filters as well as without a filter.




Figure 43: A comparison of Slice 14 with the different filter sizes applied to the Frequency map. The dataset used is interpolated with a factor 16.

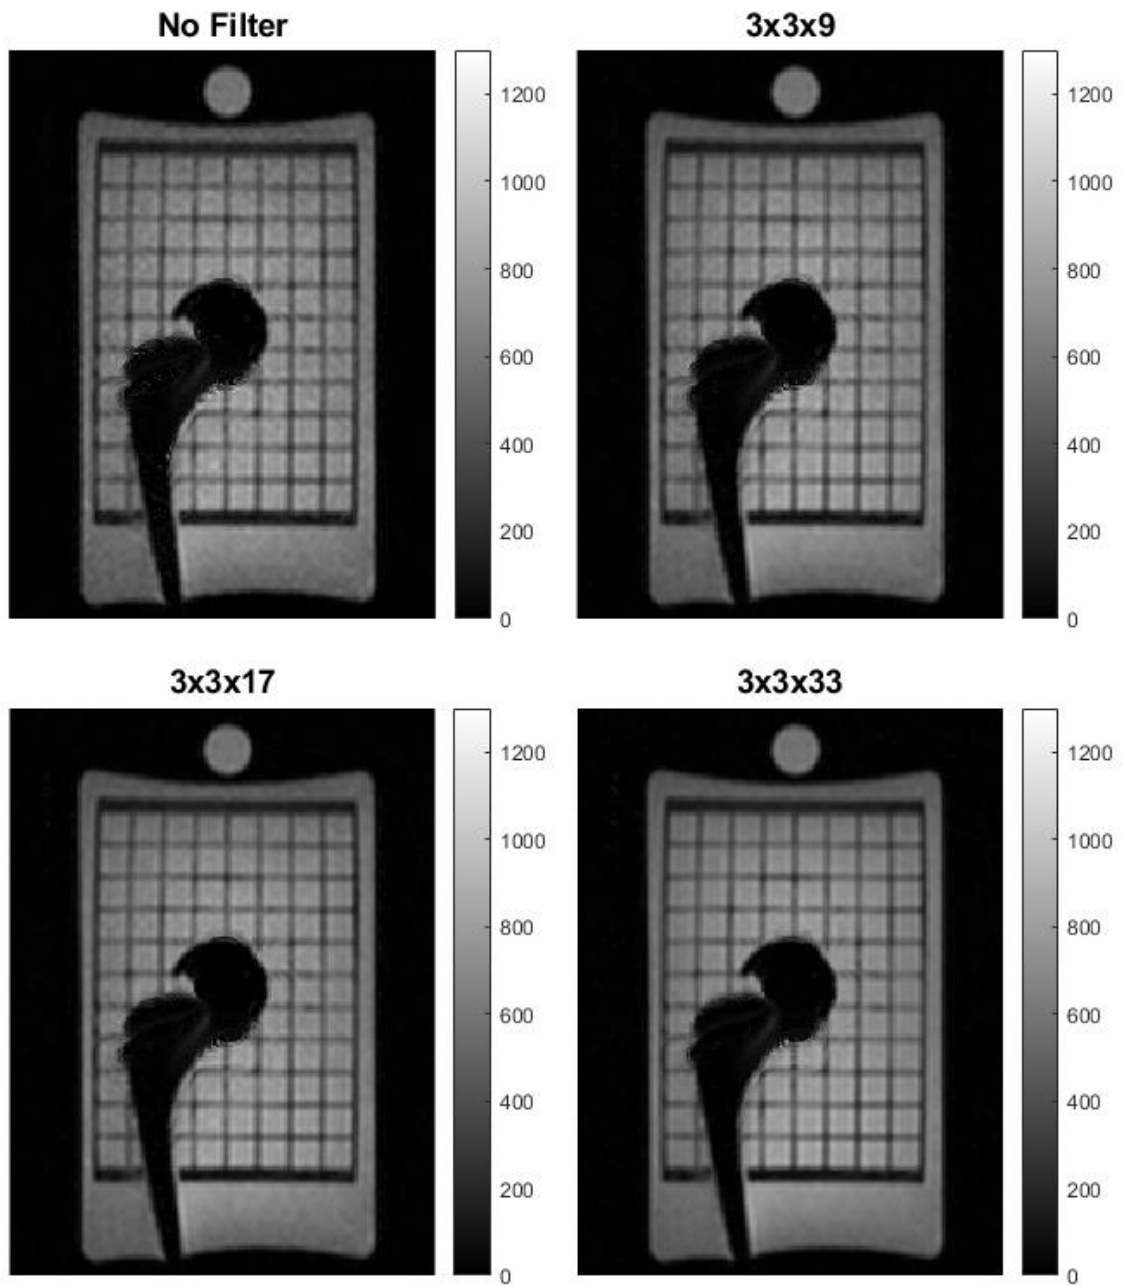


Figure 44: Comparison of the different filters for slice 16. The dataset used is interpolated with a factor 16.




Figure 45: Comparison of the different filters for slice 18. The dataset used is interpolated with a factor 16.

From these Figures (Figure 43, Figure 44 and Figure 45) it appears that the noise level decreases with each increased filter size. However, the noise reduction seems to be quite low as it is not easily noticeable.

6.2.4 Assessment of noise levels

Using the noise level sequence for Rician noise described in the method, the noise was estimated. The results are shown in Figure 46. The results for the other noise level sequence can be seen in Figure 47.




Figure 46: Estimation of Rician noise using the method by Jeny Rajan et al [16].

The Rician noise estimation indicates that more filtering of the frequency map leads to a reduced noise level. It also appears that this method is not able to distinguish noise from the grid in slices 15, 16 and 17 (16 in particular). When removing the value for the grid slice (slice 16) the average noise estimation for each filter size is shown below in Table 1.

Table 1: The average noise estimation using the Rician noise method.

No Filter	24.3226
3x3x9	17.3548
3x3x17	16.2581
3x3x33	12.5484




Figure 47: Noise level estimation of the different filters using the second method by J. Immerkær [17].

This method shows that the amount of noise is less in the unfiltered version than in the 3x3x9 and 3x3x17 datasets.

As discussed in the method, the last case for noise level estimation will be to show the average standard deviation divided by the mean of each pixel in the slice. This can be seen in Figure 48. Although everything is very close in value, it appears that the unfiltered version has the highest noise level by this estimation followed by the 3x3x9, then 3x3x17 and finally 3x3x33.




Figure 48: The average standard deviation of the pixels in the ROI of each slice divided by the mean of the pixels of the ROI of each slice.

6.3 Final Comparisons

This section will show some comparisons between the results from this project and where the project started (Morin's work). Thus, the figures below demonstrate the improvements made when comparing the final results from this project to where Morin's work left off. A full series of 16 images can be seen in Figure 60 in the appendix.

To demonstrate the improvements in both noise and artefact reduction Figure 49-Figure 52 show select comparisons between select slices 12, 16, 18, and 20 respectively. From these images, while there is still more work to be done to improve the method, there is both a noticeable reduction of the noise in the image as well as a significant reduction to the zebra-stripe artifacts.




Figure 49: Comparison of slice 12 in prior results (left) and final results (right)




Figure 50: Comparison of slice 16 in prior results (left) and final results (right)



Figure 51: Comparison of slice 18 in prior results (left) and final results (right)




Figure 52: Comparison of slice 20 in prior results (left) and final results (right)

7. Discussion

7.1 Synthetic Data: Cylinder and Sphere

In general, the results for both the cylinder and the sphere demonstrate a few things. First and foremost, the method to shift data in this way is quite far from perfect. In fact, it appears that unless the magnetic field used in the shifting method is completely uniform along the z-axis the shift cannot perfectly reconstruct the data. This is demonstrated by comparing the first cylinder with the second cylinder where even just two blank slices in the frequency map make a huge difference. There may be many causes for this, but one possibility is the inability to perfectly perform an inverse Fourier transform on the dataset. Even in the case of these synthetic datasets, which are real before sending them through the Fourier shift, the resulting data always seems to become complex values.

Another thing that seems to be the case is that the middle slices always appear to be reconstructed significantly better than the outermost slices. Again, there could be multiple causes of this. It may be the case that the slices closer to the middle simply end up less distorted than those on the edges. This seems unlikely, as the frequency field map shows that the strength of the field is like the other slices. Another possible cause could be if some sort of wrap-around is happening. That is to say that data shifted from the last slices is shifted into the first slices and vice versa.

The intent behind using these generated datasets was to attempt to mirror the problems that appeared in the reconstruction of the prosthesis phantom prior detailed in the Section Prior Work. Unfortunately, this part was only partly successful. While there are some structures in the results that appear to exhibit something like the zebra stripes, these datasets do not have the same issue with a noticeable increase in noise.

However, since the noise in the prosthetic phantom is undoubtedly induced during the Fourier shift, it seems likely that the noise is coming from elsewhere. The only other places it could originate from are the actual data set, or its frequency map. When the preprocessed data is compared to the processed data, the preprocessed data has much less visible noise. Thus, it seems unlikely that the noise is coming from the dataset. Therefore, the source of the noise increase is likely due to the frequency map for the prosthesis. Looking back at the initial measured frequency map in the Section Prior Work many of the slices appear to have a noisy signal, for example in slice 18 and 19. This could potentially be a cause to the noise in the restored images.

7.2 Prosthesis

Overall, the results for both the filtering and the interpolation are promising. It seems that both additions to the method cause a big improvement.

7.2.1 Initial Filtering

The filtering of the frequency map decreased the noise in the restored image. This seems to back up the earlier theory that a significant issue behind this method right now is that the frequency map used was not a perfect measurement. However, when using median filters like this, it may be the case that the image is a bit over-filtered. It is also quite difficult to know if this is the case, as the image itself can appear to be much better. This could be a problem and is something that should be investigated using tailormade phantoms to test for specifically this.

7.2.2 Interpolation

The interpolation of the datasets had quite a remarkable effect on the artefacts of the image, and one that was quite unexpected. The original theory behind the interpolation was that it would both allow for more through plane filtering as well as reduce the difference in shifts of magnetic field strength. It turned out that even an interpolation of only factor 2 (essentially placing an extra slice in between each slice) had a huge impact on the zebra-stripes and further interpolation increased this effect.

The result itself wasn't completely unexpected; however, the degree of success that it had was. One reason for why the interpolation worked so well is tied to why the zebra stripe artefacts appear. It's possible that the zebra stripe artefacts are a result of what is known as Gibbs artifact. Without going much into detail, the artifact is caused by using too few harmonics when reconstructing the image which causes a ringing effect around areas with especially high contrast differences [19]. It is possible that these artefacts show when using the Fourier shift method in the z-direction as well, and that by combining a series of interpolated slices this effect is averaged out.

An interesting thing to note is that different methods of interpolation were used on the frequency map and the image dataset, nearest neighbor interpolation on the data and cubic interpolation on the frequency map. The results of the interpolation seemed to indicate that a greater interpolation factor reduced the zebra-stripes more and more. This could not be tested systematically beyond a factor of 16, as the computer power (mostly RAM) and time required simply exceeded the reasonable limits when increasing the interpolation factor beyond this point. A 32x interpolation factor was the largest that was tested, and it took around 2 hours to complete, while a run-through with 16x interpolation only took around 15 minutes. Furthermore, it seemed to be unnecessary for the scope of this report to go higher than 16; however, it could be something that could be investigated in the future.

The prosthesis used for this paper was one made of stainless steel, and therefore had a very large warping effect on the surrounding volume in the phantom. This is likely the reason why the prosthesis itself cannot be correctly imaged. While the reconstruction of the phantom is quite good, it can also be seen that the resulting image is not how a hip prosthesis typically looks. If one of the datasets of the reconstructed images is examined (see ex. Figure 59-Figure

45). In Figure 53, an x-ray image of an Exeter hip implant is shown to demonstrate the actual shape of the hip implant in the image.




Figure 53: taken from article by Månsson et al. [14] with permission from the publisher. Shows an x-ray comparison of the Exeter hip implant to determine the actual outline of the prosthesis.

As can be seen when comparing Figure 53 to any other image in the results, a true depiction of this phantom may not be possible. The magnitude of the warping seems to be too large to account for in the method used for this paper. It might therefore be beneficial to attempt this method on some hip implants that have a smaller effect on the magnetic field as well.

It is important to note however, that the image in Figure 53 is taken with a higher resolution, so a direct comparison with the method presented in this paper may not be entirely appropriate.

7.2.3 Filtering combined with Interpolation and Analysis of Noise Levels

As mentioned earlier in the report, there is difficulty in determining a quantitative measure for the SNR. This was a large issue in attempting to compare the noise in the images. Furthermore, the type of noise in these images is not known, which exacerbates the issue. This is demonstrated quite clearly by the results shown in Figure 46 and Figure 47 in the results where the noise level is estimated.

Firstly, I would say there is a clear improvement the bigger the filter gets (which is also what I expected). However, judging by the figures this seems to not be the case. The first method is specifically designed to detect common MR noise, known as Rician noise. While the results shown seem to agree with my visual estimation, there are several slices where there are unexpected overlaps. The second method is more designed to detect gaussian noise. However, the results here completely disagree with the visual assessment I made earlier in that the dataset with no filter seems to be the second best. Finally, the third method, like the first method, seems to show results that align with a visual assessment of the results. Again however, the differences in noise levels seem to be rather small. To be clear it is entirely possible that the visual assessment of the results is inaccurate and that the images seen as less noisy contain noise that is difficult to visually assess. Nevertheless, a robust method to accurately determine the noise level is required.

7.2.4 Study limitations

The main problem experienced during the project is the lack of additional measured datasets, only so many conclusions can be drawn from using a single dataset. The main issue with taking additional datasets is that the measurement of magnetic fields is not something that is typically used and therefore the MR devices must be specifically adjusted to be able to do it. This problem was exacerbated due to Covid-19 which made it even more difficult to attain additional datasets as the project was done at a distance for the entire period.

7.2.5 Future improvements

There are several possible next steps mentioned earlier in this report. However, the next major step should be testing the method on different datasets. These datasets should consist of both simple phantoms of hip prosthetics like the one used in this report, but also phantoms with different structures to test specific problems this method has. An example of this could be a phantom with small quick changing structures, like small metal rods to see if they can be resolved. As mentioned earlier, this method should also be tested on phantoms with lower magnetic susceptibility, like those made of titanium.

The idea behind using the synthetic datasets is sound and should that section of the report be developed further to be more in line with reality many additional tests could be done. One such test is the addition of noise into the process to test how the method reacts to noise in a controlled environment. Furthermore, different magnitudes of magnetic field inhomogeneities could be tested to see where it performs the best.

Another important step to take would be the ability to make direct comparisons to the results of a scan using SEMAC. This would allow the noise level, scan time and the quality of the images to be directly evaluated using the SEMAC scan as a baseline. Figure 53 is a comparison, but as noted it's not entirely fair. To really examine how well the method performs it would be good to simultaneously do measurements on the same phantom using the other methods with the same or similar setups to compare.

A theory that was developed along the course of this project, and with support from the testing on the synthetic datasets, is that this method of reconstruction cannot perfectly reconstruct an image in cases where the magnetic field changes (which would be all cases where metal is involved). Should it be the case that perfect reconstruction is impossible, it doesn't mean that the method is without its merits. There are many cases where it might not be necessary to perfectly reconstruct the whole image, and if the method is good enough, there could be potential use cases.

8. Conclusion

This report has examined and attempted to improve upon the results of the method proposed by Månsson and Morin to use a measurement of the magnetic field to shift distorted slice profiles back to its origin. The project was a success in the respect that it managed to both almost completely remove the zebra-stripe artefacts that were prominent in the previous implementation, and to improve the SNR of the restored images by roughly a factor of two. However, the restored images still have a poorer SNR than the uncorrected images so there is still room for future improvement.

References

- [1] U. Berg, G. Chatziagorou, K. Gromov, N. Hailer, A. Joboryy, P. Jolbäck, J. Kärrholm, G. Limbäck-Svensson, E. Lind, H. Malchau, M. Mohaddes, E. Naucclér, J. Nåtman, (...) and P. Werner, "Svenska Höftprotesregistret: Årsrapport 2019," Svenska Höftprotesregistret, 2019.
- [2] "MRI (Magnetic Resonance Imaging)," Food & Drug Administration, 12 09 2017. [Online]. Available: <https://www.fda.gov/radiation-emitting-products/medical-imaging/mri-magnetic-resonance-imaging>. [Accessed 15 08 2021].
- [3] B. A. Hargreaves, P. W. Worters, K. B. Pauly, J. M. Pauly, K. M. Koch and G. E. Gold, "Metal-Induced Artifacts in MRI," *American Journal of Roentgenology*, vol. 197(3), pp. 547-555, 2011. DOI: 10.2214/AJR.11.7364.
- [4] W. Lu, K. B. Pauly, G. E. Gold, J. M. Pauly and B. A. Hargreaves, "SEMAC: Slice Encoding for Metal Artifact Correction in MRI," *Magnetic Resonance in Medicine*, vol. 62(1), pp. 66-76, 2009. DOI: 10.1002/mrm.21967.
- [5] V. P. Morin, G. M. Müller and S. Månsson, *Restoration of large slice profile distortions near metallic implants by frequency mapping*, Poster presented at: ISMRM, 2012.
- [6] L. Winter, L. Zilberti, M. Murbach and B. Ittermann, "MRI-Related Heating of Implants and Devices: A Review," *Journal of Magnetic Resonance Imaging*, vol. 53, no. 6, pp. 1646-1665, 2021. DOI: 10.1002/jmri.27194.
- [7] S. Currie, N. Hoggard, I. J. Craven, M. Hadjivassiliou and I. D. Wilkinson, "Understanding MRI: basic MR physics for physicians," *Postgraduate Medical Journal*, vol. 89, pp. 209-223, 2013. DOI: 10.1136/postgradmedj-2012-131342.
- [8] L. G. Hanson, "Is quantum mechanics necessary for understanding magnetic resonance?," *Concepts in Magnetic Resonance*, vol. 32A, no. 5, pp. 329-340, 2008. DOI:10.1002/cmr.a.20123
- [9] J. P. Hornak, "The Basics of MRI," Interactive Learning Software, Henrietta, NY, 2012.
- [10] B. A. Jung and M. Weigel, "Spin echo magnetic resonance imaging," *Journal of Magnetic Resonance Imaging*, vol. 37, pp. 805-817, 2013. DOI: 10.1002/jmri.24068.
- [11] K. Butts, J. M. Pauly and G. E. Gold, "Reduction of blurring in view angle tilting MRI," *Magnetic Resonance in Medicine*, vol. 53(2), no. 418-24, pp. 418-424, 2005. DOI:10.1002/mrm.20375.

- [12] R. Otazo, M. Nittka, M. Bruno, E. Raithel, C. Geppert, S. Gyftopoulos, M. Recht and L. Rybak, "Sparse-SEMAC: Rapid and Improved SEMAC Metal Implant," *Magnetic Resonance in Medicine*, vol. 78(1), pp. 79-87, 2017. DOI: 10.1002/mrm.26342.
- [13] F. Wellman, "Master of Science dissertation: Evaluation of methods for MR imaging near metallic hip prostheses," Lund University, 2011.
- [14] S. Månsson, G. M. Müller, F. Wellman, M. Nittka and B. Lundin, "Phantom based qualitative and quantitative evaluation of artifacts in MR images of metallic hip prostheses," *Physica Medica*, vol. 31, 2015. DOI:10.1016/j.ejmp.2014.12.001.
- [15] J. F. Schenck, "The role of magnetic susceptibility in magnetic resonance imaging: MRI magnetic compatibility of the first and second kinds," *Medical Physics*, vol. 23(6), pp. 815-850, 1996. DOI: 10.1118/1.597854.
- [16] J. Rajan, "MATLAB File exchange," 22 Apr 2016. [Online]. Available: <https://www.mathworks.com/matlabcentral/fileexchange/28779-noise-estimation-from-rician-noise-corrupted-images>.
- [17] J. Immerkær, "Fast Noise Variance Estimation," *Computer Vision and Image Understanding*, vol. Vol. 64, no. No. 2, 1996. DOI: 10.1006/cviu.1996.0060
- [18] T. Birdal, "Fast Noise Estimation in Images," 31 May 2012. [Online]. Available: <https://www.mathworks.com/matlabcentral/fileexchange/36941-fast-noise-estimation-in-images>. [Accessed 2021].
- [19] "Questions and Answers in MRI: Gibbs (truncation) Artefact," Elster LLC, [Online]. Available: <https://mriquestions.com/gibbs-artifact.html>. [Accessed 24 01 2022].

9. Appendix

Interpolation

2x Interpolation




Figure 54: The reconstructed image of the prosthesis using a factor 2 interpolation.

4x Interpolation

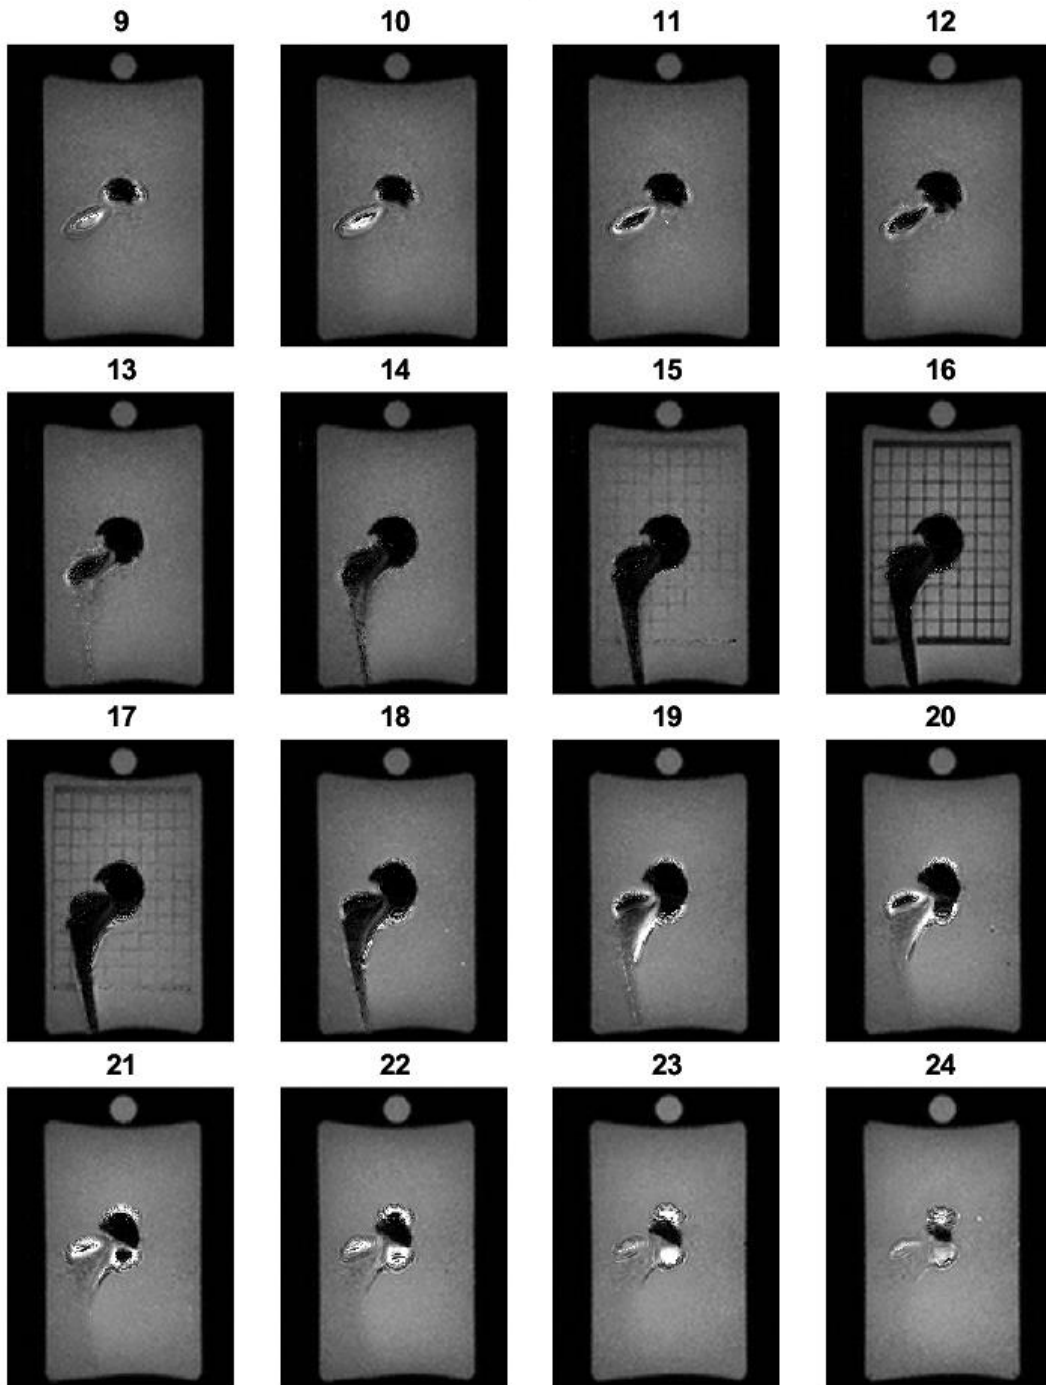


Figure 55: The reconstructed image of the prosthesis using a factor 4 interpolation

16x Interpolation




Figure 56: The reconstructed image of the prosthesis using a factor 16 interpolation

Combined interpolation and Filtering

16x Interpolation-3x3x9 Filter




Figure 57: The results from using the Fourier shift with a 16x interpolation and a median filter of size 3x3x9.

16x Interpolation-3x3x17 Filter




Figure 58: The results after using a median filter of size 3x3x17

16x Interpolation-3x3x33 Filter




Figure 59: Results for using a 3x3x33 filter.

Final data




Figure 60: the Final processed image, following both a 16x interpolation as well as a 3x3x33 filtering.