



BSc thesis APPLIED MATHEMATICS and APPLIED PHYSICS

Model based image reconstruction for low-field hand-held MRI

On imaging using field geometry and sample translations

Pieter Oppelaar

Delft, June 25, 2021

Supervisor Applied Mathematics:

Dr.ir. M.B. van Gijzen

Ir. M.L. de Leeuw den Bouter

Supervisor Applied Physics:

Dr.ir. D. Lathouwers

Acknowledgements

I would like to offer my special thanks to Dr.ir. M.B. van Gijzen and Ir. M.L. de Leeuw den Bouter for their continuous supervision, help and advice the last months. I would like to thank Dr.ir. D. Lathouwers for his role as supervisor from the physics department.

I would also like to extend my thanks to S. Tewari PhD for his insightful suggestions and of course for designing and building the scanners used in this project, as well as carrying out the measurements.

Finally, I wish to thank Dr.ir. R. van der Toorn and Dr. F.M. Vos for agreeing to be a part of the assessment committee for this bachelor thesis.

Abstract

In this report, the conversion from spin-echo signals, obtained with a low-field hand-held MRI scanner that was designed and built at the Leiden University, to images of the proton density within the sample is considered. This scanner does not make use of switchable gradient coils, but instead relies solely on the natural inhomogeneity of the field and on translations of the sample over this field for its spatial encoding. Specifically, an attempt is made to answer the question of how we can reconstruct a phantom using this kind of scanner.

This is done by deriving a signal model, discretising it and writing it as a linear least squares problem. Then, we can make use of the techniques of Conjugate Gradient for Least Squares (CGLS) with ℓ_2 -regularization and Generalized Conjugate Gradient Minimal Error (GCGME) with ℓ_1 -regularization for the difference between neighbouring pixels in order to solve this inverse problem.

Firstly, we theoretically consider combinations for magnetic field geometry and measurement strategy for their usability for image reconstruction. After this, the obtained strategies are tested in three experiments, with two magnets and two samples.

We start by doing this numerically, using a simulated phantom in combination with a measured magnetic field and the translation strategy. By doing this, we can determine if reconstruction is possible using that combination of field and strategy. Finally, the strategy is tested on real samples.

Using numerical phantoms in combination with the magnetic field and translation strategy used in the measurements, we were able to correctly reconstruct the phantoms. However, the reconstruction broke down when data from real samples was considered. A variety of possible improvements is discussed. The improvement that would have the most impact would be to design a magnet that has a less uniform gradient in the Z -direction, and instead has some locations in the xy -plane where the field falls slowly as function of Z , but quickly in other locations.

Contents

1. Introduction	1
2. Theory	2
2.1 Larmor frequency	2
2.2 The RF pulse	2
2.3 Spin-echo measurement	3
2.4 Signal model	3
3. Device	5
3.1 Device design	5
3.2 Data acquisition	5
4. Magnetic field design	7
4.1 Constant field	7
4.2 Constant linear gradient	8
4.3 Circular field	8
5. Numerical model	13
5.1 Discretisation of the integral	13
5.1.1 Incorporating translations	14
5.1.2 Bandwidth	14
5.1.3 Time-domain filtering	15
5.2 Signal model summary	15
5.3 Solving the inverse problem	16
5.4 Imaging artifacts due to data collection	16
5.4.1 Aliasing	16
5.4.2 Smearing	17
6. Results	19
6.1 Magnetic field maps	19
6.2 Phantoms	21
6.3 Experiments	22
6.4 Singular values	23
6.5 Bandwidth correction	25
6.6 Image reconstructions	27
6.6.1 Numerical phantoms	27
6.6.2 Real phantoms	29
7. Discussion	31
7.1 Improvements in signal model	31

7.2	Improvements in measurement	32
7.3	Improvements in scanner design	33
8.	Conclusion	34
	References	35

1. Introduction

After its proposal in 1973 by P.C. Lauterbur [1] (for which he received the Nobel prize in physiology or medicine in 2003), Magnetic Resonance Imaging (MRI) has grown to be one of the most widely adopted techniques used in diagnostic medical imaging. With MRI, a high resolution image can be obtained of various body parts, from the complex structures of the brain, to the softer tissues around knee or elbow joints. Now, MRI machines can be found in many hospitals, with increasing size and cost and using magnets with ever increasing magnetic field strengths.

However, MRI technology does not have to be limited to only big and complex machines with increasingly stronger and more exotic magnets. While these developments allow for high resolution imaging of the whole body, it makes MRI scanning an expensive and time consuming process. In this report, the prospect of a small, hand-held MRI-scanner is explored. While this device is no replacement for the conventional scanner design, it could prove a useful imaging apparatus for its small size and relatively low cost. The proposed use case would be scanning tissues that lie just below the skin (between 1 and 3 cm) for tumors or, if a suitable resolution can be achieved, even for scanning the nerve tissues that lie within the spine.

In this report, the problem of image reconstruction with such a hand-held scanner is considered, along with some recommendations on the design of the magnet and the data acquisition process that could make the reconstruction more accurate. The specific research question that we want to answer in this report is: "How can we reconstruct an image from the MRI-signal of a hand-held scanner?". To this end, a suitable signal model will be presented, along with some design suggestions for the magnets and measurement strategies. Then, some algorithms that can be used to invert the signal model are presented, followed by the results for both simulated phantoms and actual measurements.

This project is a collaboration between the Leiden University Medical Center (LUMC), where the prototypes are designed and built and the measurements are done, and the Delft Institute of Applied Mathematics (DIAM), where the modelling side of the project is researched. This thesis was carried out at DIAM, and therefore focuses more on the modelling side of the project.

The structure of the report is as follows. An overview of the relevant theory on MRI is given in chapter 2. The third chapter contains relevant information on the magnet design as well as the data acquisition process. A discussion on field design is presented in chapter 4. The numerical signal model is derived in the fifth chapter. Some results are shown in chapter 6. Finally, a discussion of these results and a conclusion can be found in chapters 7 and 8 respectively.

2. Theory

This section is meant to provide some insight in the relevant theory for this experiment and is by no means extensive. A more complete theory on the principles of signals for magnetic resonance imaging (MRI signals) can be found in [2].

2.1. Larmor frequency

Some elementary particles possess a property called spin. This property is a measure of a particle's intrinsic angular momentum. Composite particles, like nuclei, can also have spin, depending on their composition. When a particle has nonzero spin and is charged, it also has a magnetic dipole moment, so that it behaves like a magnetic dipole. This magnetic dipole can be influenced by applying an external magnetic field. The spin vector will then align with the magnetic field, and will precess around it. The frequency at which this precession occurs, is governed by the magnetic field strength and a particle dependent constant, called the gyromagnetic ratio. This precession frequency, also known as Larmor frequency, is given by:

$$\omega(r) = \gamma B_0(r); \quad (1)$$

where ω_0 is in rad/s, γ is the gyromagnetic ratio in rad/sT and B_0 is the magnetic field strength of the external magnetic field in T. This precession is the driving mechanism behind nuclear magnetic resonance (NMR) and magnetic resonance imaging¹

An important case of a charged particle with spin, that precesses around external magnetic fields, is the hydrogen nucleus (specifically $^1\text{H}^+$). Since hydrogen is one of the most abundant atoms, it is the ideal candidate for medical imaging, as it can be found in a multitude of tissues within the human body.

2.2. The RF pulse

When only a constant external magnetic field $B_0 = B_0 \hat{x}$ (we suppose it is a field in the x -direction) acts on a large number of hydrogen nuclei, the individual spins will either align themselves parallel or antiparallel to the field. They will however be out of phase with each other. The total magnetisation, which is the sum of the individual magnetic dipole moments of the nuclei, will therefore be aligned with the external field and be completely in the x -direction. They y and z components will cancel out due to the spins being out of phase.

The value of this macroscopic equilibrium magnetisation is given by:

$$M_0(r) = \frac{2h^2}{4k_B T} \gamma(r) B_0(r) = \frac{h^2}{4k_B T} \gamma(r) \omega(r); \quad (2)$$

with $M_0(r)$ in J/T, γ the gyromagnetic ratio in rad/sT, Planck's constant in Js/rad, Boltzmann's constant in J/K and T the temperature in K and where we have used the Larmor equation (1) and that $B_0(r)$ is oriented in the x -direction to rewrite the magnetic field to a frequency

¹These are two names for the same thing, but the term MRI is used in medical contexts.

$\rho_0(r)$.

However, when an additional magnetic field is introduced, the individual spins can be manipulated so that they end up in phase with one another, thereby creating a measurable transverse magnetisation component in the plane. This is done by generating a magnetic field that oscillates with the specific frequency. When this frequency is close to the Larmor frequency, the spins will align themselves with this second magnetic field. We call this field the B_1 -field or radio frequency (RF) pulse, since it oscillates in the radio frequency range and has a short duration, and it is given by $B_1(t) = B_{1,y} \cos(\omega_1 t) \hat{y} + B_{1,z} \sin(\omega_1 t) \hat{z}$, so that it is perpendicular to the B_0 -field.

Under the influence of the B_1 field, the individual spins end up in phase, so that after some time the total magnetisation vector has a component in the plane that is in phase with the B_1 field. When the B_1 field is switched off, the spins will precess at the Larmor frequency given by the local magnetic field. This results in a rapid decay in the component of the magnetisation vector as the spins will again become out of phase with one another. Since this changing magnetisation vector generates a changing magnetic field, it generates a voltage in a nearby coil due to Maxwell's laws of Electromagnetics and can thus be measured. From this signal, along with knowledge of the B_0 - and B_1 -fields, it is possible to reconstruct an image showing the proton density $\rho(r)$ as a function of position.

2.3. Spin-echo measurement

One way to measure the number of protons that experience a range of magnetic field strengths, is by doing a spin-echo measurement. The essential idea of this technique is to first align all magnetisations, then let them decay for a short time and finally to flip the magnetisation by 180° . This flipping makes it so that the magnetisation vectors that were pulling ahead, get flipped and are now lagging behind, and vice versa. After a short time (the same time as between the initial pulse and the flipping) the magnetisations will again be in phase and an echo can be measured. Since the receiver coil will have a finite bandwidth, only information on proton density at places experiencing a magnetic field within this bandwidth will be measured.

2.4. Signal model

This section is adapted from [3]. The signal that is measured from the spin echo routine can be mathematically described as follows. If we set the time at which the echo occurs as 0, then we know that at that time, all the magnetisations of the spins are in phase (so their magnetisations perpendicular to the z -axis are all in phase). This transverse magnetisation is proportional to the equilibrium magnetisation (equation (2)) and is therefore proportional to $\rho_0(r)$. Now, combining this with equations (2) and (1) and putting all constants into a coil response, we can write the signal model after demodulation (now the precession frequency is given by $\omega_0(r)$ via equation (1)) as:

$$S(t) = \int_{r \in D} \rho(r) \exp(-t/T_2(r)) \exp(i \omega_0(r)t) dr; \quad (3)$$

where D is the scanned spatial domain, $c(r)$ is the coil response in which we have included all constant factors, $\omega(r)$ is the frequency corresponding to the Larmor frequency in rad/s , $T_2^*(r)$ is the transverse relaxation constant in s , $\rho(r)$ is the proton spin density in m^3 , i is the imaginary unit, $\Delta\omega(r)$ is the difference between the frequency corresponding to the local magnetic field via equation (1) and the demodulation frequency in rad/s , and t is the time in s .

Since the measurement only takes a short time, the transverse relaxation factor can be omitted. The relatively small bandwidth of the measurement as compared to the total magnetic field, makes it so that $T_2^*(r)$ can be approximated as $1/\omega_{dem}$ (the demodulation frequency) and can be incorporated into the coil response. Finally, we rewrite using the Larmor equation (1) to obtain the signal model:

$$S(t) = \int_{r \in D} c(r) \rho(r) \exp(i \Delta\omega(r)t) dr; \quad (4)$$

where $\Delta\omega(r)$ is the difference between the local magnetic field strength and the magnetic field corresponding to the demodulation frequency. In chapter 4, this equation will be numerically evaluated.

3. Device

In this section, some of the physical aspects of the hand-held scanner, designed and built at the LUMC, that is used in this report are discussed. This is by no means an extensive overview, as the main focus of this project lies in the signal analysis. It is however still important to have a general idea of the design, since it influences the final signal model.

3.1. Device design

The hand-held MRI scanner consists of two main components, namely the main magnet (around $10 \times 10 \text{cm}^2$) and a radio frequency coil (RF coil). The main magnet is made up of multiple small permanent magnets, that are fixed at certain points within a larger, 3D printed enclosure. Their positions are determined using a genetic algorithm as in [4] that has as input the desired magnetic field. The fields of these smaller magnets add up to form the field, or main, static magnetic field. This field will have a large gradient in the direction perpendicular to the surface of the magnet, since it is one sided and the field falls going away from the magnet surface. The magnetic field close to the magnet can be manipulated by changing the positions of the smaller magnets.

The second component of the hand-held scanner is the RF coil. This coil is responsible for the excitation of the spins, as well as for receiving the spin echo signal. A variety of coils can be used in combination with the magnet. The choice of coil is influenced by the parameters of the coil, such as its bandwidth. The dimensions of the coil are also considered, since some coils, like a solenoid coil, will limit the size of possible phantoms. A surface coil imposes no limitations on sample size at the expense of some extra noise and a weaker signal, since it is further away from the sample. In this report, a solenoid coil is used, with the sample enclosed in it.

A third main component of most conventional MRI scanners are the gradient coils. These can change the static magnetic field to (as the name implies) add a gradient to the static magnetic field. The device used in this experiment does not have such coils. The plan is to make use of the gradient already present in the magnetic field. Since the magnetic field cannot be altered using gradient coils, a unique reconstruction must be made possible in some other way. We plan on doing this by translating the sample over the magnet, so that we in some sense vary the magnetic field.

A schematic representation of the scanner can be found in figure 1.

3.2. Data acquisition

The signal is measured at equally spaced discrete time points for some duration. We introduce t and N as measurement parameters as the time between the time points and the number of time points respectively. The data is acquired by measuring the voltage over the RF coil during the decay of the magnetisation vector. Before discretising, the signal is demodulated with the frequency f_0 , corresponding to the frequency of the RF coil. The signal is also filtered by a

Gaussian filter in the time domain, with mean = 0 and standard deviation of about 0.5ms so that noise is attenuated where no echo signal is expected (or in other words, the echo signal is amplified). Furthermore, the measurement is repeated and the average is taken in order to further improve the signal to noise ratio.

Figure 1: Schematic of the prototype hand-held scanner as viewed from the side, with the main magnet indicated in blue, the sample container in red and the RF coil in gold. All the parts are also indicated with arrows and a top-down view is added in the top left corner of the figure. It can be seen that the RF coil wraps around the sample in this prototype.

4. Magnetic field design

The goal of an MRI scanner is to make an image of the density of protons for all positions in the region of interest. For this, equation (1) can be used, which says that the frequency of the precession is a function of the external magnetic field strength at a given point. In conventional MRI, gradient coils that can be switched on and off are used to assign a unique combination of frequency and phase to each location within the field of view. Since the hand-held scanner does not make use of gradient coils, this unique encoding must be achieved in another way. This is done by making use of the inhomogeneity in the field generated by the magnet. If we were to take only one measurement, the encoding that is obtained would still not be unique. This is because the magnetic field strength will be constant on surfaces. A unique encoding can be achieved by translating the sample over the magnet. If the geometry of the planes allows this and a sufficient amount and combination of translations is done, each place in the region of interest is uniquely determined by an ordered set of magnetic field strengths.

For this to work, it is important that the planes are laid out in such a way, that a translation in fact yields new information, i.e. that points that were on the same surface before the translation, lie on different surfaces after the translation.

In order to determine if a magnetic field is usable for this kind of encoding (with translations), we must attempt to find a set of translations such that for each point in the region of interest, the resulting set of magnetic field strengths coupled to that point, is unique. In other words, there exists a finite ordered set of translations $T(r)$ such that the function $B(T(r))$ that couples an ordered set of translations to an ordered set of magnetic field strengths, is injective on the domain D , where D is the domain in space of interest. When such T exists, the field can be used to make an image using translations.

The question now arises if we can find a field, such that the function B contains as few translations as possible. In that way, the translations we do, will yield the largest amount of new information. Since D is three dimensional and the magnetic field strength is a scalar, a lower bound on the number of translations $|T|$ is 3.

It is quite difficult to directly look for a field for which a T with three translations actually exists. It is however possible to consider some simple field maps and check if they exist and how many elements it has. We do this graphically and first in two dimensions. The generalization to three dimensions is quite straightforward and the two dimensional figures are more accessible.

4.1. Constant field

We first consider a two dimensional field with no gradient at all, so that the magnetic field strength is constant over the whole domain. This field clearly is not suitable for imaging using translations, since every translation will not move any points to different magnetic field strengths (the field is constant) and no additional information is gained.

4.2. Constant linear gradient

The next field we consider is a field with a constant gradient in some direction. Another way we can represent a field, is by drawing the curves where the magnetic field strength is equal (this is done for a linear constant gradient field in figure 2 and a circular field in figure 4). We find that in the case of the constant gradient, these curves are all lines of the form

$$r_0 + s; \quad (5)$$

where r_0 is some point on the line, s is a real constant and s is a vector in the direction of the line. Furthermore, the vector s is the same for all lines. This field is also not suitable for imaging using only translations. This can be seen by simply translating a copy of figure 2 over itself (see figure 3) or mathematically. We therefore consider two arbitrary points r_1 and r_2 that lie on the same line, and an arbitrary translation t . Then, we can set, without loss of generality

$$r_0 = r_1; \quad (6)$$

and, since r_1 and r_2 lie on a line with the same field strength, we have that

$$r_2 = r_0 + s; \quad (7)$$

for some s and the same s . We now translate by t . Now, the point r_1 is translated to

$$r_1^0 = r_0 - t; \quad (8)$$

and

$$r_2^0 = r_0 - t + s; \quad (9)$$

where the minus sign is because the field is translated in the opposite direction compared to the sample. And we find that the points r_1^0 and r_2^0 again lie on a line with equal field strength. Therefore, there are no translations that yield new information and this field is also unsuitable for imaging using translations².

4.3. Circular field

The last field we consider is the field where the curves of equal field strength are concentric circles (figure 4). This field is suitable for image reconstruction using translations. This can be seen by translating a copy of figure 4 over itself (figure 5) and observing that the each red circle only crosses a black circle once, twice or not at all, but they never have completely overlapping segments (unless one considers a translation with 0, but that is trivial). With this field, it is possible to uniquely reconstruct an image after just three measurements (one measurement with the non-translated field, one with the field translated along the x -axis and one with the field translated along the y -axis).

²This field would work well if the imaging were to be done using rotations rather than translations, this can be seen by rotating a copy of the lines of equal field strength (figure 2) over itself and observing that a pair of lines intersects at most once.

Figure 2: A magnetic field with a linear constant gradient. The curves in the figure represent sets in \mathbb{R}^2 where the magnetic field strength is equal. Note that the different lines represent different field strengths.

This is one more translation than the theoretical minimum amount of translations (that would be 2, due to the two dimensional nature of the problem considered here). This is because the field has symmetry along all lines through the midpoint of the circles, and this symmetry results in an extra mirrored image, if the translations are only along one line of symmetry. Theoretically, there are fields that only require two measurements to yield a unique reconstruction. For example, a field with curves that are semicircles (if not translated along the remaining symmetry line) or quarter circles. However, since the circular field is more viable to create in practice, a magnet with a field similar to this will be used in some of the experiments. The three dimensional version of this field is the field where the surfaces (they are now two dimensional) with equal magnetic field strength have paraboloidal shapes that are rotated around the z -axis.

With this three dimensional field it is possible to achieve a unique reconstruction with a small amount of translations in the xy -plane. In practice, the image is translated more often than this, to combat noise artefacts and to minimize effects of local field inhomogeneities.

To see the effect that translating only along a symmetry line has, this was numerically simulated with the circular field and a numerical phantom drawn in figure 13. The result of this can be seen in figure 6. When, additionally, translations perpendicular to that symmetry line are done, the phantom is correctly reconstructed, see figure 7.

Figure 3: A visual representation of a translation using the lines of equal field strength. In this figure, the black lines represent the untranslated magnetic field, while the red lines represent the translated field. These lines completely overlap, therefore, this field is unsuitable for imaging using translations.

Figure 4: A magnetic field with a magnetic field strength that is radially symmetric. The curves in the figure represent sets $\{R\}$ where the magnetic field strength is equal. Note that the different lines represent different field strengths.

Figure 5: A visual representation of a translation using the curves of equal field strength. In this figure, the black curves represent the untranslated magnetic field, while the red lines represent the translated field. For any pair of a red and black curve, there are only one, two, or no intersections, making this field suitable for imaging using translations.

Figure 6: Reconstruction using GCGME (see chapter 5) when only translating along a symmetry line in the field. The phantom is reconstructed, but a mirror image is also visible.

Figure 7: Reconstruction using GCGME (see chapter 5) when translating along a symmetry line in the field and perpendicular to it. The phantom is now correctly reconstructed.

5. Numerical model

5.1. Discretisation of the integral

In this report, an attempt is made to solve the signal model equation (4) for the variable by discretising the integral. To this end, we divide the space into voxels that are small enough so that the magnetic field strength within a voxel can be assumed constant (for this, we take the magnetic field strength in the center of the voxel). Furthermore, we also assume that the proton density within the voxels is constant.

The signal model then becomes:

$$S(t) = \sum_{j=1}^M (r_j) \alpha(r_j) \exp(-i B_{r_j} t) \quad (10)$$

where M is the total number of voxels and x , y and z are the dimensions of the voxels.

Since we take discrete measurements (N samples) at specified times, we can rewrite equation 10 as:

$$S_i = \sum_{j=1}^M (r_j) a_{i,j}; \quad (11)$$

with

$$a_{i,j} = \alpha(r_j) \exp(-i B_0(r_j) t_i) \quad (12)$$

where i goes from 1 to N and j from 1 to M . Since the goal of this project is not to obtain quantitative data on the proton density, but rather to see the relative densities and try to find the shape of the phantom, the factors x , y and z are omitted from now on.

Equation (11) defines a linear system that can be written in as a matrix equation:

$$Ax = b; \quad (13)$$

where A is the matrix with elements $a_{i,j}$ as specified by equation (12), x is the vector with elements (r_j) and b is the vector with elements S_i .

In this experiment, the used phantoms were all constant in z -direction. This makes it possible to reduce the number of elements in the matrix significantly. That the phantom does not change in the z -direction, implies that all proton densities in the vector at places with the same location in the xy -plane are equal. This means that we can reduce the second dimension of A by a factor M_z , where M_z is the number of voxels in the z -direction, by simply adding the columns corresponding to the same location in the xy -plane.³

³We will keep using i as a matrix index and assign m as the index for the slices in the z -direction.

5.1.1. Incorporating translations

If one would try to solve equation (13) for x , the obtained solution would not be unique, since the scanned domain consists of surfaces where the magnetic field strength is equal. To make unique reconstruction possible, the sample is translated over the magnet multiple times. Rather than solving the matrix problem separately for all different translations, it is possible to combine the data from all the translations into a single problem.

Firstly, translating the sample only slightly changes the matrix. If we do K measurements, the matrix $A^{(k)}$ for the k th measurement will have entries given by:

$$a_{ij}^{(k)} = c(r_j) \exp(-i B(r_j - r^{(k)})t_i); \quad (14)$$

where $r^{(k)}$ is the vector by which the sample is translated, $k = 1, \dots, K$ and the minus sign is there because a translation of the sample corresponds to an inverse translation of the magnetic field.⁴

Similarly, we denote the received signal for the k th measurement as $b^{(k)}$. Then, we can write the combined matrix equation as

$$A_{\text{comb}} x = b_{\text{comb}}; \quad (15)$$

with

$$A_{\text{comb}} = \begin{pmatrix} A^{(1)} \\ A^{(2)} \\ \vdots \\ A^{(K)} \end{pmatrix} \quad (16)$$

and

$$b_{\text{comb}} = \begin{pmatrix} b^{(1)} \\ b^{(2)} \\ \vdots \\ b^{(K)} \end{pmatrix}; \quad (17)$$

The vector x remains the same. This matrix problem can then be solved using the techniques described in the next sections.

5.1.2. Bandwidth

The excitation pulses generated by the RF coil do not have an infinite bandwidth. The coil can also only receive signals from a limited bandwidth. In fact, since the magnetic field changes significantly in the domain covered by the voxels, the signal coming from some voxels will not be received by the coil. It is essential to correct for this, otherwise the model expects a signal

⁴This equation is an adaptation from equation (12). Therefore, there is still an index for all the voxels, including those in the z -direction.

from voxels that lie outside this bandwidth for a given translation. The bandwidth of the coil was determined from the frequency spectra of the measurements. After this, the matrix was tweaked to correct for this. This is done by incorporating the bandwidth in the coil sensitivity. We further assume that the coil sensitivity does not explicitly depend on the coil sensitivity is given by:

$$c(B(r)) = f_{\text{Gauss}}(B(r); 0; B_{\text{coil}}); \quad (18)$$

where $f_{\text{Gauss}}(B; 0; B_{\text{coil}})$ is the probability density function for the normal distribution with variable B , mean 0 and standard deviation B_{coil} , which is a measure of the bandwidth of the RF coil in T (can be converted to frequency by using the Larmor equation (1)).

5.1.3. Time-domain filtering

In an effort to improve the signal to noise ratio, the signal is attenuated in the time domain where no echo is expected by a Gaussian filter with a mean of 20 ms and a standard deviation of about 0.5 ms. This can be incorporated in the model quite easily, since it corresponds to a left multiplication with a diagonal matrix of the vectors \mathbf{D} and therefore also with a left multiplication with the same diagonal matrix of matrices \mathbf{A} . We will call the diagonal matrix corresponding to this filter \mathbf{D} .

5.2. Signal model summary

Since the signal model has now been built up over several chapters and sections, it is shortly summarized here for the sake of clarity.

The final signal model is given by:

$$\mathbf{A}_n \mathbf{x} = \mathbf{b}_{\text{comb}}; \quad (19)$$

with \mathbf{b}_{comb} unchanged from equation (17) and

$$\mathbf{A}_n = \begin{bmatrix} \mathbf{D} \mathbf{A}_{\text{cor}}^{(1)} \\ \mathbf{D} \mathbf{A}_{\text{cor}}^{(2)} \\ \vdots \\ \mathbf{D} \mathbf{A}_{\text{cor}}^{(K-1)} \\ \mathbf{D} \mathbf{A}_{\text{cor}}^{(K)} \end{bmatrix}; \quad (20)$$

where \mathbf{D} is the time domain filter matrix, K the total number of measurements (translations) and the elements $a_{i,j}^{(k)}$ are given by:

$$a_{i,j}^{(k;\text{cor})} = \sum_{m=1}^{M_z} f(B(r_j^{(m)})) \exp(i B(r_j^{(m)} - r^{(k)}) t_i); \quad (21)$$

where m is the index for the slice in the z -direction, going from 1 to M_z , j from 1 to $M_x M_y$ is the index for the xy -position and i is the index for the time t_i , going from 1 to N (not to

be confused with the i in the exponential, which is the imaginary unit). The function $B(r_j^{(m)})$ is the difference between the local magnetic field strength and the magnetic field strength corresponding to the demodulation frequency ω and the gyromagnetic ratio. From now on, we will refer to the matrix and vectors in equation (19) as A and b to make the future notation more compact.

5.3. Solving the inverse problem

We will attempt to solve the inverse problem, so determining x from A and b with the already known routines of Conjugate Gradient for Least Squares (CGLS) [5], [6] with possible λ -regularization and Generalized Conjugate Gradient Minimal Error (GCGME) [3] with regularization. To this end, we must rewrite the problem as a least squares problem:

$$\min_x \frac{1}{2} \|Ax - b\|_2^2 + R(x); \quad (22)$$

where λ is the regularization parameter and $R(x)$ is some kind of regularization function.

We regularize this problem in two different ways. Using CGLS with regularization using the identity matrix, so that $R(x) = \lambda \|x\|_2^2$, and using GCGME with λ -regularization, with this second algorithm, we actually carry out anisotropic total variation regularization, that is, regularization for the difference between neighbouring pixels. We do this through choosing the regularisation matrix. The details of how this is done can be found in [3]. These algorithms were provided as MATLAB code by M.L. de Leeuw den Bouter.

Due to the nature of the samples, the regularization done with the second algorithm is more suitable for the reconstruction, since neighbouring pixels tend to have similar proton density and the used form of GCGME penalizes jumps between pixels. However, the algorithm fails without regularization if the problem is ill-posed and with too much regularization, it will return a reconstruction that looks like a phantom, but is actually just a result of the regularisation. For this reason, results obtained with GCGME should be approached carefully when CGLS yields a poor result for the same problem.

5.4. Imaging artifacts due to data collection

Since the spin echo data is acquired as a finite amount of discrete samples, some artifacts may be introduced in the final reconstruction, depending on the parameters used in the data collection. Two of these effects are described in this section.

5.4.1. Aliasing

The image reconstruction does not suffer from aliasing if:

$$2f_{\max} < f_s; \quad (23)$$

where f_{\max} is the highest frequency in the signal and f_s is the sampling frequency. This can be rewritten to a relation between the absolute difference in magnetic field over the field of view B_{\max} and the time between samples Δt . We find that the image does not suffer from the effects of aliasing if:

$$\Delta t < \frac{1}{B_{\max}}; \quad (24)$$

where γ is the gyromagnetic ratio of a proton (in rad/sT). If the time between sampling is chosen larger than this, reconstructions could contain artifacts such as in figure 8. We do note that the bandwidth correction could remove some of these artifacts, when the bandwidth is small enough.

5.4.2. Smearing

The image reconstruction process resembles the computation of a discrete Fourier transform [2], where a sequence of samples in the time domain is transformed to a sequence of coefficients for N different frequencies. These frequencies are equally distributed over the interval $[0, \frac{2}{\Delta t}]$ and are therefore separated by $\frac{2}{N\Delta t}$. This imposes a limit after which adding more pixels to the reconstructed image has no effect, since the reconstruction cannot differentiate between frequencies that differ by less than this separation. Therefore, the maximum resolution that can be obtained from the data is determined by the number of samples that are collected.

This is as follows: in order to distinguish two points with a given field, there must be some orientation of the field so that these points lie at places where the difference in the magnetic field strength is large enough, so that the corresponding frequency difference is greater than $\frac{2}{N\Delta t}$. The signal corresponding to an arbitrary point in the field is of the form $\exp(i\gamma B t)$, which we can rewrite to $\exp(i\gamma B_n \Delta t)$, where n is the sample number. Signals from other points are similar, and only differ due to a different value for B . For any two points, we can denote this difference in B by ΔB , so that the difference in frequency between the signals of the two points becomes $\Delta B \gamma$ rad/s. The two points are then distinguishable if this difference is greater than $\frac{2}{N\Delta t}$. We therefore find points differing by a certain value ΔB can be distinguished if:

$$N > \frac{2}{\Delta B \Delta t}; \quad (25)$$

From this we find that increasing the resolution of the image is capped at the point where the difference between the magnetic field for all pairs of neighbouring pixels is lower than ΔB for the chosen N . If the number of samples is too low for a given field, artifacts such as in figure 9 are present in the reconstructed image. We will call this effect smearing.

⁵Another way to see this, is that the surfaces of equal field strength (as described in chapter 4) are in theory thin in theory, obtain a thickness that is given by ΔB . This makes it harder to obtain a good reconstruction.

Figure 8: Figure showing the effects of aliasing. These effects are present when the time between samples is too long. This result was obtained using a simulated phantom and eld, making the units irrelevant.

Figure 9: Figure showing the effects of smearing. These effects are present when the number of samples is too low for the translations that are carried out, in combination with the magnetic eld. This result was obtained using a simulated phantom and eld, making the units irrelevant.

6. Results

6.1. Magnetic field maps

Two different magnets were used in this experiment. Both were designed and made by S. Tewari at the Leiden University. All measurements were also obtained by him. The design of the first magnet was more focused on having a homogeneous field in the xy -plane, while having a stronger gradient in the z -direction. The second magnet was designed to have a paraboloidal field, like a simple bar magnet, so that it is closer to the three dimensional version of the circular field discussed in chapter 4. To give an idea of the field of both of the magnets, images of the middle planes of the measured magnetic field strengths are plotted in figure 10 for the first, more homogeneous magnet and in figure 11 for the second magnet, that has a paraboloidal field. From the plotted figures we can see that this is indeed the case, but that also for the second magnet, the gradient in the z -direction is still quite strong and roughly the same everywhere in the xy -plane.

(a) xy -plane at $z = 20.11\text{mm}$

(b) xz -plane at $y = 0\text{mm}$

(c) yz -plane at $x = 0\text{mm}$

Figure 10: Map showing the magnetic field of the magnet with the more homogeneous field, with in (a) the map for the xy -plane at $z = 20.11\text{mm}$, in (b) the map for the xz -plane at $y = 0\text{mm}$ and in (c) the map for the yz -plane at $x = 0\text{mm}$.

(a) xy-plane at $z = 28.22\text{mm}$

(b) xz-plane at $y = 0\text{mm}$

(c) yz-plane at $x = 0\text{mm}$

Figure 11: Map showing the magnetic field of the magnet with the paraboloidal field, with in (a) the map for the xy-plane at $z = 28.22\text{mm}$, in (b) the map for the xz-plane at $y = 0\text{mm}$ and in (c) the map for the yz-plane at $x = 0\text{mm}$.

6.2. Phantoms

For the measurements, two phantoms were created by S. Tewari at the Leiden University. The first phantom consisted of three tubes that are arranged like an equilateral triangle. These tubes were filled with oil so that they could be imaged using MRI techniques. A schematic of this phantom can be found in figure 12(a). A numerical version of this phantom was also created by M.L. de Leeuw den Bouter at DIAM. With this numerical phantom the proposed reconstructions can be tested. In this numerical version, the proton density is different in every tube. This numerical phantom can be found in figure 13. We will call this phantom the 'three-tube'-phantom or 'three-tube'-sample.

The second phantom that was created by S. Tewari consisted of a single channel in the shape of the letter 'M'. This channel was also filled with oil so that it could be measured. A schematic of this phantom can be found in figure 12(b). We will call this phantom the M-phantom or M-sample.

(a) Three tubes sample

(b) M sample

Figure 12: Schematic of the two phantoms that were imaged in the three experiments. The white area's in the schematic correspond to channels in the 3D printed phantoms that can hold oil. These white spaces are therefore the part that is imaged by the scanner. The phantom consisting of three tubes can be found in (a), and the M-phantom can be found in (b).

Figure 13: The numerical phantom used to test whether a combination of translations and magnetic field can be used for reconstruction. This phantom was created by M.L. de Leeuw den Bouter.

6.3. Experiments

With these two magnets and two samples, three sets of measurements were acquired by S. Tewari. The first set was obtained with the first magnet in combination with the 'three-tube'-phantom. In this experiment, the sample was translated over the sample, so that a measurement was obtained from 28 translations. These measurements consist of 1024 time samples with $t = 5 \mu\text{s}$ (these parameters make sure that aliasing and smearing is prevented). The demodulation frequency in this experiment was 1.079 MHz and the bandwidth of the received signal was around 20 kHz.

The second set of measurements was obtained using the paraboloidal magnet, combined with the M-sample. For this set 39 translations were done. The individual measurements again consist of 1024 time samples with $t = 5 \mu\text{s}$. The demodulation frequency in this experiment was 2.138 MHz and the bandwidth of the received signal was again around 20 kHz.

The last set of measurements was again obtained using the paraboloidal magnet, combined with the M-sample. Only now 7 translations were measured instead of 39 and all measurements consist of 512 time points instead of 1024. The t did remain $5 \mu\text{s}$. The demodulation frequency in this experiment was 2.124 MHz and the bandwidth of the received signal was around 70 kHz.

6.4. Singular values

In order to determine if a problem is ill-posed or not, the singular values of the matrices A were determined for each of the data sets. If there are singular values close to zero, the problem is ill-posed, and will require regularization in order to create an image. If a problem is very ill-posed, it could also be a sign that the translations are far from optimal in combination with the field or that the number of translations is too small to obtain a good reconstruction.

The singular value decompositions for the three experiments, along with a map of the translations that were done in order to obtain the set of measurements can be found in figure 14 for the first experiment, figure 15 for the second experiment and figure 16 for the third measurement.

We see that the first and third set of measurements have singular values that are close to zero, making the reconstruction more difficult (the lowest from the first experiment is in the order of 10^1 and the lowest singular value from the third experiment is in the order of 10^0). The lowest singular value for the second experiment lies just below 1, indicating that this experiment will require the least amount of regularisation in order to reconstruct correct images.

(a) Translations

(b) Singular value decomposition

Figure 14: Visualisation for all the translations done to obtain the first set of measurements, along with the singular value decomposition of the corresponding matrix. In the plot for the translations (a), the yellow squares indicate positions in the plane to which the center of the phantom was translated. The plot in (b) contains the singular value decomposition of the matrix A , with the singular values σ_j on a logarithmic scale and with an index on the x-axis.

(a) Translations

(b) Singular value decomposition

Figure 15: Visualisation for all the translations done to obtain the second set of measurements, along with the singular value decomposition of the corresponding matrix. In the plot for the translations (a), the yellow squares indicate positions in the plane to which the center of the phantom was translated. The plot in (b) contains the singular value decomposition of the matrix A , with the singular values σ on a logarithmic scale and with an index on the x-axis.

(a) Translations

(b) Singular value decomposition

Figure 16: Visualisation for all the translations done to obtain the third set of measurements, along with the singular value decomposition of the corresponding matrix. In the plot for the translations (a), the yellow squares indicate positions in the plane to which the center of the phantom was translated. The plot in (b) contains the singular value decomposition of the matrix A , with the singular values σ on a logarithmic scale and with an index on the x-axis.

6.5. Bandwidth correction

In order to correct for the limited bandwidth of the coil, the elements are multiplied by a weighting factor that depends on the frequency of the signals from the corresponding voxels and the bandwidth of the receiver coil. This factor is given by equation (18). The value of the factor is plotted for neighboring slices in figure 17 for the first set of measurements, figure 18 for the second set of measurements and figure 19 for the third set of measurements.

These figures are relevant, as we can visualize the evolution of the sensitive region between the slices in the z -direction using these. If the bands move significantly, then the resolution in the z -direction is too low and should be increased, since then some values miss a frequency that in fact also can originate from there. This increase can either be done by measuring or by interpolating the existing field data. If the bands are stationary, the resolution in the z -direction could be lowered for faster calculations or the resolution in the x - and y -directions could be increased for a higher resolution reconstruction. From all the figures, we see that the evolution over the slices is sufficiently slow (the figures are almost identical) in all three experiments, and could even be decreased a bit in order to aid in performance.

(a) xy -plane with $z = 20:11\text{mm}$

(b) xy -plane with $z = 20:13\text{mm}$

Figure 17: Maps showing the bandwidth correction that is applied to terms corresponding to the voxels in the slices corresponding to $z = 20:11\text{mm}$ (a) and $z = 20:13\text{mm}$ (b) for the first set of measurements. These values are given by equation (18).

(a) xy-plane with $z = 28.22\text{mm}$

(b) xy-plane with $z = 28.25\text{mm}$

Figure 18: Maps showing the bandwidth correction that is applied to term A_{sc} corresponding to the voxels in the slices corresponding to $z = 28.22\text{mm}$ (a) and $z = 28.25\text{mm}$ (b) for the second set of measurements. These values are given by equation (18).

(a) xy-plane with $z = 28.22\text{mm}$

(b) xy-plane with $z = 28.25\text{mm}$

Figure 19: Maps showing the bandwidth correction that is applied to term A_{sc} corresponding to the voxels in the slices corresponding to $z = 28.22\text{mm}$ (a) and $z = 28.25\text{mm}$ (b) for the third set of measurements. These values are given by equation (18).

6.6. Image reconstructions

6.6.1. Numerical phantoms

A good way to test if the combination of a field and translations could be used for imaging is by considering numerical phantoms. One then essentially tests whether the algorithms are capable of carrying out the inverse problem. This was done by multiplying matrix A with the numerical phantom (figure 13) to obtain the 'artificial' signal vector b . This signal is fed back into the algorithm, where it tries to reconstruct the phantom again. These reconstructions were done for all three experiments and can be found in figure 20 for the first experiment, figure 21 for the second experiment and figure 22 for the third experiment. All these reconstructions were done using CGLS and without regularization and with GCGME with regularization. From this, we see that the routine for the first set was less successful in reconstructing the phantoms, in order to get a good reconstruction, the use of GCGME with regularization was required. The routines of the second and third experiments did result in a good reconstruction, even using CGLS and without regularization. For all results, the regularization parameter is included in the caption of the figure. For the figures obtained with CGLS, a tolerance of 10^{-6} and a maximum of 500 iterations was used. For figures obtained with GCGME, the tolerances were kept at 10^{-6} while the maximum number of outer iterations was kept at 10 and the maximum number of iterations was kept at 10.

(a) CGLS

(b) GCGME

Figure 20: Results for the reconstruction of the numerical phantom using CGLS without any regularization (a) and GCGME with regularization parameter 10^{-6} (b), using the strategy from the first set of measurements.

(a) CGLS

(b) GCGME

Figure 21: Results for the reconstruction of the numerical phantom using CGLS without any regularization (a) and GCGME with regularization parameter 10^{-6} , using the strategy from the second set of measurements.

(a) CGLS

(b) GCGME

Figure 22: Results for the reconstruction of the numerical phantom using CGLS without any regularization (a) and GCGME with regularization parameter 10^{-6} (b), using the strategy from the third set of measurements.

6.6.2. Real phantoms

Finally, the combinations of magnetic field and translations were tested with the 'real' phantoms. The results of those experiments can be found in figure 23 for the first experiment, figure 24 for the second experiment and figure 25 for the last experiment. These reconstructions were unsuccessful, as no part of the phantom was reconstructed. When adding regularisation and using GCGME, the results are somewhat better, in that there is some signal from the image, although the location is not as expected, but there is a high possibility these spots are just artefacts due to the regularization and the field geometry.

(a) CGLS

(b) GCGME

Figure 23: Results for the reconstruction of the 'three-tube'-sample using CGLS without any regularization (a) and GCGME with regularization parameter 10^3 (b), using the data from the first set of measurements.

(a) CGLS

(b) GCGME

Figure 24: Results for the reconstruction of the M-sample using CGLS without any regularization (a) and GCGME with regularization parameter 10^4 (b), using the data from the second set of measurements.

(a) CGLS

(b) GCGME

Figure 25: Results for the reconstruction of the M-sample using CGLS without any regularization (a) and GCGME with regularization parameter ϵ (b), using the data from the third set of measurements.

7. Discussion

While the model can correctly reconstruct numerical phantoms, it fails when signals from real phantoms are considered. The inability to reconstruct the phantom could have several causes, of which this discussion will mention a few, along with suggestions for possible improvements in the signal model, measurement techniques and design of the scanner, which could be implemented in order to possibly obtain a system capable of reconstructing real phantoms.

7.1. Improvements in signal model

Since the signal model comprises the main part of this report, this will be the first place to look for possible improvements. In order to derive the used signal model, a number of approximations were done and assumptions were made. A logical place to start would be to look at whether or not these were justified.

One of the first assumptions is that the factor concerning T_2 relaxation can be omitted. This assumption is justified, as T_2 is much larger than the time between the first RF pulse and the echo. The factor (r) is assumed constant for the whole scanned domain, this assumption is also justified. As the bandwidth (converted to frequency using equation (1)) of the measurement is generally a factor 20 to 30 smaller than the total magnetic field strength.

Another, more significant, assumption is taken in the discretisation of the model. There, the first assumption is that the magnetic field can be approximated by a so called "center of voxel" approximation. This is only valid if the difference between neighbouring pixels is small enough. When the magnetic field for example changes significantly within a voxel, frequencies that are in fact originating from that voxel, cannot be traced back to that voxel by the model, since it assumes that that voxel only generates a signal of one distinct frequency. This can be mitigated by measuring the magnetic field on a finer grid, or, though this may result in significant errors in the magnetic field strength values, by interpolating field data measured on a coarse grid. In this report, the magnetic field was measured on a grid with a spacing of 1.5 mm between the grid points (in all dimensions), after which these results were interpolated to as much as $61 \times 61 \times 321(x \ y \ z)$. This could have resulted in a significant error between the points of the actual measurement, but that should not result in a significant error in the reconstruction, as frequencies will be moved at most 1.5 mm in every direction, and probably significantly less than that, assuming that there are no small inhomogeneities with a large amplitude within the field.

The final assumption that is made lies within the coil sensitivity. In all experiments, this sensitivity was implemented by means of a normal distribution, with the bandwidth estimated from the signal characteristics. This bandwidth could however be measured, after which one could fit some function to it and implement this function in place of (18). The second assumption in the coil sensitivity, is that it not explicitly depends on position, but only implicitly through the magnetic field strength. This assumption is justified, since the samples were enclosed in the coil and the coil was translated along with the sample.

A third assumption was made in the geometry of the phantoms, namely that they are constant in the z -direction. This assumption results in a much smaller matrix and vector, but could be problematic if this continuity is not satisfied. Since the dimensions of the phantoms are much bigger than the tolerances of a 3D printer, this assumption should be valid, but one could always omit this simplification. Then, the obvious drawback is the need for more computational power.

A different simplification in the model, that could cause the reconstruction to fail, is the assumption that the RF pulses have the same effect on at all points in the bandwidth. This might not be the case, since the magnetic field has a significant component in the z -direction at the edges of the field of view. Ideally the B_1 -field (which is also not perfectly oriented in the z -direction away from the center, due to the geometry of the coil) should be oriented perpendicular to the B_0 -field. In the center, these fields are perpendicular, but at the edges, this is not the case. Since the phantoms were kept relatively close to the center, the errors caused by this misalignment should be minimal.

7.2. Improvements in measurement

Another area where errors could arise is in the measurement of the signals. Excessive noise there could make it so that the actual data is drowned out. Some improvements that could result in less noise or in more data, are discussed in this next section.

While the measurement is repeated hundreds of times already in order to improve the signal to noise ratio, the noise could be attenuated further by repeating the measurement even more. This would however increase the time it takes to do the measurements, most likely so that fewer translations can be measured. The goal should therefore be to find a balance between these two. All experiments suggest it could be beneficial to measure fewer translations, but repeat the measurement more often, so that the final vectors are smaller in size, but also more accurate. This is because the reconstructions seemed to suffer from the noise in the signal, and not from having carried out an insufficient amount of translations, which would result in only loss of the shape of the signal, but not in no signal at all.

A second improvement that could be made in order to improve the reconstruction, would be the use of an RF coil with a larger bandwidth, so that more of the sample can be measured in a single measurement. A drawback of coils with a larger bandwidth, is that they also pick up more noise, so that it is necessary to repeat the measurement even more often. Since the reconstruction of the numerical phantom was successful for the bandwidth of the coils that were used in this experiment, more of an improvement could be made in other areas.

An obvious way to eliminate noise is to attenuate it with the use of a filter. It appeared that for the used measurements, this was not done as much as possible, as there were frequencies in the signal far outside the bandwidth of the coil. When the bandwidth of the signal is known, the signal should be low-pass filtered even before discrete samples are taken to prevent unnecessary noise in the lower frequency range. The signal was low-pass filtered in this experiment, but only digitally and after the signal was sampled. Better results could be obtained if the l-

tering were to be done by an analog filter, before the sampling takes place.

7.3. Improvements in scanner design

The last area where improvements could be made, is in the design of the scanner itself. A few ways in which the design could be altered are discussed in this section.

The gradient in the field of the used magnets was quite flat in the z -direction. While these fields were suitable for reconstructing phantoms, as can be seen from the results for the numerical phantom, it might be difficult in reality because of the relatively flat gradient in the z -direction. This flat gradient makes it so that for the signals coming from different places in the field are quite similar, since the thickness (in the z -direction) of the bandwidth only changes slightly between different places. If the gradient in this direction would vary more with position, the reconstruction might be easier.

8. Conclusion

In an effort to answer the question of how a phantom could be reconstructed using a hand-held MRI-scanner, relevant information on magnetic fields was obtained. Namely that it is essential to have the curves of magnetic field strengths not be flat planes, but have some curvature and preferably minimal symmetry.

While the results obtained using the numerical phantom would suggest that the magnetic field of the magnets in combination with the translations that were done with each of the measurements is capable of reconstructing correct images from the measured spin echo signals, the model failed to reconstruct the samples used in the measurements. Only with significant ℓ_1 -regularization for the difference between adjacent pixels, some proton density was reconstructed in the domain, but this is most likely a result of the geometry of the field, in combination with the regularization.

The combination of translations and magnetic field was shown to be suitable for image reconstruction, since the model was able to correctly reconstruct numerical phantoms. That the reconstruction fails when considering real phantoms could be because of either an approximation or assumption in the model that is invalid or because of excessive noise in the measurements. Further research is required in order to determine which of these possibilities causes the reconstructions to fail.

Ways in which the reconstruction might be improved were discussed in the previous chapter. Some of these are easier to implement than others. The alteration that potentially has the most effect is using a magnet with a gradient in the direction, that changes significantly over the imaging domain. With such a magnet, the differences between the signals from different pixels in the xy -plane would be significantly increased. Then, the reconstruction should be less influenced by noise and errors in the magnetic field strength. In addition, the coil sensitivity could be measured instead of estimated, making the model more accurate.

References

- [1] P. C. Lauterbur, “Image formation by induced local interactions: examples employing nuclear magnetic resonance,” *nature*, vol. 242, no. 5394, pp. 190–191, 1973.
- [2] Z. P. Liang and P. C. Lauterbur, *Principles of magnetic resonance imaging, a signal processing perspective*. New York: Institute of Electrical and Electronics Engineers, 2000.
- [3] M. L. de Leeuw den Bouter, M. B. van Gijzen, and R. F. Remis, “Conjugate gradient variants for ℓ_p -regularized image reconstruction in low-field MRI,” *SN Applied Sciences*, vol. 1, no. 12, pp. 1–15, 2019.
- [4] S. Tewari, T. O’Reilly, and A. Webb, “Improving the field homogeneity of fixed-and variable-diameter discrete halfbach magnet arrays for MRI via optimization of the angular magnetization distribution,” *Journal of Magnetic Resonance*, vol. 324, p. 106923, 2021.
- [5] M. R. Hestenes, E. Stiefel, *et al.*, *Methods of conjugate gradients for solving linear systems*, vol. 49. NBS Washington, DC, 1952.
- [6] Å. Björck, *Numerical methods for least squares problems*. SIAM, 1996.