Description of a Low-field MRI Scanner Based on Permanent Magnets

Merel de Leeuw den Bouter[1], Dilan Gecmen[1], Angeline Meijer[1], Danny de Gans[2], Lennart Middelplaats[2], Rob Remis[3], and Martin van Gijzen[1]

[1] Delft Institute of Applied Mathematics, Delft University of Technology
M.L.deLeeuwdenBouter-1@tudelft.nl, dbgecmen@gmail.com, angelinemeijer27@gmail.com, M.B.vanGijzen@tudelft.nl
[2] Dienst Elektronische en Mechanische Ontwikkeling, Delft University of Technology
D.H.deGans@tudelft.nl, L.N.M.Middelplaats@tudelft.nl
[3] Circuits and Systems, Delft University of Technology
R.F.Remis@tudelft.nl

Abstract. More than 6,000 infants develop hydrocephalus in East Africa every year. Magnetic Resonance Imaging is the preferred technique to diagnose hydrocephalus. In countries such as Uganda, MRI is unaffordable at even major referral hospitals. In order to provide a sustainable diagnostic tool we are developing an inexpensive and easy-to-use MRI system that yields images of sufficient quality to diagnose hydrocephalus. This paper describes our first prototype of such a scanner. We explain the lessons that we have learned from this prototype and how we used these to come up with an improved design. We also describe a dataset that has been obtained with this scanner that will be made publically available.

Keywords: Low-field MRI, Halbach array, Dataset.

1 Introduction

Hydrocephalus is a build-up of cerebrospinal fluid. The increased pressure causes the head to swell and damages brain tissue. If left untreated it may cause pain, blindness, mental disorders, and ultimately death. In industrialized countries, infant hydrocephalus is usually due to either a congenital anomaly or, in low birthweight premature infants, due to brain hemorrhages from immature blood vessels. Most of the time it is detected at an early stage, when it is relatively easy to treat the condition. Hydrocephalus in children in Uganda and other developing countries, on the other hand, is dominated by infectious causes, is seasonal, and appears to be linked to environment and farm animals living in close proximity to families [12]. It is usually detected at a late stage due to lack of access to proper health care. It is estimated that more than six thousand

 Colour and Visual Computing Symposium 2020, Gjøvik, Norway, September 16-17, 2020.

children develop hydrocephalus in East Africa in every year [14]. Many of these children remain untreated, and suffer and die as a consequence.

Hydrocephalus can be well treated if properly diagnosed. MRI is a widely used diagnostic tool for intracranial disease. Conventional MRI scanners yield high-resolution images, but are often out of reach in developing countries due to their cost, high maintenance, and requirement for cryogenic cooling. High-resolution images are, however, not necessary to diagnose hydrocephalus. A simpler device that yields lower resolution images should be sufficient for diagnostic and treatment purposes.

The aim of our research is to develop a prototype of such an inexpensive MRI scanner that is easy to install, operate and maintain, and that can provide images of sufficient quality to diagnose hydrocephalus and manage its surgical treatment. We aim at a device that costs less than 50,000 EUROs and yields images with a resolution of 5 mm^3, which is sufficient for treatment of hydrocephalus. The project is a collaboration between Leiden University Medical Center (LUMC, Netherlands), Pennsylvania State Unverisity (PSU, USA), Mbarara University of Science and Technology (MUST, Uganda), and the Delft University of Technology (TU Delft, Netherlands).

This paper describes our first prototype of a low-field MRI scanner that we have been developing over the past three years. The magnetic field inside the scanner is generated solely using inexpensive permanent magnets. We explain the design choices that we have made, and show some of the bottlenecks that we encountered. We also explain the lessons learned and how we have used these to make an improved design.

Low-field MRI has received quite some attention recently in literature. In [8] an overview is given of the challenges of low-field MRI. In [10], a low-field MRI scanner is described for infant hydrocephalus. This design uses resistive electromagnets. A paper that particularly inspired us is [4]. This paper describes a permanent magnet MRI scanner that uses rotation of the magnet for spatial encoding. In [1] a system is described that uses concentric rings of magnets to improve spatial encoding. The paper [13] proposes a spatial non-linear encoding design, based on moving magnets.

The structure of this paper is as follows. Section 2 provides relevant background on low-field MR imaging. In particular it gives the signal model. Section 3 explains the configuration of the magnets. Section 4 describes the electronics. Section 5, explains how we created a phantom which allows for a plethora of different geometric configurations. We describe the dataset of 53 different configurations or images we use in our experiments and we give some reconstruction results. In the final section we describe the lessons learned and how we used these lessons in our second protoype scanner.

2 Low-field MRI

Conventional MRI scanners use superconducting magnets to generate a strong homogeneous magnetic field, which makes the spins of the hydrogen nuclei line up

along the magnetic field. An RF-pulse at resonance frequency is applied, which make the nuclei absorb energy. After the RF-pulse stops, the energy is released which induces a tissue specific signal in a receiver. The resonance frequency depends on the magnetic field strength. To control the area that is examined, linear variations are applied to the magentic field using so-called gradient coils. The data acquired in this way can be turned into an image by applying an inverse Fourier transform.

The most expensive part of a conventional MRI system is the superconducting magnets. In our design we therefore replaced these magnets by a configuration of inexpensive off-the-shelf permanent magnets. Moreover, variations in the magnetic field can be used for spatial encoding, thus removing the necessity for gradient coils.

The signal model for low-field MR imaging is described in detail in [2]. It can be described by a linear system of equations:

$$\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e} \tag{1}$$

where, $\boldsymbol{b}$ is the measured signal polluted by noise $\boldsymbol{e}$, $\boldsymbol{A}$ is a known matrix, and $\boldsymbol{x}$ is the unknown image. The elements of $\boldsymbol{A}$ are given by

$$a_{ij} = \omega_j^2 e^{-i\Delta\omega_j t_i} \Delta x \Delta y \Delta z \tag{2}$$

where $\Delta x \Delta y \Delta z$ is the voxel size, $\Delta\omega = \omega - \omega_0$, j is the voxel number, ω denotes the resonance frequency, or Larmor frequency, and ω_0 is the demodulation frequency. The angular frequency ω depends linearly on B:

$$\omega = \gamma B, \qquad \omega_0 = \gamma B_0. \tag{3}$$

Here, γ is the gyromagnetic ratio.

Since frequencies, or equivalently magnetic field strength, are not uniquely mapped onto a location, one measurement is insufficient to make an image. In order to (partly) overcome this problem multiple measurements can be made by rotating the magnet (or the sample to be imaged). This technique was used in [4].

3 Magnet design

In this section, the configuration of rings and magnets in our MRI scanner is described. A more extensive description can be found in [9]. The magnetic hardware consists of three parts: a main magnet cylinder to generate a strong enough background field, shimming rings to increase the homogeneity of the background field, and a gradient ring to enable spatial encoding.

3.1 Magnet cylinder

The first part of the magnet consists of the magnet cylinder which is the core part of the magnet. The cylinder consists of four rings with in each ring 24 1-inch cube

magnets NdFeB N52 magnets. The residual flux density, B_r, of these magnets is between 1.45 and 1.48 T. These magnets are put into a first order Halbach configuration [6]. Since the configuration is a *Halbach* configuration, the field strength inside the rings is maximized while the field outside is approximately zero. Furthermore, a *first order* Halbach configuration is used since this yields a homogeneous magnetic field in the case of an ideal Halbach cylinder. Due to physical limitations, such as the the cylinder not being infinitely long and the magnetization of the ring being discrete with 24 distinct magnets instead of continuous, the generated magnetic field is not completely homogeneous. The four rings in the cylinder all have the same configuration and the same inner radius of 0.150 m. They are positioned 0.0246 m apart. Figure 1 shows the cylinder and the configuration of the rings. The magnitude of the magnetic field that is oriented in the z-direction is negligible.

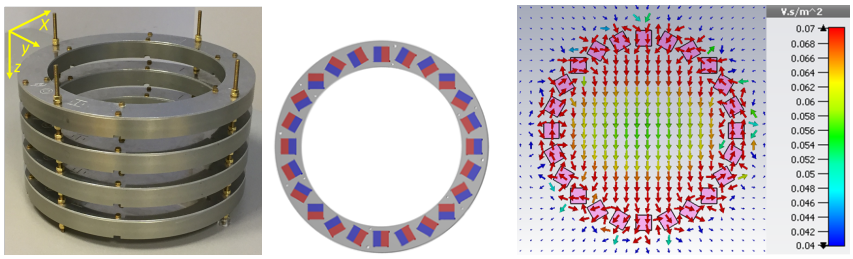

(a) Cylinder consisting of four rings.

(b) Configuration of the 24 1-inch cube magnets.

(c) Simulated magnetic field resulting from the ring in 1b.

Fig. 1: Halbach cylinder, magnets and magnetic field.

To obtain a more homogeneous field, two extra magnet rings are inserted in the cylinder. These rings are called shimming rings. Both shimming rings are of a first order Halbach configuration, with magnets removed at 0° and 180°, as can be seen in Figure 2a. The outermost ring represents a ring from the cylinder as in Figure 2, the two rings within this ring are the shimming rings. The first shimming ring consists of 20 magnets and has a radius of 0.0830 m. The second shimming ring consists of 6 magnets and has a radius of 0.060 m. The magnets that are used in these shimming rings are 12 mm NdFeB N48 cubes that have a B_r in the range of 1.37-1.42 T. Both rings are placed at a distance of 0.0827 m from the middle of the z-axis.

Besides the shimming ring, gradient rings are added in order to superimpose a gradient on top of the background field. The gradient rings are three identical rings combining 16 NdFeB N42 5 mm cube magnets. These magnets are weaker compared to the other magnets described earlier; B_r is between 1.29-1.32 T. The radii of these rings are 0.0905 m and the rings are placed 0.050 m apart from each other. The magnets in these rings are arranged in an adapted version of

a first order Halbach ring orientation. Some magnets have an orientation that is rotated over 180°, as can be seen in Figure 2b. The red line indicates the separation between the magnets in the normal Halbach configuration and the magnets that have a Halbach configuration rotated over 180°.

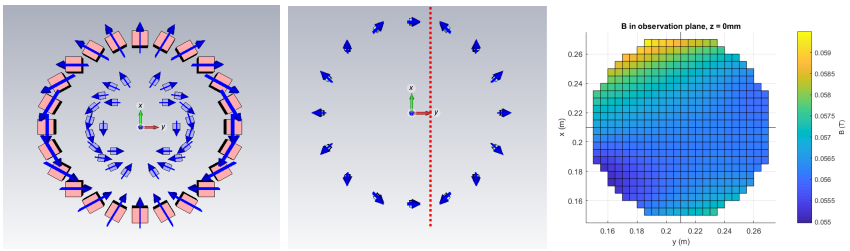

(a) A cylinder ring with the shimming rings.

(b) The gradient ring.

(c) Magnetic field measurement. The xy-plane is shown for $z = 0$.

Fig. 2: Shimming rings, gradient ring, and measured magnetic field.

The magnetic field strength B of the final configuration has been measured using a Gaussmeter Model GM 2 from Alphalab Inc. This meter can measure strong and moderate magnetic fields using a probe with an accuracy of 10^{-5} T. The measurement process is relatively easy. A COSI Measure moves the probe along a preprogrammed route through the 3D space within the cylinder. This COSI Measure is an open source multipurpose 3-axis robot with very high spacial fidelity; the robot is reliable on a spacial scale of less than a millimeter [7]. It measures the magnetic flux density in the x-, y- and z-direction every 5 mm. The results of these measurements are shown in Figure 2c.

4 Description of the electronics

The block diagram in Figure 4 shows the MRI system with each block representing an actual part of it. The PC controls the signals that are generated by the Universal Software Radio Peripheral (USRP) data acquisition. These signals, 2 RF pulses, influence the protons in the sample such that it will transmit a signal back some time after the transmitted RF pulses. This signal is received by the same coil that was transmitting just before. After amplification, filtering and limiting, the received MR signal is sampled by the USRP and transmitted to the PC for further processing.

4.1 RF amplifier

The RF excitation pulses that result in 90° and 180° rotations of the spins in the sample require a lot of power. It takes up to 1 kW in the resistance of the

coil's wire for about 20 μs with a repetition time of say 200 ms. This power is produced by the RF power amplifier (PA). The gain of the PA is about 54 dB.

4.2 T/R switch and low noise amplifier

The same coil is used for sending and receiving the MRI signals. The transmit/receive switch directs the high power signal from the power amplifier to the RF coil. The very weak signal that is picked up by the RF coil is amplified by 70 dB before entering the USRP.

Fig. 3: MRI system

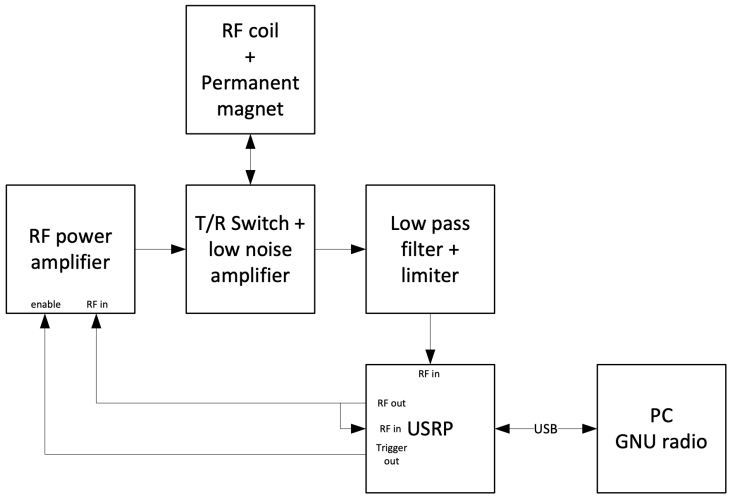

Fig. 4: Block diagram representing the MRI system

4.3 RF coil

The coil that transmits signals to the sample and receives them back from it, is a simple wire wound around a cylinder. It is accurately tuned to the frequency that corresponds to the static magnetic field. The RF frequency at which protons resonate depends on the strength of this magnetic field, see Section 2. The static magnetic field in the coil has a magnitude of about 60mT which means that the resonance frequency is approximately 2.6 MHz.

4.4 Low pass filter

This filter is located directly at the input of the USRP to prevent aliasing in the digitized sampled signal. It also limits the amplitude to prevent damage to the analogue input. The filter is a 3rd order Butterworth filter at 20 MHz.

4.5 USRP1

The data acquisition is done by the USRP1 of Ettus Research. It has two 12 bit ADC boards operating at 64 MS/s and two 14 bit DAC boards operating at 128 MS/s. Each board has two inputs or outputs. An FPGA takes care of the up- and down-conversion in frequency of the signals. This means that a fixed input value can produce an RF output signal from a DAC. Also the received MR signal is down converted to produce a low frequency complex signal. The USRP1 is combined with the GNU software defined radio for generation of the signals and processing at reception.

5 Measurements

This section describes some of the data we have acquired with our system, and also gives preliminary analysis results. A dataset has been created for public use. The samples (images) consist of variations of four geometric shapes: circles, ellipses, squares, and rectangles. In this section, we summarise how we created and developed the phantom used to generate the dataset and how the dataset was acquired. A complete description can be found in [5].

5.1 Phantoms

For the measurements an efficient phantom has been designed that can contain multiple geometric shapes in 12 different positions. The shapes consist of squares, circles, ellipses, and rectangles. A 2D illustration of the phantom and shapes is given in Figure 5.

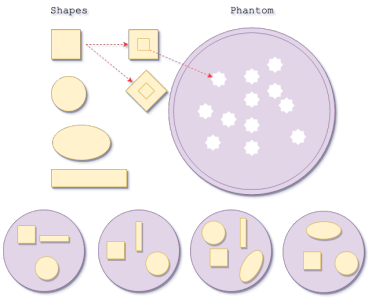

Fig. 5: The shapes are depicted in yellow, the phantom in purple, and the holes in white. A square is placed at the back of each shape that fits into the holes.

The phantom has a radius of 40 mm. Each hole in the phantom is made of two congruent squares with the same center at 45° angles. The dimensions of each square are $5 \times 5 \times 2.3$ mm. A block with the same dimensions is placed behind each shape so that each shape can be rotated by 45° in each hole. Each shape is filled with sunflower oil until 3 mm from the top.

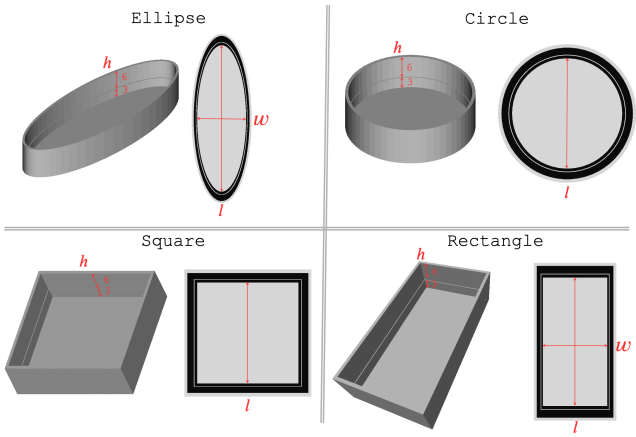

Fig. 6: The height h, length l, and width w illustrated for each shape. For each geometric shape, variations are made with the same height but with different length and width, see Table 1. This is done to make the dataset as diverse as possible.

Table 1: Length l and width w for each shape in millimeters.

	Ellipse		Circle	Square	Rectangle	
Shape	Length	Width	Length	Length	Length	Width
1.	60	20	35	40	40	20
2.	50	10	30	30	50	20
3.	40	30	25	25	30	20
4.	40	20	20	20	40	30
5.	30	20	-	-	-	-
6.	20	10	-	-	-	-

A suitable repetition time depends on the T_1 value of the liquid placed in the phantom. We use sunflower oil for the measurements for its short T_1 value of 90 ms. To avoid spilling, the total height of the shapes is 9 mm. Each shape is divided by a horizontal line at 3 mm from the top. This is illustrated in Figure 6.

SolidWorks, a computer-aided design software, is used to create 3D printable models. The phantom and shapes are printed using the 3D printer Formlabs Form 2, which uses clear resin (SLA). Using clear resin to print the shapes is very convenient as the material does not absorb sunflower oil.

The phantom is printed using the fused filament fabrication 3D printer Ultimaker 2+, which uses a continuous filament of polylactide (PLA). PLA is a compostable thermoplastic material made from renewable resources, such as sugarcane or corn starch.

5.2 Acquisition method

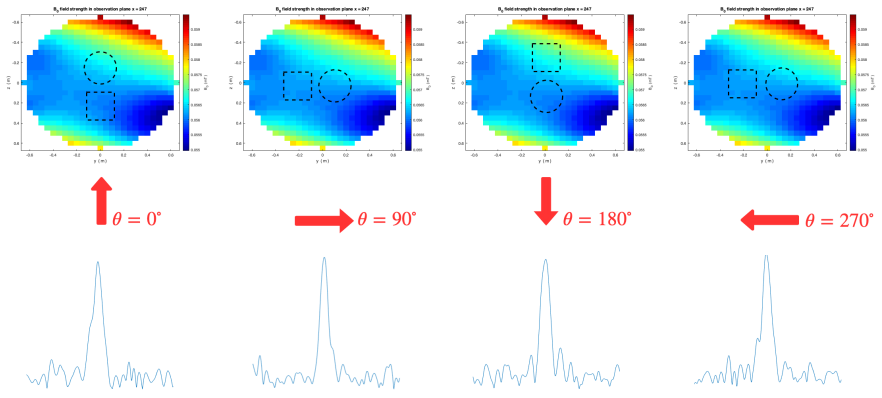

Fig. 7: Phantom containing a circle and square rotated in the field. After each rotation of 90° the signal is measured.

To improve spatial encoding the RF coil is rotated in the magnet with an angular increment θ. At each increment, the field experienced by the sample changes due to the inhomogeneity of B_0. This is illustrated in Figure 7. Starting from $\theta = 0°$ the RF coil is rotated by increments of ten degrees and after each rotation the signal is measured $(0° - 350°)$. The number of samples recorded after a spin echo sequence is set to 512. One spin echo sequence results in measured signal with a low SNR. Averaging can be applied to improve the SNR of the measurements. For each rotation, the measurements are repeated a hundred times. The hundred measurement are then summed up which results in the measured signal vector of size (1×512) for each angle θ_i. The measurements are done for 36 angles, so the signal matrix for one phantom is of size (36×512).

Of each variation of shapes in the phantom a digital image x is implemented. The image matrix x is represented by a square matrix of size (64×64) whose elements are pixel values corresponding to black or white. Black pixel values represent areas that do not give signal and white values are the areas filled with oil, which give signal. The process of creating one labelled sample is illustrated in Figure 8. Each sample $i \in \{1, \ldots, 53\}$ consists of a signal and its corresponding image (label).

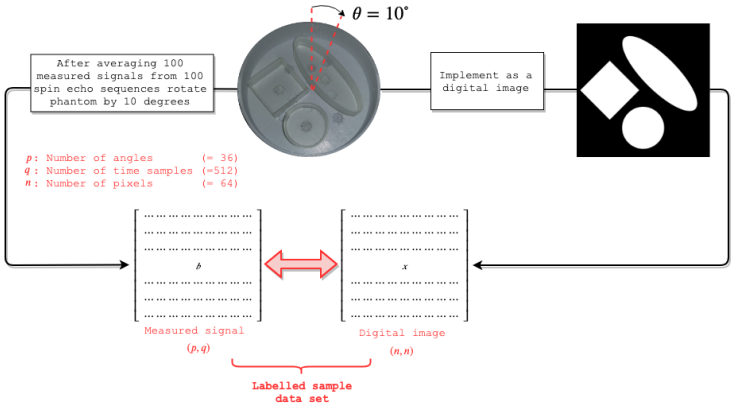

Fig. 8: Process of creating one labelled sample for the dataset. The labelled dataset can be used to apply deep learning models on so that the model can predict labels for unlabelled samples.

The data of each measurement are stored in a netCDF file, which contains all relevant acquisition parameters. All variable are in standard SI units. A total of $i = 53$ signals has been measured.

Table 2: The variables of a measurement stored in a netCDF file.
Format: netcdf4_classic

Global Attributes	
Echo time	0.01
Repetition time	0.5
RF	$2.61 \cdot 10^6$
FlipTime	1e-05
RFamplitude	36
spinEcho	100
samplingRate	$1 \cdot 10^6$
phantom	'signal_i'
angle_increment	10
Dimensions	
no_angles	36
scan_width	512
Variables	
signal_real	
Size	36×512
Dimensions	no_angles, scan_width
Datatype	double
signal_imag	
Size	36×512
Dimensions	no_angles, scan_width
Datatype	double

5.3 Model based reconstruction

Using the model described in Section 2 it is in principle possible to compute an image from the measurements. For the description of the processing techniques we refer to [2]. Unfortunately we have not been able to obtain reconstructions of sufficient quality with this approach. We assume that the main reason for the poor reconstructions is insuffiently accurate knowledge of the magnetic field. Since many locations correspond to approximately the same frequencies, small errors in the measurement for the magnetic field may result in big errors in the reconstructed images. This assumption is confirmed by the result in [3].

5.4 Data driven reconstruction

The 53 measured signals can be arranged in a matrix $\boldsymbol{M}$ of size $(pq \times i)$, where the columns of size pq correspond to the 53 measured signals. We want to see if it is possible to reconstruct an image based on a linear combination of other measured signals, which means that there must be some sort of correlation between the measured signals. This leads to a system of equations of the form:

$$\boldsymbol{M'}_{pq \times i} \boldsymbol{y}_{i \times 1} = \boldsymbol{s}_{pq \times 1} \tag{4}$$

where $\boldsymbol{y}$ is the unknown, $\boldsymbol{s}$ is one of the measured signals, and $\boldsymbol{M'}$ is the matrix $\boldsymbol{M}$ where the column corresponding to signal $\boldsymbol{s}$ is replaced by zeros. The singular value decomposition (SVD) of $\boldsymbol{M'}$ is given by:

$$\boldsymbol{M'}_{pq\times i} = \boldsymbol{U}_{pq\times i}\Sigma_{i\times i}\boldsymbol{V}^T_{i\times i} \tag{5}$$

and $\boldsymbol{M'}$ has rank $r = 52$. The least squares solution to (4) is given explicitly by:

$$\boldsymbol{y} = \sum_{j=1}^{r} \frac{\boldsymbol{v}_j \boldsymbol{u}_j^T}{\sigma_j}\boldsymbol{s} \tag{6}$$

where, $\boldsymbol{u}_j$ and $\boldsymbol{v}_j$ are the columns of $\boldsymbol{U}$ and $\boldsymbol{V}$, respectively. Now $\boldsymbol{y}$ can be used to reconstruct the signal $\boldsymbol{s}$ and the corresponding image $\boldsymbol{x}$, where the reconstruction is denoted by $\boldsymbol{s}'$ and $\boldsymbol{x}'$:

$$\begin{cases} \boldsymbol{s}' &= \sum_{j=1}^{53} y_j \boldsymbol{s}_j \\ \\ \boldsymbol{x}' &= \sum_{j=1}^{53} y_j \boldsymbol{x}_j \end{cases} \tag{7}$$

Let $\boldsymbol{s}$ and $\boldsymbol{x}$ correspond to the first measured sample, Signal_1 and Image_1. The signal and the signal corresponding to the reconstruction are shown in Figure 9. We can see that the amplitude of the signal for each angle is reconstructed reasonably well.

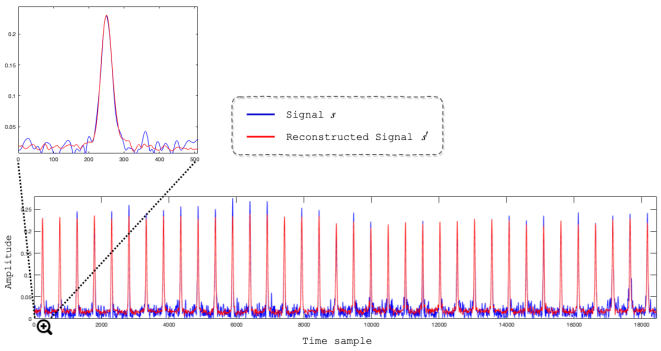

Fig. 9: The measured time signals (blue) plotted against the reconstructed signals (red) for each angle. The measured signal corresponds to the first measured sample, Sample_1.

The reconstruction of the measured signal looks promising and therefore it should also be possible to reconstruct a fairly reasonable image. The image and its reconstruction are visualized in Figure 10.

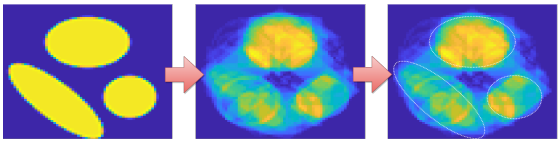

Fig. 10: Left: The image corresponds to the first measured sample, Sample_1. Middle: Linear reconstruction of the image using SVD. Right: Dotted lines around the edges of the shapes in the original image are placed in the reconstructed image for a better visualization of the reconstruction.

The samples in the measured dataset are varied. For each measurement the shapes are placed inside the phantom in a different way. From Figure 11 we can see that the images are built up of other images containing shapes in the same location. We can conclude that the signal contains information about the position of the shapes. These results give us confidence that a dataset as described above

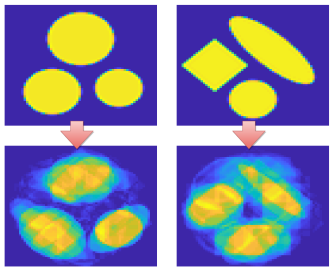

Fig. 11: The images corresponding to the seventh measured sample, Sample_7, and 46^{th} measured sample, Sample_46.

can be at the basis of a deep-learning approach to reconstruct the image, without the need to have accurate knowledge of the magnetic field. We refer to [5] for the details. This approach works well on simulated data. Unfortunately, due to COVID-19, we had to suspend our measurements and the current dataset is still too small to be used as a training set.

6 Final remarks and further research

In this paper we have described our first prototype for a low-cost MRI scanner. All parts of a conventional scanner have been redesigned and are replaced by less expensive alternatives. Both the main magnet and the gradient coils have been replaced by configurations of inexpensive off-the-shelf permanent magnets. The (hardware) cost of our design is approximately 15,000 EUROs.

Images can in principle be obtained by exploiting variations in the magnetics field. However, this requires very accurate knowledge of the magnetic field. Until now we have not been able to obtain more than rudimentary images, of much lower quality than is needed for medical purposes. As an alternative we are investigating a data-driven approach. We have acquired a dataset for this purpose and have applied some basic processing techniques to it. This dataset will be made publically available through the data repository https://researchdata.4tu.nl/en/. We aim to extend the dataset to make it large enough to be used as a training set for deep learning.

One of the main lessons learned is that encoding spatial information using field variations induced by the permanent magnets proved to be quite challenging for image reconstruction. The second prototype, which has been designed by the LUMC and is described in [11] therefore uses the conventional technique with gradient coils for spatial encoding. This design, which costs in the order of 30,000 EUROs yields images with a resolution that is well above the resolution needed for the treatment of hydrocephalus. This prototype will be sent to MUST in Uganda for further testing later this year.

Acknowledgement

This research is partly funded by TU Delft Global Initiative, by NWO WOTRO under grant W07.303.101 and by STW Open Mind under grant 15549.

The authors thank all the members of the low-field MRI team, in Uganda, the USA and in the Netherlands, for the nice collaboration and many interesting discussions. In particular we thank Andrew Webb, Tom O'Reilly, and Kirsten Koolstra of the Leiden University Medical Center for their assitance with and advice on many aspects of the work. We thank the referees for their useful comments.

References

1. Blümler, P.: Proposal for a Permanent Magnet System with a Constant Gradient Mechanically Adjustable in Direction and Strength. Concepts in Magnetic Resonance Part B: Magn. Reson. Engin. **46**(1), 41–48 (2016)
2. De Leeuw den Bouter, M.L., Van Gijzen, M.B., Remis, R.F.: Conjugate gradient variants for ℓ_p-regularized image reconstruction in low-field MRI. SN Applied Sciences **1**(12), 1736 (2019)
3. De Leeuw den Bouter, M., Koolstra, K., O'Reilly, T., Börnert, P., Webb, A., , Remis, R., Van Gijzen, M.: Joint Iterative Image Reconstruction and Field Map Estimation in Low Field MRI. In: Proceedings of the ISMRM Benelux Chapter meeting 2019 (Jan 2019)
4. Cooley, C.Z., Stockmann, J.P., Armstrong, B.D., Sarracanie, M., Lev, M.H., Rosen, M.S., Wald, L.L.: Two-dimensional imaging in a lightweight portable MRI scanner without gradient coils. Magn. Reson. in Medic. **73**(2), 872–883 (2015)
5. Geçmen, D.: Deep Learning Techniques for Low-Field MRI. Master thesis, Delft University of Technology (February 2020), online available at http://resolver.tudelft.nl/uuid:ce264a44-ddd5-45c5-96d0-c82aac0e4911

6. Halbach, K.: Design of permanent multipole magnets with oriented rare earth cobalt material. Nucl. Instr. and Meth. **169**(1), 1–10 (1980)
7. Han, H., Moritz, R., Oberacker, E., Waiczies, H., Niendorf, T., Winter, L.: Open source 3D multipurpose measurement system with submillimetre fidelity and first application in magnetic resonance. Nature Scientific Reports **7**(1), 13452 (2017)
8. Marques, J., Simonis, F., Webb, A.: Low-field MRI: An MR physics perspective. J. of Magn. Reson. Imag. **49**(6), 1499–1802 (2019)
9. Meijer, A.: Optimizing the gradient ring of a low-field MRI scanner. Bachelor thesis, Delft University of Technology (August 2019), online available at http://resolver.tudelft.nl/uuid:56a2f8d7-fb26-40f0-bb92-00271f9492df
10. Obungoloch, J., Harper, J., Consevage, S., Savukov, I., Neuberger, T., S., T., Schiff, S.: Design of a sustainable prepolarizing magnetic resonance imaging system for infant hydrocephalus. Magnetic Resonance Materials in Physics, Biology and Medicine **31**, 665–676 (2018)
11. O'Reilly, T., Teeuwisse, W.M., Webb, A.G.: Three-dimensional MRI in a homogenous 27 cm diameter bore Halbach array magnet. J. Magn. Reson. **307**, 106578 (Oct 2019)
12. Schiff, S., Ranjeva, S., Sauer, T., Warf, B.: Rainfall Drives Hydrocephalus in East Africa. J. Neurosurg. Pediatr. **10**(3), 161–7 (2012)
13. Vogel, M., Guridi, R., Su, J., Vegh, V., Reutens, D.: 3D-Spatial encoding with permanent magnets for ultra-low field magnetic resonance imaging. Nature Scientific Reports **9**(1522) (2019)
14. Warf, B.: Pediatric hydrocephalus in East Africa: Prevalence, causes, treatments, and strategies for the future. World Neurosurg. **73**, 296–300 (2010)