

Evaluating the Robustness of the Recurrent Inference Machine Based on MRI Lesion Data from Multiple Sclerosis Patients

by

Kevin Pancras

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on Monday July 5, 2021 at 13:00.

Student number:	4713311	
Project duration:	March 1, 2020 – July 5, 2021	
Thesis committee:	Prof. dr. F.M. Vos	TU Delft, supervisor
	Asst. Prof. dr. Q. Tao	TU Delft
	Asst. Prof. dr. M.W.A. Caan	Amsterdam UMC, supervisor
	ir. D. Karkalousos	Amsterdam UMC, daily supervisor

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.



Abstract

The Recurrent Inference Machine (RIM) has been developed as an alternative to the clinically used Compressed Sensing (CS) algorithm, using Deep Learning (DL). A common issue with DL networks is the generalization of the network to features that have not been trained for. In this study we evaluate the robustness of the RIM to white matter lesions in FLuid Attenuated Inversion Recovery (FLAIR) brain MRI data. We are evaluating two pre-trained RIM networks, one trained on T1 brain data and another trained on T2 knee data. This evaluation was done by comparing the two networks to CS in terms of the average relative Signal-to-Noise Ratio (SNR) and Contrast Resolution (CR) that is achieved on 15 datasets acquired from Multiple Sclerosis patients. From these comparisons it shows that the network trained on T2 knee data performs similar to CS in terms of the relative SNR, while having a higher CR. The network trained on T1 brain data has both a lower relative SNR and CR, compared to CS. The data suggest that the RIM trained on T2 knee data is robust to the inclusion of lesions in an area that the network was not trained on.

Acknowledgements

Before you lies my master's thesis, marking the end of my master's in Biomedical Engineering at the Technical University in Delft. The period in which I have worked on this project, for a duration of 1 year and 3 months, was not, as many will agree, the easiest period. With a lot of working from home and covid, my project took longer than intended, and was also in ways less social than intended. However, that does not mean that I did not have an interesting time at the Amsterdam UMC, within a great research group.

I would first like to thank my supervisor at the AUMC, Assistant Professor Matthan Caan. You are a wonderful and kind person to work with, being truly interested in the people that work in, and outside, your research group. Because of that, your group has the feeling of a family, which I hope I will be able to find, and help create, in the workplaces I will find myself in.

I would also like to thank my daily supervisor, Dimitris Karkalousos, who has taught me a lot about working with the model, and helped me get my writing to a higher level. Your expertise proved invaluable when there were problems I could not tackle on my own.

I would also like to thank my supervisor from the TU Delft, Professor Frans Vos. You were available to help, both during the stressful times, and the general course of the project. You were also one of the cornerstones of my time at the TU Delft as the head of the specialization Medical Physics.

There were also some people that helped with providing and processing the data. For that part I would like to thank Assistant Professor Hanneke Hulst, for providing the datasets, Dr. Marthijn Steenwijk and Samantha Noteboom, for segmenting the lesions from the data with their network, and I would like to thank Kai Lønning for providing me with a starting script to loading the data from the raw data format.

I would like to thank my mother for giving me the drive to achieve the things that I want to achieve. You have shown me that there is a way to get where I want to go, even if it takes longer, or requires more effort. I would also like to thank my father for always teaching me that it does not matter what I do with my life, as long as I am happy with the life that I am giving myself.

And last, but certainly not least, I would like to thank my girlfriend Vera Čaluković, without you, my life would be very boring. You help me both in keeping me company and keeping me sane, as well as being a valuable intellectual sparring partner.

Acronyms

CR Contrast Resolution.

CS Compressed Sensing.

DL Deep Learning.

FLAIR FLuid Attenuated Inversion Recovery.

FOV field-of-view.

MPRAGE Magnetization Prepared RApid Gradient Echo.

MRI Magnetic Resonance Imaging.

MS Multiple Sclerosis.

RF Radio Frequency.

RIM Recurrent Inference Machine.

SNR Signal-to-Noise Ratio.

Contents

1	Introduction	1
2	Materials and methods	4
2.1	Datasets	4
2.1.1	MS patient data	4
2.1.2	Training data	6
2.2	Preprocessing	6
2.2.1	Data reader	6
2.2.2	Coil sensitivity maps	7
2.3	Reconstruction methods	8
2.3.1	Compressed Sensing	9
2.3.2	Recurrent Inference Machine	9
2.4	Evaluation	9
3	Results	13
3.1	Visual assessment	13
3.2	Quantitative analysis	15
4	Discussion	20
4.1	Future work	21
5	Conclusion	23
A	Theory	27
A.1	Introduction	27
A.2	Background on MRI	27
A.2.1	The Inverse Problem	29
A.3	Compressed Sensing	30
A.3.1	Theory	30
A.3.2	Compressed Sensing for MRI Reconstruction	31
A.4	Machine Learning	33
A.4.1	Supervised & Unsupervised Learning	33
A.4.2	Convolutional Neural Networks	33
A.4.3	Physics-based models	34

A.5	Recurrent Inference Machine	34
A.6	The future of the RIM	37

Chapter 1

Introduction

Medical imaging is an important part of medical care, as it gives vital information about the status of the patient. For the benefit of the patient, it is best that the imaging is done in a non-invasive and non-damaging manner. One of the most popular and frequently used non-invasive imaging technique in the clinical setting today is Magnetic Resonance Imaging (MRI).

In MRI, as the name suggests, the signal is generated by the magnetic resonance of hydrogen protons. These have a specific resonance frequency depending on the magnetic field strength, making it possible to give them a magnetization [1]. Slight deviations exist in resonance frequency depending on the configuration in which a proton is embedded. Also, the magnetic susceptibility may be different depending on the environment of the proton.

These differences affect the signal that will be measured with certain scanning parameters. This is used to create different scanning sequences, that are capable of giving anatomical or functional information, such as functional MRI (fMRI) and Diffusion Weighted Imaging (DWI). In terms of anatomical variations, there are sequences like T1-weighted (which has bright fat signal) and T2-weighted (which has lower fat and higher water signals). Furthermore, it is possible to have sequences like FLuid Attenuated Inversion Recovery (FLAIR), where the signal of fluids is nulled by only recording data after the magnetization of the fluid signal has been cancelled out.

With MRI, it is not possible to directly get an image of the tissue, but it is necessary to work through the frequency domain, known as k-space. The k-space represents the Fourier-transformed MR-image, containing the harmonic frequencies that describe it in image space. This information is received through the radio frequency (RF)-receivers that can receive the frequency data of a specific subsection of k-space when a subsection of the body is magnetized.

The process of acquiring data is, however, a slow process. With general MRI only one line can be obtained at a time, with a repetition time (the time between two lines) ranging from half a second to more than 5 seconds, depending on the sequence type [25]. For 2-Dimensional imaging these acquisition times are acceptable, but for 3-Dimensional imaging there is an increasing need for




Figure 1.1: Image showing a typical blurred image due to the use of random 2D under-sampling, before reconstruction.

higher resolution and faster scanning times. The scanning times are constrained due to hardware limitations, making the data acquisition process too much of a burden for the patient.

In order to speed up the scanning times, different techniques have been developed. The current state-of-the-art technique that is used in practice, is Compressed Sensing (CS) [11]. CS makes use of the fact that when sampling less in k-space, thus acquiring less data, in a random manner, the acquired image will be evenly blurred; an example of such a blurred image can be seen in Fig. 1.1. This makes it possible to use iterative techniques for restoring an image which will be a close approximation to the one acquired from a fully-sampled k-space. As CS uses a subset of the fully sampled data, inherently the image will contain more noise and less detail than a fully-sampled one, and will therefore have a worse quality.

CS, however, is also computationally intensive and slow, regarding the image reconstruction. For accelerating the reconstruction times and achieving high quality images, research has branched out to Deep Learning (DL) [7, 18, 21]. These DL-networks focus on learning strong prior features learned from a training dataset, for the reconstruction of other datasets. The earliest models were based on Convolutional Neural Networks [20, 26], while later models made use of physics based models [17, 5]. For more information on these types of models, see Appendix A.4.

An issue with DL networks is the generalization to different types of data[19]. For addressing these issues, evaluations need to be carried out, testing the robustness of such networks under various conditions, such as with different sequences and pathologies. It has been shown before that it is possible to train

a network on one sequence type (for example T1) and use this model to reconstruct data from another sequence type, and thus different contrast [6]. Another important condition that networks need to be tested on, is on different pathologies. For the diagnosis of diseases, it is vital that tissue that has been changed due to disease can be differentiated from healthy tissue. Therefore, it is also important to evaluate DL networks on the reconstruction of data containing disease pathology.

In this thesis, we aim to show the robustness of the Recurrent Inference Machine (RIM) network [17] to image features that have not been trained for, by reconstructing prospectively under-sampled brain FLAIR image data containing white matter lesions, visible as areas with higher signal values in the white matter. For this evaluation, we will compare two pretrained RIMs to CS, where each network was trained on a different sequence and anatomy, one on T2 knee data and the other on T1 brain data.

These aforementioned lesions are an important part of the diagnosis, and the tracking of, the progression of Multiple Sclerosis (MS) [8], a neurodegenerative disease. For this purpose, it is important that the lesions are distinguishable from the areas around them, as the size of the lesions can give important information on the progression of the disease. In order to test the reconstruction quality, the lesions are segmented by a different deep learning network, and the Contrast Resolution will be used to test for the distinction between the lesions and the white matter.

As the knee data has structures in the data that have similar contrasts and tissue variations, and has a similar contrast as the data that contains the contrast rich lesion data, the expectation is that the network trained on the knee data will better reconstruct the lesion contrast within the region of the lesions.

Chapter 2

Materials and methods

2.1 Datasets

2.1.1 MS patient data

For this study, under-sampled MRI data was available from subjects of an ongoing MS study.

MS is a neurodegenerative disease marked by inflammation of the central nervous system, resulting in inflammation, or lesions, in the white and gray matter and other areas of the brain [8]. These lesions are not a unique pathology, but also develop with age [2]. That is why the diagnosis of MS done based on the way the symptoms present themselves, in combination with the presence of lesions.

The progression of MS generally happens in multiple stages. The stage that most MS patients are in when they are diagnosed, is the Relapsing-Remitting MS (RRMS) stage. In this stage, the disease is characterized by periods of worse symptoms, with the symptoms mostly improving in between these periods. This generally lasts for at least ten years, before it progresses into more aggressive forms of the disease. The data used in this thesis are of patients being in the RRMS stage of the disease.

Of the available data of this study, 68 datasets have been taken into consideration. Due to limitations that will be described later in this work (Sec. 2.1.1 and 3.2), only 15 datasets have been used for the analysis.

For each volume a T1 and a FLAIR scan is available, both being under-sampled (Sec. A.2). The under-sampling pattern used in this study is 2D and sampled as a Variable-density Poisson disk variety, in the phase and slice encoding direction. An example of the type of under-sampling pattern used is shown in Fig. 2.1. This is not the exact under-sampling pattern that was used in this study, however it is similar to the one that was used.

The data was acquired on a 3.0T Philips Ingenia scanner (Philips Healthcare, Best, The Netherlands), using a 32-channel head coil. For T1 the Magnetization Prepared RAPid Gradient Echo (MPRAGE) scheme was used. The T1 scans




Figure 2.1: Example of the type under-sampling pattern used in this study. This under-sampling pattern is not the same as was used in this study, but it is similar. The main difference is the depicted pattern has a fully sampled center, while the center of k-space was sparsely sampled in the available data. Image taken from [15].

had an isotropic voxel resolution of 1.0 mm, with a volume size of 252x240x256 voxels, and an under-sampling factor of 2.5. The FLAIR scans had an isotropic voxel resolution of 1.12 mm, with a volume size of 163x223x224 voxels, and an under-sampling factor of 7.5. Random under-sampling, using an under-sampling pattern developed by Philips Healthcare, was applied to enable sparse reconstruction. For most subjects a small and fully sampled reference scan was also available. This reference scan was scanned in the same volume as the larger datasets. During scanning, the reference scan used the same field-of-view (FOV) and orientation as the other scans. The only difference in FOV was over-sampling in the frontal axis, in the T1 datasets.

Based on the header information in the datasets, a preselection of the data was made. When working with multi-coil data, the MRI-scanners generally remove the data of individual coils that are too noisy. With the removal of noisy coils, situations arise where some coil elements are removed in the under-sampled data, but not in the reference scan, and vice versa (Fig. 2.2). This mismatch of the amount of coils, or even of different coils being classified as noisy between the under-sampled and reference scan, causes problems for the reconstructions. However, these problems are outside of the scope of this study, and thus the datasets with a mismatching number of coils have been excluded.




Figure 2.2: Image showing an example of coil mismatch. The top row is the under-sampled FLAIR data of a single slice for multiple coils, reconstructed with the zero-filling method, while the bottom row are the coil sensitivity maps that are calculated from the reference scan (see also Fig. 2.3). From left to right are a subset of the coils in the same order. Outlined in blue are the two coils of which the corresponding coil in the other dataset was removed. Outlined in green are two coils that should be above each other, but are now mismatched.

After this exclusion 18 of the 68 datasets were left.

2.1.2 Training data

For the training of the DL-networks, two different datasets were used. A T2 knee dataset, described in [3] and a T1 brain dataset. The T1 brain data were acquired using a 3D T1 MPRAGE acquisition, with a 1.0 mm isotropic voxel resolution, an FOV of 256x256x240 mm, acceleration of 10.8, and contained 32 coils. For the under-sampling of the training data, 2D Gaussian under-sampling was used with a full width at half maximum of 0.7, relative to the dimensions of the k-space. The center was also fully sampled. The fully sampled region in the center of k-space had the shape of an ellipse, with a width of the half-axes that was 2% of the sampled Gaussian region.

2.2 Preprocessing

For the preprocessing of the data, several steps have to be taken. The data has to be read from the raw data structure, a k-space mask has to be generated and the read k-space has to be normalized and cropped. Furthermore, the sensitivity maps have to be created from the smaller reference scan.

2.2.1 Data reader

The original format of the data was in .List/.Data, which is a raw data format structure used in Philips Healthcare MRI Scanners. This means that the data contain raw k-space values which need to be preprocessed. The loading is done by first reading the header from the .List file. The header contains all the necessary information for loading the data, such as the matrix size and which

bytes correspond to which k-space points. Once read, the header information is used to load the binary data in the .Data file.

Out of this loaded k-space we can extract a binary mask representing the sampling pattern, needed for the later steps. The under-sampled k-space contains data that is zero where no acquisition was done. Thus, the mask was created by including the points of a single 2D slice along the readout direction that have a value higher than zero.

The used techniques benefit from normalization of the volumes. This normalization needs to be performed in image space, in order to maintain a proper range of values, between 0 and 1. Thus, the data were transformed from k-space to image space, and back to k-space for two of the three dimensions, using the following procedure. In the end this procedure is similar to a 1D inverse Fourier transform along the readout direction, but with added normalization and cropping.

The k-space volume is transformed to image-space by the 3-Dimensional inverse Fourier transform. While the volume is in image space, the volume was cropped if oversampling occurred in the readout direction. Whether a volume has oversampling is determined from the header information in the .List file. The oversampling was cropped only in the readout-direction, as cropping the volume in the other directions would cause the k-space to change substantially, resulting in a volume that cannot be reconstructed. The volume is then normalized by dividing it by the maximum absolute value of the volume. Finally, we go back to k-space by applying the 2D Fourier transform.

However, as the Fourier transform is incapable of keeping the under-sampling pattern intact without introducing noise in k-space, the parts of k-space that had zero values before, now have non-zero values. The Fourier transform is not likely to give values of zero for frequency values when it encounters noisy data, especially after the data is normalized. These non-zero values make the reconstruction using CS impossible, since this technique makes use of the under-sampling pattern contained within the k-space. Consequently, the mask that was previously generated is used to conserve the under-sampling pattern. Having the under-sampling pattern intact then provides better reconstruction.

One important difference in the k-space obtained from the normalized image is that the 2D Fourier transform is used, rather than the 3D transform. This allows for the data to be used as individual slices, instead of as a volume. Using image slices rather than 3D image volumes is important for the applied DL method (Sec. A.5), as they were designed to operate on 2D data.

2.2.2 Coil sensitivity maps

Apart from normalizing the images, also the coil sensitivity maps needed to be estimated. The coil sensitivity maps are necessary, as they allow those techniques to reconstruct multi-coil data using localized information of the signals. When using an RF coil array to do image acquisition, every coil has a certain region where they are sensitive to signal. This sensitivity of each coil over the volume can be saved in a coil sensitivity map.




Figure 2.3: This image shows an example of the coil sensitivity maps generated during the preprocessing. The bright regions are the areas where the coils have a higher sensitivity, while the darker areas indicate a lower sensitivity.

The sensitivity map is estimated from the fully-sampled reference scan, which is needed, as the under-sampling pattern does not contain a fully-sampled center.

The sensitivity maps were estimated using the CALDIR method[14] of the BART toolbox[24]. An example of the generated coil sensitivity maps is given in Fig. 2.3. The reference scan is scanned using a lower resolution, but with the same FOV, as a high resolution is not needed for this scan, and the smaller size makes scanning times shorter. However, as the reference scan is smaller than the scans that the sensitivity maps are needed for, the sensitivity maps require resizing. Thus, the generated sensitivity maps are scaled up to the larger under-sampled volumes.

There was a difference in FOV for the T1 MPRAGE datasets, which had a larger over-sampling factor in the frontal axis. However, this change in FOV only changed the amount of background around the images, keeping the center of the image in the same place. In order to make the sensitivity maps the same size as the under-sampled volumes, the sensitivity maps are padded with zeros in k-space, which increases the volume size, without changing the FOV. For the added background regions along the frontal axis of the T1 datasets, the sensitivity maps are also padded with zeros in image-space, in order to not distort the geometry of the sensitivity map.

2.3 Reconstruction methods

For the reconstruction of data in this study, two different methods were applied. The first was the CS, and the latter the RIM, which is a DL network. For a brief overview on these methods the reader is referred to section A.3 and A.5 and the references therein. Both techniques have been used with the option to

reconstruct using a GPU (NVIDIA V100 32GB).

2.3.1 Compressed Sensing

For the CS method, the Parallel Imaging and Compressed Sensing (PICS) option of the BART toolbox was used. This technique can be used both in a slice-wise and volume-wise manner. For this study, the slice-wise manner was applied. The ℓ_1 -wavelet was used (see equation A.3), with a regularization of 0.005 and a maximum number of iterations set to 60 (empirically optimized). In order to make use of the GPU, instead of the conventional CPU, the -g flag was turned on.

2.3.2 Recurrent Inference Machine

The data were also reconstructed using the RIM. The RIM network was used in a modified version of the DIRECT toolbox [23].

With the reconstruction of data using DL networks, there is always the choice between changing the data, in order to work in the model, or training a model on data that has the same characteristics as the data that is to be reconstructed, with each having their own benefit. As the point of this thesis is to assess the robustness of the RIM to unseen data during training, the choice has fallen on bringing the data to the model and to use two different pretrained networks for the comparison. In this case it is less complicated to bring the data to the model, as the available models have already been trained on data with a specific preprocessing pipeline.

The models presented in this work, make use of the Independently Recurrent Neural Network (IndRNN) as choice of the recurrent layers [9]. The motivation for choosing this recurrent variant was the faster reconstruction time it has compared to the Gated Recurrent Unit (GRU) [6], which was originally in this paper [10]. Furthermore, this variant achieves equivalent reconstruction quality to the GRU.

The two networks were trained on different data. One was trained on T2 knee data from a public data repository [3] (an example of the data is given in Fig. 2.4), while the other was trained on T1 brain data (an example of the data is given in Fig. 2.5). Both models were trained, and reconstructed using 64 filter channels and 8 time-steps, and the ADAM optimizer [13]. The implementation of the neural networks was done in Python and using the Pytorch framework [16].

2.4 Evaluation

Evaluation was done by comparing the signal level between the white matter and the lesions as they appear in the 3D-FLAIR scans of the MS patient study. These were used to calculate the relative Signal-to-Noise Ratio (SNR) and the Contrast Resolution (CR). The SNR can give an indication of how easy it is to



Figure 2.4: Example image of the data in the knee dataset. Image taken from [3].




Figure 2.5: Example image of the data in the T1 dataset

discern the signal from the noise in the background. We would therefore prefer the SNR to be as high as possible. However, there are cases where the SNR fails to give a good estimate of the reconstruction quality. This is, for example, when an image is blurred. In the case of blurring, the SNR can still be high, while also not being useful for medical evaluation. In order to give a better approach for the quality of the reconstructions we also look at the CR, which gives a measure of the contrast between two different tissue types. The higher this is, the better they are discernable.

The lesion segmentations were performed at the Vrije Universiteit Medical Center (VUMC), using a Convolutional Neural Network (CNN) created for the segmentation of lesions in MS patients [22].

For the segmentation of the white matter, the SPM toolbox was used [4]. Contrary to the rest of the project, which was done in Python, these segmentations were made using the Matlab toolbox. For these segmentations, the first step was to create a registration of the T1 data to the FLAIR data. Then, using the built-in brain segment function in SPM, the T1 data could be used to create probability masks of the white-matter, gray-matter and CerebroSpinal Fluid (CSF) regions in the brain. Once the different probability masks were created, the white-matter probability mask could be resliced using the earlier done T1 to FLAIR data registration, while the other masks were not used.

A comparison between the different reconstruction methods was made by calculating the relative SNR and the CR.

For the calculation of the relative SNR, the white matter was used as signal region, as the segmentations were already available to determine the contrast. The white matter was used, as it should have a similar signal value for all three reconstruction methods. The noise was calculated by selecting a patch from a region-of-interest in k-space. This patch was taken in one of the corners near the top of the head. In order to get reasonable noise values from the k-space measurement, the k-space was also normalized to have a similar value range as the image-space data.

Although, usually data are normalized to the maximum value of the data, the center of k-space generally has a value that is magnitudes larger than the rest of k-space. Thus, normalizing by the maximum value gives noise values that are unreasonably low, e.g., if the value in the center of k-space is 4 magnitudes higher than the rest of k-space, the noise values will be 4 magnitudes smaller when normalizing by this highest value, and the SNR will be 4 magnitudes larger. Therefore, the maximum value of the k-space was estimated by the median value. When data is normally distributed, the median and mean value of a dataset are roughly the same, but for data that is not normally distributed this is not the case. When data is not normally distributed, the median value gives a better representation of the average [12]. Therefore, the median value is better capable of handling the outlier values that might be very small, or very large, compared to the rest of the data, leading to a better estimation of the scaling.

The CR can give a measure of the amount of contrast between two tissue types, based on the difference of the mean signals. In order to compare the




Figure 2.6: Image showing the dilation of the lesion for the calculation of the CR. Shown are the complete FLAIR image (a), the segmented lesions (b), and the complete image overlaid with the lesion segmentation in red and the used border region in the white matter in yellow(c).

signal of the lesion only with the surrounding white matter area, the lesion masks were dilated by 4 voxels, and the dilated areas that have overlap with the white matter mask were used as border region. An example of this border region is given in Fig. 2.6. As the segmented white matter is not necessarily homogenous, this gives a comparison that focuses more on the border region that is important for the segmentation of lesions. In this case, the CR is defined as

$$CR = \frac{S_A - S_B}{S_A + S_B}, \quad (2.1)$$

where S_A is the signal of the lesion and S_B is the signal in the white matter surrounding the lesion.

Chapter 3

Results

3.1 Visual assessment

A direct comparison of one slice for the three different reconstruction methods can be seen in Fig. 3.1. In this figure the root-sum-of-squares image is also shown, which is a representation of the under-sampled image using the zero-filling method, i.e., unsampled values in k-space are taken as zero. In this comparison it can be seen that the CS reconstruction contains the least amount of visible noise in the background regions. Furthermore, the two reconstructions that are the most similar, on visual inspection, are the CS reconstruction and the brain model reconstruction, as they look similar in the white matter areas next to the ventricles. On the other hand, the boundaries of the same darker colored area looks less aliased in the reconstruction using the knee model. The brain model reconstruction and CS reconstruction also look more blurred in this same area. However, the reconstruction using the knee model does have some blurred areas of its own, which can be seen in the right frontal lobe of the brain at the boundary between white and gray matter.

In Fig. 3.2, 3.3 and 3.4, the reconstructions of data from three different patients are shown. For each reconstruction method multiple different images are shown. These images are the reconstructed FLAIR image, an enlarged portion of this image showing more detail, a colored image showing the white matter and lesion segmentations in the original image, and the cutouts of both the white matter and lesion segmentations. Important to note in these figures is the enlarged portion in column B. In the images in this column, it can be seen that the T2 knee model gives reconstructions that appear sharper, with less aliased boundaries between the different tissue types. In figure 3.2 this can be seen at the boundary between the lesion and the white matter, for example, as well as at the boundary with the CSF on the right side of the image. In Fig. 3.3 this is clearly visible from the boundary of the lesion on the right, however in this case there is a blurred region in the top left of this image in row 2, at the white and gray matter boundary. This is most apparent when comparing




Figure 3.1: Image showing the aliased starting image (a) obtained using the root-sum-of-squares method on the multi-coil undersampled data, together with the reconstruction using CS (b), and the knee T2 model (c) the brain T1 model (d).




Figure 3.2: Figure showing the reconstruction of a single slice using CS (1), the knee model (2) and the brain model(3). Shown are the reconstructed image (A), a high resolution part of the image (B), the image with white matter (teal) and lesions (purple) colored (C), the white matter segmentation (D) and the lesion segmentation (E).

this region to the same area in the CS reconstruction.

3.2 Quantitative analysis

The numerical analysis of the reconstructions using the CS technique and the RIMs trained either on T1 brain data or T2 knee data has been done using the relative SNR and the CR (as defined in Sec. 2.4). Due to problems with the SPM segmentations, where white matter was not segmented correctly from the datasets, 3 datasets were excluded. This left a total of 15 datasets per reconstruction method for the analysis.

One-way repeated measures ANOVA were used to examine the effect of the reconstruction methods. Figure 3.5 shows the relative SNR between the three methods. The type of reconstruction method lead to statistically significant differences in the relative SNR ($F(2,28)=3.9667$, $p=.0304$). Further paired-

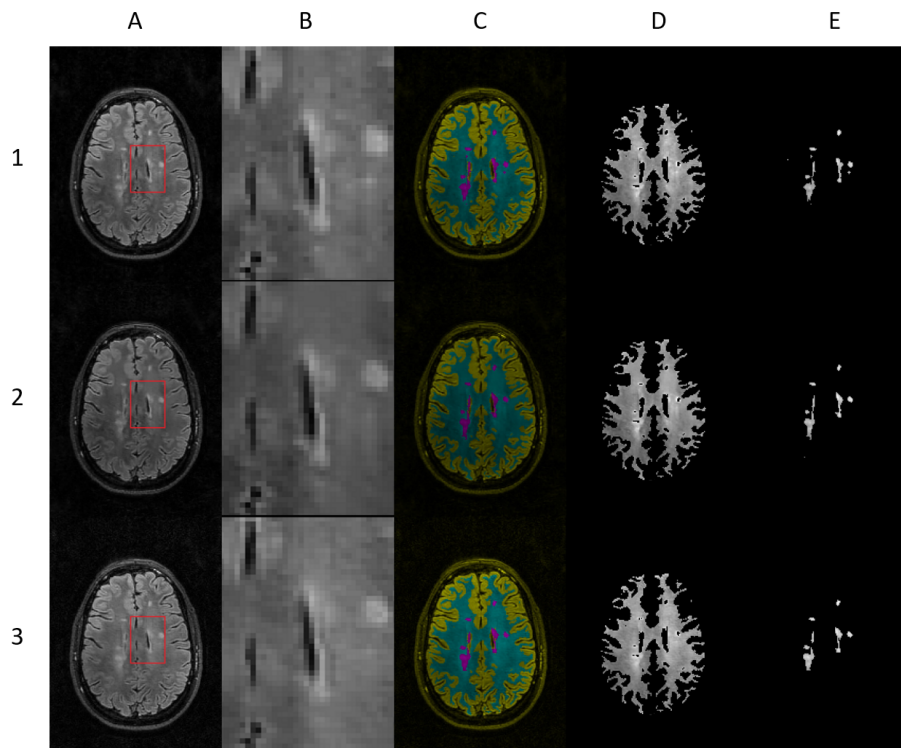


Figure 3.3: Figure showing the reconstruction of a single slice using CS (1), the knee model (2) and the brain model(3). Shown are the reconstructed image (A), a higher resolution part of the image (B), the image with white matter (teal) and lesions (purple) colorized (C), the white matter segmentation (D) and the lesion segmentation (E).




Figure 3.4: Figure showing the reconstruction of a single slice using CS (1), the knee model (2) and the brain model(3). Shown are the reconstructed image (A), a higher resolution part of the image (B), the image with white matter (teal) and lesions (purple) colorized (C), the white matter segmentation (D) and the lesion segmentation (E).




Figure 3.5: Box plot showing the difference in relative SNR between the three different reconstruction methods.

Relations	t	p
RIM Brain \rightarrow RIM Knee	-9.348	<.001
RIM Brain \rightarrow CS	-7.083	<.001
RIM Knee \rightarrow CS	5.535	<.001

Table 3.1: Outcome of the paired samples t-tests for the CR.

samples t-tests showed that there was a significant difference between the RIM trained on brains ($M = 1.323$, $SD = 0.241$) and the CS reconstruction ($M = 1.407$, $SD = 0.228$) ($t=-5.452$, $p<.001$). No significant difference was found between the RIM trained on brain data and the RIM trained on knee data ($M = 1.427$, $SD = 0.266$) ($t=-2.051$, $p=.0595$), as well as between the RIM trained on knee data and the CS reconstruction ($t=0.3178$, $p=.755$).

Figure 3.6 shows the CR between the lesions and the area around the lesion. A visual example of the area around the lesion can be seen in Fig. 2.6. The use of different reconstructions methods led to a statistically significant difference in the CR ($F(2,28)=61.1229$, $p<.001$). Further paired-samples t-tests were done to examine the statistical significance between each dataset pair. There was a statistically significant difference in the CR between all the comparisons of reconstructions (Table 3.1). This means that the T2 knee model ($M = 0.192$, $SD = 0.022$) had the highest CR, followed by CS ($M = 0.184$, $SD = 0.20$) and the T1 brain model ($M = 0.174$, $SD = 0.022$).

A different point that is important for the reconstruction in clinical use, is the inference time, or the reconstruction time of the data. For the CS technique, the average reconstruction time on the data was 533 ms per slice, while using the GPU for reconstructing. Meanwhile, the reconstruction times using the T2 knee and T1 brain models were 132 ms and 135 ms respectively. An overview

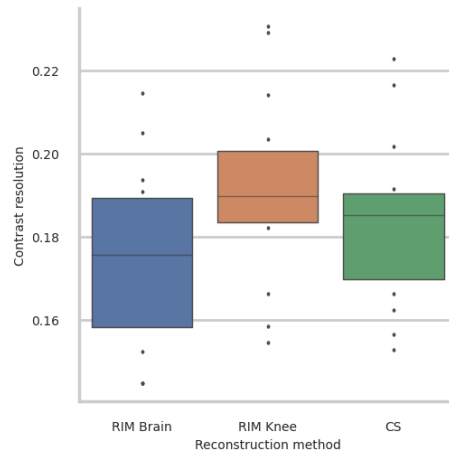


Figure 3.6: Box plot showing the difference in CR between the three different reconstruction methods.

Reconstruction method	Reconstruction time (ms per slice)
Compressed Sensing	533
T1 brain model	135
T2 knee model	131

Table 3.2: Table showing the reconstruction time per slice of the different reconstruction methods.

is also given in Table 3.2.

Chapter 4

Discussion

In chapter 3 it was shown that the relative SNR does not have a meaningful difference between most reconstruction methods, it is only between the T1 brain network and the CS reconstruction that there is a statistically significant difference. The T1 brain network has a lower relative SNR in the white matter in the FLAIR dataset than the CS reconstructions. In terms of CR, more techniques have a statistically significant difference between them, than is the case for the relative SNR. From the box plot in Fig. 3.6 it can be seen that the contrast between the average white matter signal and the average lesion signal is highest when reconstructed using the T2 knee model, and lowest when reconstructed using the T1 brain model. There are multiple factors that may aid in this difference in performance.

One factor is the contrast inherent in the data that the models were trained on. The T1 brain model was trained on T1-weighted datasets, that contain an inherently different contrast than T2-weighted datasets, such as the FLAIR dataset. The T2 weighted knee dataset was trained on T2-weighted datasets, and although it was trained on a different region of the body, might still contain valuable contrast information for the training of the network. However, in previous work the T1 brain model did perform, on average, better on cross-modality reconstruction than the T2 knee dataset, while performing slightly worse on the reconstruction of T2 knee data, when looking at the peak signal-to-noise ratio (PSNR) and the structural similarity index (SSIM) [6]. The better cross-modality performance of the T1 brain model in other studies makes it unlikely for the inherent contrast in the training data to be the only relevant factor.

A different factor that can also play a role in the contrast between white matter and lesion, is whether the model has seen similar boundaries between regions and similar intensity differences. In the T2 knee data, of which an example image can be seen in Fig. 2.4, there are more boundaries with a large contrast difference, such as between fatty tissue and muscle tissue, and between the joint fluid and the adjacent muscles. Compared to the brain model reconstructions, and in lesser extent also to the CS reconstructions, the knee model reconstructions also seems to be less aliased when inspecting the B columns of

figure 3.2, 3.3 and 3.4. This might be a result of the numerous adjacent regions, i.e., regions that have a lot of variations in voxel values, due to the structures of the muscles and more numerous tissue boundaries in the knee region, compared to the brain region.

With the improvement in CR from the T2 knee model and the comparable relative SNR overall, the RIM models were also faster in terms of reconstruction time (table 3.2), with a reconstruction time that was roughly 4 times faster. This can take the reconstruction time of a regular dataset with 256 slices, from 2 minutes and 16 seconds to 33 seconds. This also means that less computational power is needed at hospitals for the reconstruction of data, as more data can be reconstructed in a smaller time-frame.

4.1 Future work

In order to know whether the earlier made assumptions are valid or not, more research needs to be done. One suggestion would be to train a brain model containing T2 FLAIR data, and compare this model with the T1 brain model. Other experiments that could be done are to train models containing varying amounts of patients with lesion data, and see if there is an effect on the reconstruction quality of the lesions, and whether there needs to be a certain amount of lesion data in a brain model for reconstruction to be better.

For the comparison of the networks it would be useful if a fully sampled set of lesion rich data was available. This would make it possible to compare the reconstruction using the conventional PSNR and SSIM metrics.

There are other branches of future research that the data could be taken into. We have shown that the networks have shown comparable performance to the clinically accepted CS technique in terms of the relative SNR, and even giving a higher CR in the case of the knee model. One thing that could be looked in to is to use this data to train networks, where the acceleration through the under-sampling is increased by removing sampling points. This would give more insight into which technique, CS or the RIM, is more robust to high accelerations.

For future work, one thing that will need to be added in the preprocessing timeline is a way to deal with the coil problems. These coil problems were a main drawback in this study, as it made a large portion of the datasets unusable. In general, the scanner examines the SNR of the different coils, and decides to not save the data from certain coils if the SNR is too low. This process should help with the quality of the images. However, as the sensitivity maps are generated from a second smaller scan, the removal of coils with a low SNR can create a mismatch between the under-sampled data and the reference scan, when the data of different coils is thrown out, as can also be seen in Fig. 2.2. This was the main reason for most of the subject data dropping out in this study. Thus, it is necessary to implement a pipeline for recognizing these mismatched coils, and removing the coils that correspond to the thrown-out coils in the other dataset.

In the analysis of the reconstructions, three datasets were excluded. These exclusions were all due to problems with the segmentation of the white matter.

For two of the datasets, the reconstructions failed completely for all three reconstruction methods. In these reconstructions, the white matter segmentations contained most of the skull and gray matter as well, and also had one side of the head missing. In the third excluded dataset, the dark area inside the ventricles was also segmented as white matter, which then came up as an outlier in the CR estimation, as the contrast between CSF and the lesions is larger than the contrast between white matter and the lesions.

Chapter 5

Conclusion

In this work we have shown that the RIM trained on T2 knee data outperforms the CS technique in terms of CR, and the technique gives comparable results when looking at the relative SNR. The T1 brain model performed worse than the CS technique on both the relative SNR and the CR. The RIM managed to reconstruct the data 4 times faster than the CS technique. The data suggest that the RIM trained on T2 knee data is robust to the inclusion of lesions in an area that the network was not trained on.

Bibliography

- [1] R. W. Brown, Y. C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design: Second Edition*, volume 9780471720. John Wiley & Sons Ltd, Chichester, UK, 4 2014. ISBN 9781118633953. doi: 10.1002/9781118633953.
- [2] F. E. de Leeuw, J. C. de Groot, E. Achten, M. Oudkerk, L. M. Ramos, R. Heijboer, A. Hofman, J. Jolles, J. van Gijn, and M. M. Breteler. Prevalence of cerebral white matter lesions in elderly people: a population based magnetic resonance imaging study. The Rotterdam Scan Study. *Journal of neurology, neurosurgery, and psychiatry*, 70(1):9–14, 1 2001. ISSN 0022-3050 (Print). doi: 10.1136/jnnp.70.1.9.
- [3] K. Epperson, R. Rt, A. M. Sawyer, R. Rt, M. Lustig, M. Alley, M. Uecker, P. Virtue, P. Lai, and S. Vasawala. Creation of Fully Sampled MR Data Repository for Compressed Sensing of the Knee. *SMRT Conference*, 2013: 1, 2013.
- [4] K. J. Friston, J. T. Ashburner, S. Kiebel, T. E. Nichols, and W. D. Penny. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*, 2007.
- [5] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine*, 79(6): 3055–3071, 6 2018. ISSN 07403194. doi: 10.1002/mrm.26977. URL <http://doi.wiley.com/10.1002/mrm.26977>.
- [6] D. Karkalousos, K. Lønning, H. E. Hulst, S. O. Dumoulin, J.-J. Sonke, F. M. Vos, and M. W. A. Caan. Reconstructing unseen modalities and pathology with an efficient Recurrent Inference Machine. pages 1–26, 12 2020. URL <http://arxiv.org/abs/2012.07819>.
- [7] F. Knoll, T. Murrell, A. Sriram, N. Yakubova, J. Zbontar, M. Rabbat, A. Defazio, M. J. Muckley, D. K. Sodickson, C. L. Zitnick, and M. P. Recht. Advancing machine learning for MR image reconstruction with an

- open competition: Overview of the 2019 fastMRI challenge. pages 1–18, 2020. URL <http://arxiv.org/abs/2001.02518>.
- [8] H. Lassmann. Multiple sclerosis pathology. *Cold Spring Harbor Perspectives in Medicine*, 8(3):a028936, 3 2018. ISSN 21571422. doi: 10.1101/cshperspect.a028936. URL <http://perspectivesinmedicine.cshlp.org/lookup/doi/10.1101/cshperspect.a028936>.
- [9] S. Li, W. Li, C. Cook, C. Zhu, and Y. Gao. Independently Recurrent Neural Network (IndRNN): Building A Longer and Deeper RNN. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 5457–5466, 3 2018. ISBN 9781538664209. doi: 10.1109/CVPR.2018.00572. URL <https://arxiv.org/abs/1803.04831>.
- [10] K. Lønning, P. Putzky, J. J. Sonke, L. Reneman, M. W. Caan, and M. Welling. Recurrent inference machines for reconstructing heterogeneous MRI data. *Medical Image Analysis*, 53:64–78, 2019. ISSN 13618423. doi: 10.1016/j.media.2019.01.005.
- [11] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007. ISSN 07403194. doi: 10.1002/mrm.21391.
- [12] S. Manikandan. Measures of central tendency: Median and mode. *Journal of pharmacology & pharmacotherapeutics*, 2(3):214–215, 7 2011. ISSN 0976-5018 (Electronic). doi: 10.4103/0976-500X.83300.
- [13] M. Mardani, E. Gong, J. Y. Cheng, S. S. Vasanawala, G. Zaharchuk, L. Xing, and J. M. Pauly. Deep generative adversarial neural networks for compressive sensing MRI. *IEEE Transactions on Medical Imaging*, 38(1):167–179, 1 2019. ISSN 1558254X. doi: 10.1109/TMI.2018.2858752.
- [14] C. A. McKenzie, E. N. Yeh, M. A. Ohliger, M. D. Price, and D. K. Sodickson. Self-calibrating parallel imaging with automatic coil sensitivity extraction. *Magnetic Resonance in Medicine*, 47(3):529–538, 2002. ISSN 07403194. doi: 10.1002/mrm.10087.
- [15] M. H. Moghari, M. Uecker, S. Roujol, M. Sabbagh, T. Geva, and A. J. Powell. Accelerated whole-heart MR angiography using a variable-density poisson-disc undersampling pattern and compressed sensing reconstruction. *Magnetic Resonance in Medicine*, 79(2):761–769, 2018. doi: <https://doi.org/10.1002/mrm.26730>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26730>.
- [16] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In

- H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [17] P. Putzky and M. Welling. Recurrent Inference Machines for Solving Inverse Problems. (Nips), 6 2017. URL <http://arxiv.org/abs/1706.04008>.
- [18] P. Putzky, D. Karkalousos, J. Teuwen, N. Miriakov, B. Bakker, M. Caan, and M. Welling. i-RIM applied to the fastMRI challenge, 2019.
- [19] A. Raj, Y. Bresler, and B. Li. Improving Robustness of Deep-Learning-Based Image Reconstruction. *arXiv*, 2020. ISSN 23318422.
- [20] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9351:234–241, 2015. ISSN 16113349. doi: 10.1007/978-3-319-24574-4{_}28.
- [21] J. Schlemper, C. Qin, J. Duan, R. M. Summers, and K. Hammernik. $\Sigma\Sigma$ -net: Ensembled Iterative Deep Neural Networks for Accelerated Parallel MR Image Reconstruction, 2019.
- [22] M. D. Steenwijk, M. Daams, F. Barkhof, P. J. W. Pouwels, and J. J. G. Geurts. Multi-view convolutional neural networks using batch normalization outperform human raters during automatic white matter lesion segmentation, 2017. URL <https://onlinelibrary.ectrims-congress.eu/ectrims/2017/ACTRIMS-ECTRIMS2017/200729/>.
- [23] J. Teuwen, N. Moriakov, D. Karkalousos, and M. Caan. DIRECT, 2020. URL <https://github.com/directgroup/direct>.
- [24] M. Uecker, S. Rosenzweig, H. C. M. Holme, M. Blumenthal, Z. Tan, X. Wang, J. I. Tamir, and M. Lustig. Berkeley Advanced Reconstruction Toolbox (BART), 2020. URL <https://zenodo.org/record/3934312>.
- [25] D. Weishaupt, V. D. Köchli, and B. Marincek. *How does MRI work?* Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. ISBN 978-3-540-44094-9. doi: 10.1007/978-3-662-07805-1. URL <http://link.springer.com/10.1007/978-3-662-07805-1>.
- [26] G. Yang, S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, Y. Guo, and D. Firmin. DAGAN: Deep De-Aliasing Generative Adversarial Networks for Fast Compressed Sensing MRI Reconstruction. *IEEE Transactions on Medical Imaging*, 37(6):1310–1321, 6 2018. ISSN 1558254X. doi: 10.1109/TMI.2017.2785879.

Appendix A

Theory

A.1 Introduction

The major problem with MRI acquisition is that it is a lengthy process, taking multiple minutes to acquire a single dataset, with the time demand increasing when the spatial resolution of the acquired images is increased. Several techniques, such as parallel imaging [15], SENSE [12], and GRAPPA [5], exist with the capability of acquiring high quality images in low acceleration factors from multiple receiver coils. The state-of-the-art parallel-imaging technique is Compressed Sensing (CS) [10], where data are compressed through a sparsifying transform for achieving high-resolution estimations. The problem with CS is that it can be computationally expensive. In order to further reduce the reconstruction times, several techniques have been proposed facilitated by power of deep learning on learning strong prior features and having fast inference times. Such state-of-the-art techniques are the U-net [14] and the Recurrent Inference Machine [13].

A.2 Background on MRI

MRI is a technique that uses a magnetic field to magnetize protons in the tissue that is to be imaged. This magnetization is done by positioning a subject in a magnetic field. With the magnetization, the protons will get a net magnetic moment in the direction of the main magnetic field (B_0), and will start to precess around the direction of the field. This precession happens at the Larmor frequency, which is dependent on the magnetic field strength (chapter 2 in [1]). Thus, by placing a subject in the magnetic field, the protons will try to rotate around the magnetic field. This gives the protons an oscillating magnetization.

Then by giving a radio frequency (RF) pulse in a direction perpendicular to the magnetic field, it is possible to flip the direction of the magnetization of the protons by 90 degrees (chapter 3 in [1]). After a 90 degree RF pulse, the protons have been flipped to a direction perpendicular to the magnetic field




Figure A.1: Visualization of the Free Induction Decay shown as the induced voltage against time. The FID has a slope of the extremes that diminishes over time.

of the scanner, and will over time relax to their initial magnetization. Going back to the initial magnetization happens in two different ways. The directional component in the direction of the magnetic field starts growing (T1), since the protons get a net magnetization from the magnetic field. The protons that are excited in this manner each have slightly different Larmor frequencies, as the environment the protons are in is not homogenous. Due to the slightly different Larmor frequencies, the flipped magnetization in the direction perpendicular to the magnetic field will start to diminish, due to the magnetic field of the individual protons going out of phase with each other, giving the transverse decay (T2) (chapter 4 in [1]). The precession around the magnetic field, creates a small signal, due to the changing magnetic field of this precession. After the flipping of the magnetic moment of the protons, the precession will create a signal, due to the larger directional component of the induced magnetic field perpendicular to the B0 direction. When the flipped magnetic moment then relaxes back, the free induction decay (FID) can be seen (Fig. A.1). The FID signal can be measured and processed in different ways to get meaningful data, so that it can be used in further analyses.

The aim of measuring these signals is to reconstruct an image of the subject. However, the FID signal of a whole volume is not useful, but different methods have to be used to get more meaningful results. In order to obtain image data, the signals need to be encoded before being acquired by the scanner. The magnetization is encoded by giving the magnetic field a gradient in one direction, with a slightly lower magnetic field strength at one end (chapter 10 in [1]). Using this gradient it is possible to select only one slice, or a smaller region, for the flipping of the magnetic field. It is required to use gradients in different directions to create a 3D acquisition, where specific lines in k-space are imaged and stored as a volume. This volume is acquired in the frequency domain, known

as k-space, containing the individual frequency and phase components necessary for obtaining the image. In order to go from the k-space in the Fourier domain, to the image in the image domain, the inverse Fourier transform is used. With the use of multiple gradients in different directions, it is then possible to scan a whole volume over time. However, these gradients are limited in their strength and in how quickly they can be at their maximum strength, due to limitations in hardware and the adverse effects they can have on the body at high field strengths.

With MRI, it is possible to get different contrasts, depending on the setup of the acquisition (chapter 15 in [1]). Examples of these different MRI image types are T1, T2 and Fluid Attenuation Inversion Recovery (FLAIR). For the difference between T1 and T2, a different relaxation mechanism of protons is used. These properties result in inherently different contrasts, such as a brighter signal from fat in T1 images. On the other hand, the difference between T2 and FLAIR is that FLAIR is a T2 type sequence, where an RF pulse is used to first flip the signal by 180 degrees. After this initial flip there is a delay until the parallel magnetic component of fluids is non-existent. At this point, another RF pulse flips the magnetic moment by 90 degrees. Due to the inversion pulse, there will be almost no signal from fluids, giving better contrast between the fluids and the surrounding tissues.

The composition of the k-space, and thus the scanning parameters, is a direct result of the desired image quality, such as the resolution and the amount of noise, and the volume it should encapsulate (chapter 12 in [1]). For example, making the size of the k-space larger will result in a more detailed image with higher resolution, while having less distance between k-space points will result in a larger field-of-view (FOV). For medical diagnosis, it is always the desire to have data that is as detailed as possible. The time it takes to acquire the data, however, limits the possibility to obtain high resolution images.

A.2.1 The Inverse Problem

In order to understand the reconstruction of corrupted MRI data, it is important to understand the inverse problem. Let us denote x as the true image signal, and y the corrupted signal that is obtained from the scanner, we can define the model

$$y = a(x) + n, \tag{A.1}$$

where n is normally distributed noise in the image, and a is the function that takes the data from x and corrupts the data. This process of going from the true image signal to the corrupted signal is known as the forward model, as all the parameters for going from x to y are available. Solving the equation for x , on the other hand, requires knowledge on the parameters of a , which are not always available. In the case of not knowing the parameters, they have to be approximated, in order to find x . This is known as the inverse problem, where, given a set of observations, a model needs to be constructed that explain the parameters that caused these observations. One way of solving the inverse

problem is through iterations, as is done in the case of Compressed Sensing for under-sampled MRI data.

A.3 Compressed Sensing

A.3.1 Theory

When data are sparse in a certain domain, it is possible to compress them by keeping only the most important parts they contain. This sparsity in a certain domain means that in a representation of the data, the bulk of the information is comprised in a small chunk of the data. One example of such sparsity is having a vector in which there are groups of likewise values, allowing the vector to be reduced to the value, and how many times it occurs. With this example, it is beneficial to treat values close enough to each other as the same value. By treating data in that manner, it is possible to reduce the amount of data needed to store.

For MRI, the domain that is generally sparse, and thus compressible, is the wavelet domain. The wavelet domain contains the contrast of the image at different levels. In Fig. A.2 it is shown that in the wavelet domain, most of the contrast is stored at the lowest level, while the higher levels contain only little information. The sparsity in the wavelet domain allows the data to be under-sampled in k-space, while still being reconstructable. For under-sampling k-space data it is important that the under-sampling pattern is incoherent or random. If equally spaced coherent under-sampling is used, the Nyquist criterion is not adhered to, which results in the highest measurable frequency decreasing. The consequences of this are a large reduction in quality and introduction of aliasing.

Aliasing is still a problem with random under-sampling. However, in this case aliasing behaves in a noise-like manner. Pseudo-randomness can be achieved in several ways, such as skipping random lines in 2D slices and in 3D image space, and creating radial and spiral patterns (Fig. A.3). One method to quantify the incoherence of an under-sampling pattern is through the point spread function (PSF) (bottom row in Fig. A.3). From these PSF visualizations, it follows that out of four under-sampling techniques, the trajectory with the most contrast between the peak and the underlying noise is the "Random lines in 3D" method, making it the most incoherent. The limitation is that it is not random over multiple slices, as each slice will have the same under-sampling pattern. However, the sampling patterns that are possible are constrained by hardware limitations. For example, it is not possible to have too large RF gradient changes in the scanner, as these are constrained to the gradient speed and by unwanted tissue effects with too strong gradients, making it difficult to sample random points of k-space in 3D acquisition, instead of lines. Another constraint is that the distance between sequential lines needs to be minimal in order to prevent detrimental effects due to eddy currents — electrical currents due to changing magnetic fields.




Figure A.2: Example of a brain MR image (a) and the representation of the same image in the wavelet domain (b). The wavelet transformed image (b) shows the contrast gradients at different scales. From this it is also visible that the contrast within the brain image is mostly contained at a smaller scale (top left image in b). Image taken from [17].

In k-space the center is most important for defining the contrast level of the image, while the outer regions are more responsible for the level of the spatial details. Thus, it is important to create under-sampling patterns that are denser in the center of k-space than from the outer regions. If an appropriate sub-sampling pattern has been chosen, with enough incoherence and more density in the center, the image associated with the sub-sampled k-space will have diffuse aliasing artifacts. These artifacts will then manifest as noise, from which the image can be reconstructed.

A.3.2 Compressed Sensing for MRI Reconstruction

Compressed Sensing reconstruction is an optimization problem, in which a linear operator ψ is used to transform a fully-sampled image m into a sparse representation. This can be denoted as

$$\begin{aligned} & \text{minimize } \|\psi m\|_1, \\ & \text{such that } \|\mathcal{F}_u m - y\|_2 < \epsilon. \end{aligned} \tag{A.2}$$

Here, \mathcal{F}_u is the Fourier transform combined with the under-sampling pattern corresponding to the minimized linear operator, y are the k-space data from the scanner and ϵ is the threshold parameter, with a value lower than the expected noise level [10]. The second condition needs to be true, since then the frequency spectrum of the approximation and the measured data will be as similar as




Figure A.3: Example trajectories with the resulting k-space pattern and the point spread function underneath. a) random lines in 2D, b) random lines in 3D, c) radial, and d) uniform spirals. Image taken from [11].

possible, enforcing data consistency. The function to be minimized uses the ℓ_1 -norm, which is defined as

$$\|x\|_1 = \sum_i |x_i|, \quad (\text{A.3})$$

where x represents the ψ_m part of the equation. Minimization of the ℓ_1 -norm has been shown to promote sparsity [2]. An example of a sparse representation for reconstructing the data through ψ is the wavelet transform.

The general way to solve this problem is iterating over the reconstructions until the optimization criteria are satisfied. This iterative approach is based on removing the interferences in small steps until the threshold is reached. In the case of accelerated-MRI, the interferences are the added noise from the under-sampling. General MRI acquisition schemes lead to images where the main contrast is contained in the tissue boundaries, and signal being mostly homogenous between voxels of the same tissue. Because of these properties of the contrast, the volume will have a sparse representation in the wavelet domain. When MRI data are under-sampled incoherently, the contrast will still be confined to these boundaries, except for added noise. By iteratively thresholding the wavelet domain, and testing the result against the threshold, it is possible to reconstruct the data, removing the interference. For this, a balance is needed between transform sparsity and data consistency, which is reflected in the second part of eq. A.2. Both the cases of minimizing the transform sparsity too much and weighting high the data consistency, will result in adverse effects, such as blurring. Therefore, by going through iterations, sparsity is minimized, while keeping adequate data consistency.

A.4 Machine Learning

A problem with the iterative procedure of CS is that it can be time-consuming, requiring a lot of processing power. In order to mitigate this, several deep learning techniques have been proposed, with the aim of achieving fast and accurate reconstructions.

A.4.1 Supervised & Unsupervised Learning

Machine learning algorithms can be categorized regarding their learning procedure as supervised or unsupervised. Supervised algorithms learn to infer and make connections based on fully-sampled, and often labeled, training data. Facilitating the information learned from the training data, this knowledge can be applied for inference on data out of the training distribution. For reconstructing accelerated-MR images, the algorithm learns how to make connections between fully-sampled and under-sampled images, for applying this knowledge to unseen under-sampled images.

In contrary, unsupervised algorithms learn how to infer from unclassified and unlabeled data, requiring more data to learn the best approximation than supervised algorithms. This is because unsupervised learning in principle requires the algorithm to predict the right connections without a reference. In terms of accelerated-MRI reconstruction, this can be seen as an algorithm that has fully-sampled data as input in the case of supervised learning, and under-sampled data as input in the unsupervised learning case. Most networks are trained in a supervised manner, with unsupervised networks only being used when a large quantity of data is available. Especially in medical imaging, the amount of data is often limited. Because of the limited amount of data, almost all networks that are used for reconstructing under-sampled MRI data are based on supervised networks.

A.4.2 Convolutional Neural Networks

Initially, the techniques developed for the purpose of reconstructing under-sampled MRI data relied on Convolutional Neural Networks (CNN). These are networks that train on subsets of image data that are created with the use of convolution matrices on image data. Through convolutions, to the size of the selected kernels' size, and multiple hidden layers, feature maps are extracted. Those feature maps help the network to identify useful information on the image, such as edges and adjacent regions, for achieving high level of detail.

Examples of such techniques using CNNs for MRI reconstruction are the U-net [14], consisting of a series of convolutional layers, and the General Adversarial Networks (GAN) [16], where two networks are trained at the same time aiming to achieve high-resolution reconstructions by discriminating the noise from the images.

A.4.3 Physics-based models

A more sophisticated approach for reconstructing accelerated-MR images is to formulate physical properties of the problem at hand and include them into the learning procedure of the network. Therefore, the formulation of the forward model is necessary for finding the inverse transformation of mapping sparsely-sampled k-space measurements to the target image. The so-called physics-based models make use of having a target image and knowing the forward transformation to sub-sample it. We try to estimate the closest approximation with the principles of probability theory and statistics, and the Bayes Theorem for updating the network.

The Bayes' Theorem is defined by

$$P(A | B) = \frac{P(B | A)P(A)}{P(B)}, \quad (\text{A.4})$$

where A is a given hypothesis and B is some observed evidence and states that the posterior probability of having the hypothesis A be true given that evidence B is true. This is related to the conditional probability of B given A , as well as the probabilities of A and B themselves. In general, the probability of B being true is not relevant for the models. It is, therefore, possible to simplify the equation to

$$P(A | B) = P(B | A)P(A). \quad (\text{A.5})$$

Equation A.5 does not give a probability, but can lead to an estimation of the distribution through minimization of this equation. In order to maximize the performance of models for finding a solution for $P(A|B)$, it is possible to maximize the conditional probability $P(B|A)P(A)$, where usually B is the model and A the data. This maximization is called the Maximum a Posteriori (MAP) estimation.

Some examples of network that can be thought as Bayesian estimators are the Recurrent Inference Machine (RIM) [13] and the Variational Network [6]. Of these two, in this work we focus on the RIM.

A.5 Recurrent Inference Machine

The Recurrent Inference Machine was originally proposed as a general inverse problem solver [13], and later extended to reconstructing accelerated-MRI data [9]. The RIM is a supervised network that is optimized over t time-steps.

In order to find the best approximation that solves the inverse problem, the RIM makes use of the forward model. The forward model describes how data can be transformed from fully-sampled to sparsely sampled. The forward model for multi-coil under-sampling can be defined using eq. A.1 as

$$\mathbf{y}_i = \mathbf{P}\mathcal{F}\mathbf{S}_i^H \mathbf{x} + \mathbf{n}_i, \quad i = 1, \dots, c. \quad (\text{A.6})$$

Where x is the true image, y are the under-sampled data, c is the total number of coils, \mathbf{P} is the sampling mask, \mathcal{F} is the Fourier transform, \mathbf{S}_i is the sensitivity

map of coil i , with \mathbf{H} denoting the Hermitian adjoint operator, and \mathbf{n}_i represents normally distributed noise.

The quality of the approximation needs to be defined by the prior information and the fit of the approximation to the data. This is done using the MAP estimation (i.e. maximization of equation A.5), and can be defined in logarithmic space as

$$\mathbf{x}_{MAP} = \underset{x}{\operatorname{argmax}} \{ \log p(\mathbf{y}|\mathbf{x}) + \log p(\mathbf{x}) \}. \quad (\text{A.7})$$

The second part of eq. A.7 can be seen as a regularizer that depends on the model. Meanwhile, the first part can be rewritten to the log-likelihood and, given independent data that follows a Gaussian distribution, can be formulated as

$$\log p(\mathbf{y}|\mathbf{x}) = \frac{1}{\sigma^2} \sum_{i=1}^c \|\mathbf{P} \mathcal{F} \mathbf{S}_i^H \mathbf{x} - \mathbf{y}_i\|_2^2, \quad (\text{A.8})$$

where σ^2 is the variance of the data.

The RIM uses the log-likelihood gradient (llg) for updating itself at time-step τ , and perform gradient descent. The llg ($\nabla_{\mathbf{y}|\mathbf{x}_\tau}$) is defined as

$$\nabla_{\mathbf{y}|\mathbf{x}_\tau} := \frac{1}{\sigma^2} \sum_{i=1}^C \mathbf{S}_i^H \mathcal{F}^{-1} \mathbf{P}^T (\mathbf{P} \mathcal{F} \mathbf{S}_i \mathbf{x}_\tau - \mathbf{y}_i), \quad (\text{A.9})$$

where \mathbf{x}_τ is the current estimation. The iterative scheme is shaped by the following functions,

$$\begin{aligned} \mathbf{s}_0 &= 0 & \mathbf{x}_0 &= \mathcal{F}^{-1} \mathbf{y} \\ \mathbf{s}_{\tau+1} &= g_\phi(\nabla_{\mathbf{y}|\mathbf{x}_\tau}, \mathbf{x}_\tau, \mathbf{s}_\tau) & \mathbf{x}_{\tau+1} &= \mathbf{x}_\tau + h_\phi(\nabla_{\mathbf{y}|\mathbf{x}_\tau}, \mathbf{x}_\tau, \mathbf{s}_{\tau+1}), \end{aligned} \quad (\text{A.10})$$

where \mathbf{s} is the internal state at time-step τ . h is the update function, containing the core part of the network (as shown in Fig. A.4), and g is the part of the network responsible for updating the internal states. The internal states keep track of iterations, guiding the network’s gradient to the right direction for achieving convergence. One hyperparameter that needs to be fixed in advance is the number of time steps that need to be unrolled during training. The same number of time steps is applied when using a trained network.

\mathbf{s} can be thought as the result of updating the network two times, \mathbf{s}_τ^1 and \mathbf{s}_τ^2 , though the recurrent layers (Fig. A.4). As can be seen in this figure, \mathbf{s} is dependent on the type of the recurrent layer, denoted as g in eq. A.10.

The RIM does the training with the help of the recurrent layer (Fig. A.4). These recurrent layers allow the network to have a notion of memory over the time-steps taken so far. This gives the network the ability to use this memory to steer the network in more favorable directions. The complexity of the network is dependent on the type of the chosen recurrent unit. With the effect on the complexity, the inference time of the network will also change. However, having a loss in complexity does not necessarily mean that the performance will also become less [7].

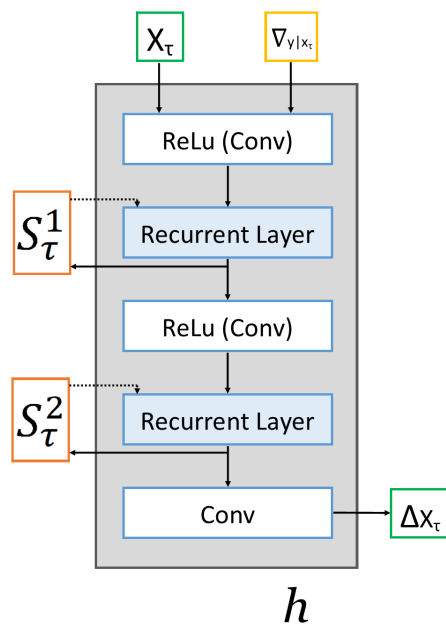


Figure A.4: Schematic view of the Recurrent Inference Machine. Image adapted from [7].

For the recurrent layers a few options have been explored, such as the Gated Recurrent Unit (GRU) and the Independently Recurrent Neural Network (IndRNN) [7].

The GRU [3] is a technique that is used in order to combat the problem that regular RNN structures have, which is the tendency for gradients to either explode or diminish over time. The gated unit solves this problem by having the ability to forget part of the previously stored memory, notably the parts of the gradient that are less useful to the performance of the network. This keeps the gradients from exploding and diminishing, and is done by introducing hyperbolic tangents and sigmoid functions, which are used in the RNN structure.

Another recurrent unit that was created to solve this problem of RNNs is the IndRNN, which also tried to solve different problems with RNN structures like the GRU [8]. These structures can have a tendency to have gradient decay over the different neuron layers because of the hyperbolic tangent and sigmoid functions that they contain. Instead of this approach, the IndRNN uses independent neurons that are connected only across layers, due to which the network is capable of processing longer sequences, tested up to 5000 time-steps. This has the effect of defeating the problem of regular RNNs, where gradients either explode or vanish, while also avoiding the gradient decay over different layers.

The convolutional layers are activated by rectifiers. The rectifier is an activation function used to saturate the output within a range. One of the most commonly used is the rectified linear unit (ReLU) [4], taking the input of the neuron and returning zero if the neuron has a value below zero, or the value of the neuron itself in the other cases.

Finally, the network aims to minimize the loss function for calculating the error between the estimation and the target image. In this work, the chosen loss function is the ℓ_1 -norm, which uses the absolute difference (Eq. A.3).

A.6 The future of the RIM

The RIM is a deep learning network implementing the physics of accelerated-MRI for reconstructing the under-sampled MRI data. It has been shown that the RIM is capable of doing these reconstructions with a higher performance, while being faster, than the state-of-the-art CS technique [7]. At the same time, more research needs to be done that explores the generalizability of the RIM network to different modalities, sampling patterns, and pathologies. The CS technique is capable of handling these variations in data, as it is generalized by design. Thus, it is necessary to also explore the generalizability of the deep learning techniques that are proposed as alternatives of CS, as it would require fewer trained networks to be used in the general workflow. Having a network that is generalized to different pathologies is also better for the treatment of patients, as it makes it less likely to miss unexpected details in the data.

Bibliography

- [1] R. W. Brown, Y. C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design: Second Edition*, volume 9780471720. John Wiley & Sons Ltd, Chichester, UK, 4 2014. ISBN 9781118633953. doi: 10.1002/9781118633953.
- [2] E. J. Candes and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 2006. ISSN 00189448. doi: 10.1109/TIT.2006.885507.
- [3] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pages 1724–1734, 2014. doi: 10.3115/v1/d14-1179.
- [4] K. Fukushima. Visual Feature Extraction by a Multilayered Network of Analog Threshold Elements. *IEEE Transactions on Systems Science and Cybernetics*, 5(4):322–333, 1969. doi: 10.1109/TSSC.1969.300225.
- [5] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6): 1202–1210, 2002. ISSN 07403194. doi: 10.1002/mrm.10171.
- [6] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine*, 79(6): 3055–3071, 6 2018. ISSN 07403194. doi: 10.1002/mrm.26977. URL <http://doi.wiley.com/10.1002/mrm.26977>.
- [7] D. Karkaloulos, K. Lønning, H. E. Hulst, S. O. Dumoulin, J.-J. Sonke, F. M. Vos, and M. W. A. Caan. Reconstructing unseen modalities and pathology with an efficient Recurrent Inference Machine. pages 1–26, 12 2020. URL <http://arxiv.org/abs/2012.07819>.

- [8] S. Li, W. Li, C. Cook, C. Zhu, and Y. Gao. Independently Recurrent Neural Network (IndRNN): Building A Longer and Deeper RNN. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 5457–5466, 3 2018. ISBN 9781538664209. doi: 10.1109/CVPR.2018.00572. URL <https://arxiv.org/abs/1803.04831>.
- [9] K. Lønning, P. Putzky, J. J. Sonke, L. Reneman, M. W. Caan, and M. Welling. Recurrent inference machines for reconstructing heterogeneous MRI data. *Medical Image Analysis*, 53:64–78, 2019. ISSN 13618423. doi: 10.1016/j.media.2019.01.005.
- [10] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007. ISSN 07403194. doi: 10.1002/mrm.21391.
- [11] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed sensing MRI: A look at how CS can improve on current imaging techniques. *IEEE Signal Processing Magazine*, 25(2):72–82, 3 2008. ISSN 10535888. doi: 10.1109/MSP.2007.914728.
- [12] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Magnetic resonance in medicine*, 42(5): 952–62, 11 1999. ISSN 0740-3194. URL <http://www.ncbi.nlm.nih.gov/pubmed/10542355>.
- [13] P. Putzky and M. Welling. Recurrent Inference Machines for Solving Inverse Problems. (Nips), 6 2017. URL <http://arxiv.org/abs/1706.04008>.
- [14] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9351:234–241, 2015. ISSN 16113349. doi: 10.1007/978-3-319-24574-4_{_}28.
- [15] D. K. Sodickson and W. J. Manning. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 38(4):591–603, 10 1997. ISSN 07403194. doi: 10.1002/mrm.1910380414.
- [16] G. Yang, S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, Y. Guo, and D. Firmin. DAGAN: Deep De-Aliasing Generative Adversarial Networks for Fast Compressed Sensing MRI Reconstruction. *IEEE Transactions on Medical Imaging*, 37(6):1310–1321, 6 2018. ISSN 1558254X. doi: 10.1109/TMI.2017.2785879.
- [17] Y. Zhang, S. Wang, Y. Huo, L. Wu, and A. Liu. Feature extraction of brain MRI by stationary wavelet transform and its applications. *Journal of Biological Systems*, 18(SPEC. ISSUE 1):115–132, 10 2010. ISSN 02183390. doi: 10.1142/S0218339010003652.