MR reconstruction of FLAIR weighted images with simulated lesions. A comparison between Compressed Sensing and a Recurrent Inference Machine.

B van Hoek[a]

[a] *University of Technology Delft*

Abstract

Background: For both hospitals and patients it would be beneficial if the scan time of MR images could be reduced. At the moment, Compressed Sensing (CS) is introduced to reduce the scan time, however, new methods are developed such as a deep learning method, called the Recurrent Inference Machine (RIM). In this study the effect of reconstructing undersampled MRI images with lesions, using the RIM and CS, was evaluated. In data of a healthy control, lesions were simulated. Evaluation is done by checking if the lesion has the correct intensity and shape after reconstruction of undersampled data.
Methods: In raw data of a healthy control lesions were simulated. To test the RIM and CS, the images with lesions where first undersampled 4x, 6x, 8x and 10x. After undersampling, the images were reconstructed with both RIM and CS. First, the peak intensity difference was measured between the reference image (with simulated lesions) and reconstructed images for both RIM and CS. Second, one lesion was undersampled ten times with different undersampling masks creating different noise, for 3 different acceleration factors (4x, 6x, 8x). These lesions were reconstructed with both RIM and CS. The maximum intensity difference between reference and reconstructed image was measured and averaged over the ten different undersampled images.
Results: In total seven different lesions were simulated in a healthy control with different intensities varying between 10% and 100% of the GM-lesion intensity in a FLAIR weighted scan. The intensities of all lesions were more accurately reconstructed with the RIM compared to CS at higher acceleration factors: the average intensity per lesions after 10 times reconstruction with RIM was more equal to the correct intensity compared to the reconstruction with CS.
Conclusion: The RIM shows robust and accurate results on data with simulated lesions. Moreover, the RIM outperformed CS on data that was more undersampled. Therefore, the RIM may be used for reconstruction of MRI data that is acquired with shorter acquisition time. And since the reconstruction time is better, it could replace CS in the future. However, before the RIM could be used, further evaluations on actual patient data are needed.

Keywords: MRI, Acceleration, Compressed Sensing, Deep Learning,, Lesion Simulation, Reconstruction, Recurrent Inference Machine, Under sampling

1. Introduction

Physicians often use magnetic resonance imaging (MRI) to image *in vivo* tissue of the human body for diagnostic, disease progress or research. MRI can distinguish differences in soft tissues, and is therefore very often used in the brain, but is also used for imaging breast, hart and blood vessels. The advantage of MRI is that it is a non invasive method to scan *in vivo* compared to Computed Tomography (CT). A disadvantage of the MRI compared to CT it that acquisition times are longer for MRI. For diagnostics, it takes around 30 minutes compared to 3 minutes for CT. During the acquisition of the image subjects should lie still in the scanner. Healthy controls sometimes have already problems with this, and it is even more difficult for patients with Parkinson's disease or small children. Moreover, planning these long scans results in higher cost for hospitals.

For cost effective purposes, as well as improvement of patient experience it is important to make the scans as short as possible without having to compromise much data quality. Research into this field has been done for many years now. Also manufacturers have been optimizing the hardware and software to get the best scan in the shortest amount of time. While the higher field strength MRI scanners were developed to improve the resolution of the images, it does not contribute to shortening in the acquisition time. Therefore, research in the field of acquisition methods and data reconstruction methods is done.

Algorithms as Compressed Sensing (CS), Sensitivity encoding (SENSE) and generalized auto-calibrating partial parallel acquisition (GRAPPA) are already used to reconstruct usable MR images from under-sampled data. Relatively new is the use of Deep Neural Networks to the field of image reconstruction. It is already proven that neural networks can reconstruct MR data. The Recurrent Inference Machine (RIM) is a recurrent neural network used for MRI reconstruction. This methods showed promising results for reconstruction. However, it is not know yet if the method could be used on clinical data that would be used for diagnosis of neurological diseases. For example, the most important pathology of multiple sclerosis (MS) are inflammatory lesions in the white matter (WM). These MS lesions are spots on the MR images that can by hyper or hypo-intense compared with the surrounding WM. For diagnostics and research the amount of lesions in the brain is important and therefore when using the RIM all the lesions should be reconstructed correctly. Therefore, we aimed in this study to compare the reconstructions of the RIM with reconstruction of compressed sensing and the reference images to concluded if the lesions are correct reconstructed by RIM.

Email address: van.hoek.bob@gmail.com (B van Hoek)
[1]Spinoza Centre, Amsterdam

2. Methods

2.1. MRI Acquisition and Reconstruction

An MRI signal is generated by the net magnetic field of the protons within the human body. A strong static magnetic field (B_0 field) polarizes net magnetic moment in the same direction and a magnetic moment parallel to the static field is resulting. Due to radio frequency excitation, called the B_1 field, a rotating magnetization $m(\vec{r})$ component is generated transversely to the B_0 field. Consequently a radio frequency signal is emitted by the rotating magnetization which is detectable by an receiver coil. This signal is dependent on many different physical attributes of the tissue [1]. For imaging the signal, depicting the spatial distribution of $m(\vec{r})$, is done in MRI. To encode the spatial distribution extra magnetic fields are superimposed on the B_0 field to create gradients. These gradients can be applied in the three Cartesian directions (G_x, G_y and G_z). This causes the rotation frequency to change linearly along the axis and the signal is distinguishable in a spatial grid and the receiver coil detects the encodes signal by the linear phase. The MRI signal equation has the form of a Fourier integral [1],

$$s(t) = \int_R m(\vec{r}) e^{-i2\pi \vec{k}(t)\cdot\vec{r}} dr, \tag{1}$$

where $k(t) \propto \int_0^t G(s)ds$. The received signal at time t is the Fourier transform of the object $m(\vec{r})$ sampled at the spatial frequency k(t)[2]. These multidimensional signals are recorded in a spatial frequency space, or k-space. A good visual example has been given by Lustig et al.[2], which is shown in figure 1.

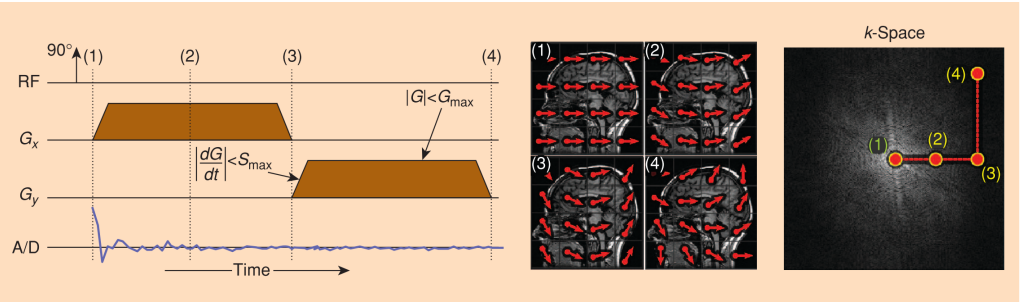

Figure 1: (Image from [2]).The MR signal directly samples the spatial frequency domain of the image. The gradients apply an linear frequency distribution, which cause frequency differences in time at different spatial positions.

After acquiring data in k-space, an inverse Fourier transform is applied to the k-space to retrieve the image. Important for a successful image reconstruction is that the sampling in k-space abides the Nyquist-Shannon sampling theorem[3]. The Nyquist-Shannon sampling theorem states that the function must be sampled at intervals of maximum length $1/2 * \text{Bandwidth}(B)$. Violating this theorem will result in aliasing artifacts. Applied to MRI, B is determined by the distance from the origin to the edge of the k-space and since B is a spatial frequency, the resolution will be limited to $1/2B$. For a given bandwidth, it is the k-space sampling rate that determines the size of the

region covered inside the scanner. Thus, if the object being scanned is larger than the field of view, the periodic properties of the Fourier transform cause the object's signal to wrap around and appear as aliasing artifacts in the image[4][5].

Encoded sampling and the freedom of choosing the path to fill up the k-space makes MRI suited for accelerated reconstruction methods. By skipping values in the k-space and zero fill them aliasing artifacts are introduced. Periodic undersampling introduces coherent aliasing artifacts, while random sampling introduces incoherent aliasing artifacts which, on the eye, look like a noise like artifact, as can be seen in figure 2.

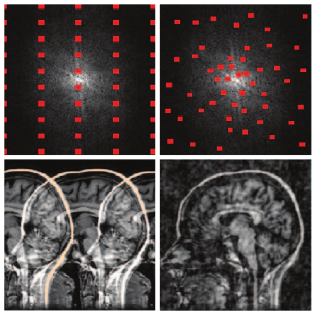

Figure 2: (Image from [2]).Aliasing artifacts. The difference between coherent (left) incoherent (right) artifacts.

In this Master thesis we used fully sampled data, which we will undersample with self made masks to impose sampling artifacts. As we will see later, incoherent noise is the favorable artifact we seek to solve with different reconstruction techniques. Therefore, we used undersampling masks with a random Gaussian like distribution. Since multiple accelerations are used for testing and training, different accelerating masks were made with different acceleration factors (figure:3). The undersampling masks are multiplied with the fully sampled k-space.

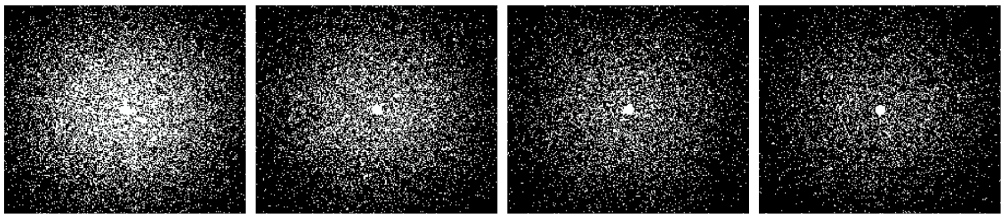

Figure 3: Undersampling masks. FLTR: 4x, 6x, 8x and 10x accelerated. White are the sampled k-lines.

2.2. Compressed Sensing

Undersampling MRI produces artifacts as seen in figure 2. A method to solve the incoherent artifacts is Compressed Sensing. CS exploits the inherent compressibility of the signal it seeks to reconstruct. Therefore, to use CS, the signal it is reconstructing needs to have an invertible sparsifying transform Ψ, such that $u = \Psi x$.

Sparsity in an image means that there are relatively few pixels with non-zero values. Images can have a sparse representation in some domain (Wavelet, finite differences). The

same holds for MRI images, for example angiograms are extremely sparse in the pixel domain. More complex images do not have a sparse representation in the pixel domain and can be transformed into, for example, the wavelet domain. In the wavelet domain an image is described in wavelets. The wavelet transform describes both frequency and location information in the image. Figure 4 is an example of a wavelet transform. The amount of zero values has increased significantly and the signal can be called sparse.

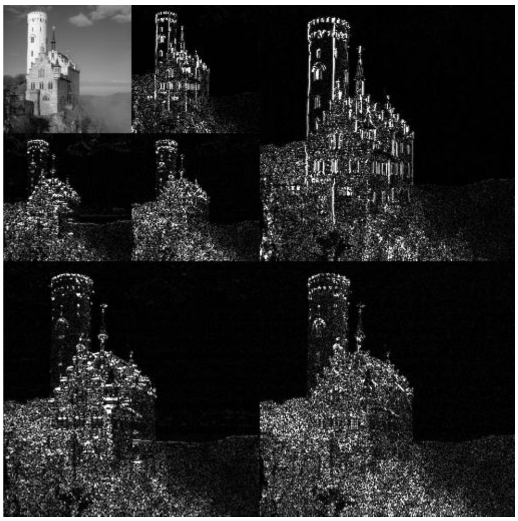

Figure 4: Example of a 2-D discrete wavelet transform.

The possibility of exploiting transform sparsity is proved by data compression in imaging. Natural images typically have inherent potential to be compressed without any visual loss of information, concluding that some information in images is abundant and probably did not have to be acquired in the first place. This inspired the theory of acquiring a compressed MRI signal while still being able to reconstruct it properly. According to the theory of compressive sampling[6][7], if an image has an sparse representation in some domain and the artifacts are incoherent, the signal can be recovered from the undersampled k-space data provided a nonlinear recovery scheme is used.

In other words, CS in MRI has three requirements: 1) The image must have a sparse representation in some domain, 2) The aliasing artifacts are incoherent and 3) a nonlinear reconstruction scheme has to be used. The scheme must enforce sparsity and data consistency of the image representation. Point 1 and point 2 were both described. Now we will look further into the nonlinear reconstruction scheme. Let the image be a vector m, let Ψ be the linear operator that transforms the vector into a sparse representation, and let F_u be the undersampled Fourier transform. The reconstruction can be done by solving the following optimization problem:

$$\underset{x}{\operatorname{argmin}} \left\{ ||F_u - y||_2^2 + \lambda \, ||\Psi m||_1 \right\} \tag{2}$$

where Ψ is the sparsifying transform to transform the image into a sparse domain, y is the actual measured k-space data and F_u is the Fourier transform for the data that was

employed. The first term enforces the data consistency via the least-squares difference. It ensures that the reconstructed image m is consistent with the data that is actually measured in k-space. The second term enforces sparsity in the appropriate transform domain. The L1-norm is used, which calculates the degree of sparsity in the image by summing the absolute value of the pixels. So this basically means that the equation finds the image m which is both sparse and also consistent with the data that has been measured in k-space. The regularizer, λ, controls the relation between the two terms. A small λ will lead to an image that is close to the measured data, relying less on the sparsity, while using a high λ will lead to an image that is relying more on the sparsity and deviate more from the collected data.[8]. Different approaches have been described in the literature for solving equation 2[9].

For this master thesis we used a CS toolbox, the "Berkeley Advanced Reconstruction Toolbox (BART)", to reconstruct the undersampled data. To use the bart toolbox, the data need to be correctly shaped, which means that the dimensions are assigned like this:

- 0 readout
- 1 phase-encoding dim 1
- 2 phase-encoding dim 2
- 3 coil dim
- 4 ESPIRiT maps

Since we don't have the sensitivity maps directly from the measurements, the bart toolbox can calculate them for us by using the "ecalib" function. This function creates sensitivity maps based on the ESPIRiT method [10] from raw k-space data. With this data the parallel-imaging compressed-sensing ('pics') function from the toolbox is used to reconstruct the data. Many variables can be given to this function. Here I will name the most important ones we considered for this study:

- l1/l2 toggle l1-wavelet or l2 regularization.
- lambda(r) regularization parameter
- R generalized regularization option

The data for input was the raw k-space (dim 0-3) and the above mentioned sensitivity maps (dim 4).

2.3. Recurrent Inference Machine

2.3.1. Forward model

For explaining the RIM, we first have to explain the forward model. Let $x \in \mathbb{C}^m$ be the true image signal and let $y_i \in \mathbb{C}^m$, be the set of sparsely sampled frequency signals measured in k-space by one of the scanner's receiver coils. The measured signal can then be described by the true image,

$$y_i = PFS_i x + n_i, \qquad i = 1, ..., c. \tag{3}$$

where P is the subsampling mask, which reduces the amount of signals measured and therefore accelerates the scanning time, F is the Fourier transform representation of the original data x, and S_i are the coil sensitivity maps. Measurements are assumed to be subjected to additive, normally distributed noise n_i. We assume that the noise can be modeled as independent and identically distributed across coils, pixels and complex components [4]. In figure 5, the forward model is illustrated.

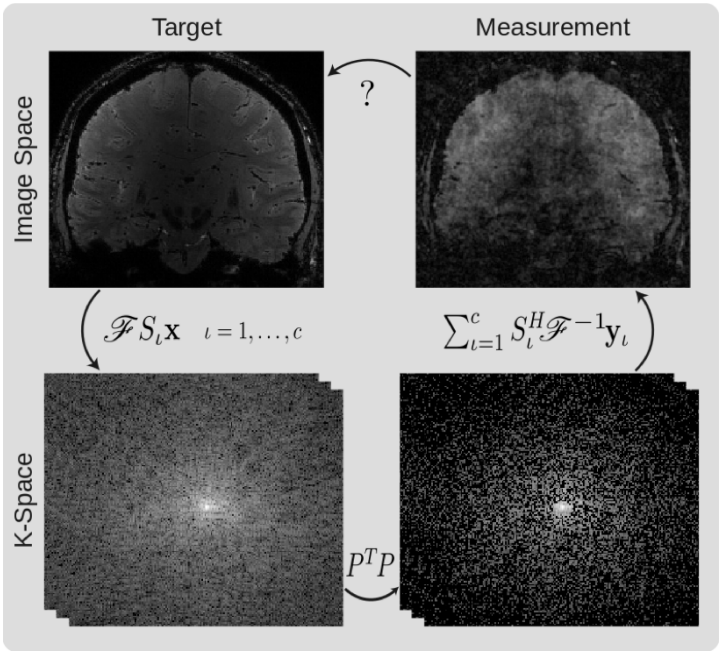

Figure 5: Image from [4]. Top left: Fully sampled MRI. Bottom left: k-space measurements which is related to image space through a coil sensitivity weighted map, followed by a Fourier transform. Bottom right: Accelerated measurements with subsample mask P. Top right: Reconstructed sparsely sampled k-space, where incoherent aliasing artifacts are visible. Reconstruction methods to find the function to go from the aliased MRI to the fully sampled high resolution MRI is the goal.

The forward model is described by going from the target to the undersampled image in figure 5. SENSE [11] reconstruction is used to reconstruct the Parallel imaging acceleration with coil sensitivity maps, which leads to the undersampled k-space, which contains incoherent aliasing artifacts when the inverse Fourier transform is taken.

Solving a function to find the inverse transform of the forward model to solve x in

$$y_i = PFS_i x + n_i, \qquad i = 1, ..., c. \tag{4}$$

is called the inverse problem. A common strategy is optimizing the maximum a posteriori (MAP) estimator given by,

$$x_{MAP} = \underset{x}{\operatorname{argmax}} \log p(y|x) + \log p(x), \tag{5}$$

which is the maximization of the sum of the log-likelihood and log-prior distributions y

and x. This can be reformulated as optimizing the regularized problem,

$$\underset{x}{\operatorname{argmin}} \left\{ \sum_{i=1}^{c} d(y_i, PFS_i x) + \lambda R(x) \right\} \tag{6}$$

where d evaluates the data consistency between the reconstruction and measurements, R is a regularizer, with regularization factor λ. Under the assumption of independent, identically and normally distributed measurement errors as seen in 3, the log-likelihood in 5, the data consistency term in 6 is given by,

$$\log p(y|x) = \frac{1}{\sigma^2} \sum_{i=1}^{c} \|PFS_i x - y_i\|_2^2 \tag{7}$$

Together with the regularization term, this equation looks the same as 2 for CS if $\|\Psi\|_1$ is the regularizer R.

2.3.2. Model design

The RIM aims to optimize equation 6 by learning an iterative scheme over t recurrent time-steps. Each time-step receives information on the current state of the reconstruction process as an input, which produces an incremental step Δx_τ to take in image space as output. The problem of evaluating the gradient of the log-prior distribution in equation 5 is solved by passing the current estimate, or external state, x_τ as an input to the network.

In figure 6 the structure of the update function and the model is shown. The model (on the left) consists of τ timesteps. The updatefunction for each timestep has as input the estimate x_τ, the internal states s_τ and the gradient of the log-likelihood. The RIM update functions (h) are responsible for producing the next internal state, which the RIM needs in order to keep track of iterations (timesteps) and modify its behaviour based on the progression of the reference procedure [4]. The update function consists of convolutional layers and gated recurrent unit cells (GRU) [12]. An GRU has the ability to be trained to keep information from many timesteps ago, without washing it through time or remove information which is irrelevant to the reconstruction. More detailled information of the networks structure can be found in [13, 4].

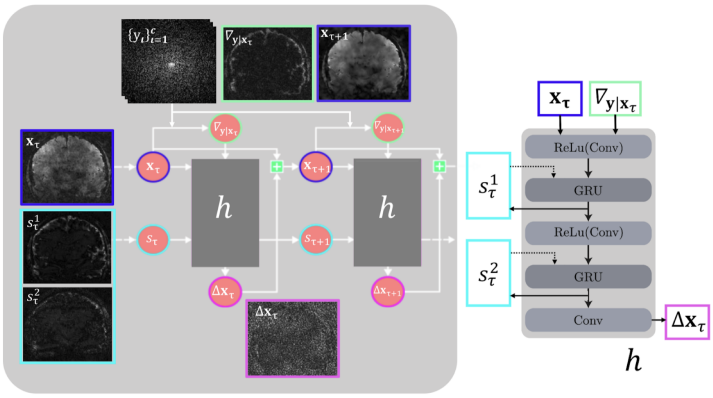

Figure 6: Image from [4] The Recurrent Inference Machine (RIM) model on the left. A detailed view on the update function h on the right. Magnitudes of intermediate internal states s1 and s2 at timestep t were averaged over features. Bold lines depict connections within a single time-step, whereas dotted lines represent recurrent connections that pass information to the next time-step.

The RIM uses the MSE as a loss function, where the estimate x_τ is evaluated against the true image x for each time-step. A weighted sum of the MSE over the time-steps t is minimized to train a good model.

2.3.3. Training and testing data

The model is trained on T_1-weighted three-dimensional magnetization prepared rapid gradient echo (MPRAGE) weighted MRI images. The images were made on a 3.0T Philips Ingenia scanner equipped with a 32-channel head coil with a resolution of 1 mm^3 and a FOV 256 x 240 mm^2, 225 slices with sagittal slice encoding direction. TFE factor 150, interval 2500 ms, inversion delay 900ms, flip angle 30° and first order shimming. The data was fully sampled with an elliptical shutter. The total scanning time was 10.8 min. In the article of Lnning et al. (2019) this same model was tested. The results showed it outperformed CS an various other DL reconstruction methods. Moreover, the article also showed that training and reconstructing on differently weighted MRI scans should not cause a problem. Since this model was tested thoroughly while trained on the above mentioned data, we adopted this model for the rest of our project and used in on T_2-weighted scans.

2.4. Data

In this project we used a T2-weighted fluid attenuated inversion recovery (FLAIR) scan of the brain, made on a 3T Philips scanner with a 32 channel head coil. The resolution was 1.1 mm^3. The scan was scanned with an acceleration of 7x. The raw data received from this scan was reconstructed with CS. This reconstructed data is now used as the target for the rest of the study. Obtaining a fully sampled scan was impossible since it would take very long and image quality would be affected by movement artifacts. The scan was made for a study into a specific medication for treating Multiple Sclerosis, called Tecfidera and approved for by the medical ethical committee. Only one participant was included, based on which multiple data containing lesions were simulated.

2.5. Lesion Simulation

To test the reconstruction methods, complex valued raw data from the MRI scanners is needed. Unfortunately, past studies within the research group do not have raw data both with and without lesions. We solved this by simulating lesions in healthy controls. Lesions were introduced by augmenting the magnitude signal of a reference standard into the raw data of the healthy control. In our somewhat simplified model for lesion simulation, we assumed that the lesions cause a higher intensity in the magnitude image and the phase would stay unchanged.

2.5.1. Intensity and implementation

When simulating a lesion, the position of the lesion and the intensity of the lesion are the two factors to be determined. Positioning of the lesions was determined by studying at other lesion maps annotated by neurologists and copying them by drawing our own lesion maps on the image of our own healthy volunteer. Most of the lesions were drawn in the deep white matter, distant from cortical gray matter. In figure 7 we see all the lesions drawn in our healthy volunteer. As for the intensity, the lesions from the Prediva study were observed and the average increase in intensity relative to the WM was calculated. We found that the a maximum intensity increase of two times occurred in lesion compared to the WM. Therefore, it was decided to scale the intensity of the lesions from 1.1 to 2.0 times the WM intensity.

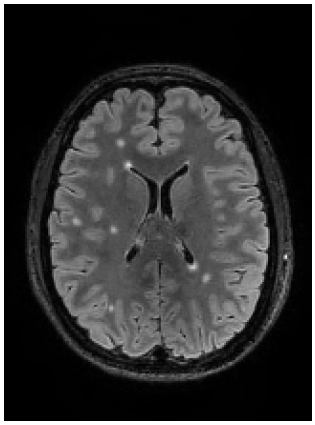

Figure 7: One slice of a healthy volunteer with all the simulated lesions at different locations (lesions at 70% of gray matter intensity).

The intensity and the position of the lesions are embedded in a mask that had the size of the actual image. To change only the magnitude of the image and keep the phase of the image unchanged, the mask was multiplied only with the magnitude of the image, see figure 8.

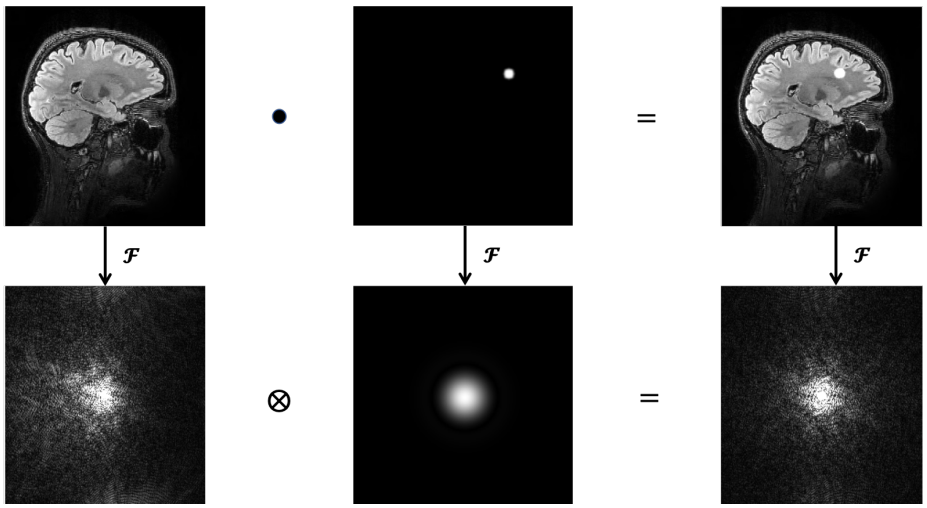

Figure 8: Simulation process of lesions in an MRI. Top row are the images in image space. Bottom row are the Fourierspace representations of the top row. Multiplying in image space is a multiplication in Fourierspace.

After insertion of the lesion, the k-space is expected to change according to the rule that a multiplication in image space, is a convolution in Fourierspace (and vise versa). which implies, that the information in k-space containing the lesion position and intensity is spread out in k-space. After undersampling, every point measured still should contain information about the lesion. The lesions were inserted in the fully sampled image and in image space. For undersampling and reconstruction, the image was transformed back into the k-space.

2.6. Design of experiment

Testing the reconstructions with simulated lesions is done with two algorithms:

- Compressed Sensing
- Recurrent Inference Machine trained on T1 weighted MRI images without lesions

The methods are tested on the accuracy and robustness in two experiments. In the first experiment 7 different lesions were simulated at different locations in one slice of a healthy control. These lesions are simulated in a range of intensities. The maximum intensity is the intensity of the gray matter and the minimum intensity is 10% of the maximum. Each image with the different lesions and intensities is accelerated with 4 different factors: 2x, 4x, 6x and 8x. All these images are reconstructed with both CS and the RIM. To compare the accuracy of reconstructing the correct contrast and shape we looked at the mean pixel intensities of a horizontal and vertical line through the lesion, we will cal this the Mean Intensity Line (MIL). These lines cross each other at the maximum intensity pixel in the lesion. From this line we will get information about intensity loss and shape differences in the lesion. In Figure 9 is an example of such a measurement.

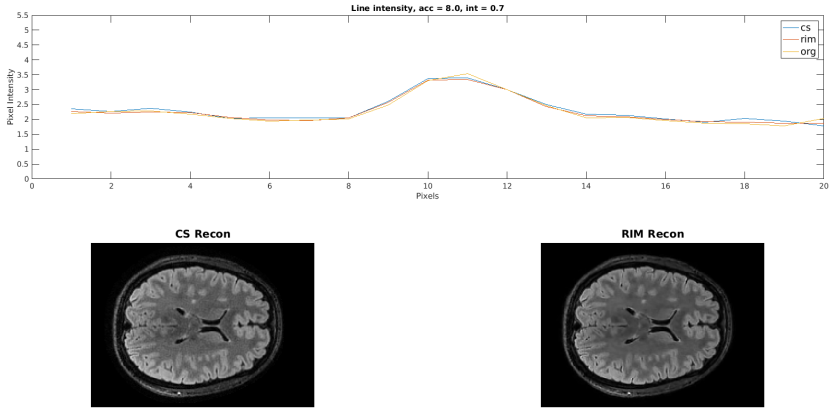

Figure 9: A plot of the mean intensity lines going through the center of the lesion. The original image with lesions, CS-constructed and RIM-constructed is plotted. Under the plot the CS reconstruction and RIM reconstruct is imaged. The red lines show the horizontal and vertical line over which the average is taken.

In the second experiment only one of the simulated lesions is used in a healthy control MRI slice. This slice is undersampled 4x, 6x, and 8x. And like the first experiment different intensities for lesion simulation are used with a maximum of 100% of the gray matter intensity. The same slice is undersampled ten times for each acceleration factor and each intensity. Again reconstruction is done with the CS and RIM for all images. For analysing the differences in contrast of the maximum value if the MIL between the reconstruction method and the reference will be used to rate the robustness of the reconstruction methods.

3. Results

In Figure 10 a visual overview is shown of the simulated and reconstructed lesions. The images are 8 times accelerated and the lesions have an intensity of 70%. In total seven lesions on different locations were simulated in healthy controls. Visual inspection shows that reconstruction with the RIM gives a smoother result in and around the lesions compared to CS, for all seven lesions compared to the original simulated lesions. There is however different smoothness visible for the different locations. When looking at the subcortical lesions (1, 2, 6 and 7), the details of the surrounding white matter seem to be more "smooth" compared to the cortical lesions (3, 4, 5). Whereas in the CS it seems to be more noisy at every lesion position. Whether this impacts the reconstruction of the lesion is not visible. As seen in Figure 10 it is not visible that the intensity and shape of lesions are different after reconstruction with both RIM and CS.

Also the plots of the intensity lines of these seven simulated lesion show that the intensity is not much different after reconstruction, see Figure 11.

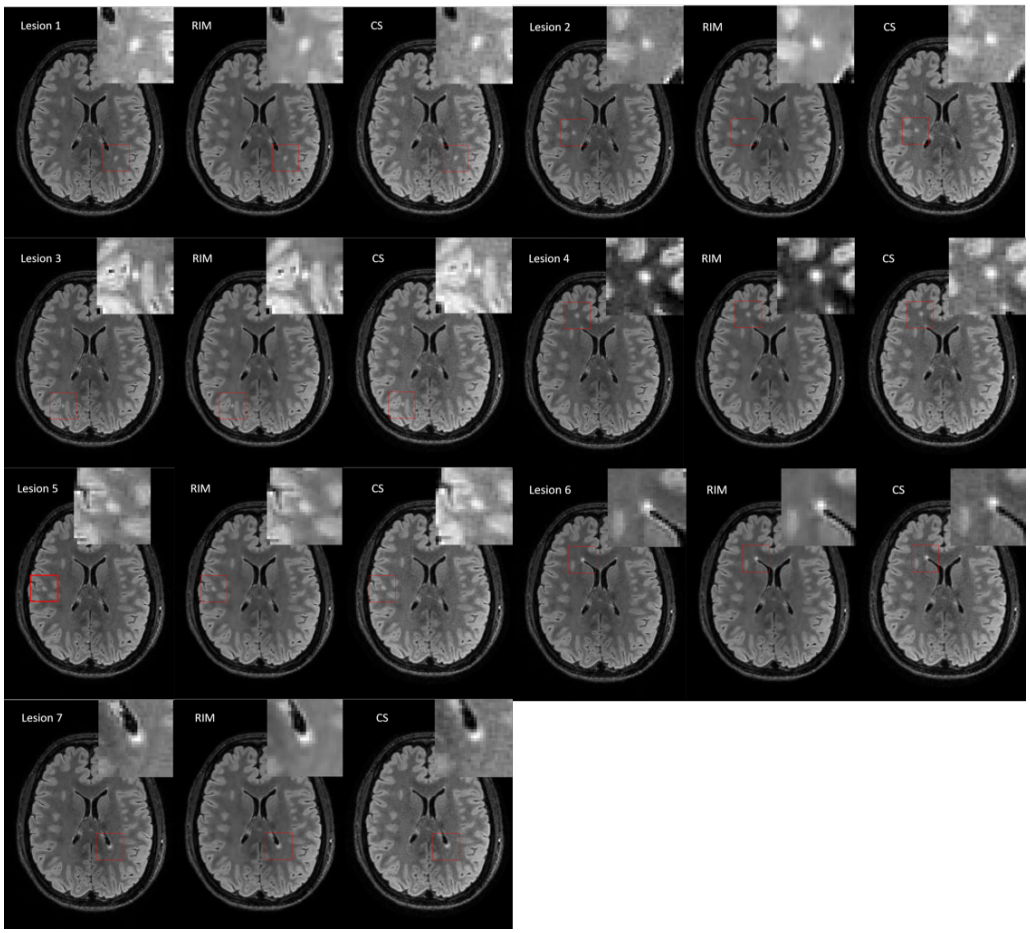

Figure 10: All seven lesions, reverence, RIM reconstructed and CS reconstructed. Lesions are zoomed in the top corner for a detailed view.

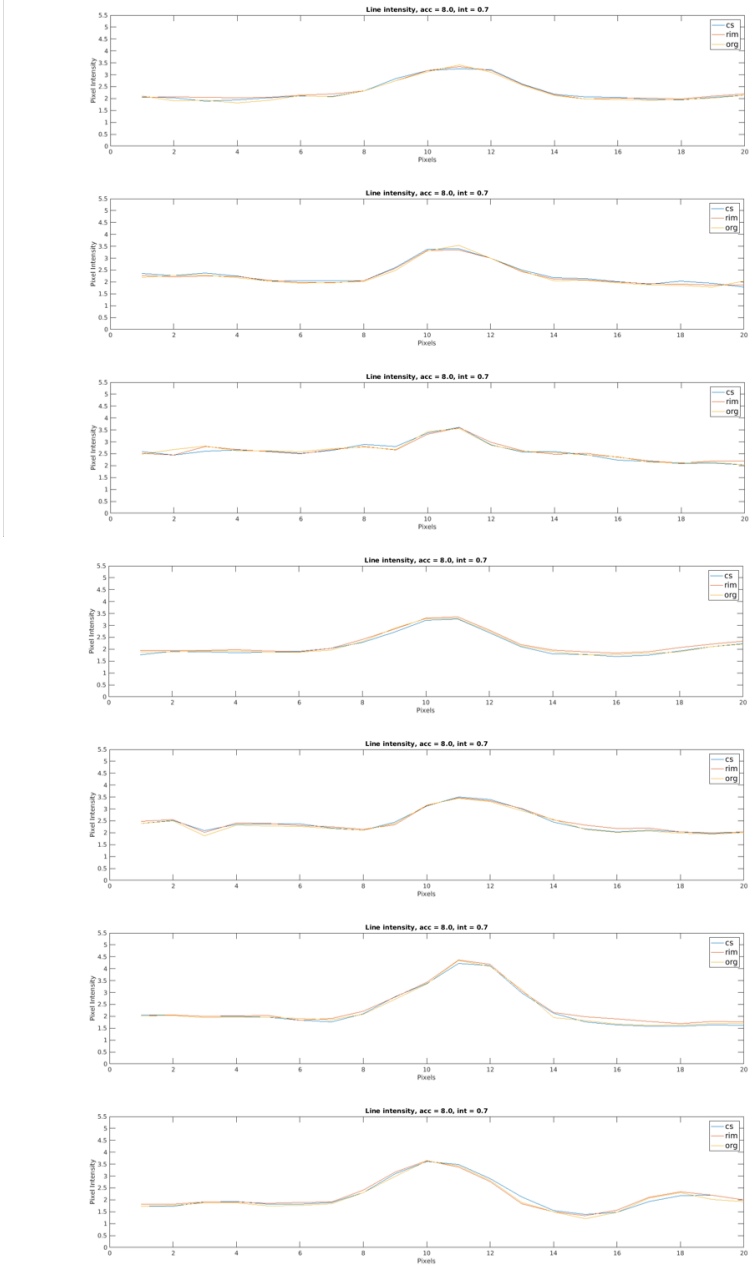

Figure 11: Seven plots of the Mean Intensity Lines going through the center of each of the seven lesions. Per plot are three lines, one of the reverence, second for the RIM and the third for the CS.

Figure 12 shows for both the RIM as CS the peak intensity difference over all lesions compared to the original simulated lesions. The different lesion intensities are measured at four different acceleration factors (2x, 4x, 6x and 8x), resulting in four different plots

per reconstruction methods and for every acceleration ten different intensities setting in a plot. The plots for the RIM are more widespread at higher acceleration factors and higher intensities compared to the CS. The plots of CS shows a more stable pattern, therefore are less affected by acceleration factor and intensity.

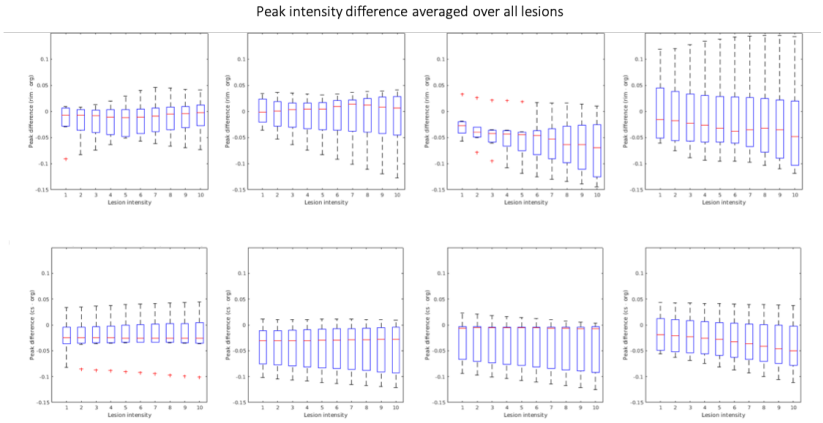

Figure 12: Boxplots of the peak intensity differences for RIM (upper four) and CS (lower four) for the four different accelerations factors. At the x-axis the lesion intensity of the reverence image is plotted (1 =10%, 10=100%). The y-axis plots the difference in peak intensity averaged over all lesions. These plots show the consistency of the reconstruction methods in terms of intensity.

3.1. Reconstruction of the MR images with one Simulated Lesion

The reconstruction of one lesion was analyzed more thoroughly. The result of the reconstruction with CS and RIM are depicted in figure 13. In this figure the performance gain of the RIM compared to the CS per acceleration factor is visible. In the three plots we see three different acceleration factors (4x, 6x and 8x). Each plot has the lesion intensity on the x-axis and the normalized contrast on the y-axis, containing three different measurements; the reference line, the CS reconstructed results and the RIM reconstructed results. The CS and RIM reconstruction are averaged over ten reconstructions with varying random masks and added noise. In the second part of the figure, under the graphs, the reconstructed lesions are shown for four different intensities, at an acceleration of 8x. The first column has the original lesion. Then all the CS results are shown followed by all the RIM results.

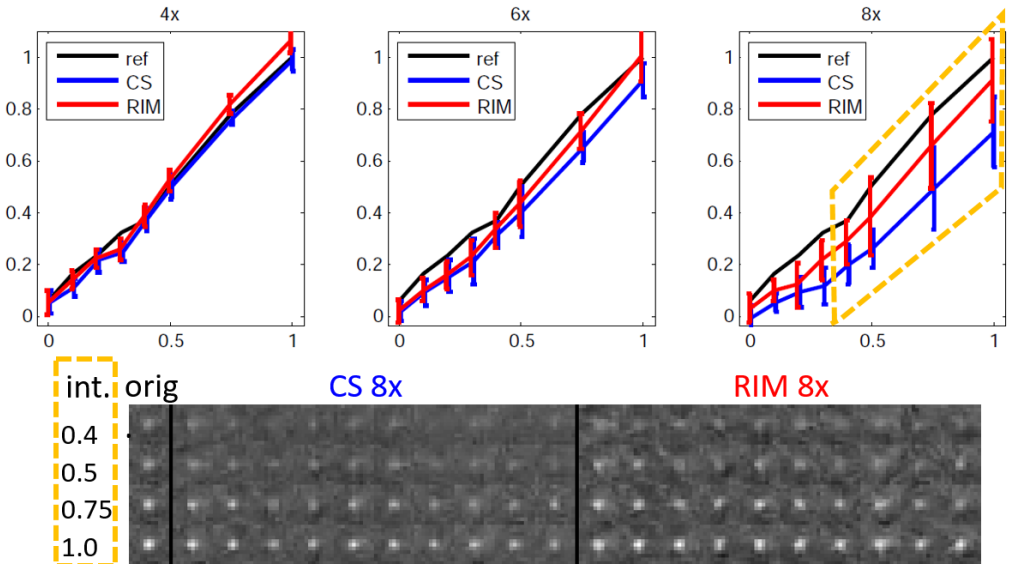

Figure 13: For three acceleration factors (4x, 6x and 8x), there were 10 reconstruction per reconstruction method. The average per lesion intensity is plotted.

4. Discussion

4.1. Overview

It would be beneficial for both the patient as the hospital if the scanning time of MR images could be reduced. Several options are addressed to shorten the acquisition time, for example undersampling the data followed by an reconstruction with CS or an other Deep learning method. However, not known yet is if the undersampling and reconstruction have influence on visible pathology in a MRI. Therefore, the aim of this thesis was to look at the gain in robustness and accuracy of reconstructing MRI data with lesions using the RIM compared to CS. To generate the data we simulated various lesions with various intensities scaled between 10% and 100% of the gray matter intensity. The data was undersampled to simulate accelerations 2x, 4x, 6x and 8x compared to fully sampling the MRI. Consequently, two different experiments were done. First the accuracy was measured and compared for both reconstruction techniques. Secondly, the robustness was measured. When looking at the results it shows us that RIM outperforms CS at higher acceleration factors, but has no significant gain at the lower acceleration factors. More specifically, the RIM appeared to be more robust in maintaining image quality at higher acceleration levels. On the other hand, the accuracy of the RIM seems to be a bit worse (Figure 12 at lower acceleration factors, although it reaches the same results as the CS at higher acceleration levels. In the next section the results will be discussed in more detail.

4.2. Performance of RIM versus CS

In the results it seems that the RIM, in general, does not underperform compared to CS. The process of the RIM to solve the reconstruction via the inference process is focused on removing aliasing artifacts in the image and not, as many other deep learning methods do, solving by imaging features. By doing so, also the noise in image is reduced as we have seen in Fig. 10. Unfortunately, noise-like details in (especially the subcortical regions) are removed during the reconstruction as well. Nevertheless, the lesions, which cover a larger spatial extent than the voxel-by-voxel noise pattern, are constructed good in terms of robustness and shape (Figure: 13) and almost no intensity is lost compared to the reference image. Comparing to the CS, we see that the same can be said for CS at lower accelerations. But starting from 6x accelerated we begin to see a drop in performance of CS compared to the RIM. The distortion of the image at higher acceleration levels seems to be high for reconstruction via CS. It might be that the sparsity term in Eq. 2 is treating too much of the aliasing artifact as noise and therefore will be removed, which leads to a lower contrast and consequently a lower intensity of the lesion.

When looking at Figure 12, we see that the RIM is not outperforming CS in terms of consistency, since the values are averaged over lesions at different positions. We suspect that the RIM might face performance differences when comparing lesions cortically or more subcortically at a brain MRI. More detailed research should be done to look into why this is the case. It could be the case that this is an artifact of the data being trained on different sort of MRI then it is tested on.

4.3. Extension of studies

This research can be extended into multiple directions. For one, the lesion simulation is a simplified model. This can be extended into a model that accounts for more variables,

like the shape of the lesion. Lesions could be segmented from other MRI scans to preserve the shape. We were not able to do this in this study because the data needs to be in RAW format, which basically means that the data still needs to be complex valued, which was not available at the time. Another extension that might be looked at for lesion simulation is how iron content in lesions might change the phase image of the MRI [14]. It is known that lesions can be detected based on iron content in the brain. Also a higher iron content at the position of the lesion might result in a difference in magnetisation and therefore result in a change in the phase image. This results in a more complex pathology then we studied in this thesis which involves only a magnitude augmentation. Such a study would reveal whether the RIM is also robust to phase modulations in addition to magnitude modulations.

An extension of this research might also be to train the RIM on data which does contain lesions. When having generated more data like described in the paragraph before, training the data with lesions becomes more easy. On one hand it might be very promising that the RIM obtains results like this while trained on a dataset which does not contains lesions, this might be a big advantage because you perhaps do not have to train the model for every pathological decease which you want to detect. On the other hand it might be very interesting to see what the performance gain or loss would be when training on data containing lesions.

Another extension could be to use different data, because in this study the lesions were simulated in only one healthy control image. The difference in data could be data of different scanners or/and data with different acquisitions parameters. CS should be robust for these different scanners and acquisitions parameters. For the RIM no studies have been done yet on pathology dataset. However, we expect that the RIM could also be robust against these changes, mainly if the RIM is retrained on a data set that includes also data of with the same pathology.

4.4. Clinical relevance

The clinical relevance of this study was to evaluate if the acquisition time of the MR images could be shortened by undersampling the image data and reconstruct is via the RIM algorithm or CS. Both RIM and CS showed promising results for using it in more clinical data. However, as said before, more testing should be done first for the RIM. As in this study only data of one scanner was used and one type of pathology more evaluation of the methods should be done.

4.5. Conclusion

Lesion simulation facilitates research because more data can be created to test, train an validate reconstruction techniques. In this study simplified WM lesions are simulated in healthy controls in raw data. But more research needs to be done to create an more detailed simulation. In this study, the simulated lesions were used to investigate whether CS and RIM were able to reconstruct the lesions accurately and robustly. Both method performed good, however, RIM is outperforming CS on terms of robustness, especially when images are more accelerated. As CS is used in clinical settings and research studies, it would be possible that the RIM could replace CS eventually. This may result in faster reconstruction, as well as in higher acceleration factors for MRI. However, follow up studies should be done to improve the RIM and more testing should be done before application in the clinical setting is warranted.

5. Acknowledgements

While writing this thesis I received lot of support and assistance in gaining knowledge in MRI and Reconstruction Techniques. I would first like to thank my supervisor from the Spinoza Centre, Matthan Caan. Your expertise and patients have been really contributing to this research and my thesis. Also my supervisor from the TU Delft, Frans Vos. You always had new and fresh ideas when I came down to show the progress. I would like to like to acknowledge Kai Lnning for his expertise on the RIM and his clear explanations of the algorithm. In addition i would like to thank Alexandra de Sitter. She helped me a lot in motivating and writing up the research in this thesis, thank you.

A. Literature: Deep learning for accelerating MRI

A.1. Introduction: Accelerated MRI

Magnetic resonance imaging (MRI), as a bio-medical imaging modality, provides images with good soft tissue contrast. Magnetic Resonance (MR) systems deploy several measuring sequences to produce a specific contrast for the intended purpose. Reducing the acquisition time for MRI is an ongoing challenge to enable imaging of certain biological processes that need to be imaged at high spatial resolution. In the current scanners some accelerated algorithms are already implemented. Some examples are Sensitivity encoding (SENSE) [11], generalized autocalibrating partial parallel acquisition (GRAPPA) [15] and compressed sensing (CS)[9]. Relatively new in the field of acceleration is the implementation of deep learning (DL) algorithms for the image reconstruction. Manufacturers have not implemented a DL algorithm for the reconstruction in the scanner software yet, however this is still an active area of research. In this literature paper we will give an general overview of accelerated MRI and highlight four methods that can be used. The first method will be CS, the three other methods will be DL based methods; the U-Net, a Variational Network (VN) and a Recurrent Inference Machine (RIM). RIM will be used as the main reconstruction method for the rest of this Master thesis.

A.1.1. Undersampling and the inverse problem.

A MRI scanner measures the data in k-space (frequency space). Sampling has to obey the Nyquist-criterion which states that the minimal amount of samples needed for a accurate reconstruction is two times the highest frequency present in the signal [3]. However, there are physical constraints to the data acquisition process, which prolong the time to produce an image. When violating the Nyquist criterion, aliasing artifacts will occur in the image which obviously is not favorable for diagnostics.

Different undersampling patterns produce different undersampling artifacts (Figure A.14). Random undersampling gives a noise-like artifact. By using de-aliasing algorithms these artifacts can be removed. This introduces solving "the inverse problem": in the context of accelerated MRI reconstruction, the forward model is a known process that describes the transformation taking the true image signal to the measured samples. What needs to be solved is the model's inverse transformation, which will produce a reconstruction of the original image. The inverse problem is often solved with CS, an accepted reconstruction method used by many MR manufactures. However, nowadays the inverse problem could also be solved with different DL algorithms. Solving the inverse problem with DL could reconstruct an image which approaches the quality of the ground truth and outperforms the CS method [4]. An other advantage of DL over CS might be an decrease in reconstruction time. Training of the model might take a long time (hours or days), but when a model is trained, it reconstructs a dataset within seconds[4].

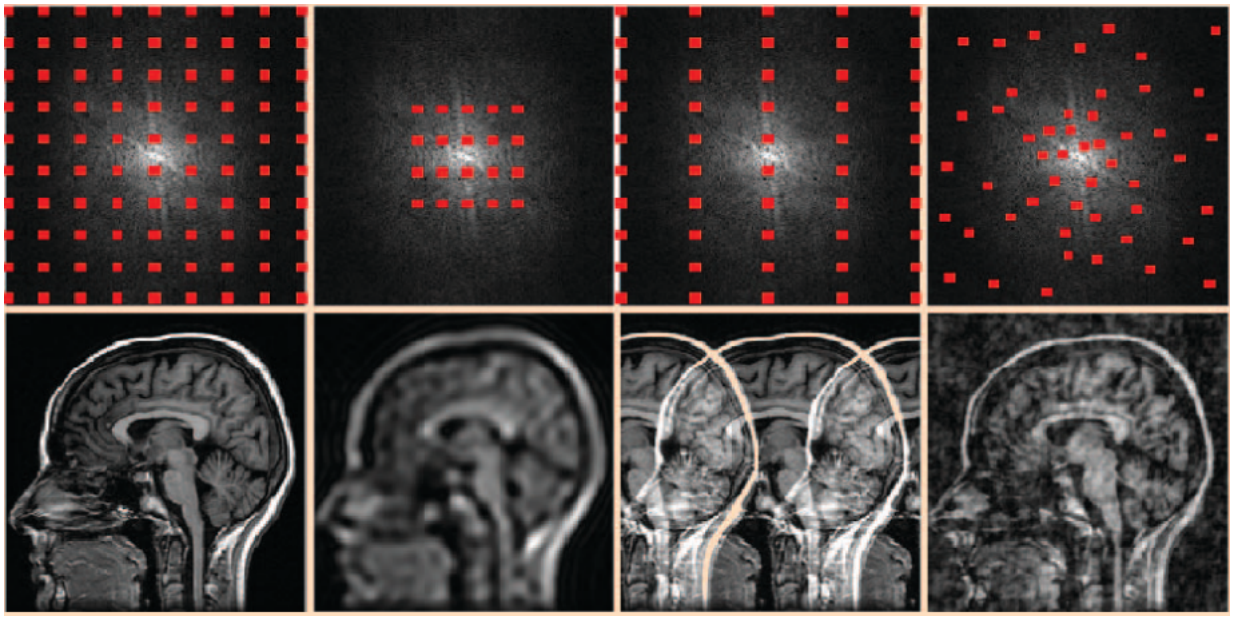

Figure A.14: (Image from [9]). Violation of the Nyquist criterion causes artifacts in linear reconstructions, which depend on the sampling pattern. Left; A fully sampled k-space. Second left; Only center of k-space sampled, but still obeying the Nyquist criterion. Second right; Sequentially undersampled. Aliasing artifacts occur, which can be solved with Sense reconstruction[10]. Right; Random undersampled. Noise like artifacts occur, which can be solved with CS and DL algorithms.

Papers about MRI reconstruction with DL are being rapidly published at the moment. In this brief literature study, it is not possible to cover all the DL methods, and therefore we will only discuss the three different DL methods which will be described A.3 and compare it to a more traditional algorithm CS.

A.2. Compressed Sensing

Compressed sensing (CS) in the literature is described as a mathematical approximation theory[6]. The reconstruction technique in CS exploits the inherent compressibility of the signal it seeks to reconstruct. Therefore, to use CS, the signal it is reconstructing needs to have a sparsifying transform Ψ, such that $u = \Psi x$. Naturally, images have a sparse representations, this feature is used a lot by image compression with file types like JPEG and MPEG[16]. Logically, it can therefore be expected that MRI images are compressible too. The sparsifying transform can vary with the application it is used for. Often the wavelet transform is used in CS reconstruction[17]. A second condition for using CS states that the aliasing artifacts of the image should be incoherent. This means a noise like artifact that should not show any distinguishable pattern that is seen in for example periodic undersampling artifacts (folding). The third and last condition needed to perform CS reconstruction is that the image should be reconstructed by a nonlinear method that enforces both sparsity of the image representation and data consistency of the reconstruction with the acquired samples[9]. The problem is formulated as:

$$\underset{x}{\operatorname{argmin}} \left\{ ||F_u - y||_2^2 + \lambda \, ||\Psi m||_1 \right\} \tag{A.1}$$

where Ψ is the sparsifying transform to transform the image into a sparse domain. An example of this transform is the wavelet transform. y is the actual measured k-space data and F_u is the Fourier transform for the data that was employed. The first term enforces the data consistency via the least-squares difference. It ensures that the reconstructed image m is consistent with the data that is actually measured in k-space. The second term enforces sparsity in the appropriate transform domain. The L1-norm is used, which calculates the degree of sparsity in the image by summing the absolute value of the pixels. So this basically means that the equation finds the image m which is both sparse and also consistent with the data that has been measured in k-space. The regularizer, λ, controls the relation between the two terms. A small λ will lead to an image that is close to the measured data, relying less on the sparsity, while using a high λ will lead to an image that is relying more on the sparsity and deviate more from the collected data.[8]. In Figure A.15[8] a schematic drawing of a simplified compressed sensing algorithm solving equation A.1 is shown. The top row shows the part which enforces sparsity and the bottom row enforces the data consistency. Other techniques can be used to solve this minimization problem as well, like convex sets or iterative soft thresholding.

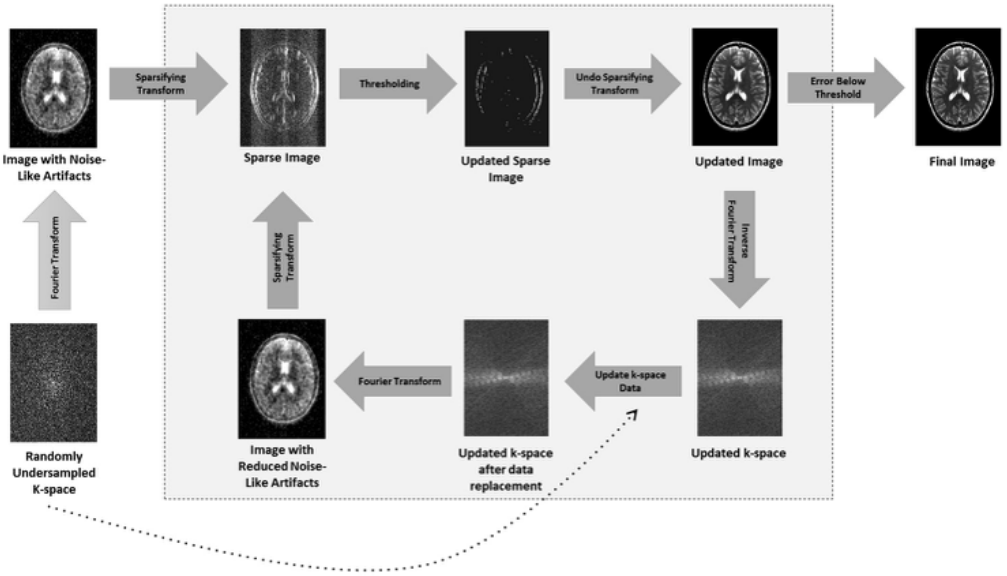

Figure A.15: (Image from [8]) Simplified CS reconstruction. Random undersampled k-space data is collected (bottom left). Linear reconstruction leads to an image which exhibits noise-like aliasing artifacts (top left). By applying a specifying transform, like finite differences (highlighting the edges), or the wavelet transform sparse representation is made. The noise-like artifacts make it possible to retain most of the significant pixels in the sparse image while removing some noise-like artifacts due to thresholding of the image. After thresholding, only the clear edges remain in the image. Some valuable edges might have disappeared due to the thresholding. An updated image is generated by transforming the threshold specified image back to image space and update the image with the reduced aliasing artifacts. Then the image is transformed back into k-space, which may contain data that is different from the collected data. To enforce data consistency, the original k-space data is reinserted into the k-space of the updated image. This updated k-space is then converted into an image. This image should have both the data from the aliased free image as the data from the measured k-space. This process is iterated until a convergence criterion is met or a number of iterations is reached.

A.2.1. Compressed Sensing and Parallel Imaging

Compressed Sensing can be used to accelerate 2D and 3D static images in the spatial transforms. It can also be used in dynamic imaging, where the temporal dimension can be accelerated[18]. Parallel imaging, like SENSE [11] or eSPIRiT, [10] can be combined with compressed sensing to accelerate even further. Acceleration factors of up to 10x are reported in the literature[19]. One application of compressed sensing reconstruction combined with parallel imaging is L1 SPIRiT [20], using the SPIRiT [21] method together with a spatial wavelet transform to sparsify the data. L1 SPIRiT is an iterative reconstruction that is typically used with 3D Cartesian datasets collected using a pseudo-random Poisson disk sampling pattern. This technique has the potential to accelerate clinical MRI scans by a factor of 8 in applications such as pediatric body imaging[22][21].

A.3. Deep Learning and MRI

Rather then using a pre-defined sparsifying transform, which is done in CS, suitable MRI prior distributions can be learned from data. A set of different possible MR-images is huge, even when restricted to a particular anatomical region, contrast mechanism or resolution[4] the variation within the image is big. From each image, a large set of corrupted images can be augmented, one for each permitted k-space sub-sampling mask. Using deep learning to find a function that maps each corruption back to the original signal, for all possible original signals, is a highly complex problem[4]. We have seen that CS does this by using compressibility and sparsity, while parallel imaging uses coil sensitivities.

If we now first consider the forward model again, the measurements can be written as:

$$y_i = PFS_i x + n_i, \quad i = 1, \ \dots \ , c \tag{A.2}$$

The image is decomposed into a set of c partial coil images through the sensitivity maps S_i. This partial image is then projected onto the Fourier domain by F. The subsample mask P reduces the measurements by discarding values in k-space which consequently accelerate the scanning time. Measurements are assumed to be subjected to additive, normally distributed noise, n_i. The inverse problem is used to solves this equation for x. Often the strategy taken is to optimize for the maximum a posteriori (MAP) estimator from statistics, given by:

$$x_{map} = \underset{x}{\operatorname{argmin}} \left\{ \log p(x|y) + \log p(x) \right\} \tag{A.3}$$

which is the maximization of the sum of the log-likelihood and log-prior distributions of y and x. This is commonly reformulated as optimizing the regularized problem [13]

$$x_{map} = \underset{x}{\operatorname{argmin}} \left\{ \sum_{i-1}^{c} (y_i, PFS_i x) d + \lambda R(x) \right\} \tag{A.4}$$

where d evaluates the data consistency and λ is the regularization factor.

DL offers the ability to learn the reconstruction parameters, thereby eliminating the need to pick a prior with some suitable sparse transformation beforehand. Another benefit is that the tuning of λ can be moved away from inference and put into the training procedure.

The combination of accelerated MRI reconstruction with deep learning is relatively new. Although, there have been several deep learning proposals for accelerated MRI reconstruction recently. We will discuss only three different methods: a U-net [23], a Variational Network (VN) [24] and a Recurrent Inference Machine(RIM)[4][13].

A.3.1. U-net

The U-net is a known DL algorithm that was first proposed for image segmentation[23]. The writers of this papers took the model of the U-net and used it for image reconstruction[25]. In contrast to CS, the image reconstruction with the U-net is not iterative. Taking the corrupted image as input, the U-net has two major components. The first component extracts features from local patches through standard CNN architecture of convolutions and max pooling layers. This produces a lot of features at the cost of global context needed for reconstruction. The second component finds a solution for the loss of global context by using unpooling layers. To obey the data consistency, the values in k-space which were measured are placed back to there original value, which they call k-space correction.

Results of the U-net were measured in terms of Structural Similarity Index Measure (SSIM) and the Mean Squared Error (MSE) and only compared to the fully sampled ground truth. The results as shown in the paper [25] displays an increase in image quality when using the U-net with k-spce correction. However, without comparison with other DL reconstruction methods or CS, no conclusion can be made whether its reconstructions had better results.

A.3.2. Variational Network

Hammernik et al.[24] introduce an efficient trainable formulation for accelerated PI-based MRI reconstruction that they call a Variational Network (VN). The model incorporates a generalized CS concept, formulated as a variational model, within a deep learning approach. The VN can reconstruct complex-valued multichannel MR data in an iterative way. The generalized concept taken from the CS is the sparsifying transform used to denoise the image. They take the Total Variation transform and extent it to a Field of Experts model[26]. This model is then substituted for the sparsifying term in formula A.1 and can be solved iteratively. Basically, one iteration of an iterative reconstruction can be related to one step in the network. Training for the VN is done on retrospectively undersampled data of musculoskeletal imaging. Evaluation is done for different acceleration factors. Using both retrospectively and prospectively undersampled clinical patient data, they investigate the applicability of the VN model approach. Evaluation is also done on clinical patient data with pathology's that were not seen in the training set. For each experiment they trained a new model but kept the models architecture fixed.

They did three main experiments. The first one was testing if the model would suffer from structured artifacts as shown in figure A.14, or if it performs better with incoherent undersampling artifacts. This they tested for acceleration factors 3x and 4x with two different contrasts. In the second experiment they evaluated the potential of the model to generalize with respect to different contrasts and orientations. In the last experiment they evaluated how the model works on prospectively accelerated data. Measures used for evaluation where mainly the MSE, Normalized Root Mean Squared Error (NRMSE) and SSIM.

The results show that VN reconstruction of the images with both structured artifacts as with incoherent artifacts performed better then CS and SENSE. It performed equally good on dictionary learning. The top row in figure A.16 shows a qualitative result of the reconstructions with structured artifacts. The table in figure A.17 shows the results. From this table the conclusion can be drawn that in all three of the experiments the VN showed better results compared to other reconstruction techniques.

Quantitative Evaluation Results in Terms of MSE, NRMSE and SSIM as well as Image Quality Reader Scores for a Clinical knee Protocol and Acceleration Factor $R = 4$ for Regular Sampling and Variable-Density Random Sampling. For the Reader Scores, we Depict the Mean Values and Standard Deviations Averaged Over Both Readers Along With the p-Value Obtained by the One-Sided Wilcoxon Signed-Rank Test. Values that Accept the Alternative Hypothesis with a Significance Level $\alpha = 0.05$ that VN Reconstructions have a Better Quality Score, are Marked as Bold.

		Regular			Random			Reader scores regular		
Data set	Method	MSE	NRMSE	SSIM in %	MSE	NRMSE	SSIM in %	Criterion	PI-CS TGV	Learning	*P*-value
Coronal PD	Zero filling	19.41 ± 4.43	0.17 ± 0.02	79.00 ± 2.36	15.83 ± 3.68	0.16 ± 0.02	80.64 ± 2.41	Artifact	3.60 ± 0.57	1.65 ± 0.07	**0.0010**
	CG SENSE	5.20 ± 0.97	0.16 ± 0.03	84.01 ± 2.21	4.26 ± 0.98	0.15 ± 0.03	85.57 ± 2.29	Sharpness/Blur	2.90 ± 0.14	2.15 ± 0.07	**0.0234**
	PI-CS TGV	2.35 ± 0.40	0.09 ± 0.02	89.80 ± 1.75	1.91 ± 0.45	0.09 ± 0.02	90.36 ± 1.79	SNR	2.60 ± 0.28	1.45 ± 0.21	**0.0078**
	Learning	**1.64 ± 0.28**	**0.08 ± 0.02**	**92.14 ± 1.68**	**1.37 ± 0.32**	**0.08 ± 0.02**	**92.86 ± 1.63**	Overall image quality	3.30 ± 0.14	2.05 ± 0.21	**0.0010**
Coronal fat-sat. PD	Zero filling	20.71 ± 4.07	0.23 ± 0.03	73.96 ± 3.04	17.69 ± 3.30	0.22 ± 0.03	75.10 ± 3.17	Artifact	3.95 ± 0.07	2.90 ± 0.42	**0.0020**
	CG SENSE	14.55 ± 1.62	0.25 ± 0.05	73.06 ± 4.62	11.79 ± 1.39	0.24 ± 0.04	74.78 ± 4.55	Sharpness/Blur	3.95 ± 0.07	3.15 ± 0.64	**0.0020**
	PI-CS TGV	7.73 ± 1.14	0.19 ± 0.04	79.19 ± 4.14	7.07 ± 1.07	0.18 ± 0.03	79.69 ± 4.09	SNR	3.75 ± 0.21	2.90 ± 0.71	**0.0049**
	Learning	**6.49 ± 0.80**	**0.17 ± 0.03**	**81.97 ± 3.60**	**5.81 ± 0.85**	**0.17 ± 0.03**	**82.47 ± 3.67**	Overall image quality	3.95 ± 0.07	3.20 ± 0.57	**0.0020**
Sagittal fat-sat. T_2	Zero filling	16.66 ± 3.14	0.19 ± 0.03	85.71 ± 2.62	17.35 ± 3.21	0.19 ± 0.03	84.91 ± 2.59	Artifact	2.90 ± 0.14	2.80 ± 0.28	0.3750
	CG SENSE	6.27 ± 1.62	0.15 ± 0.04	87.86 ± 3.08	9.55 ± 2.11	0.18 ± 0.04	85.06 ± 3.11	Sharpness/Blur	3.40 ± 0.14	2.75 ± 0.21	**0.0156**
	PI-CS TGV	3.39 ± 0.82	**0.11 ± 0.03**	91.84 ± 2.81	4.76 ± 0.95	0.13 ± 0.03	90.29 ± 2.70	SNR	3.20 ± 0.28	2.50 ± 0.28	**0.0234**
	Learning	**2.99 ± 0.68**	**0.11 ± 0.03**	**92.83 ± 2.40**	**3.92 ± 0.81**	**0.12 ± 0.03**	**91.85 ± 2.35**	Overall image quality	3.30 ± 0.28	2.75 ± 0.07	**0.0078**
Sagittal PD	Zero filling	5.17 ± 0.75	0.11 ± 0.01	87.53 ± 1.95	3.32 ± 0.51	0.09 ± 0.01	89.49 ± 1.80	Artifact	2.10 ± 0.14	2.00 ± 0.14	0.4063
	CG SENSE	0.86 ± 0.15	0.06 ± 0.02	92.74 ± 1.46	1.03 ± 0.16	0.07 ± 0.02	92.37 ± 1.48	Sharpness/Blur	2.10 ± 0.14	2.10 ± 0.14	0.6875
	PI-CS TGV	0.49 ± 0.09	0.05 ± 0.01	96.22 ± 1.17	0.64 ± 0.11	**0.05 ± 0.01**	95.47 ± 1.24	SNR	1.60 ± 0.00	1.50 ± 0.28	0.3828
	Learning	**0.44 ± 0.07**	**0.04 ± 0.01**	**96.64 ± 1.16**	**0.52 ± 0.09**	**0.05 ± 0.01**	**96.07 ± 1.17**	Overall image quality	2.20 ± 0.14	2.05 ± 0.07	0.2656
Axial fat-sat. T_2	Zero filling	44.57 ± 9.95	0.27 ± 0.02	78.52 ± 1.92	48.03 ± 11.13	0.28 ± 0.02	77.80 ± 1.98	Artifact	3.15 ± 0.07	3.10 ± 0.57	0.5000
	CG SENSE	23.75 ± 4.56	0.24 ± 0.03	80.30 ± 3.20	31.98 ± 4.88	0.27 ± 0.02	78.87 ± 2.43	Sharpness/Blur	3.05 ± 0.07	2.95 ± 0.49	0.3750
	PI-CS TGV	13.65 ± 3.78	0.18 ± 0.03	85.51 ± 3.25	15.30 ± 2.57	0.19 ± 0.02	84.93 ± 2.60	SNR	3.10 ± 0.14	2.75 ± 0.49	**0.0313**
	Learning	**10.63 ± 2.48**	**0.16 ± 0.02**	**88.46 ± 2.43**	**12.06 ± 2.13**	**0.17 ± 0.02**	**87.74 ± 2.30**	Overall image quality	3.20 ± 0.14	3.05 ± 0.49	0.2266

Figure A.16: (Table from [24]) Results of the three experiments

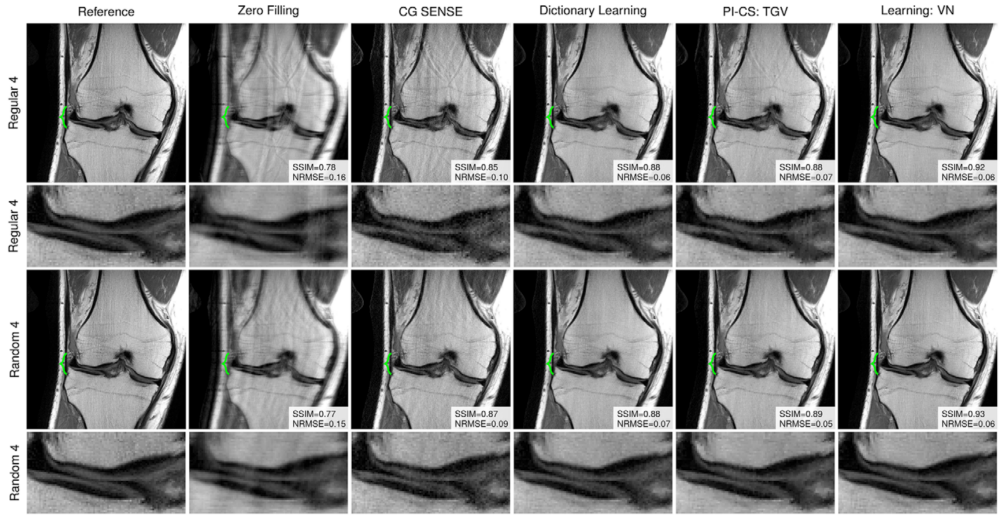

Figure A.17: (Image from [24]) Coronal PD-weighted scan with acceleration 4x of a 32-year-old male. The green bracket indicates osteoarthritis. The first and second row show reconstructions with regular Cartesian sampling, the third and forth row show reconstructions with random sampling. Zoomed views show that the learned VN reconstruction appears slightly sharper than the PI-CS TGV and dictionary learning reconstruction. The dictionary learning and VN reconstruction can significantly suppress artifacts unlike CG SENSE and PI-CS TGV. Results based on random sampling show reduced residual artifacts and slightly increased sharpness in comparison to regular sampling

A.3.3. Recurrent Inference Machine

Recurrent Inference Machines (RIM) were first proposed by Putzky and Welling (2017)[13] as an inverse problem solver. Later this is used by Lnning (2017) [4] to apply it for accelerated MRI reconstruction. This model constrains the solution space by learning an iterative process, where step-wise reassessments of the maximum a posteriori estimate lead to an incremental update that infers the inverse transform of the forward model[4].

The authors' hypothesis is that the RIM is capable of reconstructing data of multiple acceleration factors while only trained on a single acceleration factor. They verified this by comparing the performance of the models trained on a single acceleration factor with models trained on a range of acceleration factors. They also claim that the RIM has the ability to generalize structures that were unseen during training.

To test the performance of the reconstructions the SSIM [27], NRMSE and Peak Signal-to Noise Ratio (PSNR) are used as metrics. The images were also rated by a experienced neurologist in a double blinded test.

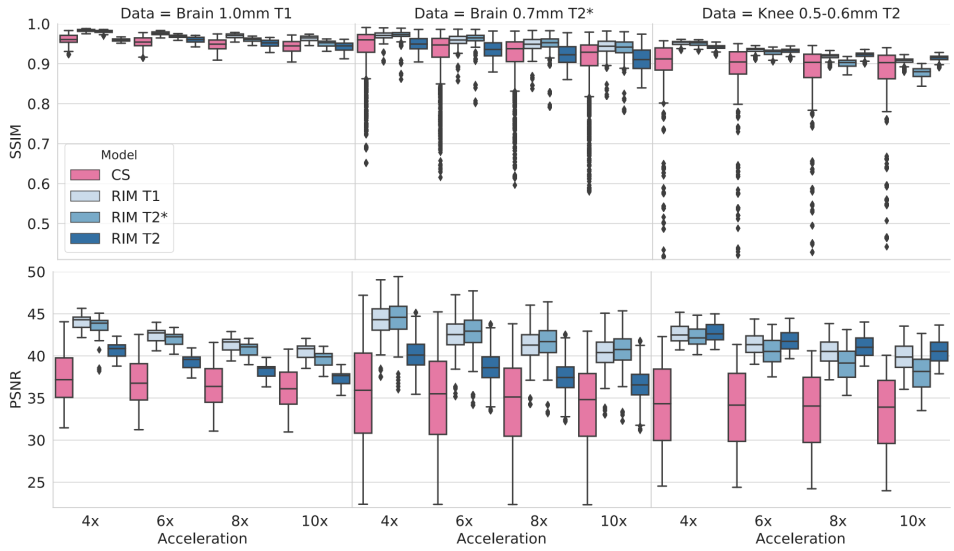

Figure A.18: (Image from [4]) Boxplots of SSIM- and PSNR-values for CS and the RIMs (both with PI) trained on all three types of data: 1.0mm T1-weighted brains, 0.7mm T2*-weighted brains, and 0.5-0.6mm T2-weighted knees. Hues indicate the model and, in the case of the RIM, the type of data trained on, whereas columns indicate the type of data evaluated on.

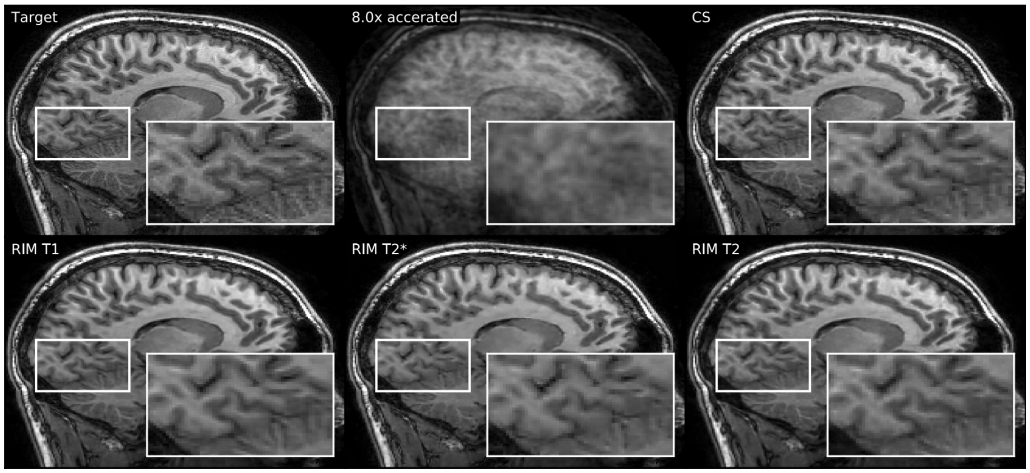

Figure A.19: (Image from [4]) Shows the 8 time accelerated images of a T1-weighted scan. Trained on T1, T2* and T2 weighted models. On the eye, there are very little differences.

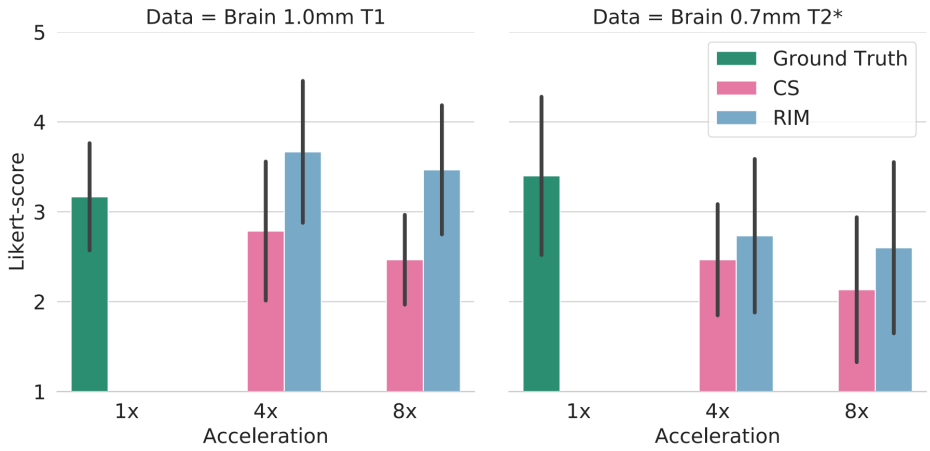

Figure A.20: (Image from [4]) The bars show the average score assigned by the neurologist.

The paper [4] shows SSIM- and NRMSE-scores for RIMs trained and cross-evaluated on three types of data, along with CS (which was optimized for this dataset). The first column of figure A.18 is illustrated qualitatively in figure A.19. It shows the 8x accelerated image from a T1-weighted brain dataset. In figure A.20 the double blinded test results from the neurologist were shown.

As we can see in the figures, the RIM outperforms the existing CS method in terms of image quality. The neurologist scores the RIM reconstructed image even higher then the fully sampled image. This might be the case because the RIM slightly smooths the fine structures, and human perception is positively sensitive for that.

A.4. Conclusion

As we have seen, the above described methods all do reconstruction in a different way. All perform well with the specific data they used in their study. The MSE and SSIM results of the studies are difficult to compare, because the studies use different data-sets and trained the models on different data. In the study on RIM and VN the methods were compared to other reconstruction methods and had better results then CS. Furthermore, as VN was the only network that was tested on pathological data (of the knee), it seemed to perform fine. Thus even though image quality seems to be excellent for the three DL algorithms in this study, more research has to be done on how well the performance is with pathological data. Moreover, a study that compares the methods should give an better interpretation of performance of the methods compared to each other and for example CS. As a result of this conclusion, there is merit in further investigation of the difference in performance between the RIM, which has promising results, and CS.

References

[1] G. Wright, Magnetic resonance imaging, IEEE Signal Processing Mag. 14 (1) (1997) 56–66.

[2] M. Lustig, D. L. Donoho, J. M. Santos, J. M. Pauly, Compressed sensing MRI, IEEE Signal Processing Magazine 25 (March 2008) (2008) 72–82. arXiv:1204.4227v1, doi:Doi 10.1109/Tit.2006.871582.
URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4472246

[3] C. E. Shannon, Communication in the presence of noise, Proceedings of the IEEE 72 (9) (1984) 1192–1201.

[4] K. Lønning, Master Thesis A Complex-Valued Recurrent Inference Machine for Accelerated MRI Reconstruction by.

[5] E. M. Haacke, R. W. Brown, M. R. Thompson, R. Venkatesan, et al., Magnetic resonance imaging: physical principles and sequence design, Vol. 82, Wiley-Liss New York:, 1999.

[6] D. L. Donoho, Compressed sensing, IEEE Transactions on Information Theory 52 (4) (2006) 1289–1306. arXiv:1204.4227v1, doi:Doi 10.1109/Tit.2006.871582.

[7] E. Candes, M. Wakin, An Introduction To Compressive Sampling, IEEE Signal Processing Magazine 25 (2) (2008) 21–30. arXiv:arXiv:1307.1360v1, doi:10.1109/MSP.2007.914731.

[8] A. C.-Y. Yang, M. Kretzler, S. Sudarski, V. Gulani, N. Seiberlich, Sparse reconstruction techniques in mri: methods, applications, and challenges to clinical adoption, Investigative radiology 51 (6) (2016) 349.

[9] M. Lustig, D. Donoho, J. M. Pauly, Sparse MRI: The application of compressed sensing for rapid MR imaging, Magnetic Resonance in Medicine 58 (6) (2007) 1182–1195. doi:10.1002/mrm.21391.

[10] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, M. Lustig, ESPIRiT-an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA, Magnetic Resonance in Medicine 71 (3) (2014) 990–1001. doi:10.1002/mrm.24751.
URL http://doi.wiley.com/10.1002/mrm.24751

[11] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, P. Boesiger, Sense: sensitivity encoding for fast mri, Magnetic resonance in medicine 42 (5) (1999) 952–962.

[12] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using rnn encoder-decoder for statistical machine translation, arXiv preprint arXiv:1406.1078.

[13] P. Putzky, M. Welling, Recurrent Inference Machines for Solving Inverse Problems (Nips). arXiv:1706.04008.
URL http://arxiv.org/abs/1706.04008

[14] S. Hametner, A. Dal Bianco, S. Trattnig, H. Lassmann, Iron related changes in ms lesions and their validity to characterize ms lesion types and dynamics with ultra-high field magnetic resonance imaging, Brain Pathology 28 (5) (2018) 743–749.

[15] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, A. Haase, Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA), Magnetic Resonance in Medicine 47 (6) (2002) 1202–1210. doi:10.1002/mrm.10171.

[16] D. Taubman, M. Marcellin, JPEG2000 image compression fundamentals, standards and practice: image compression fundamentals, standards and practice, Vol. 642, Springer Science & Business Media, 2012.

[17] M. Lustig, J. M. Pauly, SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space, Magnetic Resonance in Medicine 64 (2) (2010) 457–471. arXiv:NIHMS150003, doi:10.1002/mrm.22428.

[18] U. Gamper, P. Boesiger, S. Kozerke, Compressed sensing in dynamic MRI, Magnetic Resonance in Medicine 59 (2) (2008) 365–373. doi:10.1002/mrm.21477.
URL http://doi.wiley.com/10.1002/mrm.21477

[19] R. Otazo, D. Kim, L. Axel, D. K. Sodickson, Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion mri, Magnetic resonance in medicine 64 (3) (2010) 767–776.

[20] M. Lustig, M. Alley, S. Vasanawala, D. Donoho, J. Pauly, L1 spir-it: Autocalibrating parallel imaging compressed sensing, in: Proc Intl Soc Mag Reson Med, Vol. 17, 2009, p. 379.

[21] M. Lustig, J. M. Pauly, Spirit: iterative self-consistent parallel imaging reconstruction from arbitrary k-space, Magnetic resonance in medicine 64 (2) (2010) 457–471.

[22] M. Murphy, M. Alley, J. Demmel, K. Keutzer, S. Vasanawala, M. Lustig, Fast e-spirit compressed sensing parallel imaging mri: Scalable parallel implementation and clinically feasible runtime, IEEE transactions on medical imaging 31 (6) (2012) 1250–1262.

[23] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmenta-

tion, in: International Conference on Medical image computing and computer-assisted intervention, Springer, 2015, pp. 234–241.

[24] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, F. Knoll, Learning a variational network for reconstruction of accelerated mri data, Magnetic resonance in medicine 79 (6) (2018) 3055–3071.

[25] C. M. Hyun, H. P. Kim, S. M. Lee, S. Lee, J. K. Seo, Deep learning for undersampled MRI reconstruction, IEEE Transactions on Image Processing 26 (9) (2017) 4509–4522. arXiv:1709.02576, doi:10.1109/TIP.2017.2713099.
URL http://arxiv.org/abs/1709.02576

[26] S. Roth, M. J. Black, Fields of experts, International Journal of Computer Vision 82 (2) (2009) 205.

[27] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE transactions on image processing 13 (4) (2004) 600–612.