Application and Development of Quantitative Magnetic Resonance Imaging:

High-Resolution QMRI in an Elderly Cohort and End-to-End Quantitative Susceptibility Mapping with Deep-Learning

Elise Al

Supervisors
F. M. Vos
M. W. A. Caan
L. C. Liebrand

A thesis presented for the degree of
Master of Science

Biomedical Engineering
TU Delft
Netherlands
August 8 2022

Abstract

We present application and developments of qMRI. We reconstructed data from the Prediva Ouderen Extentie (POE) database and calculated quantitative maps. The POE database consists of 29 7 Tesla (T) whole-brain MRI scans, including both male and female participants covering the age group between 80 and 91 years old. We used this data to extend the currently existing Amsterdam Ultra-high field adult lifespan database (AHEAD), covering the entire adult life span (18-80 years). Data was acquired at a submillimeter resolution using a the MP2RAGEME sequence, resulting in complete anatomical alignment of quantitative R_1-maps, R_2^*-maps, T_1-maps, T_2^*-maps, and quantitative susceptibility mapping (QSM).

As qMRI acquisitions and reconstructions are slow, we also investigated the potential of deep learning to speed up this process using the same dataset. Conventional MRI acquisitions are already time consuming, scanning the multiple images necessary for qMRI may take even longer. Computing a quantitative susceptibility map requires several processing steps involving phase unwrapping, background phase removal, and solving an ill-posed inverse problem, which all together is computationally expensive. COSMOS is currently seen as gold standard for solving the dipole inversion, although sampling from different orientations makes it a lengthy procedure and not feasible for clinical practice. Iterative methods were developed to make high quality QSM maps out of a single orientation measurement, however they suffer from artifacts and can be computationally expensive. Deep learning may improve the QSM pipeline in terms of time, convenience and quality. Neural networks were developed to learn processing steps like phase unwrapping and dipole inversion. Especially for the dipole inversion step several methods are proposed to calculate a high quality QSM map out of single orientation measurements which could speed up the QSM pipeline. We propose a method which can be useful in accelerated end-to-end QSM calculation and highlight advances and possible points for improvement. We will conclude with a discussion of the challenges that need to be overcome to establish a fast, easy and high quality QSM pipeline.

Keywords— Quantitative magnetic resonance imaging, Quantitative susceptibility mapping, Deep learning

Contents

1 Introduction

qMRI

Magnetic Resonance Imaging (MRI) is an imaging technique which distinguishes itself from other techniques by excellent soft-tissue contrast. Quantitative MRI (qMRI) goes beyond conventional MRI since it provides specific physical parameters instead of local image contrast. Quantitative maps contain values of intrinsic tissue properties and are independent of MRI sequence parameters. The most common quantitative measures in MRI are the relaxation times T_1, T_2 and T_2^* (with their associated relaxation rates R_1, R_2 and R_2^*) and susceptibility (these will be covered in section 2.1). QMRI has proven to be better reproducible than conventional MRI [Alexander et al., 2021]. This allows monitoring changes over time more accurately making it better suited for population studies or follow-up studies. Healthy aging population studies are valuable in the clinic for early detection of abnormalities. Given the steadily increasing life expectancy and the increasing prevalence of neurodegenerative diseases, the differentiation between normal aging and early stages of neurodegeneration is crucial to identify individuals at risk of cognitive decline [Alexander et al., 2021].

The quantitative maps of tissue properties need to be calculated by fitting a parametric model to multiple images that are scanned with different specifications. While conventional MRI acquisitions are time consuming, scanning the multiple images necessary for qMRI may take even longer. Besides, the calculation of the quantitative maps can be computationally expensive. All in all, this often exceeds the acceptable time in the clinic and a lot of profit can be made by accelerating these processes. This is important for population studies, especially for subjects who cannot lie still for long, such as the elderly. In the first part of my thesis I look at the application of acceleration and motion correction (relevant for the elderly and patients) by means of a new population study (see section 4.1). In order to accelerate even further, even more advanced techniques are needed, these are explored in the second part of my thesis (see section 4.2).

POE

With healthy ageing, changes occur in the human brain. A lot is still unknown about the modifications and processes in the healthy elderly brain. Population studies with qMRI are scarce, but there is the Amsterdam Ultra-high field adult lifespan database (AHEAD) study [Alkemade et al., 2020], a population study of healthy adults (18-80 years). The AHEAD study mapped the differences in the brain that arise with age and modeled relationships between age and quantitative values like R_1, R_2^* and QSM. This has provided knowledge about anatomical and microscopical changes in certain subcortical structures. For example, iron accumulates during lifespan in the striatum. It is unknown if this will reach a limit or if it will keep increasing linearly or exponentially.

A shortcoming of the AHEAD is that the study does not contain the extreme group of elderly (over 80 years). This extreme group of elderly could contain crucial information to discriminate between linear and exponential relationships. In addition, this group can be used to find out whether certain quantitative values continue to rise or whether they may level off and reach a maximum. Here we present the Prediva Ouderen Extentie (POE) database, which can be seen an extension of the AHEAD database and contains the important information of healthy elderly between 80 and 91 years. The POE database is acquired with almost identical MRI sequence parameters and scanned on the same 7T MRI scanner as the AHEAD data, which makes it ideal for comparison. This database can complement the AHEAD study and increase our current knowledge. For these purposes, it is necessary to make quantitative images out of the acquired data.

Deep-Learning QSM

Quantitative susceptibility mapping (QSM) is an MRI postprocessing technique that quantifies the spatial distribution of magnetic susceptibility. QSM extracts the magnetic tissue susceptibility based of the phase images of multiple gradient-recalled echoes (GRE), assuming that phase shift results mainly from susceptibility-induced field inhomogeneity. Biometals and molecules such as iron, calcium, lipids and myelin are examples of common susceptibility contributors and often involved in neurodegenerative diseases.

Computing susceptibility involves several processing steps involving phase unwrapping, background phase removal, and solving an inverse problem relating the tissue phase to the underlying susceptibility distribution. Solving the inverse problem requires a dipole inversion. Due to zeros on two conical surfaces in the k-space dipole kernel this problem is ill-posed. The inverse kernel is undefined at these surfaces and noise is amplified in regions where the kernel is very small, making a simple inversion of the forward calculation impractical.

Different methods have been proposed to deal with this dipole inversion. An algorithm called COSMOS [Liu et al., 2009], which stands for Calculation of Susceptibility through Multiple Orientation Sampling, uses sampling in multiple orientations to compensate for the missing data when using a single orientation. This method is currently seen as the gold standard. A major drawback is that it requires multiple scans (at least three) in different subject orientations (requiring subject repositioning) for one susceptibility map. This means scanning time is considerably prolonged which restricts the clinical feasibility.

Iterative methods such as iterative Least squares (iLSQR) [Li et al., 2011], Morphology-Enabled Dipole Inversion (MEDI) [Liu et al., 2012] and truncated dipole inversion [Shmueli et al., 2009] were developed to restore high-quality susceptibility maps from single orientation measurements. However, these methods can be computationally intensive and time-consuming. Besides, they are prone to artifacts and require manual parameter tuning for different data sets.

Recently, deep learning has been applied to solve a variety of inverse problems as an alternative to the conventional methods. These neural networks are capable of approximating any continuous function, given enough learnable parameters and enough available data sets. A great advantage is that applying a trained network to solve an inverse problem (like the dipole inversion) is generally fast. Most of the existing methods adopted the U-Net structure, for example, DeepQSM [Bollmann et al., 2019b], QSMnet [Yoon et al., 2018] and AutoQSM [Wei et al., 2019]. Other studies of QSM calculation included xQSM, which trained an octave CNN to improve network generalization capability [Gao et al., 2021]. Previous studies successfully trained multiple deep neural networks for QSM dipole inversion. However, all of the above deep learning methods work with tissue phase maps as input, so none of them uses deep learning to reconstruct a QSM directly from raw data. Besides the dipole inversion, several other processing steps in the QSM pipeline can be done by deep learning, like de-aliasing, de-noising, masking and phase unwrapping, suggesting that deep learning might go from k-space to QSM map at once.

Inspired by the work of Zhang et al. who used a quantitative Recurrent Inference Machine (qRIM) to compute a R_2^*-map directly from raw data, we propose a method using a qRIM to reconstruct a QSM map from raw data. The main novelty is that the the QSM is calculated from raw data, so the pre-processing steps are implemented and learned in the model. Computing a high quality QSM map out of single orientation measurements with an end-to-end pipeline would make processing easier. In addition, such a deep learning pipeline has the potential to significantly speed up the process making QSM better applicable in practice. Reconstructions by deep learning already have proven to reduce the impact of noise, which may aid further preprocessing. The hypothesis is that combining deep learning reconstruction with unwrapping and fitting of a tissue phase map can be more efficient than all steps separately. Assuming this will give more accurate reconstructions and reduce uncertainties in the pre-processing steps.

2 Theory

2.1 qMRI

Quantitative MRI (qMRI) provides specific physical parameters instead of local image contrast. We will discuss the physical parameters T_1, T_2^* and QSM and how to calculate them in the next subsections.

2.1.1 T_1 Mapping

T_1 is the relaxation time for regrowth of the longitudinal magnetization M_z to its initial magnetization M_0 after an RF-pulse. This can be described by: $M_z = M_0(1 - e^{-t/T_1})$. A T_1 map can be obtained by several ways, one method is by two Gradient-Recalled Echo (GRE) image volumes acquired within a single repetition at two different inversion times, $T_{I,1}$ and $T_{I,2}$. Assuming that the recovery can be characterized by a single exponential, accurate T_1 estimates are obtained as long as the acquisition parameters are properly chosen [Metere et al., 2017].

A T_1-weighted image was made using the inversion 1 (inv_1) and the first echo of the inversion 2 (inv_{2A}) of the MP2RAGE-ME sequence. The T_1-weighted image was calculated as follows:

$$T_{1,weightedimage} = real(\frac{\overline{inv_1} \cdot inv_{2A}}{|inv_1|^2 + |inv_{2A}|^2}) \tag{1}$$

in which the bar in $\overline{inv_1}$ represents the complex conjugate and *real* means that the real part is taken. After the T_1-weighted image was made, a R_1 map was made. First the T_1-weighted image was scaled between 0 and 4095. A T_1 map is made by making use of the knowledge of the MP2RAGE sequence and its parameters (inversion times, repetition times, flip angles). The map of the corresponding R_1 relaxation rate is made by taking the reciprocal of the T_1 map, $R_1 = \frac{1}{T_1}$.

2.1.2 T_2^* Mapping

T_2^* relaxation is related to T_2 relaxation, the process by which the transverse components of magnetization (M_{xy}) decay (or dephase) caused by atomic and molecular interactions. In practice, the transverse magnetization decays much faster than predicted by atomic and molecular mechanisms resulting principally from inhomogeneities in the main magnetic field. These field inhomogeneities ΔB contribute to the relaxation rate as described below:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \gamma \Delta B \tag{2}$$

resulting in the effective spin-spin relaxation rate $\frac{1}{T_2^*}$.

T_2^* is conventionally fitted from reconstructed images of multiple echo times acquired with a multi-echo gradient echo (ME-GRE) sequence [Zhang et al., 2021]. This can be described by: $M_{xy} = M_0(e^{-t/T_2^*})$. The T_2^* map was made using the 4 echoes from the second inversion. Since the 4 echo times $3, 11.5, 20$ and $28.5ms$ were known, it was possible make a exponential fit of the data in order to retrieve the T_2^*-values. The corresponding R_2^* relaxation rate is made by taking the reciprocal of the T_2^* relaxation time, $R_2^* = \frac{1}{T_2^*}$.

2.1.3 Quantitative Susceptibility Mapping

The magnetic susceptibility is the degree that a material can be magnetized by an external magnetic field. Quantitative Susceptibility Mapping (QSM) is a post-processing technique that provides quantitative maps of tissue susceptibility [Ruetten et al., 2019]. Since the values are quantitative, the QSM maps are robust to changes in scan specifications and reproducible. It allows for quantitative analysis across subjects or between different scanners and benefits longitudinal monitoring and classification of disease processes. QSM measures relative susceptibilities to water, so less susceptibility than water means diamagnetic and higher susceptibility means paramagnetic. This makes that QSM can for example differentiate between a blood products (hyper intense) and calcifications (hypo intense).
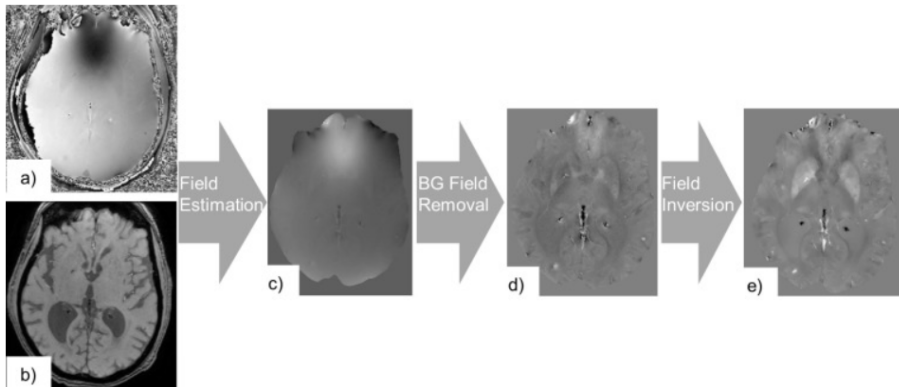

Figure 1: QSM processes the phase (a) and magnitude images (b), the field ΔB is estimated (c), contributions from background fields are removed (d) and a dipole field inversion operation calculates the susceptibility map (e) [Ruetten et al., 2019].

Magnetic field inhomogeneities cause dephasing due to imperfect static magnetic field and tissue susceptibility. By measuring the complex signal of a gradient echo scan, the susceptibility can be computed. QSM computation involves several steps to extract the tissue susceptibility from the phase information [Ruetten et al., 2019]:

1. A map of field inhomogeneities ΔB (figure 1c) is estimated from the phase data (figure 1a)

2. Field inhomogeneities generated by sources outside the region of interest (ROI) are eliminated to extract a map of field inhomogeneities generated by only susceptibility sources inside the ROI ΔB_{int} (figure 1d)

3. A dipole field inversion operation is performed on ΔB_{int} to calculate the susceptibility map χ (figure 1e)

The phase of the MR signal increases linearly with time, the slope of the phase variation (how fast the phase increases with time) is proportional to the field inhomogeneity ΔB, which can be estimated for each voxel (Figure 2).
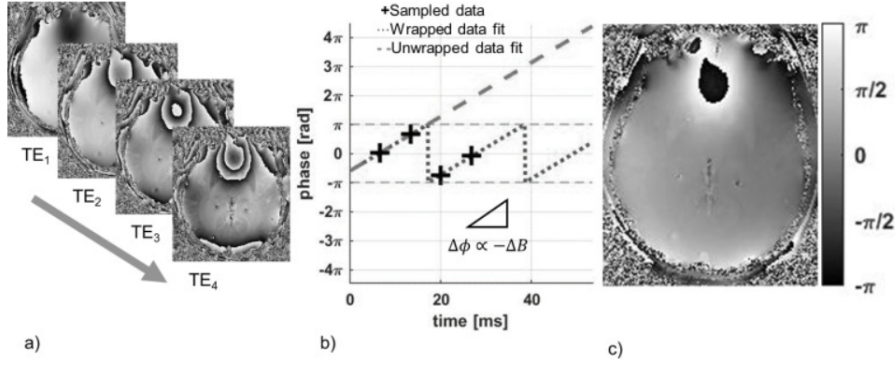
Figure 2: In order to estimate ΔB, the variation of the phase images (a) are monitored over several echo times. By fitting a linear function to the phase data, phase wraps between echoes can be estimated and removed to estimate the slope of the phase variation over time and the corresponding value of ΔB (b). This is done for every voxel in order to generate a map of temporally unwrapped ΔB values (c) [Ruetten et al., 2019].

Phase is an angle and represented in the range from $-\pi$ to π. If the phase shift goes out of that range, the phase becomes "wrapped" or "aliased". This phase wrapping can be corrected or "unwrapped" by making use of algorithms that add integer multiples (k) of 2π to the acquired phase ϕ_w [Ruetten et al., 2019]:

$$\phi_{uw} = \phi_w + k2\pi \tag{3}$$

Fortunately there are phase unwrapping methods which can easily reconstruct the original phase. There are several spatial phase unwrapping techniques which can be subdivided by Laplacian, region-growing and path-based methods. Laplacian based methods work because the complex values of the true and wrapped phase are the same. So the differentiable operator applied to the unwrapped phase can produce the same result on the wrapped phase, which can be used to solve the unwrapping problem. Note that the Laplacian unwrapping has an arbitrary background phase. For QSM this is not a big problem, because the background phase will be removed before computing the QSM. A successful method for Laplacian based phase unwrapping is HARPERELLA [Li et al., 2014]. Region growing algorithms identify discontinuities between regions and unwraps them step-by-step. As a drawback, region growing methods take some time. Path-based unwrapping is 3D voxel-by-voxel unwrapping guided by the quality of the voxel connections. There are algorithms that perform path-based unwrapping, besides there exist deep neural networks that are trained to perform path-based unwrapping, like the PHU-net [Zhou et al., 2021]. Except for the spatial phase unwrapping techniques, phase changes between echoes can be used to perform a temporal unwrapping of the data in a 1D unwrapping approach [Robinson et al., 2016]. For temporal phase unwrapping, first the aliasing between echoes is removed in the temporal domain. Phase wraps in between echoes are removed during the estimation of ΔB and a scaled version of ΔB is estimated. The remaining phase wraps are corrected by spatial phase unwrapping algorithms (Figure 3).

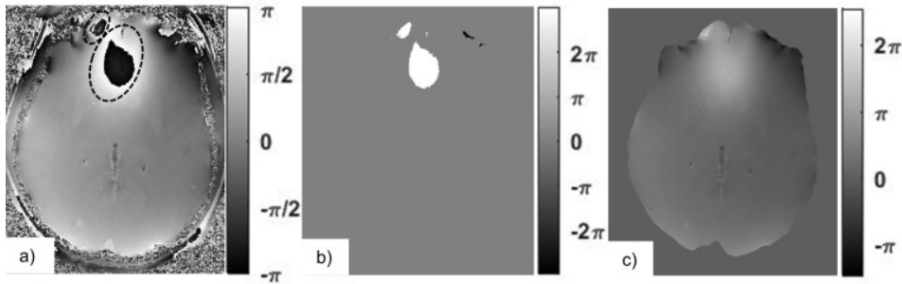Figure 3: The map of ΔB (a) includes phase wraps (dashed lines). Phase unwrapping algorithms identify these wraps and add integer multiples of 2π (b), resulting in an unwrapped field map (c) [Ruetten et al., 2019].

Temporal phase unwrapping has some disadvantages. According to the Nyquist criterion, it will only succeed in removing wraps that occur over a time greater than the inter-echo period. This will limit the

inter-echo time and constraints the acquisition. This time needs to be smaller at higher field due to the higher precession frequency of the protons which is related to the field strength. At 7T, this means that the inter-echo time needs to be smaller than 1.4 ms [Robinson et al., 2016]. In addition, a disadvantage of temporal phase unwrapping is that it is more prone to noise than spatial methods .

The next step in the QSM pipeline is the masking of the region of interest that contains an area of sufficient SNR. This seems not that important but this can affect the result quite a bit. For most background field correction methods needs to be defined what is inside and outside the object of interest. Including unreliable phase values like noise will lead to severe errors when estimating ΔB. This will propagate and lead to artifacts in the dipole inversion. In other words, masking is a trade-off between artifacts and masking out regions of interest.

The map of field inhomogeneities ΔB inside the ROI is made up from magnetized susceptibility sources located inside (B_{int}) and outside the ROI (B_{ext}). Background field removal tries to remove artifacts that are caused by susceptibility sources outside the object of interest. This background field is caused by susceptibility differences between tissue and air, shim coil fields and static field inhomogeneities. Additionally the Laplacian does contribute to the background field if Laplacian unwrapping is used.

In order to extract the susceptibility map, it is important to consider only B_{int} and remove the contributions of B_{ext}. This background field removal is performed by mathematical algorithms. The algorithms generally make assumptions about the mathematical properties of B_{ext} and B_{int} to separate them [Ruetten et al., 2019]. For example an approach that assumes B_{ext} varies at a much lower spatial frequency than B_{int} and can be removed by applying a high pass filter to ΔB.

The local field inhomogeneity can be estimated as:

$$\Delta B(k) = B_0(\chi(k)D(k)) \tag{4}$$

in which B_0 is the static magnetic field strength, $\chi(k)$ the Fourier transform of the susceptibility map and $D(k)$ the Fourier transform of the unit dipole kernel [Ruetten et al., 2019]. This equation can be solved by deconvolving B_{int} with the inverse of the dipole kernel, which corresponds to a division in k-space:

$$\chi(k) = \frac{1}{B_0}\frac{B(k)}{D(k)} \tag{5}$$

which results in an estimation of the susceptibility map [Ruetten et al., 2019]. The solution of this problem is ill-posed since at the conical surface in the frequency domain defined by $D(k) = 0$ (or $k^2 - 3k_z^2 = 0$), the inversion which is needed for susceptibility calculation is invalid. This seems to suggest that the resulting phase distribution would lack certain frequency components defined by $k^2 - 3k_z^2 = 0$ [Li et al., 2011].

Methods for QSM calculation

A lot of methods for QSM calculation were already developed, especially to handle the ill-posed dipole inversion. Each method has its own pros and cons due to the way it tries to tackle the problem. COSMOS [Liu et al., 2009] is currently used as the gold standard, it uses data from multiple (at least 3) orientations to correct from the missing data in a single orientation. A major drawback is that it requires multiple scans in one subject, this considerably prolongs the scan time and restricts the clinical feasibility. Consequently, this led to the search for faster alternatives for calculating QSM. Inverse filtering methods such as Truncated K-space Division [Shmueli et al., 2009] and the iterative methods such as iLSQR [Li et al., 2011] and MEDI [Liu et al., 2012] make QSM maps out of single orientation data. Despite the variety of methods, each of these methods suffer from artifacts or have computationally expensive calculations. Since the rise of deep learning, there is an additional way to tackle the inversion problem. The deep learning methods such as xQSM [Gao et al., 2021], AutoQSM [Wei et al., 2019], QSMnet [Yoon et al., 2018], DeepQSM [Bollmann et al., 2019b], SHARQnet [Bollmann et al., 2019a], MoDL-QSM [Feng et al., 2021] and k-QSM [He et al., 2022] try to reconstruct high-quality QSM maps out of single orientation measurements as input. For a comprehensive overview of all methods, I refer to my literature review.

All of the described methods needs preprocessed images as input in which the phase is unwrapped, the background field removed and a mask is applied. Models assume that phase wraps are accurately removed, which does not necessarily have to be the case. Further, they assume the background field is properly removed. The currently available background field corrections perform well given that the regularization parameters are carefully adjusted to a given dataset and the assumptions of the boundary conditions are not violated. Although they perform well, mistakes in regularization parameters can be made and boundary conditions could be violated resulting in errors.

Besides the preprocessing steps, the image reconstruction from raw data also has to be done and errors in this step will propagate further during QSM calculation. Noise in image reconstruction from raw data will also propagate in further steps.

There currently does not exist a method that combines all steps and compute a QSM map directly out of raw data. Despite all developments, we lack an end-to-end method able to make a QSM map direct from k-space. To develop an end-to-end method, steps like de-aliasing, de-noising, phase unwrapping, background field removal and masking all have to be implemented into one model. All these separate steps may each have their own errors. This could lead to error accumulation which will propagate during the QSM-pipeline. For this reason an end-to-end reconstruction pipeline could be promising.

2.2 qRIM

Inverse problems are a broad class of problems which can be encountered in all scientific disciplines. The goal in an inverse problem is to reconstruct a signal from observations that are subject to a known (or inferred) corruption process known as the forward model [Putzky and Willing, 2017]. This is the case for the ill-posed dipole inversion in QSM calculation.

Recurrent Inference Machines (RIM) were first proposed by Putzki and Welling as general inverse problem solvers [Putzky and Willing, 2017]. RIMs can learn inference and model parameters jointly with back-propagation through time. This make them able to learn an iterative inference algorithm without the need to explicitly specify a prior or a particular inference procedure. An RIM is an iterative map which includes the current reconstruction, a hidden memory state, and the gradient of the likelihood term which encodes information about the known generative process and measures how well the measurements are reproduced [Putzky and Willing, 2017].

The quantitative Recurrent Inference Machine (qRIM) [Zhang et al., 2021] is an iterative inverse problem solving network with a unified forward model for joint reconstruction and quantitative mapping, embedded in a RIM. The qRIM unified the reconstruction and parametric models towards a better exploitation of the correlation over the multiple images in a qMRI acquisition. This led to a better optimum of the mapping than the conventional methods with several processing steps. With the unified forward physical model, the convolutional RIM network is adopted to reconstruct directly the parameter map from the subsampled k-space.

2.2.1 Forward model

The unified forward model from the quantitative relaxation parameter map R_2^* to k-space coil signals is composed of a R_2^* estimation part and a reconstruction part. For R_2^* estimation, images of multiple echoes are acquired. With noise neglected, the measured complex-valued image $x_t \in \mathbb{C}^P$ of P voxels at echo t with echo time TE_t follows:

$$x_t = M \odot e^{-TE_t(R_2^* + iB_0)} \tag{6}$$

In which M is the net magnetization, $\odot$ represents the Hadamard product, $R_2^* \in \mathbb{R}^P$ is the parameter map of R_2^* and $B_0 \in \mathbb{R}^P$ is the off-resonance of the static magnetic field. The exponent is taken element-wise. For reconstruction of the image with echo time TE_t from the raw k-space data, the forward model is:

$$y_{t,c} = \mathcal{P}\mathcal{F}S_c x_t + n_{t,c} \tag{7}$$

In which $\mathcal{P}$ is the subsampling mask, $\mathcal{F}$ a Fourier transform, S_c the coil sensitivity maps and $n_{t,c}$ some additive noise. By combining these two equation, a unified forward model can be obtained:

$$y_{t,c} = \mathcal{P}\mathcal{F}S_c(M \odot e^{-TE_t(R_2^* + iB_0)}) + n_{t,c} \tag{8}$$

2.2.2 Problem formulation

With the previously defined unified forward model, the parameters $\Phi = \{R_2^*, M, B_0\}$ were estimated from k-space y. The problem can be solved by the maximum a posteriori (MAP) estimator. This is done by maximizing the sum of the log-likelihood and log-prior distributions of the parameters:

$$\hat{\Phi}_{MAP} = \underset{\Phi}{\operatorname{argmax}}\{logp(\boldsymbol{y}|\boldsymbol{\Phi}) + logp(\boldsymbol{\Phi})\} \tag{9}$$

The (negative) log-likelihood $logp(\boldsymbol{y}|\boldsymbol{\Phi})$ is derived as data consistency between the measured k-space y and the prediction of the estimated parameters:

$$logp(\boldsymbol{y}|\boldsymbol{\Phi}) = \frac{1}{\sigma^2}\sum_{t=1}^{T}\sum_{c=1}^{C}||\mathcal{P}\mathcal{F}S_c(M \odot e^{-TE_t R_2^* + i(\frac{2\pi}{\gamma}\phi_{tissue} - B_{ext})}) - y_{t,c}||_2^2 \tag{10}$$

The log-prior can be seen as a regularizer, acting dependent on how the prior knowledge is modelled. In deep learning based methods it can be the network trained from datasets with similar distributions as the target data [Zhang et al., 2021].

2.2.3 Architecture

The MAP estimation of the parameters as described above will be solved using the Recurrent Inference Machine (RIM). The RIM implicitly learns the log-prior $logp(\mathbf{\Phi})$ with the training data [Zhang et al., 2021]. Parameter estimation of the next step will be iteratively updated using the current estimation and hidden states. Every time step of the qRIM creates an update h to the parameter Φ. For time step $0 \leq \tau \leq T$, the update equation of the qRIM is given by:

$$\Phi_{\tau+1} = \Phi_\tau + h(\Phi_\tau, s_\tau, \nabla_{y|M_\tau}, \nabla_{y|R_\tau}) \tag{11}$$

where $\nabla_{y|M_\tau}$ and $\nabla_{y|R_\tau}$ are gradients of the log-likelihood with respect to M and R respectively. The gradient of the log-likelihood at the point of the current step encodes the data consistency information of the current estimate [Zhang et al., 2021]. The hidden state s_τ is initialized by $s_0 = 0$ and for further time steps calculated by equation:

$$s_{\tau+1} = g(\Phi_\tau, s_\tau, \nabla_{y|M_\tau}, \nabla_{y|R_\tau}) \tag{12}$$

We chose to initialize Φ_0 with the fitted parameter maps using the conventional sequential models of image reconstruction and relaxation [Zhang et al., 2021]

Figure 4: The network architecture of the qRIM [Zhang et al., 2021].

The architecture of the model is built as follows and is represented in figure 4. Each of the first two convolutional layers is followed by the Rectified Linear Unit (ReLU) activation function and interleaved with the gated recurrent units (GRU) which are used to maintain the two internal states $s = \{s_1, s_2\}$. The first convolutional layer is with 8 input channels (4 for the parameter map estimation, and another 4 for the gradients of the log-likelihood for the estimation), 128 output channels, and kernel size of 5×5, whereas the second convolutional layer is with 128 input channels, 128 output channels, and kernel size of 3×3. Each GRU is with $2 \times (128 \times 384 + 384)$ parameters. The second GRU is followed by a third convolutional layer with 128 input channels, 4 output channels, kernel size of 3×3, and without bias terms. In total the network is with 376,064 parameters.

3 Methods

3.1 Data

3.1.1 Subjects

Data is used from the Prediva Ouderen Extensie (POE), a population study of elderly. This data set consist of data from 27 healthy elderly (between 80 and 91 years old) . All participants gave written informed consent prior to inclusion in this study.

Both the AHEAD data and the POE data were acquired in the Spinoza Centre for Neuroimaging in Amsterdam. Volunteers were scanned at a 7 T scanner (Philips) containing a two-channel transmit and a 32-channel receive head coil.

3.1.2 Sequence

The data were acquired with the MP2RAGEME (Magnetization Prepared 2 Rapid Acquisition Gradient Echoes Multi Echo) sequence [Caan et al., 2019], which allows simultaneous T_1, $T_2^\star$ and QSM mapping.

The MP2RAGEME sequence is an extended MP2RAGE sequence (which has a single-echo acquisition in both inversions) and is represented in figure 5. T_1-maps can be obtained using the MP2RAGE sequence,
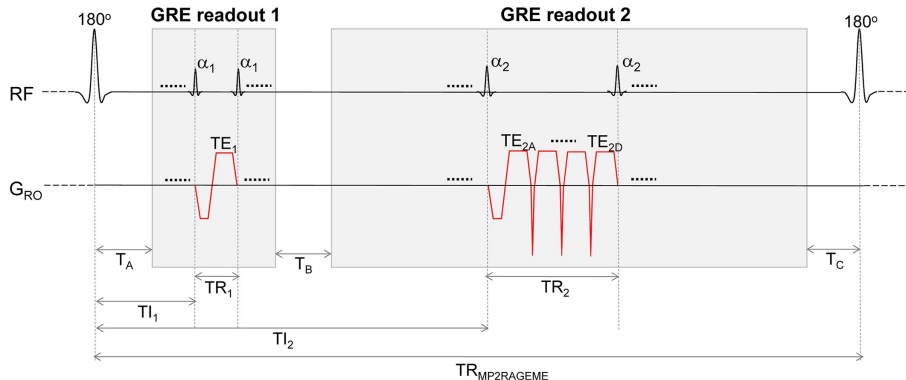

Figure 5: MP2RAGEME sequence, with the second inversion INV_2 extended to a multi-echo (ME) readout. The four echoes of INV_2 have TE's TE_{2A-D}. The TR_{GRE} differs between the first and second inversion images, in contrast with the MP2RAGE sequence [Caan et al., 2019]

which is the first part of the MP2RAGE-ME sequence. The MP2RAGE sequence is extended with a multiple echoes acquisition in the second inversion while maintaining single-echo acquisition in the first inversion. Multiple gradient echoes were acquired in the second inversion, which permits the calculation of T_2^*-maps and susceptibility maps.

The sequence parameters of the two GRE readout blocks were: Field of View (FOV) 207 mm × 205 mm × 158 mm, 225 sagittal slices, voxel size 0.7 mm isotropic, fold-over direction AP and LR. TR_{slice} = 6.8 s, inversion times $TI_1 = 0.67s$ and $TI_2 = 3.68s$, repetition times $TR_1 = 6.2ms$ and $TR_2 = 31ms$, first echo time $TE = 3.0ms$ for both readouts, for the second readout four echoes were acquired with $\Delta TE = 8.5ms$, flip angle 4° for both readouts.

The frequency encoding direction (x) was over sampled by a factor of 2 for both AHEAD and POE. However, the acceleration methods for AHEAD and POE did differ from each other. The acquisition time of the AHEAD data was accelerated by making use of GRAPPA. An acceleration factor of 2 was applied in the first phase encoding direction (y). For the POE dataset, scanning parameters were further optimized to reduce the burden on participants. The acquisition time was accelerated by making use of SENSE, the frequency encoding direction (x) was over sampled by a factor of 2 and the phase encoding directions (y and z) were under sampled by a factor of respectively 1.8 and 1.5. The under sampling factors of 1.5 and 1.8 led to an acceleration factor of $1.5 \times 1.8 = 2.7$ which was more than the acceleration of factor of 2 used for the AHEAD study. The use of these non-integer acceleration factors places special demands on the reconstruction which will be discussed in the next subsection.

3.2 Reconstruction

The acquired raw MRI data was loaded and preprocessed with MRecon. The raw SENSE k-space data was under sampled by a factor of 1.8 and 1.5 in respectively the y- and z-direction. Data were Fourier-shifted with half field-of-view in the phase encoding directions to comply with the functions of the Berkeley Advanced Reconstruction Toolbox (BART). To make this align, a Fourier shift with half FOV was applied over the phase encoding directions. This gave us images in Fourier and image domain in the right convention for BART.

Besides the under sampled raw k-space data, fully sampled low resolution coil data was acquired. The sensitivity maps were made from the coil data by running the Caldir-function of the BART toolbox (version 5.0.0) in Matlab. The Caldir-function estimated coil sensitivities from the k-space center, with a range of 90 pixels. After these sensitivity maps were made, the coil sensitivities were interpolated to the right size in image space.

The SENSE reconstruction was performed by making use of the Michigan Image Reconstruction Toolbox (MIRT). The MIRT toolbox [Fessler, 2021] is a collection of open source algorithms for image reconstruction written in Mathwork's Matlab language. By making use of the MIRT toolbox, it was possible to unfold a non-integer SENSE accelerated aliased image (see supplementary figure 38). Since the data acquisition was accelerated by non-integer factors, the acquired data points were not exactly on a grid. The image geometry had to be chosen in the same way the scan was made, with 163 and 150 trajectories in respectively y and z and with an isotropic voxel size of 0.7 mm. After the image geometry was chosen,

the sampling trajectory needed to be defined with acceleration factors of 1.8 and 1.5 in respectively y and z. A Gnufft class object was created for performing a NUFFT and different reconstruction parameters were set, the regularisation parameter was set to 1 and 10 iterations were used.

After all parameters were set and the data was in the right form and sizes, we looped over the fully sampled frequency encoding direction (x) and made slice-by-slice reconstructions for each under sampled y-z plane.

3.3 Segmentation pipeline

The data was acquired with the MP2RAGEME sequence which allowed us to make quantitative T_1, T_2^* and QSM-maps. In the segmentation pipeline we used some functions of the Nighres package, a Python package for processing of high-resolution (7T and higher) neuroimaging data [Huntenburg et al., 2018]. First, we denoised the images with a local complex-valued principal component analysis based method (`lcpca_denoising`-function in Nighres). The T_1-maps were computed using the `mp2rage_t1_mapping`-function of Nighres which computes T_1-maps by using a look-up table [P.Marques et al., 2010]. T_2^*-maps were computed by least-squares fitting in log space of the exponential signal decay over the multi-echo images of the second inversion (`flash_t2s_fitting`-function in Nighres). For QSM, phase maps were pre-processed using iHARPERELLA (integrated phase unwrapping and background phase removal using the Laplacian) of which the QSM images were computed using LSQR [Li et al., 2011]. For skull stripping we used the first echo of the second inversion. All skull information was removed, resulting in a separate image through creation of a binary mask using brain extraction protocol (BET) [Alkemade et al., 2020].

Image segmentation was performed on the T_1, T_2^* and QSM-maps by the `fuzzy_cmeans`-function of Nighres, which estimates intensity clusters with spatial smoothness and partial voluming. Image registration was done by the `embedded_antsreg_multi`-function, which runs the rigid and/or Symmetric Normalization (SyN) algorithm of ANTs and formats the output deformations into voxel coordinate mappings. This used successively rigid, affine, and non-linear transformations with high levels of regularization and mutual information as cost function [Alkemade et al., 2020]. Afterwards, the individual multimodal maps were co-registered a second time, now to the AHEAD template, using all three quantitative contrasts. Once individual images were co-registered, the corresponding delineation masks were transformed to match the AHEAD template using linear interpolation [Alkemade et al., 2020]. Finally, we used a Multi-contrast Anatomical Subcortical Structures Parcellation (MASSP) algorithm to parcellate the subcortex (`massp`-function in Nighres).

3.4 Conventional QSM calculation

Conventional QSM calculation was performed by making use of the algorithms in the STI suite package. The original phase images of the four echoes are visible in figure 20 and contain phase wraps and contributions from the background.
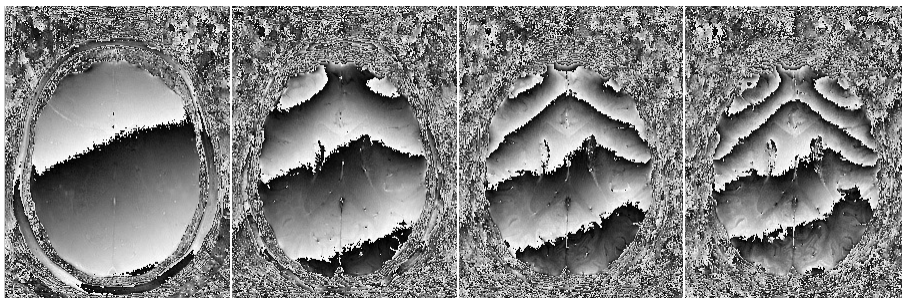

Figure 6: The phase reconstructions of a slice for 4 different echoes

Integrated phase unwrapping and background phase removal can be done using the Laplacian. The Laplacian can be derived from the wrapped phase directly by making use of sine and cosine functions as described in the equation below:

$$\nabla^2\phi = cos\phi\nabla^2 sin\phi - sin\phi\nabla^2 cos\phi \tag{13}$$

The Laplacian is used to calculate the phase only originated from the tissue susceptibility differences. Solving the equation written above yields the unwrapped phase that is free of contributions from sources outside the FOV.

11

The relationship between the Laplacian of the phase image and the underlying magnetic susceptibility distribution is given by [Li et al., 2014]:

$$\nabla^2 \phi = \gamma \cdot TE \cdot \mu_0 H_0 (\nabla^2 \chi / 3 - \partial^2 \chi / \partial z^2) \tag{14}$$

in which γ is the gyromagnetic ratio, TE the echo time, μ_0 the magnetic permeability in vacuum and H_0 is the magnetic field strength. In this equation the local elements $(\nabla^2/3 - \partial^2/\partial z^2)\chi$ can be regarded as sources that generate the tissue phase obeying the principle of superposition [Li et al., 2014]. Solving this equation gives the susceptibility map.

The Laplacian equation and the previous equation can be solved in spatial frequency domain if the phase is available everywhere within the FOV. This approach is fast as it takes advantage of the Fast Fourier Transform (FFT) algorithm. However phase measurements are typically not available outside the tissue. To mitigate this, the equations can be solved with boundary conditions set at the tissue boundaries, but as a consequence the FFT algorithms can no longer be applied. Defining those boundary conditions is hard and solving the partial differential equations would be computationally intensive.

To take advantage of the FFT algorithm, the phase outside the tissue has to be determined. The spherical mean value property of harmonic functions can be applied to estimate the phase Laplacian outside the FOV (similarly as is done in SHARP). For this we need to define the interior (I) and boundary (O) regions of the tissue and the external surrounding part (E) in the FOV. Boundary region O is the set of tissue voxels next to the boundaries that are within a distance of the radius of the spherical mean value filter. The phase Laplacian within the region of E shall be the solution of the following minimization problem:

$$\min_{\nabla^2 \phi_E} ||S(\nabla^2 \phi_E) + S(\nabla^2 \phi_{I+O}) - \overline{S(\nabla^2 \phi_I)}||_2 \tag{15}$$

where S represents the spherical mean operator. This equation states that when all background sources are removed, the only sources of phase reside in the trustable region. Once the phase outside the tissue ϕ_E is determined, the phase in the whole FOV can be calculated as follows:

$$\nabla^2 \phi_{FOV} = \nabla^2 \phi_{I+O} + \nabla^2 \phi_E \tag{16}$$

The background removed and unwrapped phase can be obtained using the following FFT-based inverse Laplacian:

$$\phi = -FT^{-1}[FT(\nabla^2 \phi_{FOV})] \tag{17}$$

in which FT represents the Fourier transform. This method is referred to as HARPERELLA which stand for HARmonic PhasE REmoval with LapLAcian. It is emphasized that HARPERELLA achieves both phase unwrapping and background phase removal in a single integrated procedure purely based on the Laplacian operator [Li et al., 2014]. The STI suite package in Matlab contains the HARPERELLA functions and other functions for QSM reconstruction which are used in order to generate QSM maps of the different echoes. The echo time of the first echo is very short, this means that there will be a relatively low SNR. This is the reason that the first echo is excluded from and the susceptibility is calculated for the second, third and fourth echo. This will give us 3 susceptibility images as shown in figure 7.
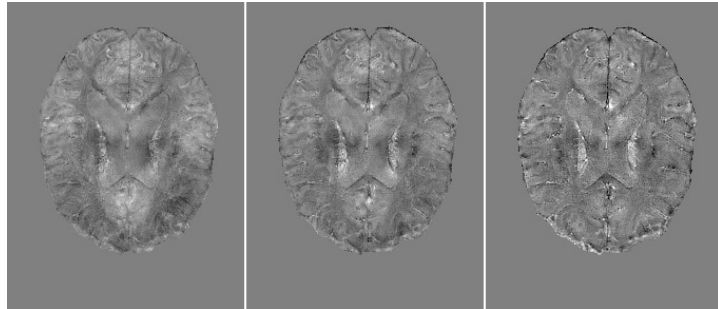

Figure 7: The susceptibility images of echo 2, 3 and 4.

The images of the different echo-times are combined by taking the mean over the three images. This will result in a final QSM map as in figure 8.
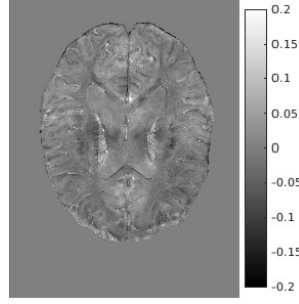
12

Figure 8: A final QSM map.

3.5 qRIM

3.5.1 Forward Model

Regarding our QSM problem, we modified the forward model as defined by Zhang et al. in equation 8.

$$y_{t,c} = \mathcal{PFS}_c(M \odot e^{-TE_t(R_2^* + iB_0)}) + n_{t,c} \tag{18}$$

We were interested in the changes in the phase over the echo times. This information was found in the imaginary part of the exponential $TE_t B_0$ and represented the perturbations in the phase. These perturbations consisted of changes due to external field inhomogeneities B_{ext} and changes due to local field perturbations. We assumed the changes due to local field perturbations originated from differences in susceptibility in the tissue and referred to it as B_{int}. In order to make the forward model suitable for QSM calculations, we replaced the imaginary part of the exponential $TE_t B_0$ by $B_{ext} + B_{int}$ and the forward model looks like:

$$y_{t,c} = \mathcal{PFS}_c(M \odot e^{-TE_t R_2^* - i(B_{ext} + B_{int})}) + n_{t,c} \tag{19}$$

The step from a B_{int} map to a tissue phase map ϕ_{tissue} was easily made since there was a linear relationship between the two $B_{int} = -\frac{2\pi}{\gamma}\phi_{tissue}$. This resulted in the following unified forward model;

$$y_{t,c} = \mathcal{PFS}_c(M \odot e^{-TE_t R_2^* + i(\frac{2\pi}{\gamma}\phi_{tissue} - B_{ext})}) + n_{t,c} \tag{20}$$

This unified forward model facilitated mapping from our parameters of interest, R_2^*, M and ϕ_{tissue} to k-space signals $y_{c,t}$.

3.5.2 Problem formulation

With the previously defined unified forward model, we aimed to estimate the parameters $\Phi = \{R_2^*, M, \phi_{tissue}, B_{int}\}$ from k-space y. The problem can be solved by the maximum a posteriori (MAP) estimator. This is done by maximizing the sum of the log-likelihood and log-prior distributions of the parameters in the same way as descibed by Zhang et al. [Zhang et al., 2021]:

$$\hat{\Phi}_{MAP} = \underset{\Phi}{\operatorname{argmax}}\{logp(\boldsymbol{y}|\boldsymbol{\Phi}) + logp(\boldsymbol{\Phi})\} \tag{21}$$

The (negative) log-likelihood $logp(\boldsymbol{y}|\boldsymbol{\Phi})$ is derived as data consistency between the measured k-space y and the prediction of the estimated parameters:

$$logp(\boldsymbol{y}|\boldsymbol{\Phi}) = \frac{1}{\sigma^2}\sum_{t=1}^{T}\sum_{c=1}^{C}||\mathcal{PFS}_c(M \odot e^{-TE_t R_2^* + i(\frac{2\pi}{\gamma}\phi_{tissue} - B_{ext})}) - y_{t,c}||_2^2 \tag{22}$$

The log-prior can be seen as a regularizer, acting dependent on how the prior knowledge is modelled. In deep learning based methods it can be the network trained from datasets with similar distributions as the target data [Zhang et al., 2021].

4 Results

The results section is divided into two parts: in subsection 4.1 we extend the AHEAD dataset with the quantitative data of the POE data and in subsection 4.2 we will show the results of the qRIM.

4.1 POE

The results of present study using POE data were plotted alongside the results of the AHEAD dataset in each figure for direct comparison. Alkemade et al. [Alkemade et al., 2020] have shown the relationships between age and quantitative values like R_1, R_2^* and QSM. These relationships were calculated for the inter quartile range (IQR) and the median values of the quantitative values. As an addition, the median of the volume is given with respect to age. We focused on three regions of interest which may be interesting regarding healthy ageing. The regions of interest were the striatum, red nuclei and substantia nigra. We will discuss the results of the different nuclei in the next 3 subsections.

4.1.1 Striatum

The results of the striatum are added to the AHEAD data together with the model based on the AHEAD data and displayed in figure 9. The volume and the median R_1 and R_2^* values of the POE database are
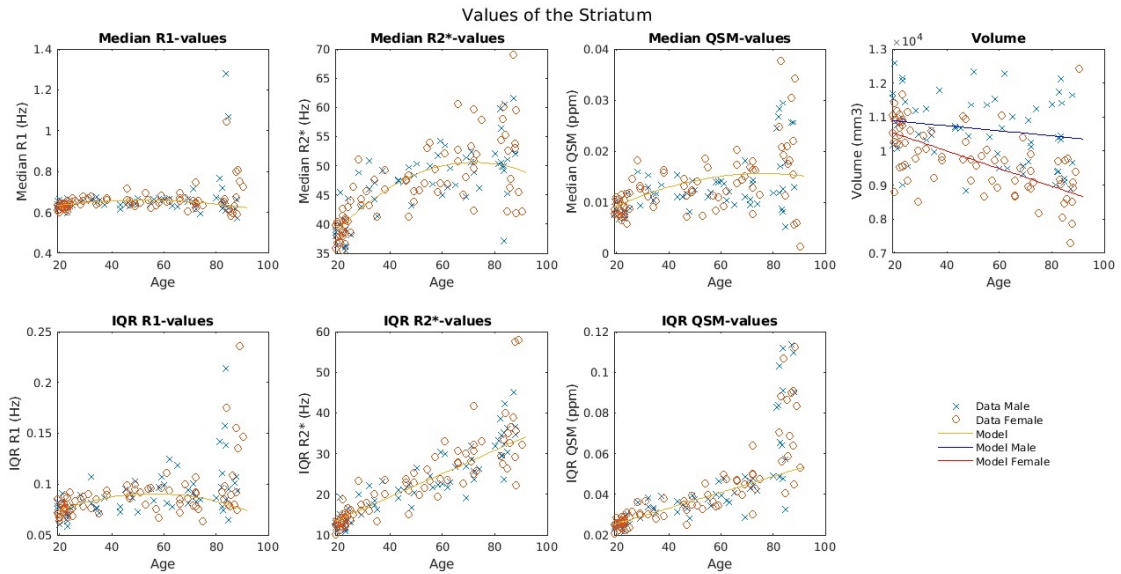
Figure 9: The calculated quantitative values of the POE data from the striatum were added to the already existing values from the AHEAD data with their calculated models.

in line with the AHEAD data and the model based on the AHEAD data. Regarding the IQR of the R_1 and QSM values, there is more spread in the values of the POE data. The QSM values of the POE data have also more spread in the median values. We made histograms of all subjects to determine whether the segmentations were accurate (see supplementary figures 26, 27 and 28). We expected the histograms to have a Gaussian shape with one peak. If the histograms showed a deviant shape, we excluded the data. We plotted the remaining data and made a model based on the information of the AHEAD data and the POE data combined in figure 35.

4.1.2 Substantia Nigra

The results of the substantia nigra are added to the AHEAD data together with the model based on the AHEAD data and displayed in figure 11. The median R_1 and R_2^* values of the POE database seem to be in line with the AHEAD data and the model based on the AHEAD data. The volume seems to become smaller than expected by the AHEAD model. The IQR and median of the QSM-values show different results. We made histograms of all subjects to determine whether the segmentations were accurate (see supplementary figures 29, 30 and 31). We expected the histograms to have a Gaussian shape with one peak. If the histograms showed a deviant shape, we excluded the data. We plotted the remaining data and made a model based on the information of the AHEAD data and the POE data combined in figure 36.
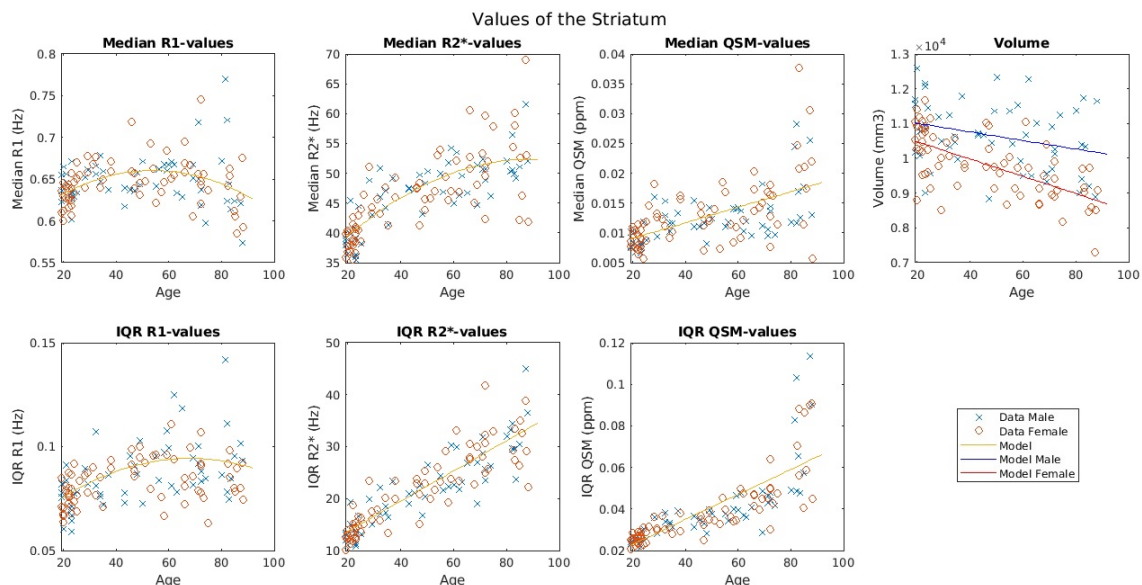
Figure 10: After excluding probably wrong segmentations, the the models for the quantitative values over age were calculated again, based on the AHEAD data and the POE data.
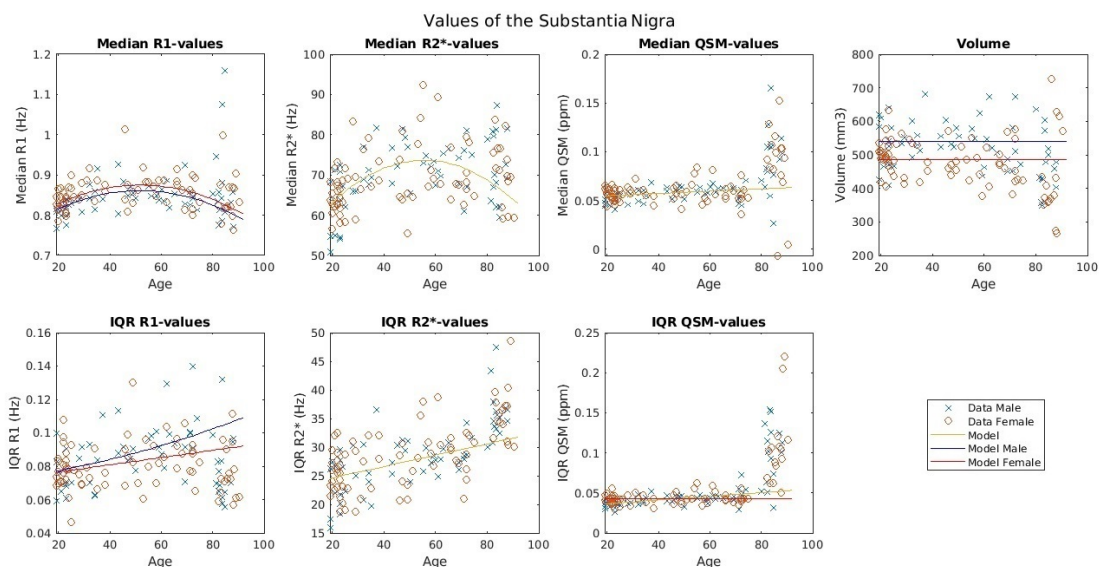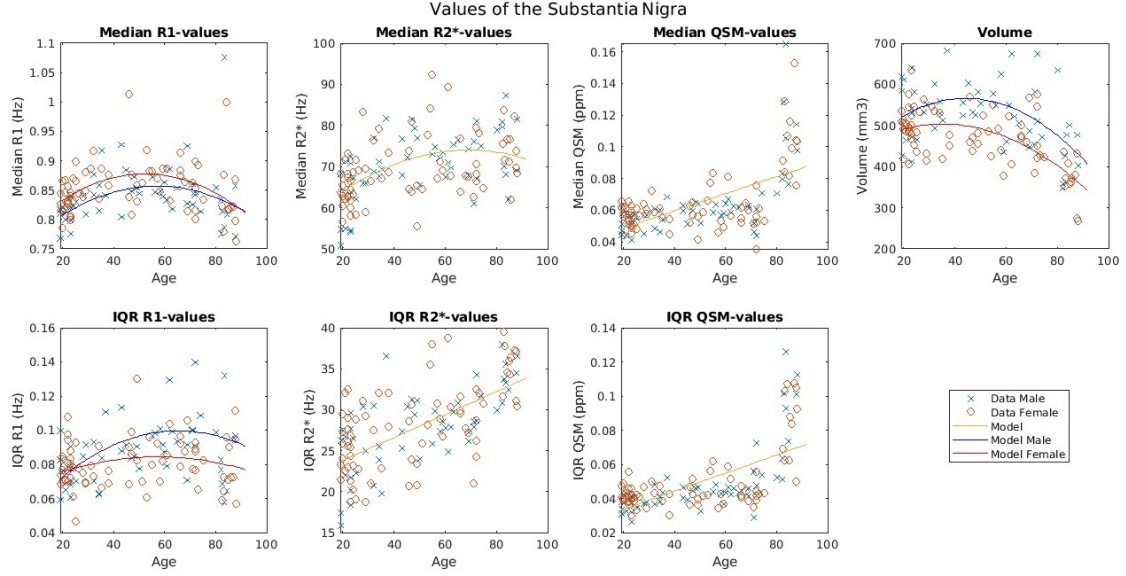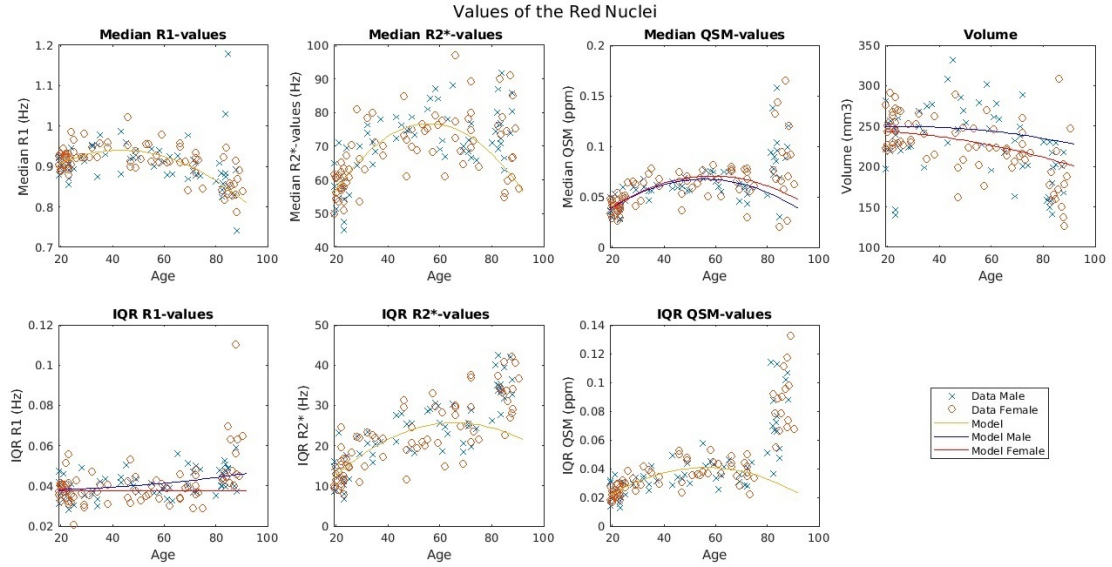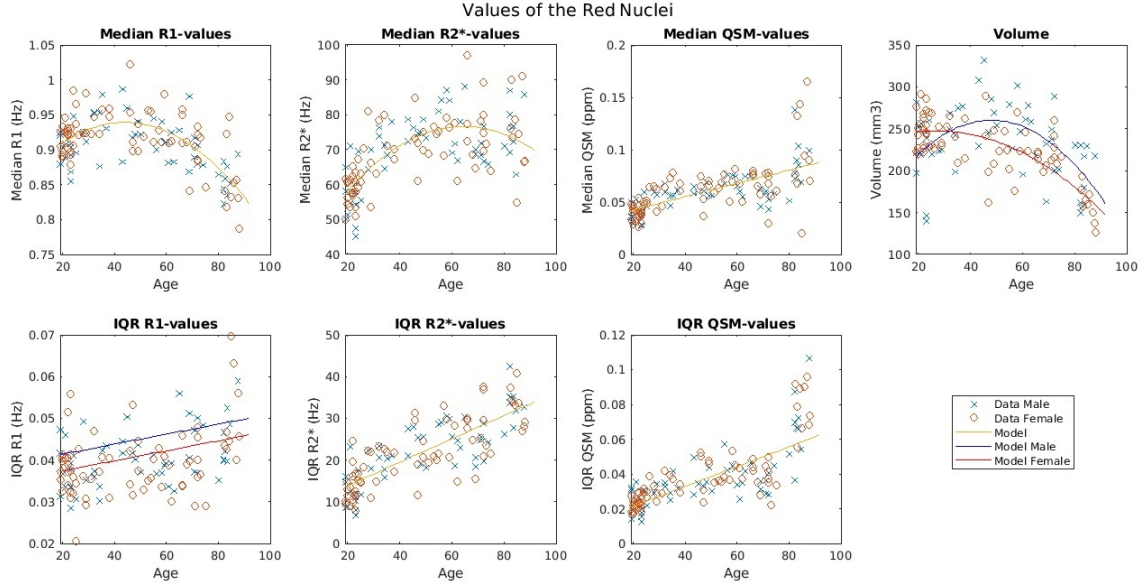

Figure 11: The calculated quantitative values of the POE data from the substantia nigra were added to the already existing values from the AHEAD data with their calculated models.

4.1.3 Red Nuclei

The results of the red nuclei are added to the AHEAD data together with the model based on the AHEAD data and displayed in figure 13. The median and the IQR of the R_1 values of the POE database are in line with the model based on the AHEAD data, but do contain more outliers than for the AHEAD database. The median and IQR if the R_2^* values do not deviate much from the data of the AHEAD. However, the values do deviate from the models based on the AHEAD dataset. The median values of the QSM contain a wide range of values and the IQR is also a lot higher than expected. Regarding the volume of the segmentations, the volumes correspond to the other data points, but are lower than expected by the AHEAD model. We made histograms of all subjects to determine whether the segmentations were accurate (see supplementary figures 32, 33 and 34). We expected the histograms to have a gaussian shape with one peak. If the histograms showed a deviant shape, we excluded the data. We plotted the remaining

Figure 12: After excluding probably wrong segmentations, the the models for the quantitative values over age were calculated again, based on the AHEAD data and the POE data.


Figure 13: The calculated quantitative values of the POE data from the red nuclei were added to the already existing values from the AHEAD data with their calculated models.

data and made a model based on the information of the AHEAD data and the POE data combined in figure 37.

4.2 qRIM

We built upon a previously developed qRIM [Zhang et al., 2021] to do QSM. The data structure of the POE-data was modified to suit a pre-trained qRIM provided by the authors. The initial estimations are shown in figure 15. The pre-trained model did not improve on the initial estimations, as can be seen in figure 16.

The R_2^*-model [Zhang et al., 2021] was able to calculate a B_0-map which represented the differences in the phase. Out of this B_0-map we wanted to extract a tissue phase map. Our aim was to extract the tissue contribution from the B_0 map, which would provide us with B_int. This estimate could be used to

16

Figure 14: After excluding probably wrong segmentations, the the models for the quantitative values over age were calculated again, based on the AHEAD data and the POE data.
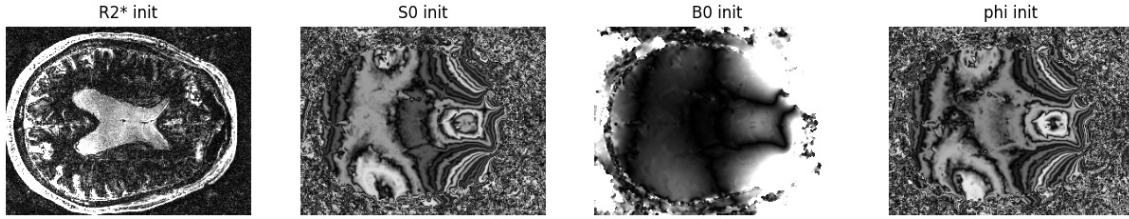
Figure 15: The initialization images R_2^*, S_0, B_0 and ϕ_0 created by the model of Zhang et al.

create an initialization image since B_int and tissue phase are linearly related. Since the pre-trained model did not work on our data, we used a reconstruction on the data included in the original work to generate a possible B_{int} initialization image. The B_0-map contains contributions of the B_{int} with high-frequency changes and contribution of the low-frequency varying B_{ext}. We assumed that this low-frequency B_{ext}-field can be seen as some sort of biasfield. For this reason, we used a bias field correction to generate an initial estimate for B_int. We used the N4 Bias field correction [Tustison et al., 2010] from the SimpleITK toolkit [Beare et al., 2018]. The images before and after biasfield correction are shown in figure 17.

5 Discussion

5.1 POE

General

The POE results show similarities and differences compared to the expected trend based on the AHEAD study. The volumes for all nuclei match the trend and show no remarkable outliers or deviations. This indicates that the segmentations have a correct size. All three nuclei showed an age-related decline of their volume. Age-related volumetric decline has previously been hypothesized to be the result of loss of neuronal bodies, decreases in dendritic arborization, synaptic pruning, refinement of circuits, and/or loss of connections [Tullo et al., 2019].

Besides the volume, the medians of the R_1-values roughly correspond to the model based on the AHEAD study. All nuclei show a decrease in R_1-values above a certain age. Previous work observed an age-related increase of T_1-values (so a decrease in R_1-values) across the subcortical areas, most likely indicating a loss of myelin [Keuken et al., 2017]. This myelin loss could exlplain our age-related decrease in R_1 as well.
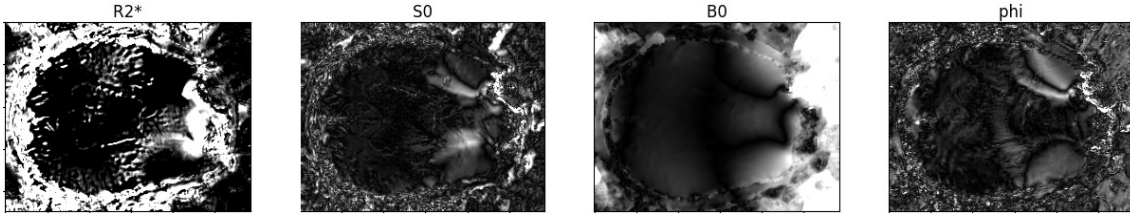
17

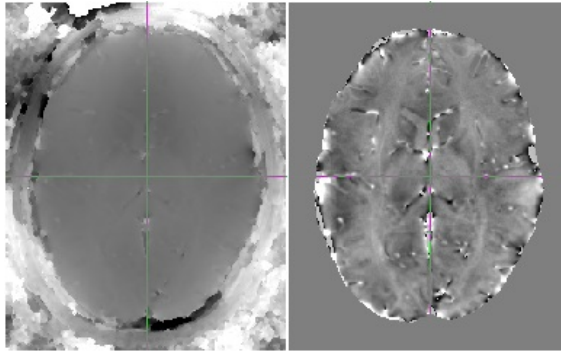Figure 16: The reconstruction made by the qRIM of Zhang et al.

Figure 17: On the left is shown the B_0-map as reconstructed by the model of Zhang et al. and on the right is a possible initialization image of B_int after applying a biasfield correction. The scale has not yet been determined, it is calculated as phase term in the model with internal scaling, but still needs to be calibrated to make it in (m)T.

The medians of the R_2^*-values also correspond to the previously made models based on the AHEAD data. Our results show an age-related increase in R_2^*- and QSM-values for the striatum, substantia nigra and the red nuclei. A decrease of T_2^* combined with an increase of QSM is indicating an increase of iron [Keuken et al., 2017]. This means that iron accumulation in these brain regions could explain these age-related changes. We will discuss each nuclei separately in the next sections.

Striatum

We made histograms because some outliers led to the suspicion that some segmentations still contain multiple structures. The histograms of the R_1, R_2 and QSM-values of the striatum can be found in the supplementary materials in figure 26, 27 and 28. For 8 subjects it can be seen that there exist two peaks in the histograms of the R_1-values. This probably means that a surrounding structure is uncorrectly labeled as striatum. For almost all of these subjects, deviations from the expected Gaussian shape are also visible in the histograms of the R_2 and QSM values.

We excluded these subjects to reduce the effect of potentially incorrect segmentations and models were made based on the AHEAD and the POE dataset. Figure 35 in the supplementary materials shows how the model relate to eachother. The models for the volume, R_1 and R_2^*-values (both median and IQR) correspond very well to the models of the AHEAD study. The volume of the striatum decreases linearly with age, with the age-related shrinkage being greater in females. This is as expected, since previous studies showed a similar age- and gender-related dependence [Brabec et al., 2003] [Keuken et al., 2017]. The median of the R_1-values of the striatum showed a decrease. A decrease of R_1 corresponds to an increase of T_1, as previous research has shown [Keuken et al., 2017]. An increase in T_1-values is thought to reflect a decrease of myelinisation [Callaghan et al., 2014]. The R_2^*-values showed an increase as a function of age. Previous work also showed a significant increase in effective transverse relaxation rate (R_2^*) as a function of age due to iron accumulation [Callaghan et al., 2014]. Iron accumulation turned out to be a predictor for the volumetric shrinkage of the striatum and the amount of accumulated iron is a risk factor for neural and cognitive decline in normal aging [Daugherty et al., 2015]. The accumulation of iron also has an effect on the QSM map since iron causes an increase of magnetic sucseptibility. The median values of the QSM show an increase for the POE data, but there are differences. The median and IQR of the QSM-values do differ from the AHEAD data, since there is much more spread in the data and the median values appear in a broader range. This causes that our model fit is influenced by these values and

differs from the AHEAD model. Possible causes of this wider spread will be discussed in the subsection Discrepancies with respect to AHEAD.

Substantia nigra

We made histograms because some outliers led to the suspicion that some segmentations still contain multiple structures. The histograms of the R_1, R_2 and QSM-values of the substantia nigra can be found in the supplementary materials in figure 29, 30 and 31. The histograms of the R_1 values show an expected shape, all the histograms are narrow Gaussian-shaped peaks. Some R_2^* histograms have a shape that differs somehow from the Gaussian shape. The QSM histograms contain two peaks for 8 subjects, which may be due to a wrong segmentation. All subjects with a different shape of the R_2^* histograms also have a different shape of QSM histogram.

We excluded these subjects to reduce the effect of potentially incorrect segmentations after visually checking that the outliers indeed contain multiple structures. We made models based on the AHEAD and the POE dataset (figure 12). Figure 36 in the supplementary materials shows how our model relates to AHEAD-model. First note that our volumes are smaller than expected with the constant model of the AHEAD study. This makes that for the total data a quadratic fit which decreases with age is more optimal. Previous work has shown that from the age around 45 years, neurons decline by about 1% per year [Cabello et al., 2002]. Which could be a possible cause of the decline in the volume of the elderly.

The median of the R_1 values of the substantia nigra showed a decrease in age corresponding to the AHEAD model. This is possibly due to a decrease of myelinisation causing an increase in T_1 values [Callaghan et al., 2014]. The median of the R_2^*-values are a little higher than expected by the AHEAD model, so are the IQR-values. Our results show an increase in R_2^* with age, possibly due to iron accumulation [Callaghan et al., 2014]. While slight iron accumulation seems to appear in healthy elderly, abnormal iron accumulation around the substantia nigra is a diagnostic indicator of Parkinsonism [Lee et al., 2021]. Iron accumulation also causes an increase in QSM-value, which can be seen in the data. Although we see an increase in the median of the QSM-values, there are some differences. The median and IQR of the QSM-values do differ from the AHEAD data, which makes our model differ from the AHEAD model. The differences in IQR may be partly due to the smaller volume, since segmented volumes are smaller for the POE data, it can lead to abnormalities in the IQR more quickly. More possible reasons for the deviations for the QSM values will be discussed in the subsection Discrepancies with respect to AHEAD.

Red nuclei

We made histograms because some outliers led to the suspicion that some segmentations still contain multiple structures. The histograms of the R_1, R_2 and QSM-values of the red nuclei can be found in the supplementary materials in figures 32, 33 and 34. The histograms of the R_1 and R_2 values have a good shape, for 4 subjects the histogram of the R_2^* seems to have a small additional peak. The QSM histogram confirm that the segmentation were not completely accurate, since for 9 subjects a second peak is visible in the QSM histogram. This peak does contain some negative values, which is not as expected since iron accumulations will cause positive QSM values in the red nuclei.

We excluded these subjects to reduce the effect of potentially incorrect segmentations after visually checking that the outliers indeed contain multiple structures. We made models based on the AHEAD and the POE dataset (figure 12). Figure 37 in the supplementary materials shows how the models relate to each other. First note that our volumes are smaller than expected with the constant model of the AHEAD study. This makes that for the total data a steeper (compared to the AHEAD model) quadratic fit is more optimal. The median R_1-values are almost exactly as expected by the AHEAD model. The median R_1-values decay with age, possibly indicating a loss of myelin or incomplete remyelination [Keuken et al., 2017]. The median R_2^* values are slightly higher than expected by the AHEAD model, our results do not show a sharp decline in the extreme age group of the POE. This is not remarkable, as this extreme age group is missing from the AHEAD data. In addition to the increase in R_2^* values with aging, there is also an age-related increase in QSM values. A decreased T_2^* and increased magnetic susceptibility is converging evidence for a potential increase of iron concentration with age within the red nuclei, reported in a previous study [Keuken et al., 2017]. The median QSM-values were higher than expected which results in a different model. Regarding the IQR of the quantitative values, the IQR of the R_2^* and the QSM-values turn out to be higher than expected by the AHEAD model, which will be discussed in the next subsection. By adding the POE data, IQR of the R_2^*-values show a linearly increasing relation instead of a quadratic relation based on the AHEAD study. This may be caused by iron accumulation in the red nuclei during lifespan.

Discrepancies with respect to AHEAD

A remarkable difference between the POE and AHEAD data is in the spread. Overall, the IQR of the POE data is considerably higher than seen in the AHEAD data. This can be caused by several reasons. The used pipeline for the calculation of the quantitative maps and the segmentation is the same for the POE as for the AHEAD data. As a result, we do not expect any differences through the segmentation pipeline and we can exclude this as a possible cause.

A possible cause of the higher IQR can be the motion correction. A motion correction was performed during the reconstruction of the POE data, while this was not the case for the AHEAD data. The estimated motion parameters of the POE data can be found in supplementary figure 25. The amount of motion in mm and rotation in degrees are represented over time. Almost all elderly moved less than 3 mm in a 14-minute scan. Although this seems small, a shift of 2.1 mm already represents a shift of 3 voxels since we scanned at ultra-high resolution.

This motion correction may lead to a sharper image and make the process more sensitive for high and low values which will be blurred without the motion correction. The motion correction will not be very important for segmentation purposes, since it is in the level of details. However, this mechanism could cause more spread in IQR values. We visually compared our reconstructions before and after motion correction. Despite there being no major difference, the motion-corrected image showed more fine detail and looked slightly sharper. Although previous reports showed limited motion artifacts for the AHEAD study, which did not result in noticeable image degradation [Alkemade et al., 2020]. The fact that the AHEAD study did not use any motion correction could lead to differences with the POE dataset. We did not have insight to the motion parameter of the AHEAD study to compare the amount of motion in the AHEAD study with the POE study. Motion correction may especially be important for detection of micro structures like microlesions [Abdel-Fahim et al., 2014].

Second, the higher IQR can be caused by the reconstruction method. The AHEAD data used the reconstructed Nifti-files from the scanner as input for the pipeline. Since the raw POE data was not yet reconstructed, we first had to reconstruct the raw data before putting it into the pipeline. The reconstruction from the scanner is visually optimized and uses unknown filters to perfect the reconstructed image. Since these are unknown, we will get a reconstruction for our POE data without these additional filters, while the AHEAD data is visually optimized by these filters. A possible explanation of the higher IQR for the POE is that the filters have a smoothing effect on the data or have some kind of additional noise suppression. A smoother image with less noise and smoothed should contain fewer peaks and valleys, which could result in a lower IQR for the AHEAD data.

A final reason for a difference in IQR could be that it is harder to perform the segmentations on the extreme group of elderly. The atlas we used [Alkemade et al., 2020] was developed for an average adult brain. It might not represent the anatomy of this elderly cohort, for example due to atrophy. This could lead to mismatches in the segmentations. We looked into the volume and the histograms of the values to get an estimate if the segmentations were performed well. An unexpected deviation in volume could mean that the segmentation is too small or too large. This means that surrounding structures were given the wrong label, and values of other structures are incorrect taken into a different structure. We looked into the histogram values of the structures to see if the shape is as expected. We expected a Gaussian shape in the values. If the histograms contain two peaks, this might indicate that other structures are allocated to the wrong structure.

5.2 qRIM

The pre-trained qRIM failed to make good reconstructions on our data, while it performed excellently in previous research [Zhang et al., 2021]. This may be due to the use of a different dataset. The initialization images of R_2^* and B_0 appear to be sufficient, while the S_0 and ϕ are not as expected. It seems that phase wraps appear over the image. The torch-functions in Python did not yield the expected results, so probably torch expects the data in a different format which means the POE data needs to be preprocessed differently.

For all maps, R_2^*, S_0, B_0 and ϕ, the reconstructions do not look as expected. This may be caused by incorrect image initialization. Remarkably, despite the initialization images of R_2^* and B_0 look as expected, the reconstructions of the qRIM do not look correct. Possibly, data scaling and normalization used for the pre-training was different from the one used in our dataset. We suggest to retrain the qRIM or to convert it to a more generalized model.

5.3 Limitations

This study has some limitations regarding the data. The data of the AHEAD is biased for the age range between 18 and 30 years. Every decade contains 12 or 13 participants, while the first decade contains 42

participants. This will affect the models because the subjects are not evenly spread over the ages.

The POE data were acquired from healthy elderly. The dataset could be biased due to (healthy) reaching old age. Not everyone will reach the age of 90 years at all, especially not healthy. These elderly were fit enough to participate in a population study, which is not self-evident for people of that age. There is a plausible chance that participants will be skewed towards a particular subpopulation. For example, people with on average more or less iron accumulation in some deep brain nuclei, slowing down the ageing process. If these people may have slower brain aging due to underlying mechanisms, it may result in brains that are exceptionally good for their age (for example, brains that look 10 years younger). This may mean that the subjects do not have an average brain for their age category. At the moment little is known about healthy aging. Despite this data may contribute to the current knowlegde of the healthy elderly brain, more knowledge should be gained about it in the future.

The used POE data were acquired in a single orientation. For the QSM reconstructions, we used the iterative least squares method and the functions of HARPERELLA. Although the QSM maps calculated by this method are better than methods such as truncated dipole inversion, this is not an optimal solution regarding QSM maps. In order to optimize a network, one has to train the network in the future with high quality QSM map as can be achieved with COSMOS, which is currently seen as the gold standard. Multiple orientation data were not available of the POE data, so this was the highest achievable for this dataset.

The POE data was accelerated by SENSE and under sampled by non-integer acceleration factors in the phase encoding directions. The non-uniform fast Fourier transform is used to properly reconstruct the images. With this method, there is no intermediate step in which coil images can be reconstructed. The qRIM is an inverse problem solver that solves by filling the gaps in k-space. This means that a k-space mask had to be defined in order to let the qRIM solve it. Besides, the k-space needs to be defined for all the coils separately. Since these intermediate steps are skipped in NUFFT reconstruction, we needed to make coil images retrospectively and add noise to them. The model could be improved by making use of fully sampled k-space data, which is probably not feasible due to the long scan time. If it has to be accelerated with under sampling, then preferably in a way that the model can handle the raw data, so that there is no need to under sample retrospectively and add noise yourself.

5.4 Future work

Although the proposed idea is a step closer to an end-to-end QSM pipeline, there is still a lot of work to do. Starting with getting the model working for our dataset. It would be an improvement if the model works for all data types and not for one particular data set. Besides, it would be an addition if the model can handle non-integer acceleration factors. Data scaling and normalization should be modified in order to let the pretrained qRIM properly reconstruct our dataset. We suggest to retrain the qRIM or convert it to a more generalized model, so this problem does not have to be tackeled again for subsequent datasets.

Once the model works, it should be retrained to reconstruct quantitative tissue phase images. For example, by making use of our proposed idea to use a biasfield correction as initial guess. The estimated B_{int} value after the biasfield correction must be calibrated so it is quantitative and a quantitative tissue phase image can be obtained from it. From that point on, the ill-posed dipole inversion needs to be implemented in the forward model. By completely unifying the reconstruction and QSM calculation we could reach a better exploitation of the correlation over the multiple images in a qMRI acquisition and potentially a better optimum of the mapping than the conventional solution with sequential image reconstruction and parameter fitting.

Further, fully sampled data scanned in multiple orientations can give higher quality reconstructions and higher quality COSMOS QSM-maps. This will be better training data and will lead to better estimations from the model.

Finally, the model should be trained on more subjects with more variety in age range. The elderly brain can not be seen as an average brain since this an extreme population. This extreme group of people have for example higher QSM values in several brain regions like the red nuclei and the striatum, possibly due to iron accumulation. Also, some region in the brain have extreme volumes since some regions like the lateral ventricles increase with age due to atrophy. By solely training for this extreme age group, the model could be overtrained for the changes in the elderly brain. To make the model more robust to age, the training data should consist subjects of all ages.

6 Conclusion

6.1 POE

In this study, we have expanded the high-resolution quantitative MRI benchmarks of the human lifespan as published by AHEAD by a decade. We found that relaxation parameters R_1 and R_2^* were mostly following the expected trends. With the additional information of the POE data, only slight changes in the models were necessary. However, QSM parameters did differ more for the POE data which we attribute to the applied motion correction and the different reconstruction method.

6.2 qRIM

We proposed an idea to extend the currently existing qRIM towards an end-to-end QSM pipeline. The qRIM was not able to reconstruct decent R_2^*, S_0, B_0 and ϕ maps. A working model, could be extended with a reconstruction of a quantitative tissue phase image, which is one step closer to the QSM map. However, this requires that the model works properly and will be retrained to reconstruct the tissue phase data.

References

[Abdel-Fahim et al., 2014] Abdel-Fahim, R., Mistry, N., Mougin, O., Blazejewska, A., Pitiot, A., Retkute, R., Gowland, P., and Evangelou, N. (2014). Improved detection of focal cortical lesions using 7t magnetisation transfer imaging in patients with multiple sclerosis. *Elsevier*.

[Alexander et al., 2021] Alexander, S., Ulrike, N., Pavel, H., Annemarie, R., Michelle, M., Simon, B., Sven, M., Felix, R., Helmuth, S., Marlies, W., Elke, H., Ralf, D., and René-Maxime, G. (2021). Multiparametric quantitative mri in neurological diseases. *Frontiers in Neurology*, 12.

[Alkemade et al., 2020] Alkemade, A., Mulder, M. J., Groot, J. M., Isaacs, B. R., van Berendonk, N., Lute, N., Isherwood, S. J., Bazin, P.-L., and Forstmann, B. U. (2020). The amsterdam ultra-high field adult lifespan database (ahead): A freely available multimodal 7 tesla submillimeter magnetic resonance imaging database. *Neuroimage*.

[Bazin et al., 2020] Bazin, P.-L., Nijsse, H. E., van der Zwaag, W., Gallichan, D., Alkemade, A., Vos, F. M., Forstmann, B. U., and Caan, M. W. (2020). Sharpness in motion corrected quantitative imaging at 7t. *NeuroImage*, 222:117–227.

[Beare et al., 2018] Beare, R., Lowekamp, B., and Yaniv, Z. (2018). Image segmentation, registration and characterization in r with simpleitk. *Journal of Statistical Software*, 86.

[Bollmann et al., 2019a] Bollmann, S., Kristensen, M. H., Larsen, M. S., Olsen, M. V., Pedersen, M. J., Østergaard, L. R., O'Brien, K., Langkammer, C., Fazlollahi, A., and Barth, M. (2019a). Sharqnet – sophisticated harmonic artifact reduction in quantitative susceptibility mapping using a deep convolutional neural network. *Zeitschrift für Medizinische Physik*, 29:139–149.

[Bollmann et al., 2019b] Bollmann, S., Rasmussen, K. G. B., Kristensen, M., Blendal, R. G., Østergaard, L. R., Plocharski, M., O'Brien, K., Langkammer, C., Janke, A., and Barth, M. (2019b). Deepqsm - using deep learning to solve the dipole inversion for quantitative susceptibility mapping. *Neuroimage*, 195:373–383.

[Brabec et al., 2003] Brabec, J., Krásený, J., and Petrovický, P. (2003). Volumetry of striatum and pallidum in man–anatomy, cytoarchitecture, connections, mri and aging. *Sb Lek*.

[Caan et al., 2019] Caan, M., Bazin, P., Marques, J., de Hollander, G., Dumoulin, S., and van der Zwaag, W. (2019). Mp2rageme: T1, t2*, and qsm mapping in one sequence at 7 tesla. *Human brain mapping*, 40.

[Cabello et al., 2002] Cabello, C., Thune, J., Pakkenberg, H., and Pakkenberg, B. (2002). Ageing of substantia nigra in humans: cell loss may be compensated by hypertrophy. *Neuropathol Appl Neurobiol*.

[Callaghan et al., 2014] Callaghan, M. F., Freund, P., Draganski, B., Anderson, E., Cappelletti, M., Chowdhury, R., Diedrichsen, J., FitzGerald, T. H., Smittenaar, P., Helms, G., Lutti, A., and Weiskopf, N. (2014). Widespread age-related differences in the human brain microstructure revealed by quantitative magnetic resonance imaging. *Neurobiol Aging*.

[Daugherty et al., 2015] Daugherty, A. M., Haacke, E. M., and Raz, N. (2015). Articles, behavioral/cognitive striatal iron content predicts its shrinkage and changes in verbal working memory after two years in healthy adults. *Journal of Neuroscience*.

[Feng et al., 2021] Feng, R., Zhao, J., Wang, H., Yang, B., Feng, J., Shi, Y., Zhang, M., Liu, C., Zhang, Y., Zhuang, J., and Wei, H. (2021). Modl-qsm: Model-based deep learning for quantitative susceptibility mapping. *NeuroImage*, 240.

[Fessler, 2021] Fessler, J. A. (2021). Michigan image reconstruction toolbox.

[Gao et al., 2021] Gao, Y., Zhu, X., Moffat, B. A., Glarin, R., Wilman, A. H., Pike, G. B., Crozier, S., Liu, F., and Sun, H. (2021). xqsm: quantitative susceptibility mapping with octave convolutional and noise-regularized neural networks. *NMR Biomedicine*, 34.

[He et al., 2022] He, J., Wang, L., Cao, Y., Wang, R., and Zhu, Y. (2022). Learn less, infer more: Learning in the fourier domain for quantitative susceptibility mapping. *Neuroimage*.

[Huntenburg et al., 2018] Huntenburg, J. M., Steele, C. J., and Bazin, P.-L. (2018). Nighres: processing tools for high-resolution neuroimaging. *GigaScience*, 7.

[Karkalousos et al., 2021] Karkalousos, D., Noteboom, S., Hulst, H. E., Vos, F., and Caan, M. (2021). Assessment of data consistency through cascades of independently recurrent inference machines for fast and robust accelerated mri reconstruction.

[Keuken et al., 2017] Keuken, M. C., Bazin, P., Backhouse, K., Beekhuizen, S., Himmer, L., Kandola, A., Lafeber, J., Prochazkova, L., Trutti, A., Schäfer, A., Turner, R., and Forstmann, B. (2017). Effects of aging on t_1, t_2^*, and qsm mri values in the subcortex. *Brain Struct Funct*.

[Lee et al., 2021] Lee, H., Lee, M. J., Kim, E.-J., Huh, G. Y., Lee, J.-H., and Cho, H. J. (2021). Iron accumulation in the oculomotor nerve of the progressive supranuclear palsy brain. *Sci Rep*, 11.

[Li et al., 2014] Li, W., Avram, A. V., Wu, B., Xiao, X., and Liu, C. (2014). Integrated laplacian-based phase unwrapping and background phase removal for quantitative susceptibility mapping. *NMR in Biomedicine*, 27:219–227.

[Li et al., 2011] Li, W., Wu, B., and Lui, C. (2011). Quantitative susceptibility mapping of human brain reflects spatial variation in tissue composition. *NeuroImage*, 55:1645–1656.

[Liu et al., 2012] Liu, J., Liu, T., de Rochefort, L., Ledoux, J., Khalidova, I., Chend, W., Tsiouris, A. J., Wisnieff, C., Spincemaille, P., Prince, M. R., and Wang, Y. (2012). Morphology enabled dipole inversion for quantitative susceptibility mapping using structural consistency between the magnitude image and the susceptibility map. *NeuroImage*, 59:2560–2568.

[Liu et al., 2009] Liu, T., Spincemaille, P., de Rochefort, L., Kressler, B., and Wang, Y. (2009). Calculation of susceptibility through multiple orientation sampling (cosmos): A method for conditioning the inverse problem from measured magnetic field map to susceptibility source image in mri. *Magnetic Resonance in Medicine*, 61:196–204.

[Metere et al., 2017] Metere, R., Kober, T., Moller, H. E., and Schafer, A. (2017). Simultaneous quantitative mri mapping of t1, t2 and magnetic susceptibility with multi echo mp2rage. *PLOS ONE*, 12.

[P.Marques et al., 2010] P.Marques, J., Kober, T., Krueger, G., van der Zwaag, W., de Moortel, P.-F. V., and Gruetter, R. (2010). Mp2rage, a self bias-field corrected sequence for improved segmentation and t1-mapping at high field. *NeuroImage*, 49.

[Putzky and Willing, 2017] Putzky, P. and Willing, M. (2017). Recurrent inference machines for accelerated mri reconstruction.

[Robinson et al., 2016] Robinson, S. D., Bredies, K., Khabipova, D., Dymerska, B., Marques, J. P., and Schweser, F. (2016). An illustrated comparison of processing methods for mr phase imaging and qsm: combining array coil signals and phase unwrapping. *NMR in Biomedicine*.

[Ruetten et al., 2019] Ruetten, P. P., Gillard, J. H., and Graves, M. J. (2019). Introduction to quantitative susceptibility mapping and susceptibility weighted imaging. *The British journal of radiology*, 92.

[Shmueli et al., 2009] Shmueli, K., de Zwart, J. A., van Gelderen, P., Li, T.-Q., Dodd, S. J., and Duyn, J. H. (2009). Magnetic susceptibility mapping of brain tissue in vivo using mri phase data. *Magnetic Resonance in Medicine*, 62:1510–1522.

[Tullo et al., 2019] Tullo, S., Patel, R., Devenyi, G. A., Salaciak, A., Bedford, S. A., Farzin, S., Wlodarski, N., Tardif, C. L., Breitner, J. C. S., and Chakravarty, M. M. (2019). Mr-based age-related effects on the striatum, globus pallidus, and thalamus in healthy individuals across the adult lifespan. *Hum Brain Mapp*.

[Tustison et al., 2010] Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., and Gee, J. C. (2010). N4itk: Improved n3 bias correction. *IEEE Trans Med Imaging*, 29.

[Wei et al., 2019] Wei, H., Cao, S., Zhang, Y., Guan, X., Yan, F., Yeom, K. W., and Liu, C. (2019). Learning-based single-step quantitative susceptibility mapping reconstruction without brain extraction. *NeuroImage*, 201.

[Yoon et al., 2018] Yoon, J., Gong, E., Chatnuntawech, I., Bilgic, B., Lee, J., Jung, W., Ko, J., Jung, H., Setsompop, K., Zaharchuk, G., Kim, E. Y., Pauly, J., and Lee, J. (2018). Quantitative susceptibility mapping using deep neural network: Qsmnet. *Neuroimage*, 179:199–206.

[Zhang et al., 2021] Zhang, C., Poot, D. H. J., Bazin, P., Coolen, B. F., Vrenken, H., Forstmann, B. U., and Caan, M. W. A. (2021). A unified model for reconstruction and r2* mapping of accelerated 7t data using the quantitative recurrent inference machine. *IEEE*.

[Zhou et al., 2021] Zhou, H., Cheng, C., Peng, H., Liang, D., Liu, X., Zheng, H., and Zou, C. (2021). The phu-net: A robust phase unwrapping method for mri based on deep learning. *Magnetic Resonance in Medicine*, 86:3321–3333.

7 Supplementary material

7.1 CIRIM

In addition to the method with the qRIM, another idea we worked on was to modify a currently existing CIRIM [Karkalousos et al., 2021]. CIRIM stands for Cascades of Independently Recurrent Inference Machines and is capable of reconstructing an image from raw k-space data.

7.1.1 Preprocessing

The CIRIM was possible to reconstruct under sampled k-space data as long as there existed a sampling mask of the sampling pattern. For non-integer under sampling factors there did not exist a sampling mask since the sampled points did not fall on a grid. In order to work with the CIRIM we retrospectively under sampled the reconstructed POE data with acceleration factors of 2, 3 and 4 in an equidistant way. Several files were needed and saved: the fully sampled k-space, under sampled k-space, the under sampling mask, the phase mask, the tissue phase image. All the files were saved for each echo separately and converted into 3 different orientations (axial, coronal and sagittal).

7.1.2 Architecture

The CIRIM consisted of sequentially connected RIM blocks, see figure 18 [Karkalousos et al., 2021].
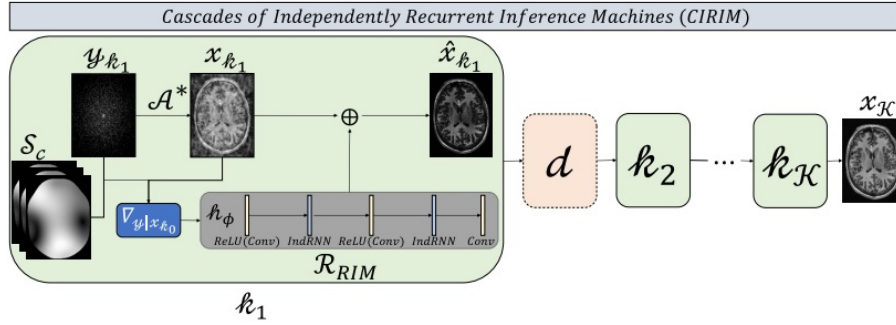
Figure 18: The architecture of of the CIRIM in which the RIM is used as a regularizer. The prediction after the first cascade will be the input to the next cascade, while an (optional) additional data consistency step can be performed through an explicitly formulated term. After all cascades the network will return the final prediction [Karkalousos et al., 2021]

The update equation of the RIM was formulated by:

$$x_{\tau+1} = x_\tau + h_\phi(\nabla_{y|x_\tau}, x_\tau) \tag{23}$$

which meant that a next prediction was an optimization of the previous prediction with optimizer h with its own set of parameters ϕ. The cascades of RIM blocks in the CIRIM allowed building a deep RNN without vanishing or exploding gradients issues and further evaluated the update equation originating of the RIM through k cascades. In this way the RIM acted as regularizer ($\mathcal{R}_{RIM}$), the update equation of the CIRIM are given by:

$$\hat{x}_{k+1} = x_k + \lambda\mathcal{R}_{RIM_k}(x_k) \tag{24}$$

for cascade k, with $1 \leq k \leq \mathcal{K}$. The Independently Recurrent Neural Networks (IndRNN) was used as recurrent unit for the RIM as a more efficient unit for balancing the network's complexity while increasing the number of trainable parameters through the cascades. Optionally, a data consistency term could be added which can applied k-space corrections. If this data consistency term was included, the update equation will be:

$$\hat{x}_{k+1} = \mathcal{A}^*(y_k - d(y_k - y) - \mathcal{A}(\hat{x}_{k+1})) \tag{25}$$

In which $\mathcal{A}$ is the forward operator and $\mathcal{A}^*$ the adjoint operator or backward operator.

7.1.3 Loss functions

For calculating the loss, the tissue phase images derived from the complex-valued estimations $\hat{\phi_{tissue}}$ could be compared against the reference ϕ_{tissue} calculated from the fully sampled image. The l_1-norm can be

chosen as loss function and represents the sum of the absolute difference, given by:

$$\mathcal{L}^{l_1}(\hat{\phi}_{tissue}) = |\hat{\phi}_{tissue} - \phi_{tissue}| \tag{26}$$

Since the CIRIM contained multiple cascades, that had to be taken into account. The l_1-loss was weighted depending on the number of iterations and averaged over the cascades, to emphasize the predictions of the later iterations. The loss was then formulated as:

$$\mathcal{L}^{l_1}(\hat{\phi}_{tissue}) = \frac{\sum_{i=1}^{c}(\frac{1}{q\mathcal{T}}\sum_{\tau=1}^{\mathcal{T}} w_\tau |\hat{\phi}_{tissue} - \phi_{tissue}|)}{\mathcal{K}} \tag{27}$$

7.1.4 Forward model

The process of accelerating the MRI acquisition can be described through a forward model. Let the true image be denoted by $\mathbf{x} \in \mathbb{C}^n$, with $n = n_x \times n_y$ and let $y \in {}^m$ with $m << n$, be the set of the sparsely sampled data in k-space. The forward model describes how the measured data are obtained from an underlying reference image. For the i-th coil of c receiver coils, the forward model is formulated as:

$$y_i = \mathcal{P} \cdot \mathcal{F} \cdot (S_0 \cdot x, ..., S_c \cdot x) + \sigma_i, i = 1, .., c \tag{28}$$

in which $\mathcal{P}$ is a subsampling mask, $\mathcal{F}$ denotes the Fourier transform, S_i are the coil sensitivity maps transforming $\mathbf{x}$ into $\mathbf{x}_c$ coil images and σ_i denotes the measured noise in the i-th coil.

Our idea was to extend this currently existing forward model in order to get tissue phase images. In order to this, first we wanted to get the phase images from the reconstructions. The phase images can be obtained by taking the angle of true image $\mathbf{x}$ (which is complex).

$$\phi = angle(x) = tan^{-1}\frac{imag(x)}{real(x)} \tag{29}$$

In which ϕ is the phase image, $imag$ is the imaginary part, $real$ is the real part and $\mathbf{x}$ the complex reconstructed image. The other way around, the true image consisted of a magnitude and the phase part:

$$x = abs(x) \cdot e^{i \cdot angle(x)} = abs(x) \cdot e^{i\phi} \tag{30}$$

Integrated phase unwrapping and background phase removal can be done using the Laplacian. The Laplacian can be derived from the wrapped phase directly by making use of sine and cosine functions as described in the equation below:

$$\nabla^2\phi = cos\phi\nabla^2 sin\phi - sin\phi\nabla^2 cos\phi \tag{31}$$

The Laplacian was used to calculate the phase only originated from the tissue susceptibility differences. Solving the equation written above yields the unwrapped phase that is free of contributions from sources outside the FOV.

The Laplacian equation can be solved in spatial frequency domain (by assuming boundary conditions at the edges of the FOV) if the phase is available everywhere within the FOV. However phase measurements are typically not available outside the tissue, so the phase outside the tissue has to be determined. This is usually done by the sperical mean value, but we will train the model to solve this by itself. Once the phase outside the tissue $\phi_{surrounding}$ is determined, the phase in the whole FOV can be calculated as follows:

$$\nabla^2\phi_{FOV} = \nabla^2\phi_{tissue} + \nabla^2\phi_{surrounding} \tag{32}$$

Our idea was to train the RIM to solve for the $\phi_{surrounding}$ part and we wanted to give a mask as input. The background removed and unwrapped phase can be obtained using the following FFT-based inverse Laplacian:

$$\phi = -\mathcal{F}^{-1}[\mathcal{F}(\nabla^2\phi_{FOV})/4\pi^2 k^2] \tag{33}$$

in which FT represents the Fourier transform. Combining the above equations, this will lead through the following equation:

$$\phi = -\mathcal{F}^{-1}[\mathcal{F}(\nabla^2 tan^{-1}\frac{imag(I)}{real(I)})/4\pi^2 k^2] \tag{34}$$

26

7.1.5 Inverse problem

When solving the inverse problem, the goal is to map the sparsely sampled k-space measurements to a high resolution tissue phase image in which aliasing, the background field and phase wraps are removed. The actual image x can be reconstructed by making use of this relationship:

$$x = \sum_{i=1}^{c} \mathcal{S}_i^H \cdot (\mathcal{F}^{-1} \cdot \mathcal{P}^T \cdot y_i) \tag{35}$$

We know that the phase images ϕ can be made with the following equation:

$$\phi = angle(x) = tan^{-1} \frac{imag(x)}{real(x)} \tag{36}$$

From this wrapped phase image we can get to an unwrapped and background field removed tissue phase ϕ_{tissue}:

$$\phi_{tissue} = -\mathcal{F}^{-1}[\mathcal{F}(\nabla^2 \phi / 4\pi^2 k^2] \tag{37}$$

when combining all of the above equations, one can expressed this into one equation:

$$\phi_{tissue} = -\mathcal{F}^{-1}[\frac{\mathcal{F}(\nabla^2 angle(\sum_{i=1}^{c} \mathcal{S}_i^H \cdot (\mathcal{F}^{-1} \cdot \mathcal{P}^T \cdot y_i)}{4\pi^2 k^2})] \tag{38}$$

This is schematically shown in figure 19 below.
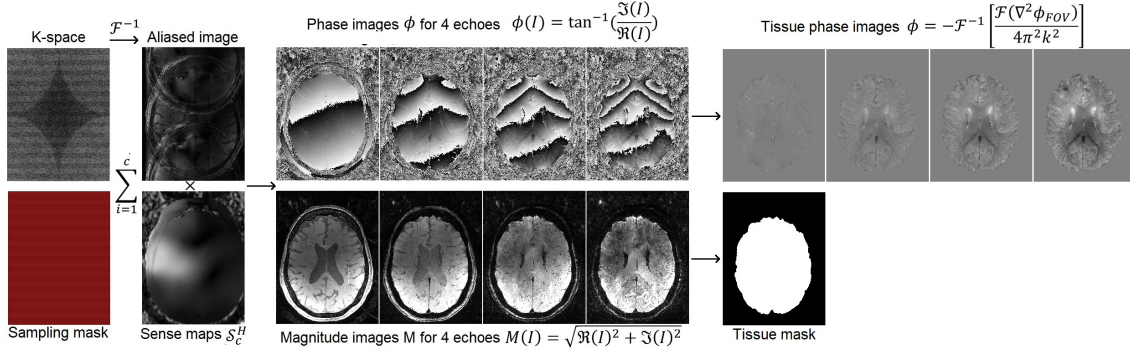
Figure 19: The total inverse problem

The inverse transformation of restoring the tissue phase image from the set of the sparsely sampled k-space measurements can be found through the Maximum A Posteriori (MAP) estimator, given by

$$\phi_{MAP} = argmax_\phi[logp(y|\phi) + logp(\phi)] \tag{39}$$

which is the maximization of the sum of the log-likelihood and the log-prior distribution of y and ϕ respectively. The log-likelihood expresses the log probability that k-space data y are obtained given an image ϕ yielding a data fidelity term derived from the posterior $p(y|\phi)$. The log-prior distribution regularizes the solution by representing an MR-image's most likely appearance

7.1.6 Training

The network can be trained to get an unwrapped and background field removed phase image. This is a necessary step in the QSM pipeline. The phase images can be obtained by taking the angle of the reconstruction (which is complex).

$$\Psi = angle(I) = tan^{-1} \frac{imag(I)}{real(I)} \tag{40}$$

In which Ψ is the phase image, $imag$ is the imaginary part, $real$ is the real part and I the complex reconstructed image. The original phase images of the 4 echoes are visible in figure 20.

These phase images need to be unwrapped, the unwrapping process can be written as the following inverse problem:

$$\Psi = mod(\phi, 2\pi) \tag{41}$$
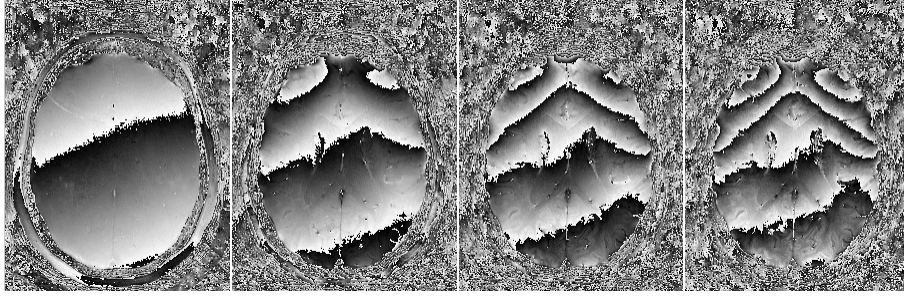
Figure 20: The phase reconstructions of a slice for 4 different echoes

In which Ψ is the phase image with phase wraps and ϕ is the unwrapped phase image. In order to learn this phase unwrapping, we need priors for the network to train on. These priors are made by Laplacian based phase unwrapping. The Laplacian based phase unwrapping is performed with the STI suite package in in Matlab which consist an algorithm for phase unwrapping and background phase removal using the Laplacian operator [Li et al., 2014]. This phase unwrapping leaves us with the following phase unwrapped images, see figure 21.
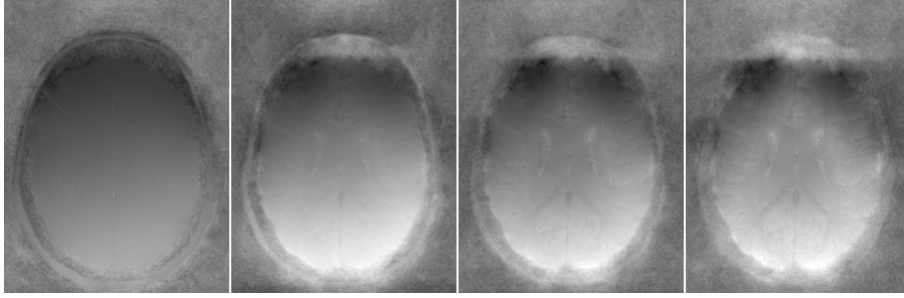
Figure 21: Unwrapped phase images

Since the phase outside the region of interest needs to be removed, a masking is applied. The masking is done based on the magnitude image of the second echo. Morphological operations like erosions and dilations are applied in order to close possible holes in the mask (e.g. due to veins). The mask is multiplied with the phase unwrapped image, which results in a phase unwrapped image of only the region of interest, see figure 22.
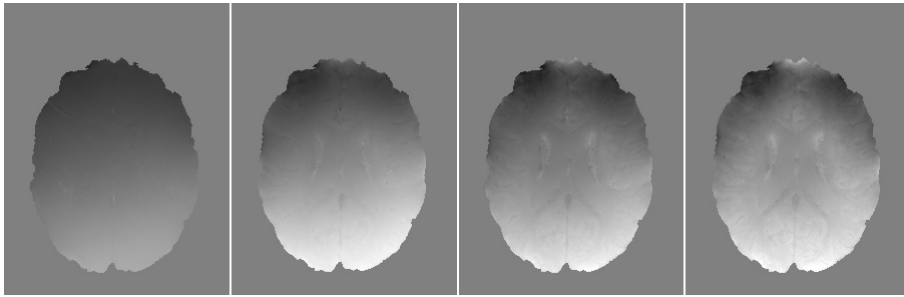
Figure 22: Unwrapped phase images of the 4 echoes after applying a mask

For QSM it is important that the differences in phase only originate from susceptibility differences inside the tissue. This means that external contributions should be removed. External contributions are the background field and this can be removed by making use of the Laplacian. The Laplacian can be derived from the wrapped phase directly by making use of sine and cosine functions as described in the equation below:

$$\nabla^2\phi = cos\phi\nabla^2 sin\phi - sin\phi\nabla^2 cos\phi \tag{42}$$

The Laplacian was calculated by making use of the functions available in the STI suite package and is shown in the figure 23 below.
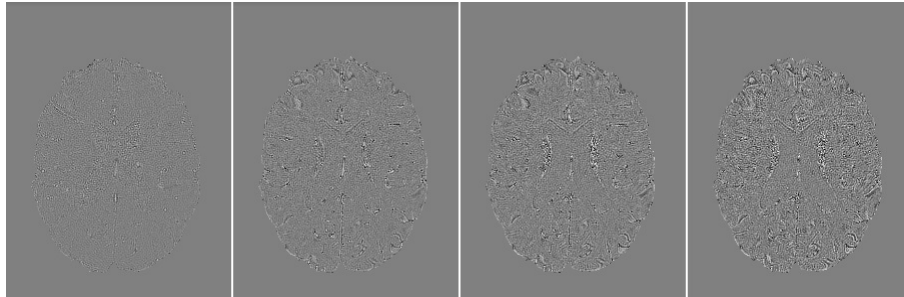
Figure 23: The calculated Laplacian of the phase images of the 4 echoes.

The Laplacian is used to calculate the phase only originated from the tissue susceptibility differences. Solving the equation written above yields the unwrapped phase that is free of contributions from sources outside the FOV. This is referred to as the tissue phase and is shown in figure 24.
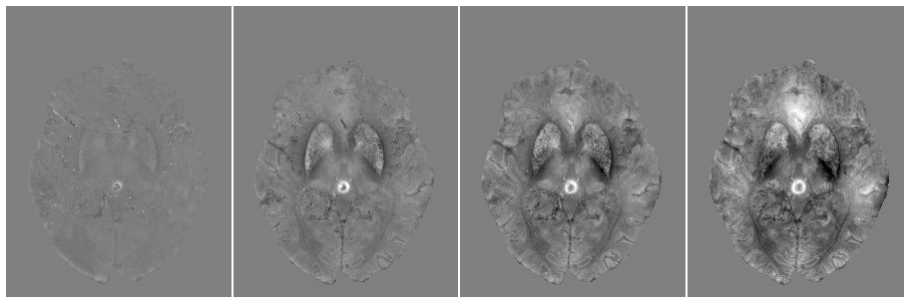

Figure 24: The tissue phase images of the 4 echoes.

This can be used for network training in future work, but regarding time I didn't tried this method but focused on de qRIM.

7.2 Motion analysis

A 3D Echo Planar Imaging (3D-EPI) Fatnav was inserted into the sequence after the two GRE readout blocks in the MP2RAGEME sequence. Fat was selectively excited using a fat-excitation pulse, which results in a sparse image where mainly the subcutaneous fat layer appears bright, while the rest of the structures are dark. This sparse fat signal circumscribing the brain can be used to determine rigid brain movement over time. According to several points of the FatNav image, one can get insight in and correct retrospectively for the motion made during the acquisition. K-space weighted motion metrics were estimated through rigid realignment of a series of FatNav images. For the realignment of image i to some reference image, we can define the rigid transformation matrix T^i:

$$T^i = \begin{bmatrix} \mathbf{R}^i & \mathbf{t}^i \\ 0 & 1 \end{bmatrix}$$

given the rotation matrix $\mathbf{R}^i$ and the translation vector $\mathbf{t}^i = [t_x^i t_y^i t_z^i]^T$
[Bazin et al., 2020]. For each voxel in the image, the displacement vector according to the FatNav's is calculated through $d^i(\overrightarrow{x}) = (T^i)^{-1}\overrightarrow{x} - \overrightarrow{x}$
[Bazin et al., 2020]. By making use of translation and rotation parameters it is possible to correct for motion, this operation takes place in k-space and results in an image with an improved resolution.

Motion correction and reconstruction were performed in Retro-MoCo-Box 1 a Matlab (The MathWorks, Inc., Natick, MA) toolbox for retrospective motion-correction of 3D MRI k-space data. FatNav's were rigidly realigned with six Degrees Of Freedom using Statistical Parametric Mapping software (SPM12). The motion parameters of the 11 center FatNav's correspond to the readout of the center of k-space and were averaged and used as anatomical reference point during reconstruction. We used the rotation matrices to get insight in the movement during acquisition and made this visible in plots (see figure 25).
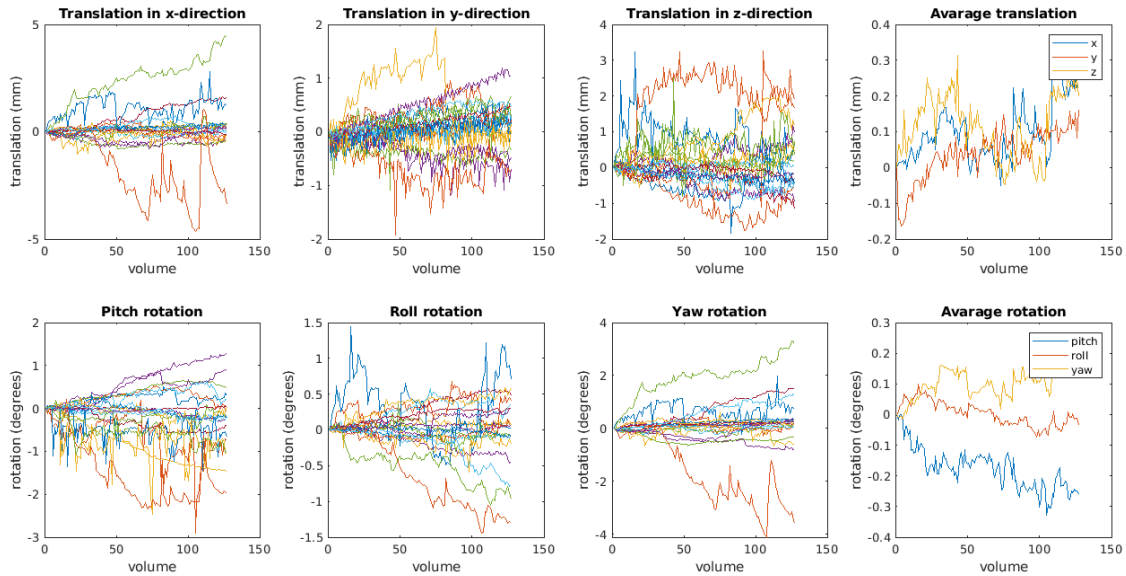
Figure 25: Translation and rotation parameters for all subjects computed through rigid realignment.

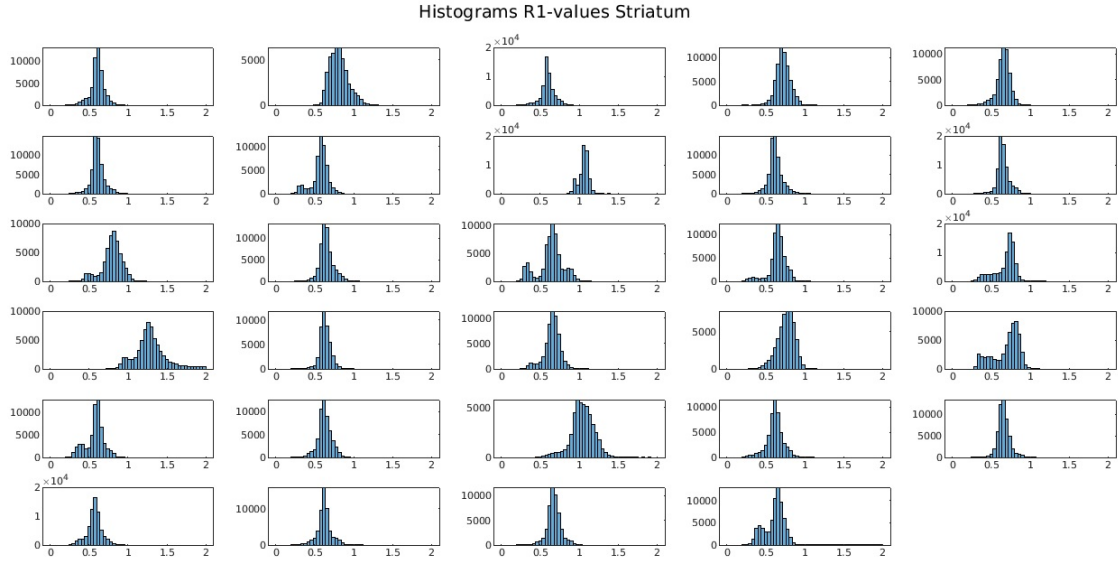## 7.3 Supplementary figures

Histograms

Figure 26: The histograms of the R_1-values of the Striatum. The horizontal axes show the R_1-values in Hz and the vertical axes the number of voxels.
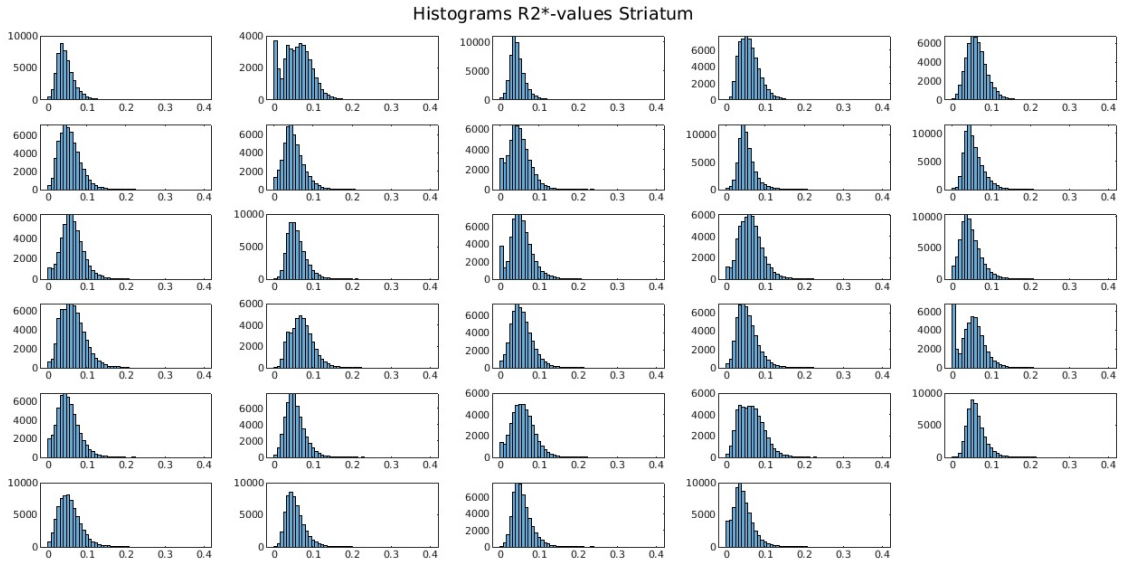

Figure 27: The histograms of the R_2^*-values of the Striatum. The horizontal axes show the R_2^*-values in Hz and the vertical axes the number of voxels.
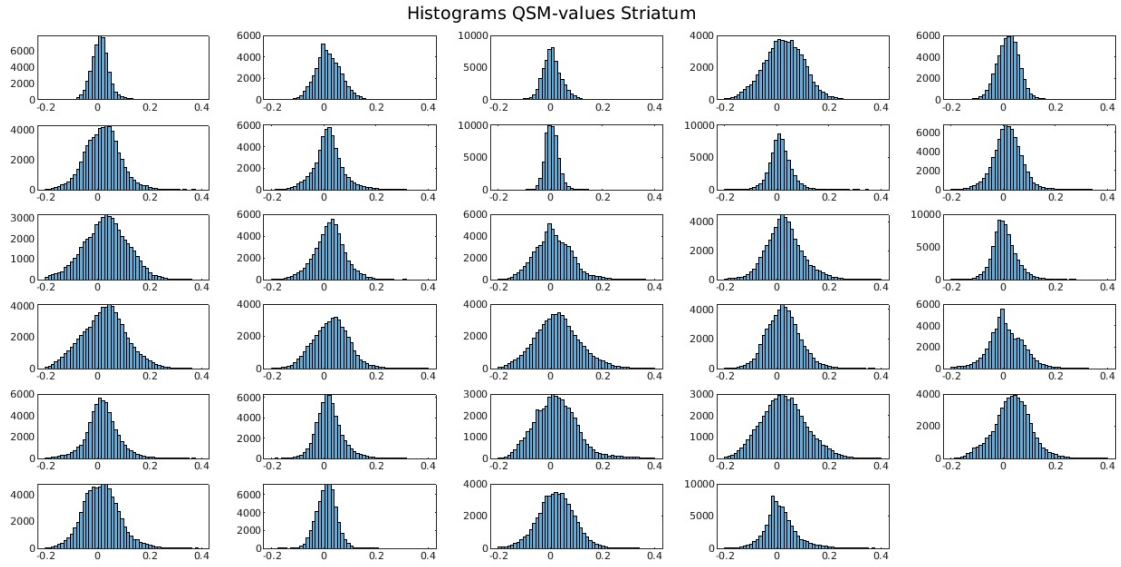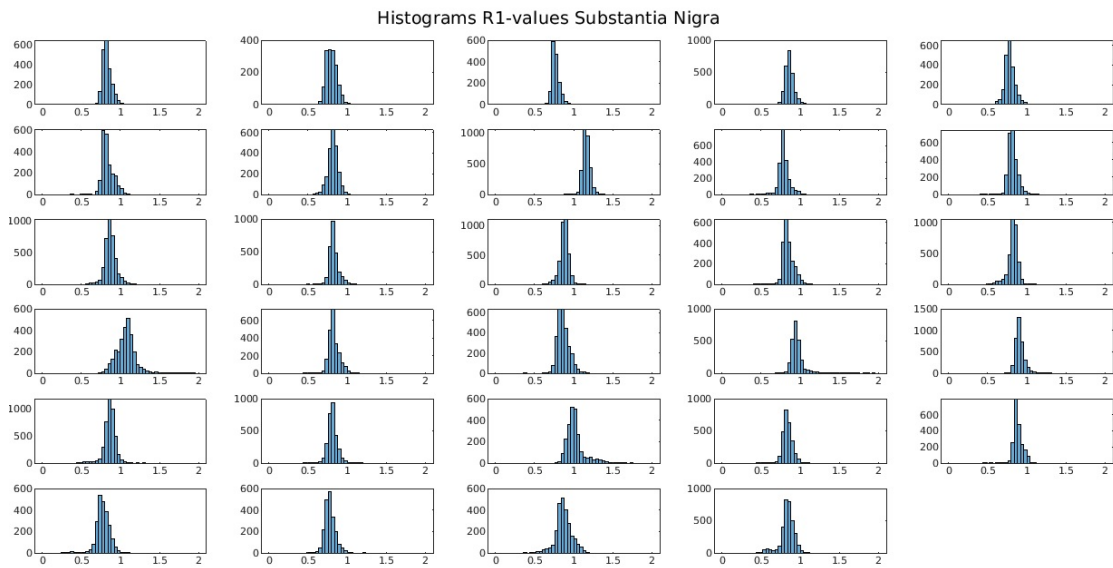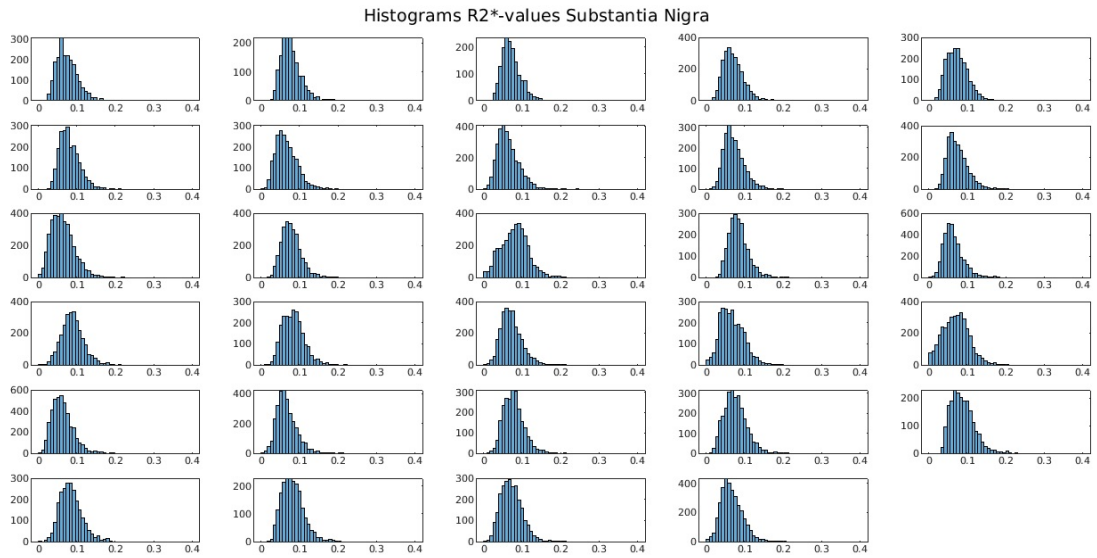
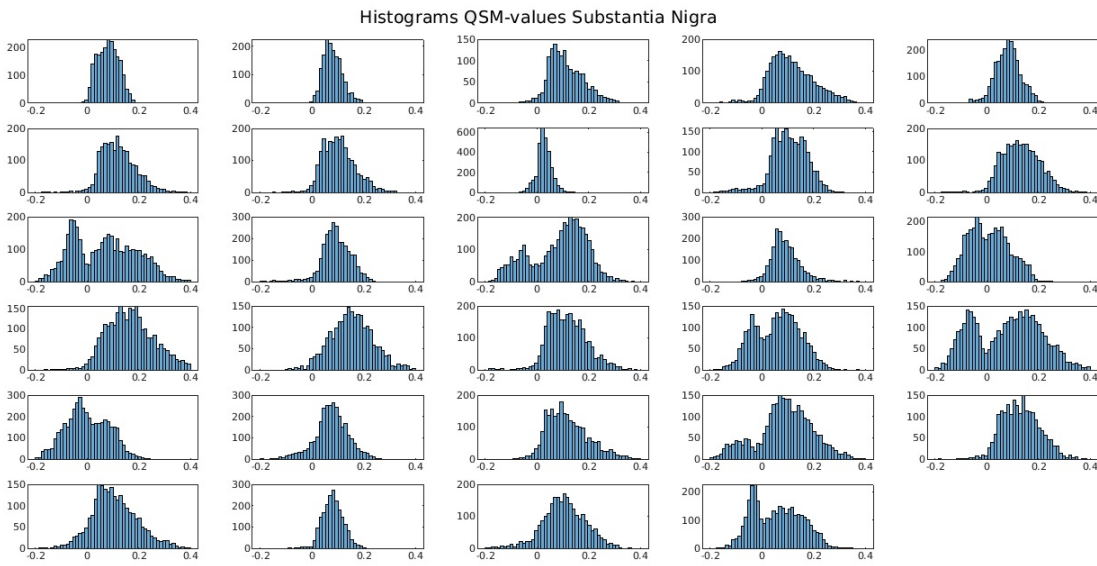Figure 28: The histograms of the QSM-values of the Striatum. The horizontal axes show the QSM-values in ppm and the vertical axes the number of voxels.

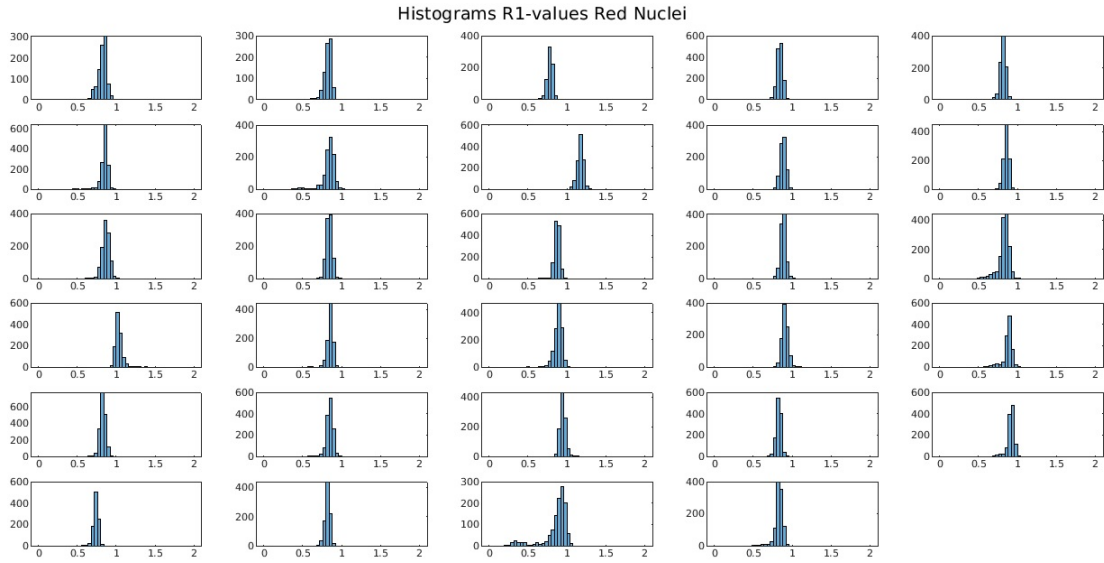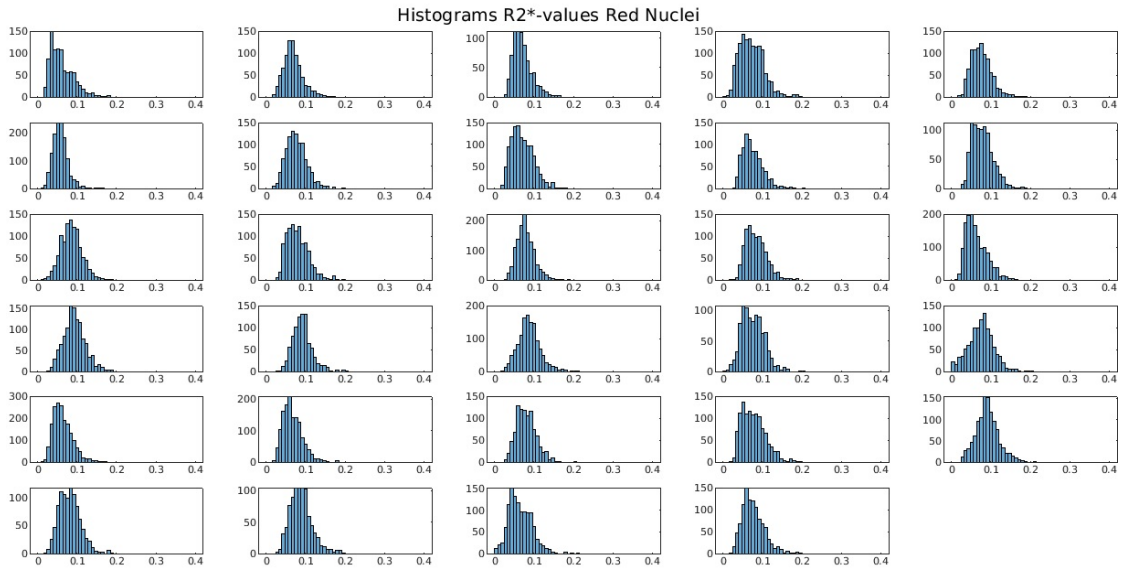Figure 29: The histograms of the R_1-values of the Substantia Nigra. The horizontal axes show the R_1-values in Hz and the vertical axes the number of voxels.

Figure 30: The histograms of the R_2^*-values of the Substantia Nigra. The horizontal axes show the R_2^*-values in Hz and the vertical axes the number of voxels.
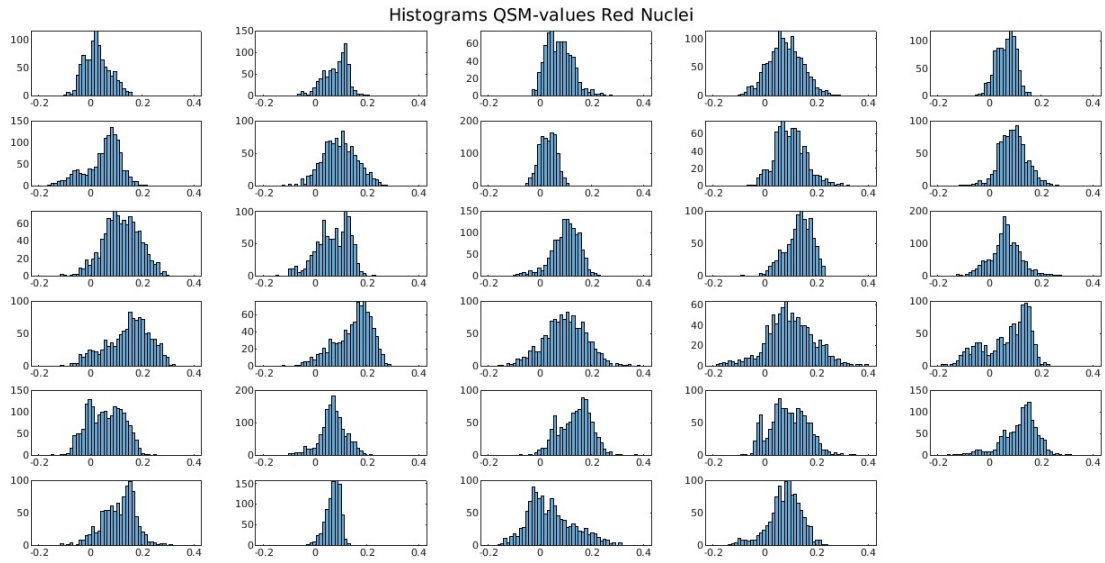

Figure 31: The histograms of the QSM-values of the Substantia Nigra. The horizontal axes show the QSM-values in ppm and the vertical axes the number of voxels.

Figure 32: The histograms of the R_1-values of the Red Nuclei. The horizontal axes show the R_1-values in Hz and the vertical axes the number of voxels.

Figure 33: The histograms of the R_2^*-values of the Red Nuclei. The horizontal axes show the R_2^*-values in Hz and the vertical axes the number of voxels.

Figure 34: The histograms of the QSM-values of the Red Nuclei. The horizontal axes show the QSM-values in Hz and the vertical axes the number of voxels.
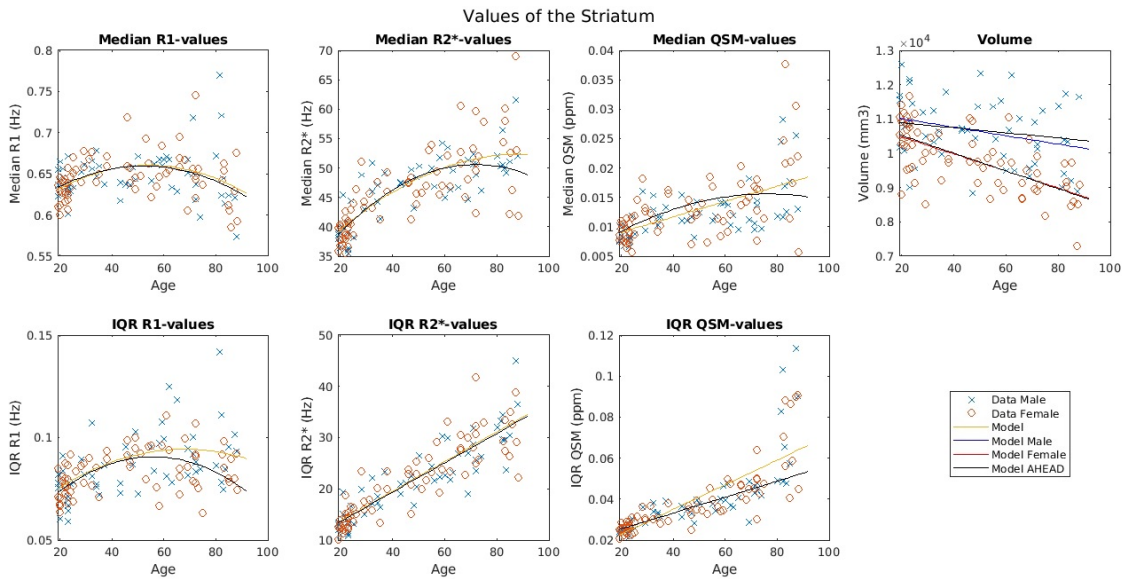
Results with both models

Figure 35: The results of the Striatum plotted together with the models based on the all data and the models based on solely the AHEAD dataset.
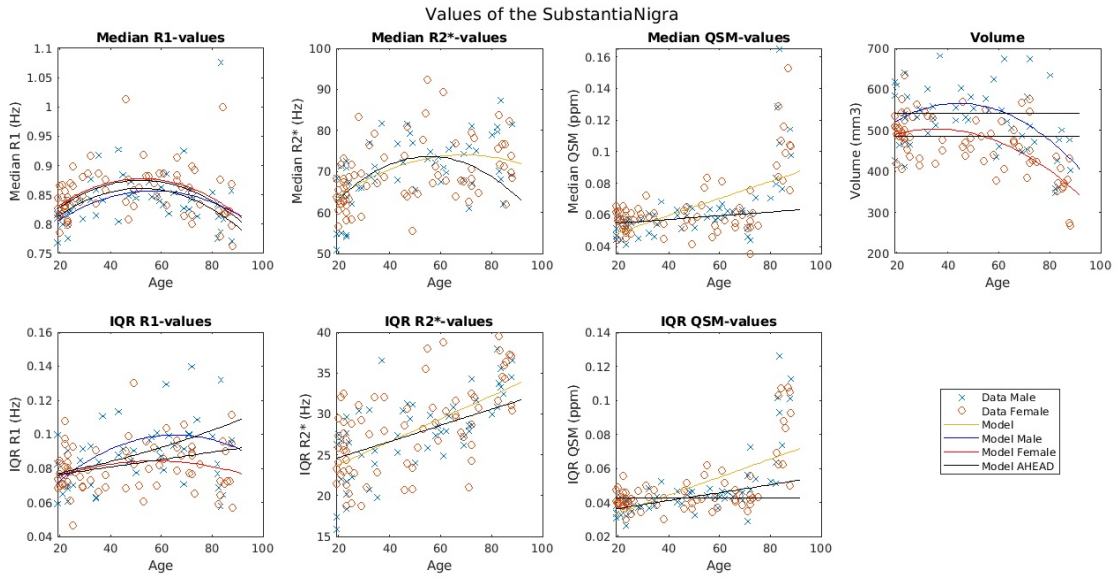
Figure 36: The results of the Substantia Nigra plotted together with the models based on the all data and the models based on solely the AHEAD dataset.
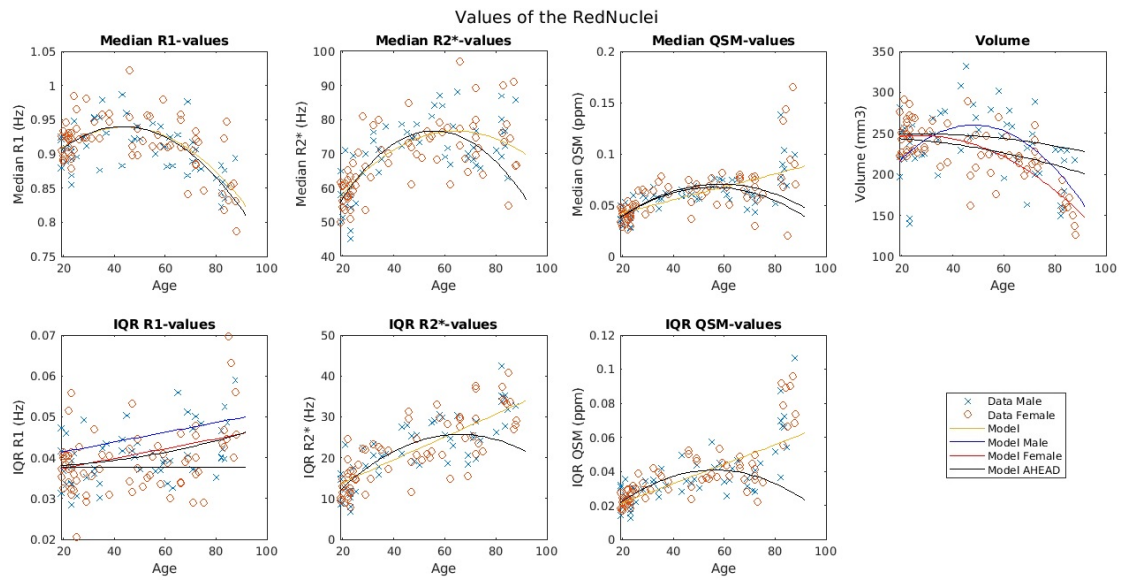

Figure 37: The results of the Red Nuclei plotted together with the models based on the all data and the models based on solely the AHEAD dataset.

Figure 38: A y-z slice shows an aliased field-of-view with non-integer acceleration factors of 1.8 in y and 1.5 in z.