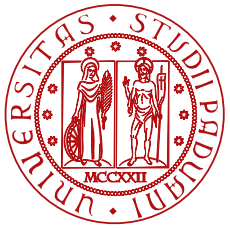

Università degli Studi di Padova

Dipartimento di Ingegneria dell'Informazione

Master Thesis in Bioengineering

Assessment of a Neural Network-Based Subspace MRI Reconstruction Method for Myocardial T1 Mapping Using Inversion-Recovery Radial FLASH

Master Candidate

Chiara Fantinato

Student ID 2039066

Supervisor

Prof. Alessandra Bertoldo

University of Padova

Co-supervisors

Prof. Martin Uecker, M.Sc. Moritz Blumenthal

Technical University of Graz

Academic Year
2022/2023

Abstract

In cardiovascular MRI, myocardial T_1 mapping provides an imaging biomarker for the non-invasive characterization of the myocardial tissue, with the potential to replace invasive biopsy for the diagnosis of several pathological heart muscle conditions such as fibrosis, iron overload, or amyloid infiltration. Over the last few years, deep learning has become increasingly appealing for image reconstruction to improve upon the commonly employed user-dependent regularization terms by automatically learning image properties from the training dataset. This thesis investigates a novel neural network-based subspace MRI reconstruction method for myocardial T_1 mapping, which uses a single-shot inversion-recovery radial FLASH sequence. The neural network utilized in this study is NLINV-Net, which draws inspiration from the NLINV image reconstruction technique. NLINV-Net addresses the nonlinear inverse problem for parallel imaging by unrolling the iteratively regularized Gauss-Newton method and incorporating neural network-based regularization terms into the process. It learned in a self-supervised fashion, i.e., without a reference, correlations between the individual parameters encoded with the FLASH sequence, and, consequently, a well-tuned regularization. NLINV-Net outperformed NLINV in terms of T_1 precision and generated high-quality T_1 maps. The T_1 maps computed using NLINV-Net were comparable to the ones obtained using another baseline method, which combines parallel imaging and compressed sensing using the ℓ_1-Wavelet regularization when solving the linear inverse problem for parallel imaging. In this case, the advantage of NLINV-Net is that it removes the subjective regularization parameter tuning that comes with the forenamed benchmark method. Thus, it provides an excellent basis for myocardial T_1 mapping using a single-shot inversion-recovery radial FLASH sequence.

Sommario

La mappatura T_1 del miocardio si è affermata come un promettente biomarker per la caratterizzazione non invasiva del muscolo cardiaco nell'ambito della risonanza magnetica cardiovascolare. Questo approccio ha il potenziale di sostituire la biopsia nella diagnosi di diverse condizioni patologiche del miocardio, come la fibrosi, l'accumulo di ferro o amiloidosi. Negli ultimi anni, il deep learning ha suscitato un crescente interesse per la ricostruzione delle immagini, portando a notevoli miglioramenti rispetto alle tecniche che richiedono la predefinizione dei parametri di regolarizzazione da parte dell'operatore, rendendo così il processo parzialmente soggettivo. Il miglioramento è reso possibile grazie alla capacità delle reti neurali di apprendere automaticamente le proprietà presenti nelle immagini del dataset utilizzato per il training. La presente tesi si focalizza sull'analisi di un nuovo metodo di ricostruzione subspaziale delle immagini di risonanza magnetica basato su reti neurali per la mappatura T_1 del miocardio, che utilizza una sequenza chiamata single-shot inversion-recovery radial FLASH. È stata impiegata una rete neurale nota come NLINV-Net, la quale trae ispirazione dalla tecnica di ricostruzione delle immagini NLINV. NLINV-Net risolve il problema inverso non lineare per il parallel imaging srotolando l'iteratively regularized Gauss-Newton method e incorporando nel processo termini di regolarizzazione basati su reti neurali. La rete neurale ha appreso le correlazioni esistenti tra i singoli parametri codificati dalla sequenza FLASH in modo auto-supervisionato, ovvero senza richiedere un riferimento esterno. NLINV-Net ha dimostrato di superare NLINV per la precisione dei valori T_1, producendo mappe T_1 di alta qualità. Le mappe ottenute con NLINV-Net sono paragonabili a quelle ottenute con un altro metodo di riferimento, che combina parallel imaging e compressed sensing utilizzando la regolarizzazione ℓ_1-Wavelet nella risoluzione del problema lineare inverso per il parallel imaging. Il vantaggio di NLINV-Net rispetto al suddetto metodo di appoggio è quello di sbarazzarsi della predefinizione dei parametri di regolarizzazione da parte dell'operatore. In questo modo, NLINV-Net fornisce una solida base per la mappatura T_1 del miocardio utilizzando la sequenza single-shot inversion-recovery radial FLASH.

Contents

List of Figures

List of Tables

List of Acronyms

T_1 spin-lattice relaxation time

T_1^* effective spin-lattice relaxation time

T_2 spin-spin relaxation time

T_2^* effective spin-spin relaxation time

Adam Adaptive moment estimation

ANN Artificial Neural Network

ANOVA Analysis Of Variance

BART Berkeley Advanced Reconstruction Toolbox

CG Conjugate Gradient

CG-SENSE Conjugate Gradient SENSE

CNN Convolutional Neural Network

CV Coefficient of Variation

DFT Discrete Fourier Transform

DC Data Consistency

ECG electrocardiogram

FID Free Induction Decay

FFT Fast Fourier Transform

FISTA Fast Iterative Shrinkage-Thresholding Algorithm

FLASH Fast Low Angle SHot

FOV Field Of View

IRGNM Iteratively Regularized Gauss-Newton Method

MoDL MOdel-based reconstruction using Deep Learned priors

MRI Magnetic Resonance Imaging

MSE Mean Square Error

NaN Not a Number

NLINV NonLinear Inversion

NMSE Normalized Mean Square Error

NMR Nuclear Magnetic Resonance

nuFFT non-uniform Fast Fourier Transform

PCA Principal Component Analysis

PICS Parallel Imaging Compressed Sensing

ReLU Rectified Linear Unit

ResNet Residual neural Network

RMSE Root Mean Square Error

RF Radio Frequency

ROI Region Of Interest

SD Standard Deviation

SENSE SENSitivity Encoding

SGD Stochastic Gradient Descent

SVD Singular Value Decomposition

Chapter 1

Introduction

1.1 Background

Cardiovascular Magnetic Resonance Imaging (MRI) is an imaging technique that offers a non-invasive tool for the accurate assessment of both the function and structure of the cardiovascular system. It stands out from other imaging techniques, such as X-ray Computed Tomography (CT) or Positron Emission Tomography (PET), by employing neither ionizing radiation nor radioactive substances but static and time-varying magnetic fields. A primary challenge associated with MRI, especially in cardiovascular MRI, is the extended acquisition time, which can cause patient discomfort and heighten the risk of motion artifacts. A breakthrough in addressing this issue occurred in 1986 when Haase et al. introduced the Fast Low Angle SHot (FLASH) sequence [1]. This innovation harnessed Radio Frequency (RF) pulses with small flip angles to preserve most of the longitudinal magnetization during excitation, eliminating the requirement for waiting periods between consecutive experiments. Another notable advancement was the development of parallel imaging in 1997 by Sodickson et al. [2]. This technique involves the utilization of multiple receiver coils for signal acquisition, resulting in a significant acceleration of acquisition times, often by a factor of two or more [2]. A milestone in the quest for rapid MRI occurred in 2007 when Lustig et al. [3] and Block et al. [4] introduced compressed sensing in the field of MRI. It enables image reconstruction from signals acquired at sub-Nyquist sampling rates, leading to substantial reductions in scan times. Modern MRI techniques combine parallel imaging and compressed sensing by incorporating sparsifying regularization terms when solving the inverse problem associated with parallel

imaging. Another approach developed to enhance imaging speed is non-Cartesian MRI. Notably, the adoption of radial sampling has garnered significant attention due to its inherent robustness against motion artifacts and its partial tolerance to undersampling.

In the context of cardiac MRI, myocardial T_1 mapping provides an imaging biomarker for the non-invasive characterization of the myocardial tissue with the potential to replace invasive biopsy for several pathological heart muscle conditions such as fibrosis, iron overload, or amyloid infiltration [5]. Quantitative T_1 mapping typically involves an appropriate magnetization preparation pulse, e.g., an inversion pulse, followed by the acquisition of a predefined number of images of the relaxation period. After data acquisition and image reconstruction, the relaxation model is fit to the images to derive the parameter maps. The basic requirements for myocardial T_1 mapping comprehend speed, practically robustness, T_1 accuracy, and sufficiently high spatial resolution, and a method ensuring these needs was proposed by Wang et al. in 2016 [6]: myocardial T_1 mapping can be accomplished by combining single-shot inversion-recovery radial FLASH sequence with breath-hold and finger pulse triggering, i.e., data are acquired within a breath-hold, and the inversion pulse is triggered to the early diastolic phase with the use of finger pulse signal, iterative image reconstruction by NonLinear Inversion (NLINV), as well as T_1 fitting with automated deletion of systolic frames based on the finger pulse triggering [6]. Instead of solving the image reconstruction problem for as many images as the number of time points, it is possible to model the physical laws that govern the MRI signal generation and extract quantitative maps of the underlying physical parameters directly from the measured k-space data without intermediate image reconstruction by fitting the model to the data through a process called model-based MRI reconstruction [7]. Furthermore, subspace-constrained MRI reconstruction [7] aims to reduce the dimensionality of the MRI reconstruction problem as the model-based MRI reconstruction does but preserves linearity through the projection of the data onto a subspace. Finally, over the last few years, deep learning has become even more appealing for image reconstruction to improve upon the previously user-dependent regularization terms by automatically learning image properties from the training dataset [8].

1.2 Research Question

The main focus of this thesis is to develop and investigate a new neural network-based MRI reconstruction method for myocardial T_1 mapping based on self-supervised subspace deep learning. The chosen neural network is NLINV-Net [9]. Through its self-supervised learning approach, NLINV-Net can autonomously learn how to reconstruct subspace images jointly with coil sensitivity maps from undersampled radial cardiac data, which do not possess a ground truth reference. The aim is to enhance the image quality of parameter maps compared to NLINV method results and, if possible, improve over techniques incorporating hand-crafted regularization terms.

1.3 Structure

As far as the organization of this thesis is concerned, Chapter 2 introduces fundamental MRI concepts and advanced image reconstruction techniques, including parallel imaging, compressed sensing, and, for quantitative MRI, model-based and subspace-constrained reconstruction. It also explores cardiovascular MRI, focusing on myocardial T_1 mapping. Furthermore, it delves into the basics of deep learning and its integration into image reconstruction, representing the core investigation in this thesis. Chapter 3 outlines the materials and methods exploited in this study. Chapter 4 presents the experimental results. Chapter 5 delves into a detailed discussion and analysis of the results, highlighting the strengths, limitations, and potential future improvements. Chapter 6 presents the conclusions drawn from the current study and highlights potential outlooks for future research.

Chapter 2

Theoretical Background

This chapter introduces fundamental MRI concepts and advanced image reconstruction techniques, including parallel imaging, compressed sensing, and, for quantitative MRI, model-based and subspace-constrained reconstruction. It also explores cardiovascular MRI, focusing on myocardial T_1 mapping. Furthermore, it delves into the basics of deep learning and its integration into image reconstruction, representing the core investigation in this thesis.

2.1 Basics of Magnetic Resonance Imaging

MRI is a widely used non-invasive imaging technique to study the anatomy and function of humans and animals. In contrast to other imaging techniques, such as X-ray Computed Tomography (CT) or Positron Emission Tomography (PET), MRI stands apart by employing neither ionizing radiation nor radioactive substances but static and time-varying magnetic fields. The manipulation of these fields and their timing determines the possibility of acquiring images with different contrasts, such as those based on signal relaxation times T_1 and T_2. For comprehensive information about MRI, it is recommended to refer to the textbook [10].

2.1.1 Nuclear Magnetic Resonance

Nuclear Magnetic Resonance (NMR) is the basis of MRI. It was first described by Rabi in 1938 [11], and then extended by Bloch and Purcell in 1946 [12][13]. NMR is the physical phenomenon in which nuclei in a strong static magnetic field are perturbed by a weak oscillating magnetic field and respond by inducing an

electric current with a frequency related to the strength of the static magnetic field applied in a nearby coil.

Nuclei with odd number of protons and/or neutrons, such as ^{1}H, ^{13}C, ^{19}F, ^{23}Na, and ^{31}P, carry an intrinsic angular momentum $\boldsymbol{I}$, called spin. The magnetic field that arises from the spin of a nucleus is called nuclear magnetic dipole moment, and is related to the spin by

$$\boldsymbol{\mu} = \gamma \boldsymbol{I} \tag{2.1}$$

where γ is the gyromagnetic ratio. The magnitude of $\boldsymbol{\mu}$ is certain and given by

$$|\boldsymbol{\mu}| = \gamma \hbar \sqrt{I(I+1)} \tag{2.2}$$

where $\hbar$ is the Planck's constant divided by 2π, and I is the spin quantum number. By contrast, the direction of $\boldsymbol{\mu}$ is random due to random thermal motion, resulting in zero macroscopic magnetization $\boldsymbol{M}$, where $\boldsymbol{M} = \sum_i \boldsymbol{\mu}_i$. In the presence of a strong static external magnetic field $\boldsymbol{B}_0 = B_0 \boldsymbol{e}_z$, the spins experience the Zeeman effect, which for the hydrogen protons means that they align in two energy states given by

$$E_\uparrow = -\gamma \frac{\hbar}{2} B_0 \quad E_\downarrow = \gamma \frac{\hbar}{2} B_0 \tag{2.3}$$

where $\uparrow$ denotes the parallel and $\downarrow$ the anti-parallel alignment of $\boldsymbol{\mu}$ with $\boldsymbol{B}_0$. According to the Boltzmann statistics, the distribution of the spins in the two population stick to

$$\frac{N_\uparrow}{N_\downarrow} = \exp\left(-\frac{\Delta E}{K_b T}\right) \tag{2.4}$$

where K_b is the Boltzmann constant and T is the temperature. Therefore, the presence of $\boldsymbol{B}_0$ generates an observable macroscopic magnetization $\boldsymbol{M}$, which points along the positive direction of $\boldsymbol{B}_0$ and has a magnitude equal to

$$M_0 = |\boldsymbol{M}| = \rho \frac{\gamma^2 \hbar^2}{4 K_b T} B_0 \tag{2.5}$$

where ρ is the proton density. M_0 is called equilibrium magnetization.

The nuclear magnetic dipole moment $\boldsymbol{\mu}$ precesses around $\boldsymbol{B}_0$ with an angu-

lar frequency equal to

$$\omega_0 = \gamma B_0 \ , \tag{2.6}$$

and the precession is clockwise if observed against the positive direction of $\boldsymbol{B}_0$. In pulsed NMR experiments, typical values of ω_0, known as Larmor frequency, are in the RF band, i.e., in the MHz range.

Excitation

The application of a weak, time-varying magnetic field $\boldsymbol{B}_1(t)$ with frequency ω_0 and perpendicular to $\boldsymbol{e}_z$ rotates the macroscopic magnetization $\boldsymbol{M}$ by an angle α with respect to $\boldsymbol{e}_z$, with $\boldsymbol{M}$ experiencing a torque perpendicular to the direction of $\boldsymbol{M}$ according to

$$\frac{\mathrm{d}\boldsymbol{M}(t)}{\mathrm{d}t} = \gamma \boldsymbol{M}(t) \times \boldsymbol{B}(t) \tag{2.7}$$

where

$$\boldsymbol{B}(t) = \boldsymbol{B}_0 + \boldsymbol{B}_1(t) = \begin{pmatrix} 0 \\ 0 \\ B_0 \end{pmatrix} + \begin{pmatrix} B_1(t)\sin(\omega_0 t) \\ B_1(t)\cos(\omega_0 t) \\ 0 \end{pmatrix} \ . \tag{2.8}$$

During the excitation, $\boldsymbol{M}$ is tilted towards the xy-plane where it processes with ω_0. The rotation of $\boldsymbol{M}$ is a spiraling-downward motion due to the continuing precession about the direction of $\boldsymbol{B}_0$. The flip angle α is proportional to the integral over the envelope $B_1(t)$ of the pulse. $\boldsymbol{B}_1(t)$ is also known as RF pulse because of the frequency band of ω_0. A pulse which rotates the magnetization vector by a certain flip angle α will be called an α-pulse (e.g., a 90°-pulse) in the following. The physical phenomenon underlying the aforementioned excitation is called resonance.

Relaxation

Once the RF pulse is turned off, the system tends toward its equilibrium state and $\boldsymbol{M}$ realigns with $\boldsymbol{B}_0$, releasing electromagnetic energy in the so-called relaxation process. This process can be described by the Bloch equations:

$$\frac{\mathrm{d}\boldsymbol{M}(t)}{\mathrm{d}t} = \gamma \boldsymbol{M}(t) \times \boldsymbol{B}(t) + \begin{pmatrix} -\frac{1}{T_2} M_x(t) \\ -\frac{1}{T_2} M_y(t) \\ \frac{M_0 - M_z(t)}{T_1} \end{pmatrix} \ . \tag{2.9}$$

T_1 represents the time constant governing the spin-lattice relaxation. This process involves the absorption of released energy by the lattice and pertains to the relaxation of the longitudinal magnetization M_z. T_2 is the time constant associated with the spin-spin relaxation. This mechanism involves the absorption of the released energy by other spins and pertains to the relaxation of the transversal magnetization. T_1 and T_2 are specific characteristics of the sampled material, allowing for the distinction of different tissues. To simplify the analysis, the two components of the transversal magnetization M_x and M_y are usually combined into one single complex-valued quantity:

$$M_\perp(t) = M_x(t) + iM_y(t) \ . \tag{2.10}$$

In the absence of $\boldsymbol{B}_1(t)$, the solution of Equation 2.9 with respect to $M_\perp(t)$, where $M_\perp$ satisfies the differential equation $\frac{\mathrm{d}M_\perp(t)}{\mathrm{d}t} = -\left(i\gamma B_z(t) + \frac{1}{T_2}\right) M_\perp(t)$, and M_z is given by

$$M_\perp(t) = M_\perp(0)e^{-t/T_2}e^{-i\gamma\int_0^t B_z(\tau)d\tau} \tag{2.11}$$

$$= M_\perp(0)e^{-t/T_2}e^{-i\omega_0 t} \tag{2.12}$$

$$M_z(t) = M_0 + (M_z(0) - M_0)e^{-t/T_1} \tag{2.13}$$

where $M_\perp(0)$ and $M_z(0)$ are the transversal and longitudinal magnetization immediately after turning off the RF pulse.

Free Induction Decay

The application of an RF pulse $\boldsymbol{B}_1(t)$ in the presence of a static magnetic field $\boldsymbol{B}_0$ results in a nonzero, time-varying transversal magnetization $M_\perp(t)$ precessing with Larmor frequency ω_0 around $\boldsymbol{B}_0$. $M_\perp(t)$ can be captured by placing a receiver coil in the xy-plane: according to the Faraday-Lens's law of induction, the time rate of change of the magnetic flux of $M_\perp(t)$ will induce in the coil an electromotive force, which is called Free Induction Decay (FID) signal. The FID signal $s(t)$ oscillates at the Larmor frequency ω_0 and can be modeled by the transversal magnetization $M_\perp(t)$. Theoretically, the Larmor frequency ω_0 is uniform across all positions within $\boldsymbol{B}_0$. In practice, external field inhomogeneities lead to fluctuations in the Larmor frequency as position shifts. Consequently, dephasing accelerates, and the transversal relaxation occurs more rapidly than what T_2 describes. The signal decays with an effective spin-spin relaxation time T_2^* smaller than T_2.

2.1.2 Spatial Encoding

The FID signal generated by the application of an RF pulse $\boldsymbol{B}_1(t)$ in the presence of a static magnetic field $\boldsymbol{B}_0$ carry any information concerning the position of the spins generating the signal. To inject spatial information into the acquired signal, magnetic field gradients come into play. By adding gradients to $\boldsymbol{B}_0$, the resulting magnetic field varies in space:

$$\boldsymbol{B}(\boldsymbol{r},t) = \boldsymbol{B}_0 + \boldsymbol{G}(t)\cdot\boldsymbol{r} = \begin{pmatrix}0\\0\\B_0\end{pmatrix} + \begin{pmatrix}G_x\\G_y\\G_z\end{pmatrix}\cdot\begin{pmatrix}x\\y\\z\end{pmatrix} \tag{2.14}$$

Consequently, an additional phase term is generated on the transversal magnetization:

$$M_\perp(\boldsymbol{r},t) = M_\perp(\boldsymbol{r},0)e^{-t/T_2^*(\boldsymbol{r})}e^{-i(\omega_0 t+\gamma\int_0^t \boldsymbol{G}(\tau)d\tau\cdot\boldsymbol{r})} \tag{2.15}$$

$$= M_\perp(\boldsymbol{r},0)e^{-t/T_2^*(\boldsymbol{r})}e^{-i(\omega_0 t+2\pi\boldsymbol{k}(t)\cdot\boldsymbol{r})}\ . \tag{2.16}$$

Here, to describe that the time evolution of the gradient induced spatial phase variations, the k-space trajectory is defined as

$$\boldsymbol{k}(t) = \frac{\gamma}{2\pi}\int_0^t \boldsymbol{G}(\tau)d\tau\ , \tag{2.17}$$

and the accumulated phase due to the application of the gradients is

$$\phi(\boldsymbol{r},t) = \gamma\int_0^t \boldsymbol{G}(\tau)d\tau\cdot\boldsymbol{r} = 2\pi\boldsymbol{k}(t)\cdot\boldsymbol{r}\ . \tag{2.18}$$

The signal detected by the receiver coil can be modeled by the integral

$$\int_V M_\perp(\boldsymbol{r},t)d\boldsymbol{r}\ . \tag{2.19}$$

Removing the term $e^{-i\omega_0 t}$ by quadrature demodulation, omitting the relaxation term $e^{-t/T_2^*(\boldsymbol{r})}$, and assuming that the initial transversal magnetization $M_\perp(\boldsymbol{r},0)$ is proportional to the spin density ρ, the signal is then proportional to the Fourier transform of ρ, sampled on a k-space trajectory $\boldsymbol{k}(t)$:

$$s(\boldsymbol{k}(t)) \propto \int_V \rho(\boldsymbol{r})e^{-i2\pi\boldsymbol{k}(t)\cdot\boldsymbol{r}}d\boldsymbol{r} \tag{2.20}$$

or

$$s(k_x(t), k_y(t), k_z(t)) \propto \int\int\int \rho(x, y, z) e^{-i2\pi(k_x(t)x + k_y(t)y + k_z(t)z)} dx\ dy\ dz \quad (2.21)$$

where

$$k_x(t) = \frac{\gamma}{2\pi} \int_0^t G_x(\tau) d\tau \quad k_y(t) = \frac{\gamma}{2\pi} \int_0^t G_y(\tau) d\tau \quad k_z(t) = \frac{\gamma}{2\pi} \int_0^t G_z(\tau) d\tau\ . \quad (2.22)$$

Equation 2.22 shows that, by independently varying the gradients G_x, G_y, and G_z, it is possible to sweep over the 3D k-space. Then, the challenge is to manipulate the application of the gradients in such a way that the 3D k-space can be sampled sufficiently. This is required in order that the reconstructed image $\hat{\rho}(\boldsymbol{r})$ be an accurate estimate of $\rho(\boldsymbol{r})$, where $\hat{\rho}(\boldsymbol{r})$ is the inverse Discrete Fourier Transform (DFT) of the measured data. According to the Nyquist-Shannon sampling theorem, to avoid aliasing the maximum distance of samples in k-space Δk and the desired Field Of View (FOV) of the object should satisfy

$$\Delta k \leq \frac{1}{\text{FOV}}\ . \quad (2.23)$$

An acquisition is "fully sampled" if Equation 2.23 holds. To spare acquisition time, it is generally desirable to perform undersampling acquisitions, acquiring fewer data than what the Nyquist-Shannon condition demands. Section 2.2 presents the concepts of parallel imaging and compressed sensing, which circumvent the Nyquist-Shannon condition and allow for suitable reconstructions even for high undersampling factors.

2.1.3 Slice Selection

Slice selection aims to excite only a slice of the sample. A gradient is applied along one of the axes, e.g., z-axes, resulting in a Larmor frequency varying linearly along the gradient direction:

$$\omega(z) = \gamma(B_0 + G_z \cdot z) = \omega_0 + \gamma G_z \cdot z\ . \quad (2.24)$$

As a result, by applying an RF pulse with a bandwidth $\Delta\omega$, only the spins with the corresponding Larmor frequencies will be excited. Once a slice has been selected, the remaining 2D plane can be spatially encoded.

The slice selected gradient induces dephasing in the slice select direction across the slice thickness Δz, and the procedure to correct for this phase accumulation is to have the gradient reversed after the RF pulse is turned off, such that the spins realign in the transversal plane at the end of the reversed gradient lobe.

2.1.4 k-Space Trajectories

In principle, it is possible to generate any k-space trajectory by switching gradients along time according to Equation 2.17. An example of 2D Cartesian trajectory is shown in Figure 2.1a. k-space lines are collected line-by-line sequentially to fill in a Cartesian grid. An example of 2D radial sampling trajectory is shown in Figure 2.1b. k-space lines, known also as a radial spoke or simply spoke, are still collected line-by-line sequentially, but in a star-shaped pattern. Thus, each spoke is defined by a specific angle and passes through the center of k-space. It is accomplished by setting the gradients G_x and G_y according to

$$G_x = G_{max} \cdot \cos(\theta) \quad G_y = G_{max} \cdot \sin(\theta) \tag{2.25}$$

where G_{max} denotes the amplitude required for sampling the central k-space row ($k_y = 0$), and θ is the desired angle of the spoke. To date, the most used rotation scheme is golden-angle radial sampling, in which spokes are rotated by a so-called "golden angle" [14]. Golden-angle based strategy covers the entire k-space without repeating any angles and allows for arbitrary binning of spokes to form one k-space frame [14]. A novel extension of the standard golden-angle radial sampling is the tiny golden angle radial sampling [15].

In the case of Cartesian sampling, every sample lies on the Cartesian grid, and this allows for simple reconstruction with a Fast Fourier Transform (FFT), i.e., an efficient algorithm for calculating the DFT, when the Nyquist-Shannon condition (Equation 2.23) is fulfilled. As for radial sampling, the samples do not necessarily lie on the Cartesian grid, preventing direct FFT reconstruction. Reconstructing radial k-space data involves a process called "gridding" before FFT, where radial k-space data are interpolated onto a Cartesian grid. The entire reconstruction process is also known as non-uniform Fast Fourier Transform (nuFFT) [16]. Furthermore, Cartesian sampling is insensitive to gradient imperfections, while radial sampling is heavily affected by inaccuracies of gradient timing, which

can result in deviation of actually sampled k-space locations from the theoretical sampling positions (gradient delay effect). As a result, Cartesian sampling is the most common trajectory employed in clinical practice. However, radial sampling combines the advantages of line-by-line scanning with better undersampling behavior and is more robust to motion. Thanks to significant advancements in modern MRI scanners and methods for gradient delay compensation, e.g., the RING method [17], the overall performance of radial MRI has been significantly improved. As a result, radial sampling now serves as a valuable alternative to Cartesian sampling when enhanced motion robustness is required [18].

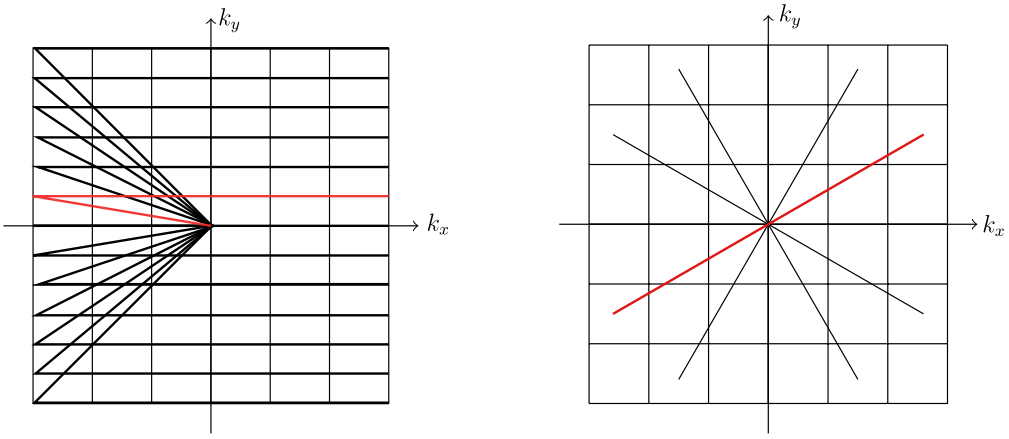

(a) Cartesian FLASH. (b) Cartesian trajectory.

Figure 2.1: Schematic representation of Cartesian (a) and radial (b) k-space trajectory. A schematic of the gradients used to acquire the red k-space lines can be found in Figure 2.2.

2.1.5 Pulse Sequence

Up until now, the focus has centered solely on the FID signal. However, the realm of MRI encompasses a broader spectrum of signals. Within this context, pulse sequences take center stage. A pulse sequence refers to a series of RF pulses and magnetic field gradients designed to generate a meaningful MRI signal, allowing for the reconstruction of an image with a specific contrast. The convention will be to take the two imaging dimensions in the xy-plane, implying that the slice select gradient is along the z-axis [19].

Cartesian Gradient Echo

A slice select gradient G_z is applied along the z-axis during the RF pulse for a total time of τ_{RF}. Since the slice selected gradient induces dephasing in the slice select direction across the slice thickness Δz, the gradient is reversed immediately after the RF pulse is turned off for a time of $\tau_{RF}/2$. At the end of the rephase lobe of the slice select gradient, all the transversal magnetization components within the slice are in phase with a common accumulated phase value $\phi = 0$. For a boxcar gradient, the signal immediately following the slice selection is given by

$$s(\tau_{RF}) \propto \int \int \left(\int_{z_0 - \frac{\Delta z}{2}}^{z_0 + \frac{\Delta z}{2}} \rho(x, y, z) dz \right) dx \; dy \; . \tag{2.26}$$

After the slice select gradient is turned off, a phase encoding gradient G_y is applied along the y-axis for a time of τ_y, and during this interval the magnetization accumulates a y-dependent phase. The y-axis determines the so-called "phase encoding" direction. Assuming for the moment that no other gradient is applied during this period and that the applied gradient is a boxcar gradient, the signal immediately following the phase encoding is given by

$$s(\tau_{RF} + \tau_y) \propto \int \left(\int \left(\int_{z_0 - \frac{\Delta z}{2}}^{z_0 + \frac{\Delta z}{2}} \rho(x, y, z) dz \right) e^{-i2\pi k_y(G_y) y} dy \right) dx \tag{2.27}$$

where

$$k_y(G_y) = \frac{\gamma G_y \tau_y}{2\pi} \; . \tag{2.28}$$

To gather information about the y dependence of the spin density, the independent variable G_y will be varied in a step-like fashion with step sizes ΔG_y. This will be a repeated experiment where τ_y stays the same in every cycle.

Finally, a read gradient G_x is applied along the x-axis, which determines the so-called "frequency encoding" direction. G_x has a first negative dephasing lobe followed by a second positive read (or rephasing) lobe, during which the signal is measured. The second lobe aims to generate an echo, where the time echo T_E corresponds to that time during the read lobe where the evolved area under the second lobe just cancels the area of the first lobe. In terms of time-shifted variable $t' = t - T_E$, the signal during the application of the read gradient is given

by

$$s(t', G_y) \propto \int \left(\int \left(\int_{z_0 - \frac{\Delta z}{2}}^{z_0 + \frac{\Delta z}{2}} \rho(x, y, z) dz \right) e^{-i2\pi k_y(G_y) y} dy \right) e^{-2\pi k_x(t') x} dx \qquad (2.29)$$

where

$$k_x(t') = \frac{\gamma G_x t'}{2\pi} \ . \qquad (2.30)$$

Equation 2.29 shows that the signal obtained by measuring the received electromotive force over a period of time T_s, in the presence of a gradient echo structure in the read direction after it is phase encoded by a fixed value of G_y, gives a line in the k-space set of 2D Fourier transform values of the 2D spin density for the selected slice. Thus, stepped changes in G_y produces a series of sampled lines in a coverage of the 2D k-space pertaining to the 2D Fourier transform.

To shorten the echo time, the rephase lobe of the slice select gradient, the phase encoding gradient table, and the dephase lobe of the read gradient are all switched on at the same time. The gradient echo is the base of the FLASH sequence.

FLASH Sequence

The FLASH sequence [1] takes advantage of RF pulses with small flip angles. Using RF pulses with small flip angles, most of the longitudinal magnetization remains unaffected, eliminating the need for waiting periods in between successive experiments and reducing the acquisition time. The FLASH sequence is based on the acquisition of the FID in the form of a gradient echo. A slice select gradient is applied during the RF pulse. After termination of the RF pulse the slice select gradient is inverted for proper refocusing of the transversal magnetization. The in-plane spatial discrimination may be achieved by applying a fixed read gradient and a perpendicular phase encoding gradient of variable strength. The read gradient is inverted prior to the data acquisition period, leading to a gradient echo. Immediately after acquisition of the data the experiment is repeated with a repetition time T_R, i.e., the time between successive RF pulses, given by the time needed for slice selection and data acquisition. Thus, the duration of the entire imaging experiment is reduced by the same factor as conventional repetition times of the order of 1 s are reduced to about 10-20 ms. After application of the first 20-40 excitation pulses, the spin system reaches a steady-state where the loss of longi-

tudinal magnetization by excitation is compensated for by transversal relaxation during the imaging sequence. Figure 2.2 reports the pulse diagram for a Cartesian and a radial FLASH sequence.

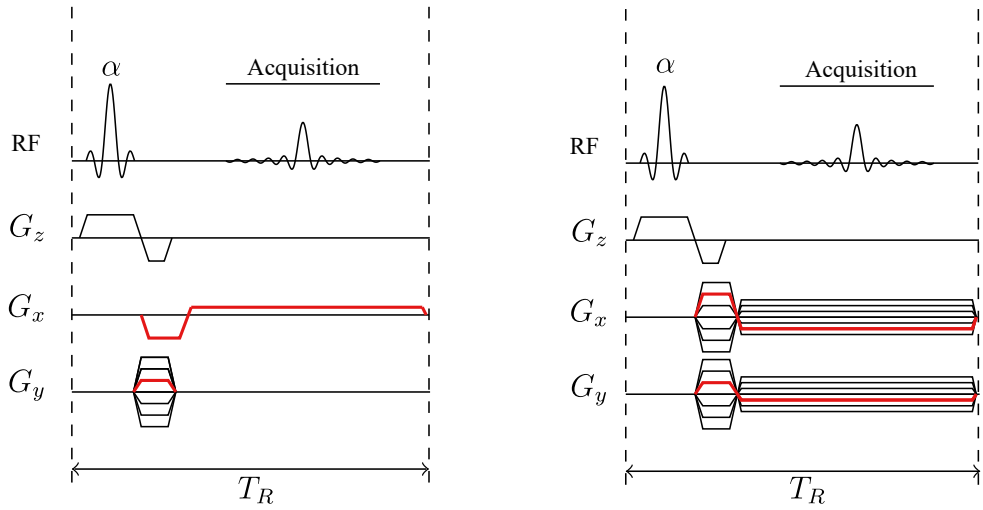

(a) Cartesian FLASH. (b) Radial FLASH.

Figure 2.2: Pulse diagram for a Cartesian (a) and a radial (b) FLASH sequence. The FLASH sequence takes advantage of RF pulses with small flip angles and is based on the acquisition of the FID in the form of a gradient echo. A slice select gradient G_z is applied during the RF pulse. After termination of the RF pulse the G_z is inverted for proper refocusing of the transversal magnetization. The in-plane spatial discrimination may be achieved by applying a fixed read gradient G_x and a perpendicular phase encoding gradient G_y of variable strength. G_x is inverted prior to the data acquisition period leading to a gradient echo. Immediately after acquisition of the data the experiment is repeated with a repetition time given by the time needed for slice selection and data acquisition. The red color highlights the gradient strength used to acquire a k-space line, which for the Cartesian FLASH sequence is schematically depicted in Figure 2.1a, while for the radial FLASH sequence it is represented in Figure 2.1b.

Inversion-Recovery

Inversion-recovery is useful in highlighting differences in T_1 behavior. An inversion-recovery pulse sequence consists of two parts. First, a 90°-pulse (inversion pulse) is applied to invert the magnetization. Second, after a waiting time known as inversion time T_I, a self-contained pulse sequence, e.g., gradient echo, is played out.

2.1.6 MRI Reconstruction as an Inverse Problem

The fundamental idea of MRI reconstruction originates from Equation 2.20 which formalizes the Fourier relation between the MRI signal and the image content: MRI reconstruction refers to the process of recovering an image from the acquired k-space data through inverse Fourier transform. The mathematical relationship mapping the image content into the acquired k-space data is known as the forward model. Consequently, the MRI reconstruction problem is also known as forward problem. Since the number of unknowns in the image, i.e., the number of pixels, is larger than the number of acquired k-space points, the forward problem is underdetermined. For this reason, one may address it in the opposite direction , i.e., given the image, recover the k-space data [4]. Let $\boldsymbol{m} = \rho$ be the unknown image content, $\boldsymbol{y}$ the acquired k-space data, and $\boldsymbol{A}$ the linear forward model that maps the image content $\boldsymbol{m}$ into the k-space data $\boldsymbol{y}$ and includes the a matrix of Fourier coefficients $\boldsymbol{\mathcal{F}}$ and the sampling pattern $\boldsymbol{\mathcal{P}}$, i.e., the projection onto a certain k-space trajectory. Omitting noise, the forward problem may be written as

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{m} \quad \boldsymbol{A} = \boldsymbol{\mathcal{P}}\boldsymbol{\mathcal{F}} \; . \tag{2.31}$$

For fully sampled Cartesian data, $\boldsymbol{A}$ reduces simply to $\boldsymbol{\mathcal{F}}$. Instead of trying to directly invert Equation 2.31, it is more convenient to iteratively estimate the image content $\boldsymbol{m}$ from k-space data $\boldsymbol{y}$ by solving the optimization problem

$$\hat{\boldsymbol{m}} = \arg\min_{\boldsymbol{m}} \underbrace{\|A\boldsymbol{m} - \boldsymbol{y}\|_2^2}_{\text{Data Consistency}} + \underbrace{\lambda \mathcal{R}(\boldsymbol{m})}_{\text{Regularization}} \tag{2.32}$$

where the optimization function is composed of a least-square data consistency term as well as an additional regularization term. This is because the problem is not only underdeterminded but also ill-posed, i.e., small variations in the acquired data can result in large variations in the estimated image, and the noise amplification decreases by adding a regularization term into the inversion. λ is the regularization parameter controlling the balance of noise reduction and the preservation of image details. More details on the MRI modelling are given in 2.2.

2.2 Advanced MRI Reconstruction Techniques

The acquisition of MRI data often requires long acquisition times, leading to patient discomfort and increased susceptibility to motion artifacts. To overcome this challenge, various advanced image reconstruction techniques have been proposed.

2.2.1 Parallel Imaging

Parallel imaging is a technique used in MRI that enables faster data acquisition by simultaneously collecting data from multiple receiver coils. Each coil refers to a distinct spatial sensitivity profile, which serves as an additional spatial encoding function. This characteristic allows for the reduction of scan time by subsampling k-space and reconstructing images exploiting sensitivity information. The MRI signal obtained for multiple receiver coils is given by

$$s_j(\boldsymbol{k}(t)) \propto \int \rho(\boldsymbol{r}) c_j(\boldsymbol{r}) e^{i2\pi \boldsymbol{k}(t)\cdot \boldsymbol{r}} d\boldsymbol{r} \quad j = 1, ..., N \ . \tag{2.33}$$

Here, ρ denotes the proton density, c_j the complex-valued coil sensitivity maps, $\boldsymbol{k}(t)$ the chosen trajectory, and N the number of receiver coils.

Parallel Imaging as a Linear Inverse Problem

SENSitivity Encoding (SENSE) [20] poses the parallel imaging reconstruction problem as a linear inverse problem. Let $\boldsymbol{m}$ be the unknown image content, and $\boldsymbol{y}$ acquired k-space data. The linear forward model that maps the image content $\boldsymbol{m}$ into the k-space data $\boldsymbol{y}$ is:

$$\boldsymbol{A} = \boldsymbol{\mathcal{P}}\boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{C}} \ . \tag{2.34}$$

Here, $\boldsymbol{\mathcal{C}}$ is the multiplication with the coil sensitivity maps, $\boldsymbol{\mathcal{F}}$ is the 2D DFT, and $\mathcal{P}$ is the projection to the sampling pattern. Figure 2.3 explains visually the linear SENSE forward model. The coil sensitivity maps can be determined from autocalibration signal lines, e.g., using the ESPIRiT algorithm [21]. As an alternative, they can be computed directly from the acquired k-space data using the NLINV method, which will be presented in the next section. The forward model $\boldsymbol{A}$ is used to formulate the following regularized linear inverse problem, which can be solved for the image content $\boldsymbol{m}$ given the k-space data $\boldsymbol{y}$:

$$\hat{\boldsymbol{m}} = \arg\min_{\boldsymbol{m}} \|\boldsymbol{A}\boldsymbol{m} - \boldsymbol{y}\|_2^2 + \lambda \mathcal{R}(\boldsymbol{m}) \ . \tag{2.35}$$

Here, $\mathcal{R}(\cdot)$ is the regularization term, and λ is the regularization parameter controlling the balance of noise reduction and the preservation of image details. When a quadratic regularization is used, Equation 2.35 can be solved with very efficient algorithms such as the Conjugate Gradient (CG) method, because the problem is not only ill-posed, but also very large. This is known as CG-SENSE.

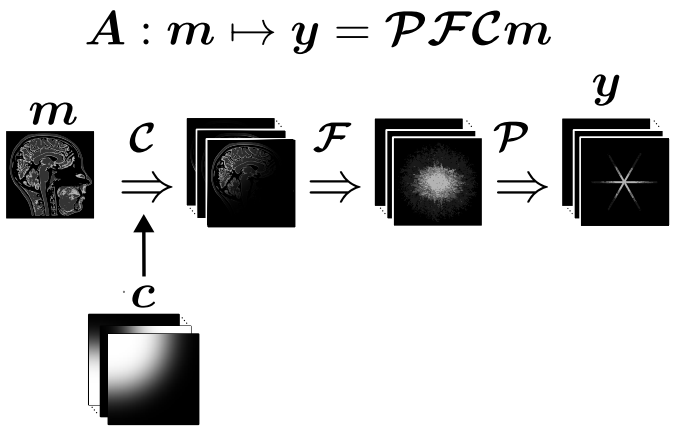

Figure 2.3: The linear SENSE forward model. The linear SENSE forward model $\boldsymbol{A}$ maps the unknown image content $\boldsymbol{m}$ into the acquired k-space data $\boldsymbol{y}$. It can be factorized into the multiplication with the coil sensitivity maps $\boldsymbol{\mathcal{C}}$, the 2D DFT $\boldsymbol{\mathcal{F}}$, and the projection to the sampling pattern $\boldsymbol{\mathcal{P}}$. $\odot$ denotes the Hadamard product, which is also known as the element-wise product. $\boldsymbol{c}$ denotes the coil sensitivity maps. Figure adapted from [7].

Parallel Imaging as a Nonlinear Inverse Problem

NLINV [22] poses the parallel imaging reconstruction problem as a nonlinear inverse problem. It extends the capabilities of the SENSE method by enabling joint estimation of both the image content and the coil sensitivity maps directly from the acquired k-space data, eliminating the need for a separate calibration step and leading to better coil sensitivity maps. Thus, both the image content and coil sensitivity maps are unknown, and the inverse problem that NLINV attempt to solve is nonlinear. Let $\boldsymbol{m}$ be the unknown image content, $\boldsymbol{c}$ the unknown coil sensitivity maps, and $\boldsymbol{y}$ the acquired k-space data. The nonlinear forward model that maps the image content $\boldsymbol{m}$ and the coil sensitivity maps $\boldsymbol{c}$ into the k-space data $\boldsymbol{y}$ is:

$$F(\boldsymbol{x}) = F\begin{pmatrix}\boldsymbol{m}\\ \boldsymbol{c}\end{pmatrix} = \boldsymbol{\mathcal{P}}\boldsymbol{\mathcal{F}}(\boldsymbol{m} \odot \boldsymbol{c}) \ . \tag{2.36}$$

Here, $\odot$ denotes the Hadamard product, which is also known as the element-wise product, $\boldsymbol{\mathcal{F}}$ is the 2D DFT, and $\boldsymbol{\mathcal{P}}$ is the projection to the sampling pattern. Figure 2.4 explains visually the NLINV forward model, i.e., the nonlinear SENSE forward model. The forward model F is used to formulate the following regularized nonlinear inverse problem, which can be solved for the image content $\boldsymbol{m}$ and the coil sensitivity maps $\boldsymbol{c}$ given the k-space data $\boldsymbol{y}$:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \|F(\boldsymbol{x}) - \boldsymbol{y}\|_2^2 + \lambda \mathcal{R}(\boldsymbol{m}) \ . \tag{2.37}$$

Here, $\mathcal{R}(\cdot)$ is the regularization term, and λ is the regularization parameter. Equation 2.37 can be solved using the Iteratively Regularized Gauss-Newton Method (IRGNM) [23], which iteratively approximates the nonlinear minimization problem to the following linear minimization subproblem:

$$\boldsymbol{x}^{n+1} = \arg\min_{\boldsymbol{x}} \|J_F(\boldsymbol{x}^n)(\boldsymbol{x} - \boldsymbol{x}^n) - (\boldsymbol{y} - F(\boldsymbol{x}^n))\|_2^2 + \alpha^n \|\boldsymbol{m}\|_2^2 + \alpha^n \|W\boldsymbol{c}\|_2^2 \ . \tag{2.38}$$

Here, $J_F(\boldsymbol{x}^n)$ denotes the Jacobian of F at the point $\boldsymbol{x}^n$. The regularization parameters are always chosen to be of the form of $\alpha_n = \alpha_0 q^n$ with the same $q \in (0, 1)$, usually $q = \frac{1}{2}$. For more details, see [22].

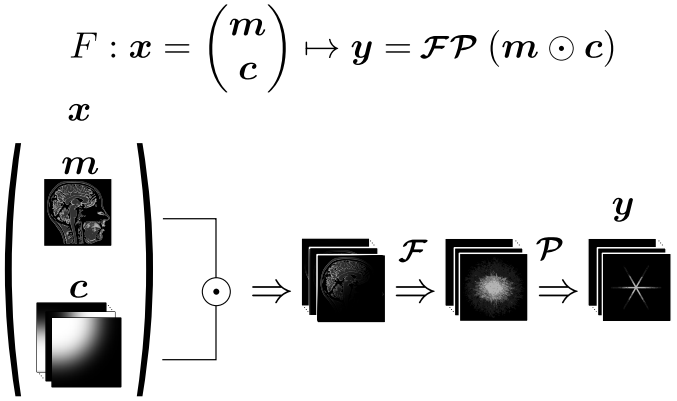

Figure 2.4: The NLINV forward model, i.e., the nonlinear SENSE forward model. F which maps the unknown image content $\boldsymbol{m}$ and the coil sensitivity maps $\boldsymbol{c}$ into the acquired k-space data $\boldsymbol{y}$. It can be factorized into the Fourier transform $\mathcal{F}$, and the projection to the sampling pattern $\boldsymbol{\mathcal{P}}$. $\odot$ denotes the Hadamard product, which is also known as the element-wise product. Figure adapted from [7].

2.2.2 Compressed Sensing

Compressed Sensing [3] [4] is a technique used in MRI to expedite data acquisition by acquiring significantly fewer k-space samples compared to what is traditionally required according to the Nyquist-Shannon sampling theorem (Equation 2.23). The compressed sensing approach requires that: (i) the desired image exhibits a sparse representation in a known transform domain, i.e., it is compressible; (ii) the aliasing artifacts due to k-space undersampling be incoherent in that transform domain; (iii) a nonlinear reconstruction be used to enforce both sparsity of the image reconstruction and consistency with the acquired data.

Advanced MRI reconstruction methods integrate parallel imaging and compressed sensing by using sparsifying regularization terms when solving the inverse problem for parallel imaging. Let $\boldsymbol{m}$ be the unknown image content, $\boldsymbol{y}$ the acquired k-space data, and $\boldsymbol{A}$ the linear SENSE forward operator. The forward model $\boldsymbol{A}$ is used to formulate the following regularized linear inverse problem which can be solved for the image content $\boldsymbol{m}$ given the k-space data $\boldsymbol{y}$ [3]:

$$\hat{\boldsymbol{m}} = \arg\min_{\boldsymbol{m}} \|A\boldsymbol{m} - \boldsymbol{y}\|_2^2 + \lambda \|\Psi \boldsymbol{m}\|_1 \ . \tag{2.39}$$

Here, Ψ is the linear operator that transforms from pixel representation into a sparse representation, and λ is the regularization parameter. Thus, among all solutions which are consistent with the acquired data, Equation 2.39 finds a solution compressible by the transform Ψ. An example of a sparsifying transform is the Wavelet transform.

2.3 Quantitative MRI

While conventional MRI aims primarily at creating specific anatomical contrast, e.g., T_1-weighted contrast, quantitative MRI aims to provide quantitative measurements of specific physical parameters, e.g., T_1, that, ideally, are directly comparable across imaging sites, time points, and different MRI acquisition methods. In other words, a conventional T_1-weighted image acquired on the same patient but on scanners produced by different vendors may show different values for the same tissue, and a simple calibration by rescaling the image intensities would usually not resolve this issue because the intensity values typically depend nonlinearly on the

settings in the imaging protocol and the MRI parameters. A perfect quantitative T_1 map would indicate the same value at both scanners, facilitating standardized diagnosis and multicenter or longitudinal studies [24]. This section presents an advanced MRI reconstruction method for quantitative MRI, i.e., model-based MRI reconstruction, together with subspace-constrained MRI reconstruction, which the present work exploits to perform quantitative T_1 mapping, that is the final topic of the section.

2.3.1 Model-Based MRI Reconstruction

Model-based reconstruction [7] explicitly models the physical laws that govern the MRI signal generation. Quantitative maps of the underlying physical parameters can then be extracted directly from the measured k-space data without intermediate image reconstruction by fitting the model to the data. The specific signal model depends on the applied sequence protocol and specifies which parameters can be estimated. Often, an analytical model can be derived from the simulation of the Bloch equations.

Let $\boldsymbol{p}$ be the unknown model parameters, and $\boldsymbol{y}$ the acquired k-space data. The nonlinear forward model that maps the model parameters $\boldsymbol{p}$ into the k-space data $\boldsymbol{y}$ is given by

$$F(\boldsymbol{p}) = \boldsymbol{\mathcal{P}}\boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{C}}\mathcal{M}(\boldsymbol{p}) \ . \tag{2.40}$$

Here, $\mathcal{M}$ is the signal model, $\boldsymbol{\mathcal{C}}$ is the multiplication with the coil sensitivity maps, $\boldsymbol{\mathcal{F}}$ is the 2D DFT, and $\boldsymbol{\mathcal{P}}$ is the projection to the sampling pattern. Figure 2.5 explain visually the nonlinear model-based reconstruction forward model.

Model-based MRI reconstruction addresses a fundamental limitation of the classical reconstruction methods, which involve solving the image reconstruction problem for as many images as the number of time points. This limitation becomes particularly challenging when the k-space is highly undersampled to save time, leading to limited measurements available for each time point and making reconstruction difficult without imposing structure. Model-based MRI reconstruction takes a different approach by directly solving the image reconstruction problem for the parameter maps instead of reconstructing each time point individually, achieving a significant reduction of the problem's dimensionality.

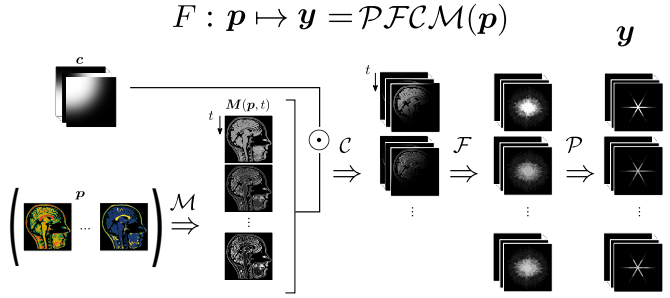

Figure 2.5: The nonlinear model-based reconstruction forward model. The nonlinear model-based reconstruction forward model F maps the unknown parameters $\boldsymbol{p}$ into the acquired k-space data $\boldsymbol{y}$. It can be factorized into the signal model $\mathcal{M}$, the multiplication with the coil sensitivity maps $\boldsymbol{\mathcal{C}}$, the 2D DFT $\boldsymbol{\mathcal{F}}$, and the projection to the sampling pattern $\boldsymbol{\mathcal{P}}$. $\odot$ denotes the Hadamard product, which is also known as the element-wise product. $\boldsymbol{c}$ denotes the coil sensitivity maps. Figure adapted from [7].

2.3.2 Subspace-Constrained MRI Reconstruction

Subspace-constrained MRI reconstruction [7] aims to reduce the dimensionality of the MRI reconstruction problem as the model-based MRI reconstruction does but preserves linearity, i.e., it approximates the MRI signal generation and maintains a linear forward model. For quantitative MRI, the idea is to approximate the signal evolutions in a linear, lower dimensional subspace, reconstruct in the subspace using SENSE, and fit the model to the subspace reconstructions. The subspace basis is computed as follows:

- A dictionary of representative signal curves is generated by evaluating the signal model over a range of model parameters, and the signal curves are summarized into a $m \times n$ matrix $\boldsymbol{X}$.

- The Singular Value Decomposition (SVD) of the matrix $\boldsymbol{X}$ is computed:

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T \tag{2.41}$$

  Here, $\boldsymbol{U}$ is a $m \times m$ orthogonal matrix, i.e., $\boldsymbol{U}\boldsymbol{U}^T = \boldsymbol{I}$ and its columns are called left-singular vectors, $\boldsymbol{\Sigma}$ is a $m \times n$ rectangular diagonal matrix and its diagonal entries of $\boldsymbol{\Sigma}$ are called singular values, $\boldsymbol{V}$ is a $n \times n$ orthogonal matrix, i.e., $\boldsymbol{V}\boldsymbol{V}^T = \boldsymbol{I}$ and its columns are called right-singular vectors.

- The first K left-singular vectors are taken as the subspace basis:

$$\boldsymbol{\mathcal{B}}_K = \boldsymbol{U}[:, 1:K] \tag{2.42}$$

Given the subspace basis $\boldsymbol{\mathcal{B}}_K$, the signal model can be interpreted as a linear combination of the K basis functions:

$$\mathcal{M}(\boldsymbol{r}, t) \approx \boldsymbol{\mathcal{B}}_K(t) \cdot \boldsymbol{a}_K(\boldsymbol{r}) = \sum_{k=1}^{K} \boldsymbol{\mathcal{B}}_k(t) \cdot \boldsymbol{a}_k(\boldsymbol{r}) \ . \tag{2.43}$$

Here, $\boldsymbol{r}$ is the position, t is the time, $\boldsymbol{a}_K$ is the matrix summarizing the K coefficient maps, and $\boldsymbol{\mathcal{B}}_K$ is the matrix summarizing the K basis functions. In the following, $\boldsymbol{a}$ will refer to $\boldsymbol{a}_K$ and $\boldsymbol{\mathcal{B}}$ will refer to to $\boldsymbol{\mathcal{B}_K}$.

Let $\boldsymbol{a}$ be the unknown coefficient maps, and $\boldsymbol{y}$ the acquired k-space data. The linear forward model that maps the coefficient maps $\boldsymbol{a}$ into the k-space data $\boldsymbol{y}$ is:

$$\boldsymbol{A} = \boldsymbol{\mathcal{P}\mathcal{F}\mathcal{C}\mathcal{B}} \ . \tag{2.44}$$

Here, $\boldsymbol{\mathcal{B}}$ is the multiplication of the coefficient maps with the subspace basis, $\boldsymbol{\mathcal{C}}$ is the multiplication with the coil sensitivity maps, $\boldsymbol{\mathcal{F}}$ is the 2D DFT, and $\boldsymbol{\mathcal{P}}$ is the projection to the sampling pattern.

Subspace-constrained reconstruction has two main advantages with respect to the conventional nonlinear model-based reconstruction. Firstly, it is computationally more efficient. Since the subspace-constrained forward operator is linear, $\boldsymbol{\mathcal{P}\mathcal{F}\mathcal{C}\mathcal{B}} = \boldsymbol{\mathcal{P}\mathcal{B}\mathcal{F}\mathcal{C}}$. For this reason, it is possible to compute the Fourier transform K times, where K is the number of basis functions and of coefficient maps, instead of computing it T times, where T is the number of time points. Secondly, having more degrees of freedom, which are equal to the number of basis functions, can be beneficial as it allows for enhanced expressiveness.

Finally, it is worth noting that not all the signal models are low dimensional, meaning that this approach could not work for any signal model. For this reason, it is fundamental to check if the considered data does live in a low-dimensional space, otherwise it will not make sense to proceed with the subspace-constrained reconstruction. For instance, an inversion-recovery signal can be well represented

through singular value decomposition, while a multi-gradient-echo signal cannot. For more information, see [7].

2.3.3 Quantitative T_1 Mapping

This study focuses on quantitative T_1 mapping, which is a quantitative MRI technique used to quantitatively assess the T_1 relaxation times of tissues and display them on a parametric map. The basis of several T_1 mapping methods is the Look-Locker approach.

Inversion-Recovery Look-Locker T_1 mapping

The inversion-recovery Look-Locker sequence comprises an inversion pulse with a spoiler gradient followed by a FLASH sequence to monitor the T_1 relaxation. It may be shown that an effective relaxation time T_1^* shorter than T_1 is observed, and that the magnetization approaches a saturation value M_0^* lower than the equilibrium value M_0 [25]. M_0^* is called steady-state magnetization. After spin inversion, the longitudinal magnetization relaxes according to

$$\mathcal{M}(\boldsymbol{r}, t) = M_0^*(\boldsymbol{r}) - (M_0^*(\boldsymbol{r}) + M_0(\boldsymbol{r})) \cdot e^{-t \cdot R_1^*(\boldsymbol{r})} \tag{2.45}$$

with

$$R_1^*(\boldsymbol{r}) = \frac{1}{T_1^*(\boldsymbol{r})} = \frac{1}{T_1(\boldsymbol{r})} - \frac{1}{T_R} \cdot \log(\alpha) \ , \tag{2.46}$$

and

$$M_0^*(\boldsymbol{r}) = M_0(\boldsymbol{r}) \cdot \frac{T_1^*(\boldsymbol{r})}{T_1(\boldsymbol{r})} \ . \tag{2.47}$$

Here, $\boldsymbol{r}$ denotes the position, t the inversion time, M_0^* the steady-state magnetization, M_0 the equilibrium magnetization, R_1^* the effective T_1 relaxation rate, T_R the repetition time, and α the flip angle. Since knowing the exact flip angle α is impossible, computing T_1 from Equation 2.46 is not trivial. Instead, if a three-parameter fit according to $\mathcal{M}(\boldsymbol{r}, t) = A(\boldsymbol{r}) + B(\boldsymbol{r}) \cdot e^{-t \cdot R_1^*(\boldsymbol{r})}$ is performed, T_1 can be calculated as

$$T_1(\boldsymbol{r}) = -\frac{1}{R_1^*(\boldsymbol{r})} \left(1 + \frac{B(\boldsymbol{r})}{A(\boldsymbol{r})}\right) = \frac{M_0(\boldsymbol{r})}{M_0^*(\boldsymbol{r})} \cdot T_1^*(\boldsymbol{r}) \ . \tag{2.48}$$

Noteworthy, the delay T_D between the inversion and the start of the acquisition would introduce a systematic error in the T_1 evaluation [25]. This delay can amount

to several milliseconds due to the duration of the inversion pulse and the subsequent spoiler gradient. During T_D, there is a T_1 relaxation, so the longitudinal magnetization at the start of the acquisition is $-M_0'$ with

$$M_0'(\boldsymbol{r}) = \left(M_0(\boldsymbol{r}) - 2M_0(\boldsymbol{r}) \cdot e^{-\frac{T_D}{T_1(\boldsymbol{r})}}\right) \approx M_0(\boldsymbol{r}) \cdot \left(1 - 2 \cdot \frac{T_D}{T_1(\boldsymbol{r})}\right) \tag{2.49}$$

Consequently, T_1 evaluation according to 2.47 yields an erroneous value T_1' given by

$$T_1'(\boldsymbol{r}) = T_1(\boldsymbol{r})^* \cdot \frac{M_0^*(\boldsymbol{r})}{M_0(\boldsymbol{r})} = T_1(\boldsymbol{r}) \cdot \frac{M_0'(\boldsymbol{r})}{M_0(\boldsymbol{r})} = T_1(\boldsymbol{r}) - 2 \cdot T_D \tag{2.50}$$

Thus, T_1 values have to be corrected by adding $2 \cdot T_D$.

2.4 Cardiovascular MRI

Cardiovascular MRI is an incredibly potent imaging technique that offers a non-invasive means for the accurate assessment of both the function and structure of the cardiovascular system.

2.4.1 The Cardiovascular System

The cardiovascular system consists of the heart, blood vessels, and blood. The heart (Figure 2.6) is a muscular organ that acts as the central pump of the cardiovascular system. It is composed of three layers. The outer layer is the epicardium. The middle and thicker layer is the myocardium, which is made up of cardiac muscle tissue and is responsible for the heart's pumping action. The inner layer is the endocardium. It has four chambers: two atria and two ventricles. The heart contracts rhythmically to pump oxygen-rich blood to the body's tissues and oxygen-poor blood to the lungs for reoxygenation. Blood vessels form a network of tubes that carry blood. Blood is a fluid consisting of red blood cells transporting oxygen, white blood cells composing the immune system, platelets assisting in clotting, and plasma carrying nutrients and waste products. The cardiac cycle (Figure 2.7) refers to the sequence of events that occur during one complete heartbeat, as the heart contracts and relaxes to pump blood throughout the circulatory system. It can be divided into two main phases: one during which the heart muscle relaxes and refills with blood, called diastole, following a period of robust contraction and pumping of blood, called systole. Assuming a healthy heart and a typical rate of 70

to 75 beats per minute, each cardiac cycle, or heartbeat, takes about 0.8 seconds to complete the cycle. For comprehensive information about the cardiovascular system, it is recommended to refer to the textbook [26].

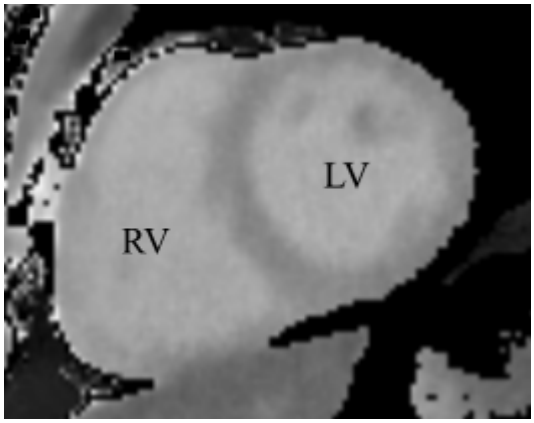

Figure 2.6: Overview of the heart. LV stands for left ventricle, while RV indicates the right ventricle. These two ventricles are distinct chambers separated by the interventricular septum.

2.4.2 Myocardial T_1 Mapping

In the context of cardiovascular MRI, myocardial T_1 mapping provide an imaging biomarker for non-invasive characterization of the myocardial tissue with the potential to replace invasive biopsy for several pathological heart muscle conditions such as fibrosis, iron overload or amyloid infiltration [5]. Wang et al. in 2016 [6] demonstrated that myocardial T_1 mapping can be accomplished by combining single-shot inversion-recovery radial FLASH sequence with breath-hold and finger pulse triggering, i.e., data are acquired within a breath-hold and the inversion pulse is triggered to the early diastolic phase with the use of finger pulse signal, iterative image reconstruction by NLINV, as well as T_1 fitting with automated deletion of systolic frames based on the finger pulse triggering [6]. This method meets the basic requirements for cardiac T_1 mapping, which comprehend T_1 accuracy, speed, practically robustness, and sufficiently high spatial resolution, and is the basis of the myocardial T_1 mapping method presented in this thesis.

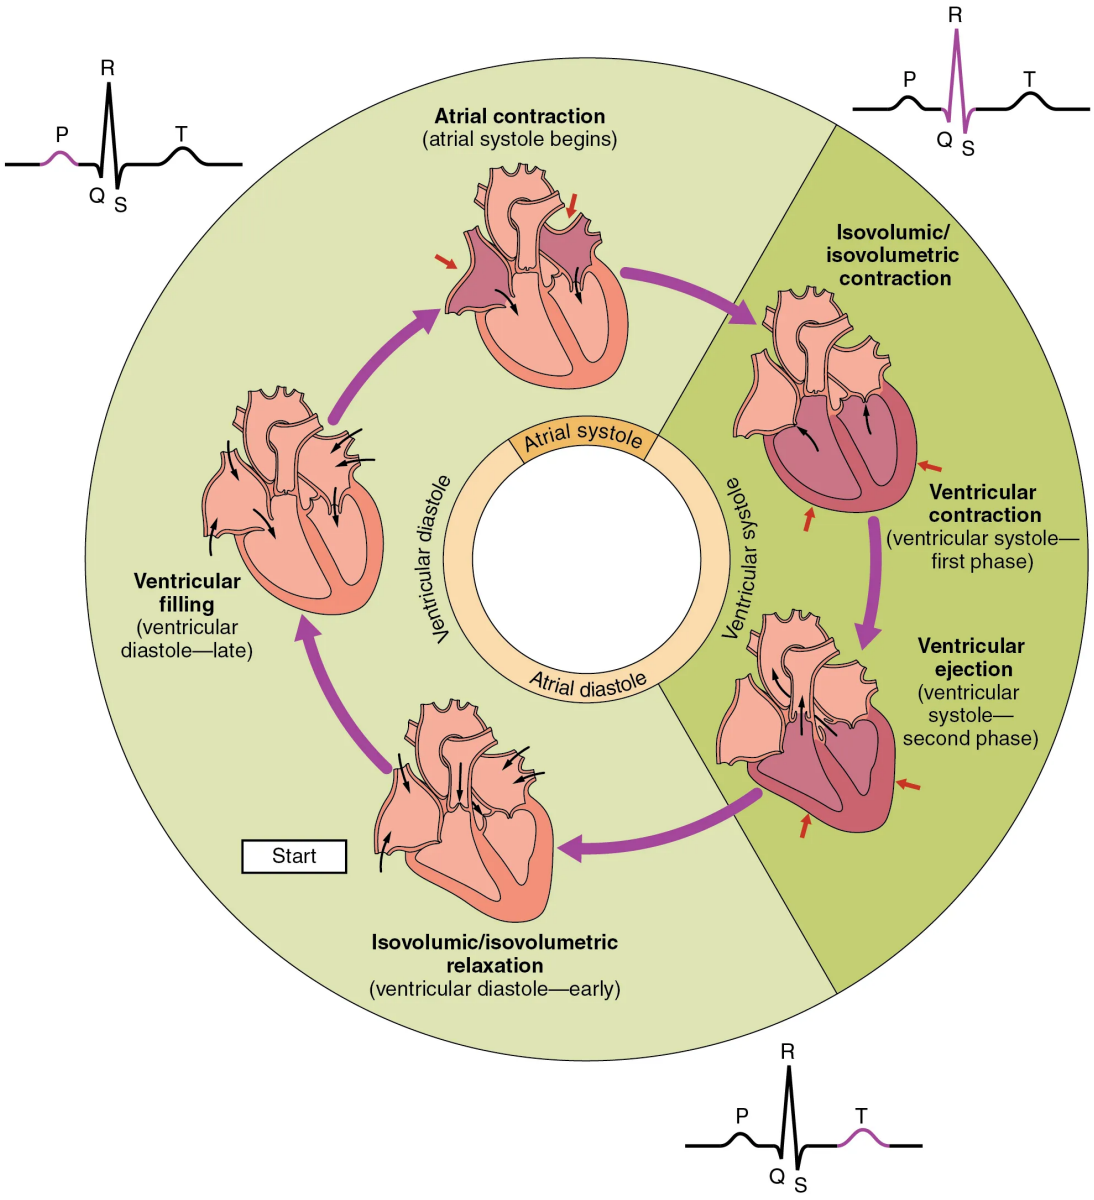

Figure 2.7: Overview of the Cardiac Cycle. The cardiac cycle begins with atrial systole and progresses to ventricular systole, atrial diastole, and ventricular diastole, when the cycle begins again. Correlations to the electrocardiogram (ECG) are highlighted. Figure extracted from [26].

2.5 Deep Learning

The human brain is a complex organ that serves as the command center for the body. It consists of billions of specialized cells called neurons that communicate with each other through electrical and chemical signals. These neurons form intri-

cate networks, allowing information to be processed and transmitted throughout the brain and to different parts of the body. Artificial Neural Network (ANN)s are information processing systems inspired by biological neural networks consisting of interconnected artificial neurons, which communicate by sending signals to one another. Each connection between neurons is associated with a parameter called connection weight or simply weight. Artificial neurons are mathematical models that capture the basic aspects of the computation carried out at the single neuron level. In mathematical terms [8], an ANN is a nonlinear function $F(\cdot)$ mapping the input data $\boldsymbol{x}$ and the weights $\boldsymbol{w}$ to an output $\boldsymbol{y} = F(\boldsymbol{x}; \boldsymbol{w})$. Training an ANN corresponds to fitting the ANN to a training dataset by minimizing some suitable error function E:

$$\hat{\boldsymbol{w}} = \arg\min_{\boldsymbol{w}} E(\boldsymbol{w}) \quad E(\boldsymbol{w}) = \sum_{n=1}^{N} E_n(\boldsymbol{w}) \quad E_n(\boldsymbol{w}) = E(\boldsymbol{y}_n, F(\boldsymbol{x}_n; \boldsymbol{w})) \ . \tag{2.51}$$

Here, N is the number of training samples. Usually, ANNs are constructed from small building blocks such as fully connected layers, convolutional layers, or activation functions. Backpropagation is used to compute the gradients of the loss needed for gradient-based optimization algorithms such as stochastic gradient descent or Adam. In the following, ANN and neural network will refer to the same concept.

2.5.1 Fully Connected Layers

Fully connected or dense layers layers consist of a set of neurons where each neuron is connected to every neuron in the previous layer. Let $\boldsymbol{x}$ be the input data, and $\boldsymbol{w}$ the weights. The output $\boldsymbol{y}$ is obtained by applying an activation function $f(\cdot)$ pointwise to the matrix multiplication $\boldsymbol{w} \cdot \boldsymbol{x}$:

$$\boldsymbol{y} = f(\boldsymbol{w} \cdot \boldsymbol{x}) \ . \tag{2.52}$$

2.5.2 Convolutional Layers

Convolutional layers consist of a set of filters or kernels that convolve across the input data. If the input data is an image, this means that the filter slide over the image, spatially convolving the kernel with each image patch. Let $\boldsymbol{I}$ be the input image and $\boldsymbol{K}$ the kernel with dimension $k_1 \times k_2$. The output $S(i, j)$ of the

convolution operation when the kernel $\boldsymbol{K}$ is in position (i, j) is given by

$$S(i,j) = (\boldsymbol{I} \otimes \boldsymbol{K})_{ij} = \sum_{m=0}^{k_1-1} \sum_{n=0}^{k_2-1} \boldsymbol{I}(i+m, j+n)\boldsymbol{K}(m,n) \ . \tag{2.53}$$

A Convolutional Neural Network (CNN) is a type of ANN that includes at least one convolutional layer in its architecture, which means that it employs convolution in place of matrix multiplication in at least one layer. It is particularly effective in processing data with a known grid-like topology, such as sequences (grid over time) or images (grid over space).

2.5.3 Activation Functions

The activation function determines the firing intensity of the neuron. The input of the activation function is called activation. It is given by the weighted sum of the outputs coming from the neurons in the previous layer. There are several alternatives for activation functions, e.g. linear activation function, sigmoid function, hyperbolic tangent, rectified linear unit.

Linear Activation Function

The linear activation function simply passes the input data through unchanged:

$$\boldsymbol{y} = f(\boldsymbol{w} \cdot \boldsymbol{x}) = \boldsymbol{w} \cdot \boldsymbol{x} \ . \tag{2.54}$$

Stacking together multiple layers with linear activation function results in a linear model. Linear models have limited representation power and can only capture simple linear relationships in the data. Introducing nonlinear activation functions is crucial to enable the neural network to learn and represent more complex and nonlinear relationships. Nonetheless, using a linear activation function as the last layer of a deep neural network can be entirely suitable if the previous layers are nonlinear.

Sigmoid Function

The sigmoid function is an S-shaped, bounded with output range [0,1], differentiable, and monotonic function:

$$\sigma(c) = \frac{1}{1+e^{-c}} \ . \tag{2.55}$$

Since the output range is [0,1], this function is useful when predicting probabilities. The sigmoid function has two main problems. Firstly, the output of the sigmoid neuron saturates at either tail of 0 or 1, and the gradient at these regions is almost zero. During backpropagation (see 2.5.5), the gradient of the activation function, i.e., the local gradient, is one of the factors involved in the computation of the error gradient with respect to the weight. Therefore, if the local gradient is tiny, the resulting gradient will also be little, and almost no signal will flow through the neuron. This phenomenon is known as the vanishing gradients problem. To counter this problem, one must pay extra attention when initializing the weights of sigmoid neurons to prevent saturation. For example, if the initial weights are too large, most neurons will become saturated, and the network will barely learn. Secondly, sigmoid outputs are not zero-centered, which is undesirable since neurons in later layers of processing in a neural network would be receiving data that is not zero-centered, leading to zig-zagging dynamics in weight updates during gradient-based optimization (see 2.5.4): if the data coming to a neuron is always positive, then the gradient on the weights will become during training either all positive or all negative. While this is an inconvenience, it has less severe consequences than the saturated activation problem, as it often is mitigated by updating a large batch of data.

Hyperbolic Tangent

Hyperbolic tangent is an S-shaped, bounded with output range [-1,1], differentiable, and monotonic function:

$$\tanh(c) = \frac{e^{c} - e^{-c}}{e^{c} + e^{-c}} \ . \tag{2.56}$$

Since tanh has saturated regions, similar to the sigmoid function, the use of tanh as an activation function leads to the vanishing gradients problem. However, since the output range is [-1,1], tanh outputs are zero-centered, leading to balanced dynamics in the gradient updates. For this reason, tanh is always preferable with respect to the sigmoid function. Its output is simply the scaled version of the sigmoid output:

$$\tanh(c) = 2\sigma(2c) - 1 \ . \tag{2.57}$$

Rectified Linear Unit

Rectified Linear Unit (ReLU) function it is given by

$$\text{ReLU}(c) = \max(0, c) \ . \tag{2.58}$$

While sigmoid and tanh neurons involve expensive operations, e.g., exponentials, ReLU can be implemented by simply thresholding a matrix of activations at zero. Unfortunately, the ReLU units can irreversibly die during training: a large gradient flowing through a ReLU neuron could cause the weights to update such that the neuron will never activate on any data point again. If this happens, then the gradient flowing through the unit will forever be zero from that point on. A proper setting of the learning rate can solve this problem.

2.5.4 Gradient-Based Optimization Algorithms

A gradient-based optimization algorithm is an iterative method used to find the minimum or maximum of an objective function based on the calculation and utilization of the gradient, which is the vector that indicates the direction of the steepest descent or ascent of the objective function at a specific point. The general idea behind gradient-based optimization is to iteratively update the parameters of a model in the direction of the negative gradient (for minimization) or the positive gradient (for maximization) to converge toward the optimal solution. In a deep learning framework, the model is a neural network. The basic steps involved in a gradient-based optimization algorithm are the initialization of the parameters with some initial values, the computation of the gradient of the objective function with respect to each parameter, the update of the parameters by taking a step in the opposite direction of the gradient (this step is determined by the learning rate, which controls the step size of the parameter update). The gradient computation and parameter update steps are repeated until a convergence criterion is met. This criterion could be a specific number of iterations, reaching a certain threshold for the objective function, or observing small changes in the parameters. Different gradient-based optimization algorithms may have variations in how they update the parameters or adjust the learning rate. Some common gradient-based optimization algorithms include gradient descent, Stochastic Gradient Descent (SGD), mini-batch gradient descent, and Adam. The backpropagation algorithm can be used for gradient computation.

Gradient Descent

Gradient descent is gradient-based optimization algorithm based on the gradients computed from all the training examples at once. The update rule is given by

$$w_{ij}(t+1) = w_{ij}(t) - \eta \frac{\partial E(\boldsymbol{w})}{\partial w_{ij}} \ . \tag{2.59}$$

Here, $w_{ij}(t)$ is the weight of the link between neuron i and neuron j at step t, $w_{ij}(t+1)$ is the weight of the link between neuron i and neuron j at step $t+1$, η is the learning rate, and $E(\boldsymbol{w})$ is the error function computed from all the N training examples as defined in Equation 2.51.

Stochastic Gradient Descent

SGD is a gradient-based optimization algorithm based on the gradients computed from a single training example at a time. Thus, it replaces the actual gradient calculated from the entire data set with an estimate of the gradient calculated from a randomly selected subset of the data. Using the SGD algorithm, instead of taking a smooth descent towards the minimum as for the gradient descent algorithm, the weight vector tends to jitter around the surface of the error function. The result is a random walk, i.e., the path could locally deviate from the good direction. The update rule is given by

$$w_{ij}(t+1) = w_{ij}(t) - \eta \frac{\partial E_n(\boldsymbol{w})}{\partial w_{ij}} \ . \tag{2.60}$$

Here, $E_n(\boldsymbol{w})$ is the error function computed from the n-th training example as defined in Equation 2.51.

Mini-Batch Gradient Descent

Mini-batch gradient descent is a gradient-based optimization algorithm based on the gradients computed from a mini-batch, i.e., a small number of training examples, at a time. The update rule is given by

$$w_{ij}(t+1) = w_{ij}(t) - \eta \frac{\partial E_{\text{batch}_k}(\boldsymbol{w})}{\partial w_{ij}} \ . \tag{2.61}$$

Here, E_{batch_k} is the error computed from the k-th mini-batch.

Adaptive Moment Estimation

Adaptive moment estimation (Adam) [27] is a stochastic gradient-based optimization algorithm that computes individual learning rates for each parameter from adaptive estimates of the first and second moments of the gradients, which are used to update the parameters in a way that is invariant to diagonal rescaling of the gradients. This makes the algorithm well-suited for problems with sparse or noisy gradients, as it can adapt to the local geometry of the objective function and converge more quickly to a good solution. Additionally, Adam is computationally efficient and has low memory requirements, making it a popular choice for large-scale optimization problems in machine learning and other fields. Overall, the algorithm is a powerful and versatile tool for optimizing a wide range of objective functions, and has been shown to outperform other methods in many empirical tests.

2.5.5 Backpropagation

Backpropagation refers to the algorithm used to compute the gradients of the error function with respect to the weights by propagating the error gradients from the output layer to the input layer. The first step of the backpropagation algorithm is the forward pass. The input data is fed into the neural network, and the computations propagate forward through the layers. Each neuron of the current layer applies a nonlinear activation function to the weighted sum of the outputs coming from the neurons in the previous layer and passes its output to the next layer:

$$y_j = f(a_j) \quad a_j = \sum_k w_{kj} y_k \ . \tag{2.62}$$

Here, y_j is the output of the neuron j belonging to the current layer, $f(\cdot)$ is the nonlinear activation function, a_j is the activation of neuron j, i.e., the input of the nonlinear activation function given by the weighted sum of the outputs coming from the neurons in the previous layer, k counts for the neurons in the previous layer, then, y_k are the outputs of the previous layer, and w_{kj} are the weights connecting the current neuron j with the neurons in the previous layer.

The second step is the error computation at the output layer. The error is computed by comparing the predicted output of the neural network with the desired output. The specific error metric depends on the task at hand, e.g., the

Mean Square Error (MSE) can be used for regression. In the context of gradient descent, where the error is computed across all the training examples at once, the MSE is given by

$$E(\boldsymbol{w}) = \frac{1}{N}\sum_{n=1}^{N} \|\boldsymbol{y}_n - \boldsymbol{t}_n\|^2 \tag{2.63}$$

where N counts for the training examples, $\boldsymbol{y}_n$ is the vector summarizing the outputs at the last layer, i.e., the predicted output of the neural network, for the n-th training example, $\boldsymbol{t}_n$ is the vector summarizing the desired outputs at the last layer for the n-th training example.

The third step is the backward pass. The error gradients with respect to the weights are computed at the output layer, and then they are propagated backward through the layers of the network using the same connections that spread the computations forward. Always in the context of gradient descent, the computation of the error gradients with respect to the weights is done using the chain rule twice:

$$\frac{\partial E(\boldsymbol{w})}{\partial w_{ij}} = \frac{\partial E(\boldsymbol{w})}{\partial y_j}\frac{\partial y_j}{\partial w_{ij}} = \frac{\partial E(\boldsymbol{w})}{\partial y_j}\frac{\partial y_j}{\partial a_j}\frac{\partial a_j}{\partial w_{ij}} \tag{2.64}$$

with

$$\frac{\partial a_j}{\partial w_{ij}} = y_i \quad \frac{\partial y_j}{\partial a_j} = \frac{\partial f(a_j)}{\partial a_j}\ . \tag{2.65}$$

Here, $E(\boldsymbol{w})$ is the error function computed for all the training examples, w_{ij} is the weight of the link between neuron i and neuron j, a_j is the activation of neuron j, y_j is the output of the nonlinear activation function for neuron j. $\frac{\partial f(a_j)}{\partial a_j}$ is called local gradient. $\frac{\partial E(\boldsymbol{w})}{\partial y_j}$ is straightforward to compute if the neuron j is in the output layer because, in this case, the $E(\boldsymbol{w})$ depends directly of y_j. If the neuron j is in a hidden layer, the computation is less obvious, and a recursive expression for the gradient is obtained:

$$\frac{\partial E(\boldsymbol{w})}{\partial y_j} = \sum_{\ell}\frac{\partial E(\boldsymbol{w})}{\partial a_\ell}\frac{\partial a_\ell}{\partial y_j} = \sum_{\ell}\frac{\partial E(\boldsymbol{w})}{\partial y_\ell}\frac{\partial y_\ell}{\partial a_\ell}\frac{\partial a_\ell}{\partial y_j} = \sum_{\ell}\frac{\partial E(\boldsymbol{w})}{\partial y_\ell}\frac{\partial y_\ell}{\partial a_\ell}w_{j\ell} \tag{2.66}$$

Here, ℓ counts for the neurons receiving input from neuron j, i.e., the neurons of the next layer. Therefore, at each layer, the gradients with respect to the outputs $\frac{\partial E(\boldsymbol{w})}{\partial y_j}$ can be calculated if all the gradients with respect to the outputs of the next layer $\frac{\partial E(\boldsymbol{w})}{\partial y_\ell}$ are known.

2.5.6 Batch Normalization

Batch normalization [28] is a technique used in deep learning to by normalizing layer inputs for each training mini-batch. It addresses the issue of the internal covariate shift problem, which refers to the change in the distribution of layer inputs during the training process. The basic idea of batch normalization is to normalize the input to a layer by subtracting the mini-batch mean and dividing by the mini-batch standard deviation. This ensures that the inputs have zero mean and unit variance, which helps in stabilizing and accelerating the training process.

2.5.7 Residual Learning

Conventional learning aims to learn the desired underlying mapping from the input to the output vector. Residual learning aims to learn the residual mapping, representing the difference between desired underlying mapping and the input vector, instead of directly learning the desired underlying mapping. A residual block can be expressed through Equations 2.67 and 2.68, as presented in [29]. A Residual neural Network (ResNet) is a type of ANN where residual blocks are stacked together to form the overall architecture of the network.

$$\boldsymbol{y}_\ell = \boldsymbol{x}_\ell + F(\boldsymbol{x}_\ell, \boldsymbol{w}_\ell) \tag{2.67}$$

$$\boldsymbol{x}_{\ell+1} = f(\boldsymbol{y}_\ell) \tag{2.68}$$

Here, $\boldsymbol{x}_\ell$ and $\boldsymbol{x}_{\ell+1}$ are input and output of the ℓ-th residual unit, $\boldsymbol{w}_\ell = \{\boldsymbol{w}_{\ell,k} \mid 1 \leq k \leq K\}$ is the set of weights associated with the ℓ-th residual unit, K is the number of layers in the ℓ-th residual unit, $F(\cdot)$ is a residual mapping, and $f(\cdot)$ is a ReLU function.

Residual Learning for Image Denoising

Image denoising aims to recover a clean image x from a noisy observation $\boldsymbol{y} = \boldsymbol{x}+\boldsymbol{v}$, where $\boldsymbol{v}$ is the measurement noise. One common assumption is that $\boldsymbol{v}$ is additive white Gaussian noise with standard deviation σ. Various methods have been developed to address image denoising, ranging from classical filtering techniques to more advanced algorithms based on statistical models and machine learning. For instance, residual learning can be used for separating noise from noisy observation. The author of [30] proposed denoising CNN (DnCNN) designed to predict the

residual image (i.e., the difference between the noisy observation and the latent clean image) rather than directly outputting the denoised image. In other words, the proposed DnCNN implicitly removes the latent clean image with the operations in the hidden layers. The batch normalization technique is further introduced to stabilize and enhance the training performance of DnCNN. It turns out that residual learning and batch normalization can benefit from each other, and their integration is effective in speeding up the training and boosting the denoising performance.

2.5.8 Self-Supervised Learning

Self-supervised learning refers to a type of machine learning approach where a model learns from unlabelled data by creating surrogate supervisory signals from the data itself. In self-supervised learning, the model is trained to predict or reconstruct missing or corrupted parts of the input data, effectively creating its own supervision without the need for explicit labels. By leveraging the inherent structure or patterns in the unlabelled data, self-supervised learning enables models to learn useful representations and capture meaningful features.

2.6 Deep Learning-Based MRI Reconstruction Techniques

Deep learning-based MRI reconstruction algorithms have emerged as a cutting-edge trend in accelerated MRI research. Some of these techniques, e.g., [9] [31] [32] [33], leverage the concept of unrolling an iterative reconstruction algorithm for a predefined number of iterations. The unrolled network alternates between Data Consistency (DC) and regularization, in which the regularization is implemented through an ANN.

2.6.1 Model-Based Reconstruction using Deep Learned Priors

Let $\boldsymbol{m}$ be the unknown image content, $\boldsymbol{y}$ the acquired k-space data, and $\boldsymbol{A}$ the linear forward model that maps the image content $\boldsymbol{m}$ into the k-space data $\boldsymbol{y}$. MOdel-based reconstruction using Deep Learned priors (MoDL) [33] method uses

the forward model $\boldsymbol{A}$ to formulate the following regularized linear inverse problem which can be solved for the image content $\boldsymbol{m}$ given the k-space data $\boldsymbol{y}$:

$$\hat{\boldsymbol{m}} = \arg\min_{\boldsymbol{m}} \|\boldsymbol{A}\boldsymbol{m} - \boldsymbol{y}\|_2^2 + \lambda \|\mathcal{N}_{\boldsymbol{w}}(\boldsymbol{m})\|_2^2 \ . \tag{2.69}$$

Here, $\mathcal{N}_{\boldsymbol{w}}(\boldsymbol{m})$ is a learned CNN estimator of noise and alias patterns, which depends on the learned parameters $\boldsymbol{w}$. The first term is the DC term, and the second term is the regularization term multiplied by the regularization parameter λ. $\mathcal{N}_{\boldsymbol{w}}(\boldsymbol{m})$ can be expressed as:

$$\mathcal{N}_{\boldsymbol{w}}(\boldsymbol{m}) = \boldsymbol{m} - \mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m}) \ . \tag{2.70}$$

Here, $\mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m})$ is the denoised version of $\boldsymbol{m}$. The use of the CNN-based prior $\|\mathcal{N}_{\boldsymbol{w}}(\boldsymbol{m})\|_2^2$, which gives high values when $\boldsymbol{m}$ is contaminated with noise and alias patterns, results in solutions that are data-consistent and are minimally contaminated by noise and alias pattern. The substitution of Equation 2.70 in Equation 2.69 results in Equation 2.71.

$$\hat{\boldsymbol{m}} = \arg\min_{\boldsymbol{m}} \|\boldsymbol{A}\boldsymbol{m} - \boldsymbol{y}\|_2^2 + \lambda \|\boldsymbol{m} - \mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m})\|_2^2 \ . \tag{2.71}$$

The optimization problem in Equation 2.71 is solved recursively as follows. $\mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m}^n + \Delta\boldsymbol{m})$ can be approximated using Taylor series as

$$\mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m}^n + \Delta\boldsymbol{m}) \approx \mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m}_n) + J^n \Delta\boldsymbol{m} \ . \tag{2.72}$$

Here, n counts for the iterations, and J^n is the Jacobian matrix. Setting $\boldsymbol{m}^n + \Delta\boldsymbol{m} = \boldsymbol{m}$, the regularization term can be approximated as

$$\|\boldsymbol{m} - \mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m})\|_2^2 \approx \|\boldsymbol{m} - \mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m}_n)\|_2^2 + J^n \Delta\boldsymbol{m} \ . \tag{2.73}$$

The Taylor approximation holds only if $\Delta\boldsymbol{m} \to 0$, and under this condition the resulting alternating algorithm that approximates Equation 2.71 is given by

$$\boldsymbol{z}^n = \mathcal{D}_{\boldsymbol{w}}(\boldsymbol{m}^n) \tag{2.74}$$

$$\boldsymbol{m}^{n+1} = \arg\min_{\boldsymbol{m}} \|A\boldsymbol{m} - \boldsymbol{y}\|_2^2 + \lambda \|\boldsymbol{m} - \boldsymbol{z}^n\|_2^2 \ . \tag{2.75}$$

Equation 2.74 refers to the denoising block, and Equation 2.75 refers to the DC block. The sub-problem described in Equation 2.75 can be solved as

$$\boldsymbol{m}^{n+1} = (\boldsymbol{A}^H\boldsymbol{A} + \lambda\boldsymbol{I})^{-1}(\boldsymbol{A}^H\boldsymbol{m} + \lambda\boldsymbol{z}^n) \ . \tag{2.76}$$

The algorithm is initialized with $\boldsymbol{z}^0 = 0$. Once the number of iterations is fixed, the update rules can be viewed as an unrolled linear CNN, whose weights at different iterations are shared. Since the operator $(\boldsymbol{A}^H\boldsymbol{A}+\lambda\boldsymbol{I})$ is not analytically invertible for complex forward model such as in the case of multi-channel MRI, Equation 2.76 is solved using CG optimization. This means that the unrolled linear CNN will have sub-blocks consisting of numerical optimization layers.

2.6.2 NLINV-Net

NLINV-Net [9] is a deep learning-based MRI reconstruction method proposed to jointly estimate the image content and the coil sensitivity maps of radial cardiac MRI data. NLINV-Net unrolls the IRGNM to solve the nonlinear SENSE inverse problem, i.e., Equation 2.36, and the resulting alternating algorithm is given by

$$\boldsymbol{m}^n_{\text{ref}} = \text{Net}(\boldsymbol{m}^n) \tag{2.77}$$

$$\begin{aligned}\boldsymbol{x}^{n+1} = \arg\min_{\boldsymbol{x}} &\|J_F(\boldsymbol{x}^n)(\boldsymbol{x}-\boldsymbol{x}^n) - (\boldsymbol{y} - F(\boldsymbol{x}^n))\|_2^2 \\ &+ (\alpha^n + \lambda)\|\boldsymbol{m} - \boldsymbol{m}^n_{\text{ref}}\|_2^2 + \alpha^n\|W\boldsymbol{c}\|_2^2 \ .\end{aligned} \tag{2.78}$$

Equation 2.77 refers to the denoising block, which consist of a neural network. Equation 2.78 refers to the DC block and is solved using the CG optimization. In the absence of fully sampled data and a ground truth reference, NLINV-Net can be trained using a self-supervised approach [31]. The acquired k-space data $\boldsymbol{y}$ is randomly divided spoke-wise into two disjunct subsets $\boldsymbol{\Theta}$ and $\boldsymbol{\Lambda}$ according to 8:2. $\boldsymbol{\Theta}$ is used by NLINV-Net to estimate the image content and the coil sensitivity maps. From that, $\boldsymbol{\Lambda}$ is predicted by the NLINV forward model presented in Equation 2.36 with the respective pattern. The weights of NLINV-Net, which are shared across iterations, are optimized to minimize the difference between the predicted and actual $\boldsymbol{\Lambda}$. Figure 2.8 depicts a possible architecture for NLINV-Net and the aforementioned self-supervised learning strategy.

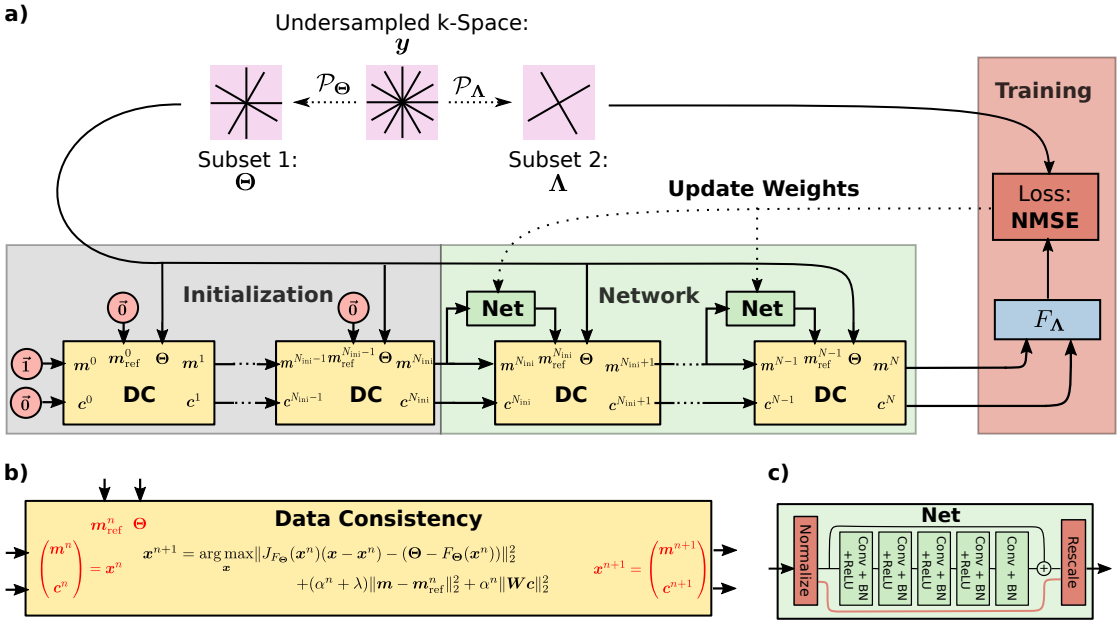

Figure 2.8: NLINV-Net and self-supervised learning. (a) Schematic representation of the self-supervised training strategy. The acquired k-space data $\boldsymbol{y}$ is randomly divided spoke-wise into two disjunct subsets $\boldsymbol{\Theta}$ and $\boldsymbol{\Lambda}$ according to 8:2. $\boldsymbol{\Theta}$ is used by NLINV-Net to estimate the image content and the coil sensitivity maps. From that, $\boldsymbol{\Lambda}$ is predicted by the NLINV forward model with the respective pattern. The weights of NLINV-Net, which are shared across iterations, are optimized to minimize the NMSE of the predicted and actual $\boldsymbol{\Lambda}$. (b) The DC block of NLINV-Net. It is implemented using the CG algorithm. (c) A possible denoising block for NLINV-Net. It is a 5-layer ResNet, where each layer consists of a convolutional layer (conv) with 64 filters followed by batch normalization (BN) and ReLU acting independently on the real and imaginary parts. The input of the denoising block is normalized to a maximum magnitude of one, and its output is rescaled to the original value. Figure adapted from [9].

2.7 Berkeley Advanced Reconstruction Toolbox

All the image reconstruction-related computations in this work were done using the Berkeley Advanced Reconstruction Toolbox (BART). It is a free and open-source image-reconstruction framework for computational MRI developed by the research groups of Martin Uecker (Graz University of Technology), Jon Tamir (UT Austin), and Michael Lustig (UC Berkeley). It consists of a programming library and a toolbox of command-line programs. The library provides common operations on multi-dimensional arrays, Fourier and wavelet transforms, as well as

generic implementations of iterative optimization algorithms. The command-line tools provide direct access to basic operations on multi-dimensional arrays as well as efficient implementations of many calibration and reconstruction algorithms for parallel imaging and compressed sensing [34]. The command structure follows `bart + command + options + input / output`.

Chapter 3

Materials and Methods

This chapter details the materials and methods used to perform subspace-constrained myocardial T_1 mapping. The general procedure for inversion-recovery Look-Locker T_1 mapping involves several key steps: (i) data acquisition using an inversion-recovery FLASH sequence, which is characterized by a specific signal model depending on specific parameters; (ii) subspace basis computation; (iii) subspace-constrained MRI reconstruction to estimate the subspace images and the coil sensitivity maps; (iv) pixel-wise fitting of signal model to the reconstructed images in the subspace to compute the model parameters; (v) Look-Locker computation of the T_1 values based on the estimated model parameters.

3.1 Data Acquisition

The MRI data used in this work was acquired by the research group of Martin Uecker at the University Medical Center Göttingen, Germany with funding of the DZHK on a Siemens Skyra 3T scanner using a 32-channel body coil from thirty-four volunteers without known illness after obtaining written informed consent and with approval of the local ethics committee. The chosen acquisition scheme for myocardial T_1 mapping is illustrated in Figure 3.1, which refers to a single-shot inversion-recovery radial FLASH sequence with breath-hold and finger pulse triggering: data acquisition starts with a non-selective adiabatic 180° inversion pulse, which is triggered to the early diastolic phase using a finger pulse signal. After inversion, the signal is continuously acquired within almost 4 s breath-hold using the FLASH readout to monitor the T_1 relaxation recovery. Only data from the diastolic phase will be retrospectively selected, based on the finger pulse triggering, to

avoid the motion effects of the systolic phase (see 3.2). Myocardial T_1 maps were acquired at a nominal in-plane resolution of $1.0 \times 1.0\text{mm}^2$ and 6mm section thickness, using a FOV= $364 \times 364\text{mm}^2$ in combination with a resolution of 364 complex data points per radial spoke. Other parameters were $T_R/T_E/T_D = 2.18/1.34/15.12\text{ms}$, flip angle 6°. A 9-th tiny golden angle sampling scheme was employed to sample the data continuously during inversion-recovery.

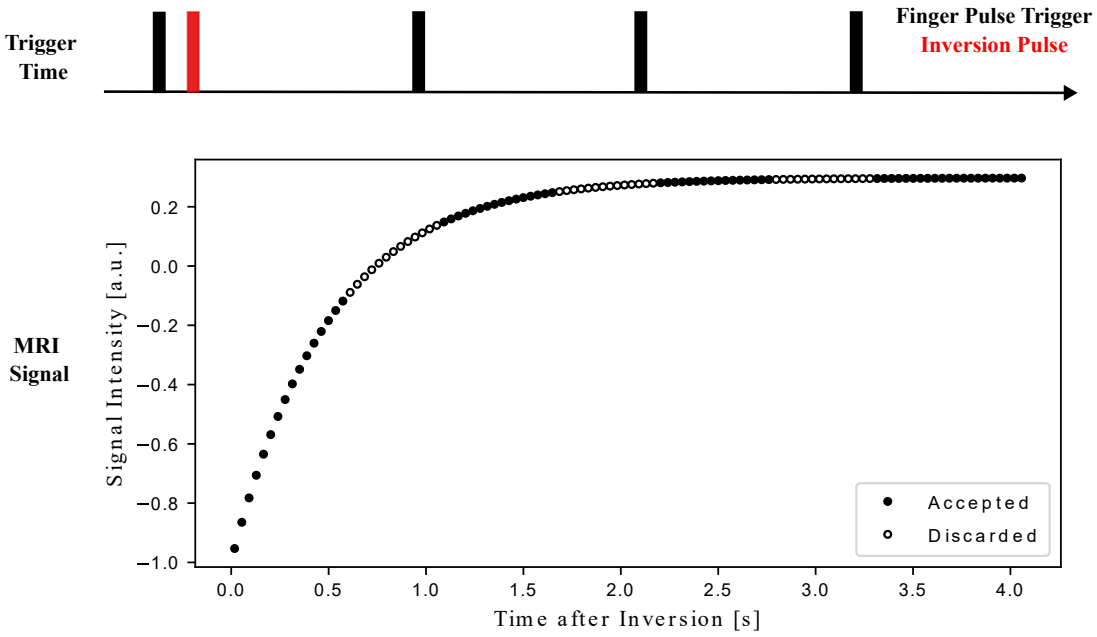

Figure 3.1: Single-shot inversion-recovery radial FLASH sequence with breath-hold and finger pulse triggering. The data acquisition starts with a non-selective inversion pulse, which is triggered to the early diastolic phase using a finger pulse signal. After inversion, the signal is continuously acquired within 4 s breath-hold using the FLASH readout. Only data from the diastolic phase will be retrospectively selected, based on the finger pulse triggering, to avoid the motion effects of the systolic phase. Finger pulse trigger = black bars. Inversion pulse = red bar. Figure adapted from [6].

3.2 Data Preparation

Since binning spokes from neighborhoods with a proper size helps to reduce the computation time as long as the T_1 accuracy is not compromised [35], and golden-angle-based strategy covers the entire k-space without repeating any angles and allows for arbitrary binning of spokes to form one k-space frame [14], 17 spokes

were binned to create one k-space frame resulting in a temporal resolution of almost 36 ms and 110 frames in the present study. As shown in Figure 3.2, k-Space trajectories consistent with the acquired k-space data were generated using the `bart traj` command, by specifying 364 readout samples `-x364`, 1870 (17×110) spokes `-y1870`, radial sampling `-r`, 9-th tiny golden angle `-G -s9`, double-angle `-D`, and 1 turn `-t1`, i,e, single-shot acquisition. Data was corrected for gradient delays. Gradient delays were estimated using the `bart estdelay` command, which implements the RING method [17]. The first six dimensions for the k-space data and k-space trajectories were assigned as reported in Tables 3.1 and 3.2. Gradient delays were computed considering only the last 270 repetitions because the magnetization signal follows the single-shot inversion-recovery FLASH curve and only the last time points are in a steady-state condition. Data was compressed using the `bart cc` tool into 10 virtual coils through Principal Component Analysis (PCA) to reduce the size of the dataset and, consequently, the computational complexity. Only data from the diastolic phase was retrospectively selected using subject-specific masks $\boldsymbol{\mathcal{D}}$ created specifically to handle the motion effects during the systolic phase. The subject-specific masks were designed to have a value of one for diastolic time points and a value of zero for systolic time points.

Dimension	**Usage for k-space data**
0	Not used, set to 1
1	Readout dimension
2	Number of binned spokes per frame
3	Number of coils
4	Not used, set to 1
5	Number of frames

Table 3.1: The first six dimensions for the k-space data.

Dimension	**Usage for k-space trajectories**
0	Number of dimension, set to 3
1	Readout dimension
2	Number of binned spokes per frame
3	Not used, set to 1
4	Not used, set to 1
5	Number of frames

Table 3.2: The first six dimensions for the k-space trajectories.

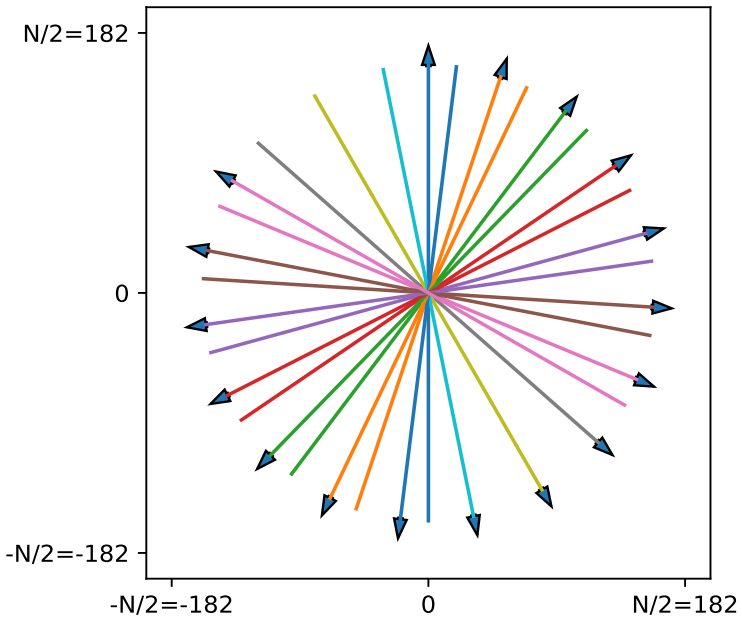

Figure 3.2: k-Space trajectory. k-Space trajectories consistent with the acquired k-space data were generated using the `bart traj` command, by specifying 364 readout samples `-x364`, 1870 (17×110) spokes `-y1870`, radial sampling `-r`, 9-th tiny golden angle `-G -s9`, double-angle `-D`, and 1 turn `-t1`, i,e, single-shot acquisition.

3.3 Subspace Basis Computation

After spin inversion, the longitudinal magnetization relaxes according to Equation 2.45, that was,

$$\mathcal{M}(\boldsymbol{r}, t) = M_0^*(\boldsymbol{r}) - (M_0^*(\boldsymbol{r}) + M_0(\boldsymbol{r})) \cdot e^{-t \cdot R_1^*(\boldsymbol{r})} \ .$$

Here, $\boldsymbol{r}$ is the position, t is the inversion time, M_0^* is the steady-state magnetization, M_0 is the equilibrium magnetization, and R_1^* is the effective T_1 relaxation rate. Based on this equation and on the method to create a subset basis explained in 2.3.2, a dictionary of 100000 representative signal curves was generated using 1000 different R_1^* values ranging from 5 to 5000 ms, combined with 100 M_0^* values ranging from $0.01 \cdot M_0$ to M_0. Since the simulated curves were highly correlated, they can be well represented by a small number of basis functions $\boldsymbol{\mathcal{B}}$, which were

computed by applying SVD and shared across subjects. Figure 3.3 shows a subset of the signal curves composing the simulated signal curves dictionary, the first four basis functions, and the first four basis functions masked to select data from the diastolic phase only. Notably, the basis functions are the same for all subjects, while the masks to discard systolic time points are subject-specific.

3.4 Subspace-Constrained MRI Reconstruction

Given the k-space data $\boldsymbol{y}$, the coefficient maps $\boldsymbol{a}$ were estimated jointly with the coil sensitivity maps $\boldsymbol{c}$ using subspace NLINV-Net (see 3.4.2), which this thesis proposes as a new self-supervised subspace deep learning method for myocardial T_1 mapping to automatically learn a regularization term and increase image quality while decreasing computational efficiency.

3.4.1 Baseline Methods

To evaluate the effectiveness of the proposed neural network-based approach for subspace-constrained myocardial T_1 mapping, plain NLINV (see 3.4.1) and the ℓ_1-Wavelet Parallel Imaging Compressed Sensing (PICS) methods (see 3.4.1) were used as baseline approaches.

Subspace NLINV

Let $\boldsymbol{a}$ be the unknown coefficient maps, $\boldsymbol{c}$ the unknown coil sensitivity maps, and $\boldsymbol{y}$ the acquired k-space data. The nonlinear forward model that maps the coefficient maps $\boldsymbol{a}$ and the coil sensitivity maps $\boldsymbol{c}$ into the k-space data $\boldsymbol{y}$ is

$$F(\boldsymbol{x}) = F\begin{pmatrix}\boldsymbol{a}\\ \boldsymbol{c}\end{pmatrix} = \boldsymbol{\mathcal{P}}\boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{D}}\boldsymbol{\mathcal{B}}(\boldsymbol{a} \odot \boldsymbol{c}) \ . \tag{3.1}$$

Here, $\boldsymbol{\mathcal{B}}$ is the multiplication of the coefficient maps with the subspace basis, $\boldsymbol{\mathcal{D}}$ is the multiplication of the signal evolution with the subject-specific mask selecting only data from the diastolic phase, $\boldsymbol{\mathcal{P}}$ is the sampling pattern, $\boldsymbol{\mathcal{F}}$ is the 2D DFT. The forward model F is used to formulate the following regularized nonlinear inverse problem, which can be solved for the coefficient maps $\boldsymbol{a}$ and the coil sensitivity maps $\boldsymbol{c}$ given the k-space data $\boldsymbol{y}$:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \|F(\boldsymbol{x}) - \boldsymbol{y}\|_2^2 + \lambda\mathcal{R}(\boldsymbol{m}) \ . \tag{3.2}$$

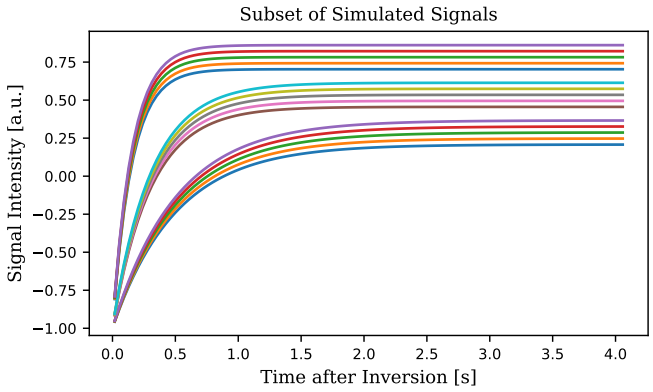

(a) A subset of the 100000 representative signal curves generated using 1000 different R_1^* values ranging from 5 to 5000 ms, combined with 100 M_0^* values ranging from $0.01 \cdot M_0$ to M_0

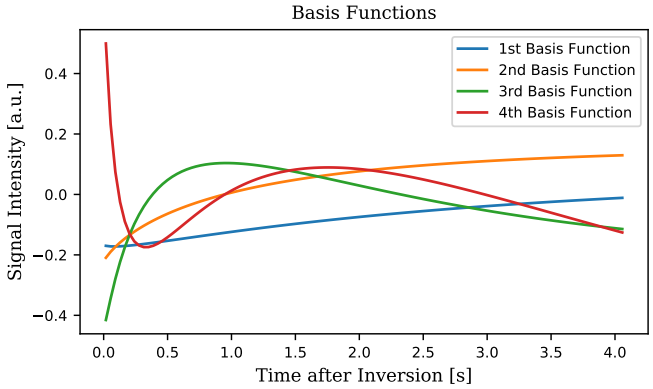

(b) The first four basis functions obtained by applying singular value decomposition to the signal curves dictionary.

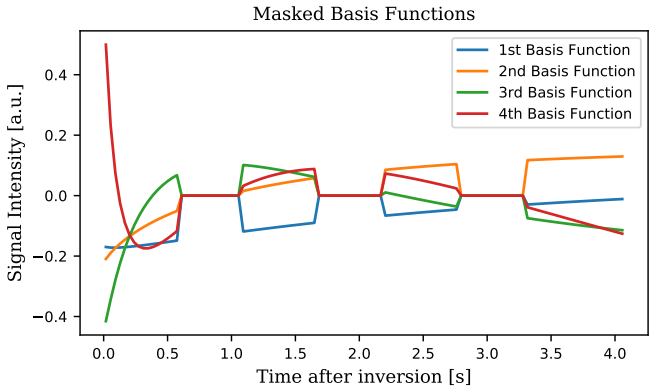

(c) The basis functions masked to select data from the diastolic phase only and deal with motion effects.

Figure 3.3: Basis function computation.

Here, $\mathcal{R}(\cdot)$ is the regularization term, and λ is the regularization parameter. The regularized nonlinear inverse problem is solved using the IRGNM, which itera-

tively approximates the nonlinear problem to the following linear minimization sub-problem:

$$\boldsymbol{x}^{n+1} = \arg\min_{\boldsymbol{x}} \|J_F(\boldsymbol{x}^n)(\boldsymbol{x} - \boldsymbol{x}^n) - (\boldsymbol{y} - F(\boldsymbol{x}^n))\|_2^2 + \alpha^n \|\boldsymbol{a}\|_2^2 + \alpha^n \|W\boldsymbol{c}\|_2^2 \ . \quad (3.3)$$

The NLINV reconstruction was performed in BART using the `bart nlinv` tool and specifying 9 iterations `-i9`.

Subspace ℓ_1-Wavelet Parallel Imaging Compressed Sensing

The ℓ_1-Wavelet PICS method combines parallel imaging and compressed sensing by using ℓ_1-wavelet regularization when solving the linear inverse problem for parallel imaging, i.e., the SENSE problem. Let $\boldsymbol{a}$ be the unknown coefficient maps, $\boldsymbol{c}$ the coil sensitivity maps estimated by NLINV (see 3.4.1), and $\boldsymbol{y}$ the acquired k-space data. The linear forward model that maps the coefficient maps $\boldsymbol{a}$ into the k-space data $\boldsymbol{y}$ is

$$\boldsymbol{A} = \boldsymbol{\mathcal{P}}\boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{C}}\boldsymbol{\mathcal{D}}\boldsymbol{\mathcal{B}} \ . \quad (3.4)$$

Here, $\boldsymbol{\mathcal{B}}$ is the multiplication of the coefficient maps with the subspace basis, $\boldsymbol{\mathcal{D}}$ is the multiplication of the signal evolution with the mask selecting data only from the diastolic phase, $\boldsymbol{\mathcal{C}}$ is the multiplication with the coil sensitivity maps, $\boldsymbol{\mathcal{P}}$ is the sampling pattern, $\boldsymbol{\mathcal{F}}$ is the 2D DFT. The forward model $\boldsymbol{A}$ is used to formulate the following regularized linear inverse problem which can be solved for the coefficient maps $\boldsymbol{a}$ given the k-space data $\boldsymbol{y}$:

$$\hat{\boldsymbol{a}} = \arg\min_{\boldsymbol{a}} \|\boldsymbol{A}\boldsymbol{a} - \boldsymbol{y}\|_2^2 + \lambda \|\Psi \boldsymbol{a}\|_1 \ . \quad (3.5)$$

Here, the sparsifying transform Ψ is a Wavelet transform, and λ is the regularization parameter to be tuned to balance the preservation of fine details and the residual noise. The regularized linear inverse problem is solved using the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA). The PICS reconstruction was performed in BART using the `bart pics` tool and specifying 9 iterations `-i9` and ℓ_1-Wavelet regularization `-R W`. Note that this wavelet regularization for compressed sensing could also be included in the NLINV pipeline, but so far this variant has not been implemented in BART.

3.4.2 Subspace NLINV-Net

NLINV-Net was used to jointly estimate the coefficient maps $\boldsymbol{a}$ and the coil sensitivity maps $\boldsymbol{c}$ of the radial cardiac MRI data $\boldsymbol{y}$. NLINV-Net unrolls the IRGNM to solve the nonlinear SENSE inverse problem, i.e., Equation 3.2, for 9 iterations but uses no network regularization, i.e., plain NLINV, in the first 6 iterations. The resulting alternating algorithm is given by

$$\boldsymbol{a}_{\text{ref}}^{n} = \text{Net}(\boldsymbol{a}^{n}) \tag{3.6}$$

$$\begin{aligned} \boldsymbol{x}^{n+1} = \arg\min_{\boldsymbol{x}} & \|J_F(\boldsymbol{x}^n)(\boldsymbol{x}-\boldsymbol{x}^n) - (\boldsymbol{y} - F(\boldsymbol{x}^n))\|_2^2 \\ & + (\alpha^n + \lambda)\|\boldsymbol{a} - \boldsymbol{a}_{\text{ref}}^n\|_2^2 + \alpha^n \|W\boldsymbol{c}\|_2^2 \ . \end{aligned} \tag{3.7}$$

Equation 3.6 refers to the denoising block. Equation 3.7 refers to the DC block and is solved using CG optimization. The NLINV-Net reconstruction was performed in BART using the `bart nlinvnet` tool, which has two modes: training and inference. The output of NLINV-Net in the training mode consists of the weights, which are provided afterward in the inference mode to obtain a reconstruction. The dataset was split into 24 training subjects and 10 test subjects. The training subjects were used for training and learning the weights of NLINV-Net, while the test subjects were used for inference and evaluation.

Denoising Block Architecture

The proposed architecture for the denoising block is a ResNet made of 5 layers. Each layer consists of a convolutional layer (Conv) with 32 filters followed by ReLU acting independently on the real and imaginary parts. The last layer does not have the ReLU to avoid truncating the negative part of the learned noise patterns. The ResNet has one input/output channel per coefficient map. Following the residual learning strategy, the learned noise, i.e., the output of the last layer of the ResNet, is added to the input of the denoising block to get the reconstructed image as the output of the denoising block. The input of the denoising block is normalized to a maximum magnitude of one, and its output is rescaled to the original value.

Training

The training step aimed to learn the weights of NLINV-Net. The k-space data, the k-space trajectories, and the subject-specific masks for discarding systolic data

of the 24 training subjects were stacked along the 15-th dimension, as the `bart nlinvnet` command in the training mode requires that independent datasets, i.e., datasets of different subjects, should be stacked along the fiftieth dimension. Then, the stacked variables and the basis functions, where the basis functions are the same for all the subjects, were provided as input to NLINV-Net. In the absence of fully sampled data and a ground truth reference, NLINV-Net was trained using the self-supervised strategy presented in 2.6.2, with the difference that here, for subspace-constrained reconstruction, the objects of estimation are the coefficient maps and not the image content. Figure 3.4 shows the proposed architecture for subspace NLINV-Net and the self-supervised learning strategy in the case of subspace-constrained reconstruction. The weights of NLINV-Net, which are shared across iterations, are optimized to minimize the MSE of the predicted and actual $\boldsymbol{\Lambda}$ as follows:

$$\hat{\boldsymbol{w}} = \arg\min_{\boldsymbol{w}} \sum_{i=1}^{N} \text{MSE}(\mathcal{P}_{\boldsymbol{\Lambda},i} \cdot \boldsymbol{y}_i; F_{\boldsymbol{\Lambda},i}(\text{NLINV-Net}(\mathcal{P}_{\boldsymbol{\Theta},i} \cdot \boldsymbol{y}_i, \boldsymbol{w}))) \ . \tag{3.8}$$

Here, $\mathcal{P}_{\boldsymbol{\Lambda},i}$ is the sampling pattern to get the k-space subset $\boldsymbol{\Lambda}_i$, $\mathcal{P}_{\boldsymbol{\Theta},i}$ is the sampling pattern to get the k-space subset $\boldsymbol{\Theta}_i$, NLINV-Net$(\mathcal{P}_{\boldsymbol{\Theta},i} \cdot \boldsymbol{y}_i, \boldsymbol{w})$ is the output of the last iteration, and $F_{\boldsymbol{\Lambda},i}$ is the forward model including the sampling pattern $\mathcal{P}_{\boldsymbol{\Lambda},i}$. The number of complex-valued trainable parameters were 30,085, the batch size was 12, the learning rate was 0.001, the network was trained for 500, 1000, and 2000 epochs.

Inference

The inference step aimed to reconstruct the coefficient and coil sensitivity maps using the weights learned during the training step. It involved the 10 test subjects. For each test subject, the basis functions, the k-space data, the k-space trajectory, the subject-specific mask for discarding systolic data, and the learned weights were provided as input to the trained NLINV-Net.

3.5 T_1 Quantitation

After reconstruction, the coefficient maps $\boldsymbol{a}$ were fitted pixel-wise to the signal model in the subspace using the Gauss-Newton algorithm through the `bart`

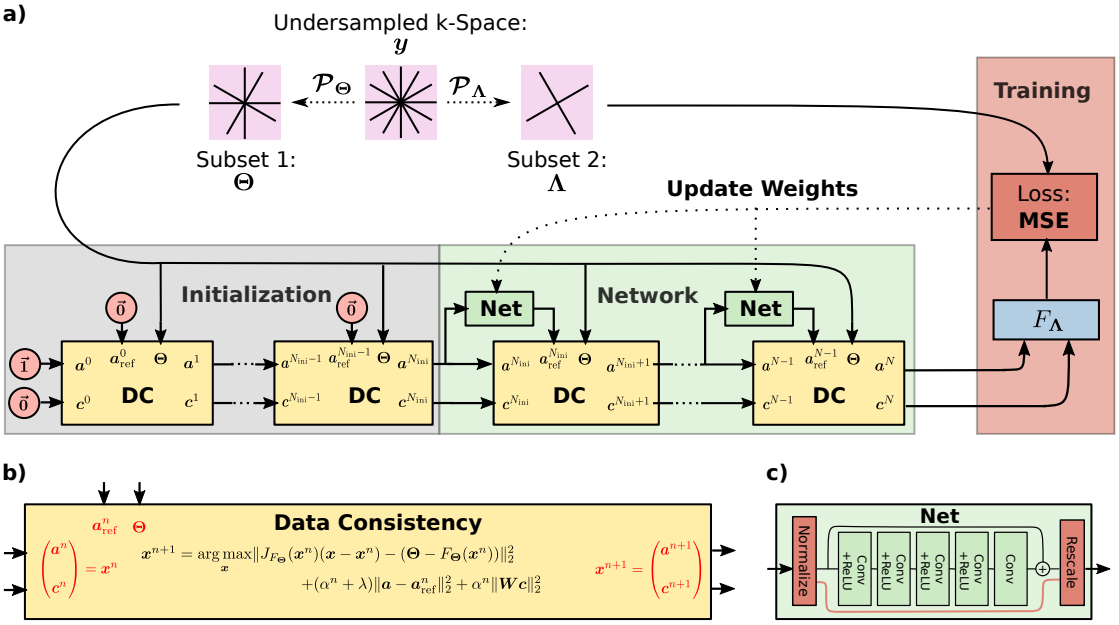

Figure 3.4: Subspace NLINV-Net and self-supervised learning. (a) Schematic representation of the subspace self-supervised training strategy. The acquired k-space data $\boldsymbol{y}$ is randomly divided spoke-wise into two disjunct subsets $\boldsymbol{\Theta}$ and $\boldsymbol{\Lambda}$ according to 8:2. $\boldsymbol{\Theta}$ is used by NLINV-Net to estimate the coefficient maps and the coil sensitivity maps. From that, $\boldsymbol{\Lambda}$ is predicted by the NLINV forward model with the respective pattern. The weights of NLINV-Net are optimized to minimize the MSE of the predicted and actual $\boldsymbol{\Lambda}$. (b) The DC block of NLINV-Net. It is implemented using the CG algorithm. (c) The proposed denoising block for subspace NLINV-Net. It is a ResNet made of 5 layers. Each layer consists of a convolutional layer (conv) with 32 filters followed ReLU acting independently on real and imaginary part. The last layer does not have the ReLU to avoid truncating the negative part of the learned noise patterns. The ResNet has one input/output channel per coefficient map. Following the residual learning strategy, the learned noise, i.e., the output of the last layer of the ResNet, is added to the input of the denoising block to get the reconstructed image as the output of the denoising block. The input of the denoising block is normalized to a maximum magnitude of one, and its output is rescaled to the original value. Figure adapted from [9].

`mobafit` command:

$$\hat{\boldsymbol{p}} = \begin{pmatrix} M_0^* \\ M_0 \\ R_1^* \end{pmatrix} = \arg\min_{\boldsymbol{p}} \left\| \boldsymbol{a} - \mathcal{B}^H \mathcal{M}(\boldsymbol{p}) \right\|_2^2 . \tag{3.9}$$

Here, $(\cdot)^H$ denotes the conjugate transpose. The coefficient maps were normalized so that the steady-state magnetization was scaled to one, allowing consistent initialization to the same value and improving the convergence of the parameter estimation algorithm.

After the estimation of the model parameters M_0^*, M_0, and R_1^*, T_1 was computed using the `bart looklocker` which apply the corrected Look-Locker equation:

$$T_1 = \frac{M_0}{M_0^* \cdot R_1^*} - 2 \cdot T_D \ . \tag{3.10}$$

Here, T_D is the delay between the inversion and the start of the acquisition.

3.6 Data Analysis

The quality of myocardial T_1 maps was evaluated on the ten test subjects using a combination of qualitative and quantitative evaluations, even though quantitative assessment has been challenging due to the absence of any ground truth reference. However, significant insights emerged from the computation of several quantitative metrics and the comparative analysis between the proposed neural network-based approach and the baseline methods. The quantitative metrics computed were the Coefficient of Variation (CV) within a specific Region Of Interest (ROI) to quantify the precision of T_1 estimation and the edge sharpness to quantify the preservation of fine details.

3.6.1 ROI analysis

For the assessment of the myocardial T_1 values, the interventricular septum was carefully selected to exclude the blood pool using Python. Excluding the blood pool from the evaluation process holds significant importance. This necessity arises from the dynamic nature of blood circulation during diastole within the ventricles, which results in image intensities that deviate from the expected signal model. This deviation ultimately hinders the reliable estimation of blood T_1 relaxation times. The interventricular septum ROIs of the ten test subject are shown in Figure 3.5. Similar to [35], the precision of the T_1 estimation was evaluated through the CV

withing the selected ROI:

$$\mathrm{CV}_{\mathrm{ROI}} = \frac{\mathrm{SD}_{\mathrm{ROI}}}{\mathrm{mean}_{\mathrm{ROI}}} \cdot 100\% \ . \tag{3.11}$$

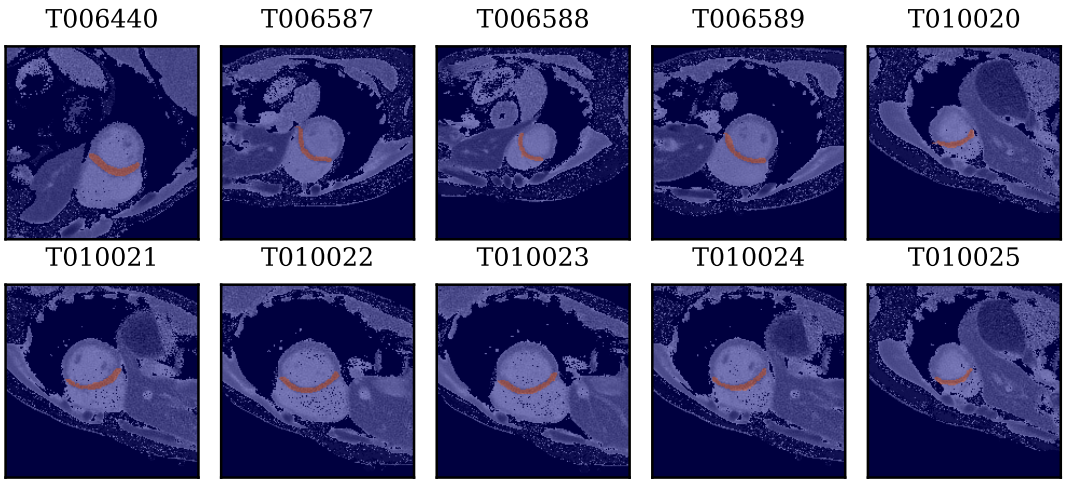

Figure 3.5: Interventricular septum ROIs. For the assessment of the myocardial T_1 values, the interventricular septum was carefully selected to exclude the blood pool using Python.

3.6.2 Edge Sharpness Analysis

T_1 edge sharpness was quantitatively measured as follows [36]: (i) a binary edge map was generated using the Canny Edge Detector (CED), and one edge-of-interest (EOI) was manually selected using a graphical user interface in Matlab; (ii) at each pixel of the EOI, an intensity profile was read along the horizontal direction; (iii) the intensity profiles corresponding to all EOI pixels were individually fitted with a sigmoid function characterized by four parameters, including one that represents the edge sharpness:

$$y(\rho, a_0, a_1, a_2, s) = \frac{a_1}{1 + 10^{s(a_0+\rho)}} + a_2 \ . \tag{3.12}$$

Here, ρ is the intensity profile, a_0 determines the central location, a_1 determines the vertical range, a_2 defines the vertical offset, and s quantifies the growth rate or sharpness of the sigmoid. The higher $|s|$, the sharper the edges. The above

nonlinear least square fitting was performed in MATLAB using the Levenberg-Marquardt algorithm. Since there are multiple intensity profiles, the sharpness of each image is represented by the distribution of s.

3.6.3 Statistical Analysis

Differences among the various methods for myocardial T_1 mapping were statistically tested using one-way ANOVA ($\alpha = 0.05$), which is a test that is used to find out whether there exists a statistically significant difference between the mean values of more than one group. In addition, t-tests were performed to test one-to-one differences. Multiple comparison correction [37] [38] was applied for controlling False Discovery Rate (FDR) and for estimating corrected significant p-values for each subject.

3.7 Implementation

All the image reconstruction-related computations were done in BART on a server system equipped with two AMD EPYC 7662 processors, 1008 GB of RAM and four NVIDIA A100-SXM-80GB GPUs from which one was used. Data analysis was performed using Python 3.9.2, except for the edge sharpness analysis, which was carried out in Matlab R2022a.

Chapter 4

Results

4.1 Comparison of Different Basis Dimensions

Figure 4.1 shows myocardial T_1 maps estimated using a variable number of principal components. A larger number of components results in noisier T_1 maps. Figure 4.1a shows that simple NLINV reconstruction fails for some datasets when 20 components are used, providing in those cases over-smoothed coefficient maps, as can be observed in Figure 4.2. Table 4.1 reports the execution time of NLINV and ℓ_1-Wavelet PICS reconstruction in seconds as a function of the number of components, showing a significant increase in computational effort when increasing the size of the basis. Considering both the precision of the T_1 maps and the computational effort, the superior result of using three, four, or five components compared to using ten or twenty components is evident. However, the results of three, four, or five components are difficult to distinguish. Figure 4.3 depicts the graph of the cumulative explained variance and indicates that most of the variance, i.e., the 98.56%, is already explained by the first four components. Therefore, based on this evidence and suggestions from the literature [7], the first four principal components were used as subspace basis.

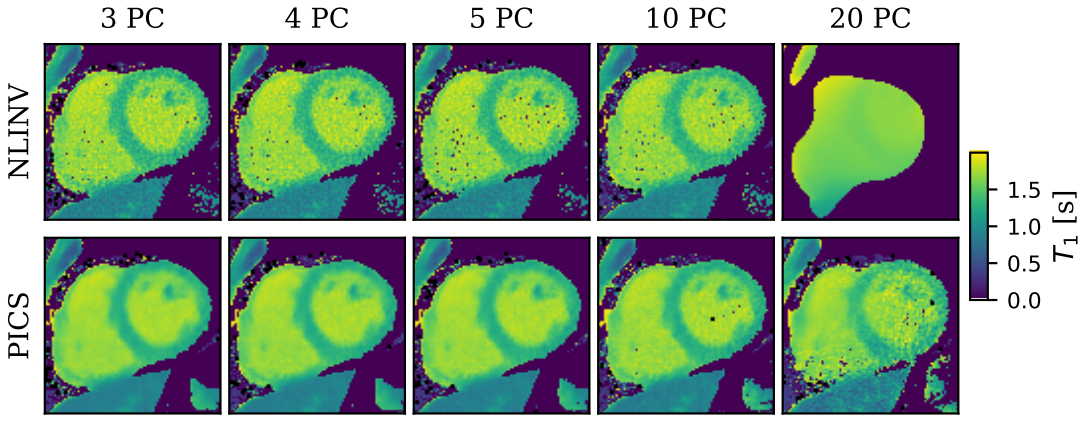

(a) First test dataset.

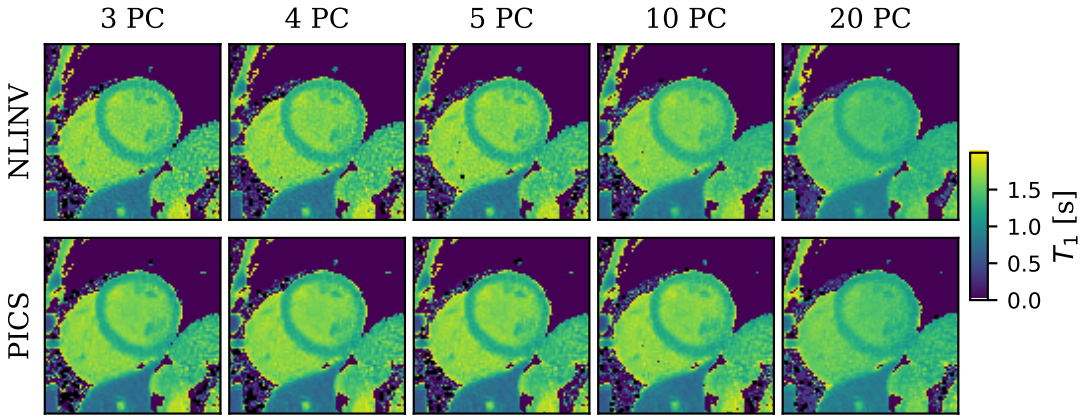

(b) Second test dataset.

Figure 4.1: NLINV and ℓ_1-Wavelet PICS myocardial T_1 maps estimated using a variable number of principal components. A higher number of components results in noisier T_1 maps. Simple NLINV reconstruction fails for the first test dataset when using twenty components, providing over-smoothed coefficient maps, as shown in Figure 4.2. Black pixels represent Not a Number (NaN) values.

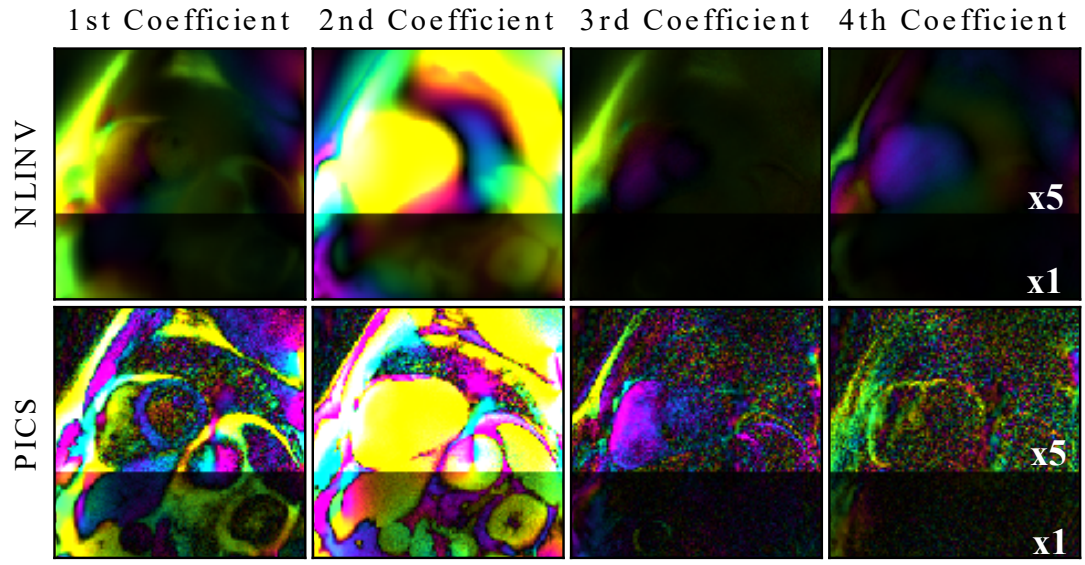

Figure 4.2: NLINV and ℓ_1-Wavelet PICS coefficient maps reconstructed using twenty components. Simple NLINV reconstruction fails for the first test dataset when using twenty components, providing over-smoothed coefficient maps. The color encodes the phase.

Number of Components	**Execution Time** [s]	
	NLINV	**ℓ_1-Wavelet PICS**
3	4	17
4	5	19
5	5	19
10	6	25
20	10	57

Table 4.1: NLINV and ℓ_1-Wavelet PICS execution time in seconds as a function of the number of components. There is a significant increase in computational effort when increasing the basis dimension.

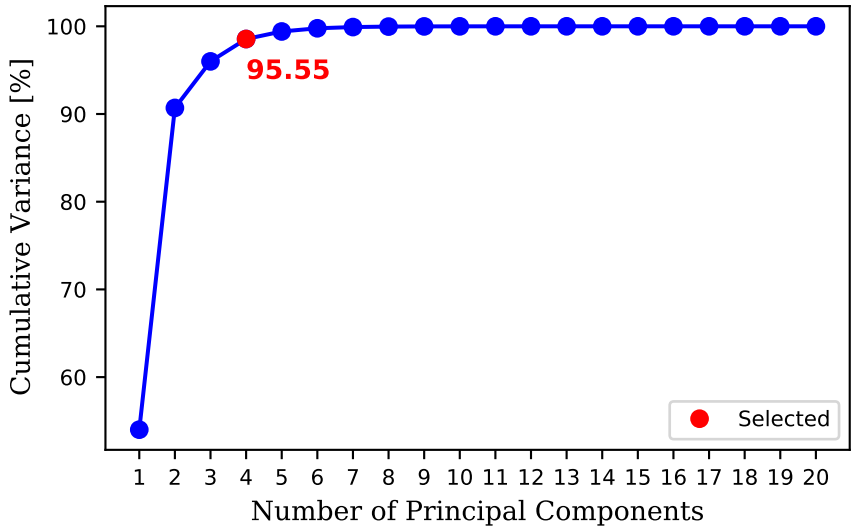

Figure 4.3: Cumulative Explained Variance. Most of the variance, i.e., the 98.56%, is already explained by the first four components. The red dot marks the selected basis dimension. Connecting lines serve to guide the eye.

4.2 Comparison of Different Regularization Parameters

Figure 4.4 and Figure 4.5 demonstrate the influence of the regularization parameter λ used in the ℓ_1-Wavelet PICS reconstruction on CV and edge sharpness. Low values of λ increase noise in the myocardial T_1 maps, i.e., higher CV, while high values of λ lead to blurring, i.e., lower edge sharpness. The CV was evaluated in the interventricular septum. The edge sharpness of each image is represented by the distribution of the edge sharpness of the multiple sigmoid functions fitted to the selected T_1 intensity profiles. Figure 4.6 reports the fitting results of a representative T_1 intensity profile for the first test dataset and shows a good agreement, i.e., low RMSE, between the profiles and the fitted sigmoid curves. This evidence holds to all the datasets and all the selected profiles. Table 4.2 reports the mean $\pm$ SD of the CV computed in the interventricular septum across test subjects. The precision of the T_1 values increases for higher λ values. Once the CV values were identified as normally distributed by the Lilliefors test, t-tests revealed that (i) the CV values for $\lambda = 1.5 \cdot 10^{-4}$ are significantly higher than the CV values for all the other λ values ($p < 0.005$), (ii) the CV values for $\lambda = 6.0 \cdot 10^{-4}$ are significantly higher than the CV values for the highest λ value, i.e., $\lambda = 1.5 \cdot 10^{-3}$ ($p < 0.015$), (iii) no significant difference between the CV values for two highest λ values, i.e., $\lambda = 1.0 \cdot 10^{-3}$ and $\lambda = 1.5 \cdot 10^{-3}$ ($p < 0.395$). Since there is no denoising improvement from $\lambda = 1.0 \cdot 10^{-3}$ to $\lambda = 1.5 \cdot 10^{-3}$ but, as can be observed in the last column of Figure 4.4 and Figure 4.5, the edge sharpness keeps decreasing, $\lambda = 1.0 \cdot 10^{-3}$ was chosen to balance between noise reduction and preservation of image details.

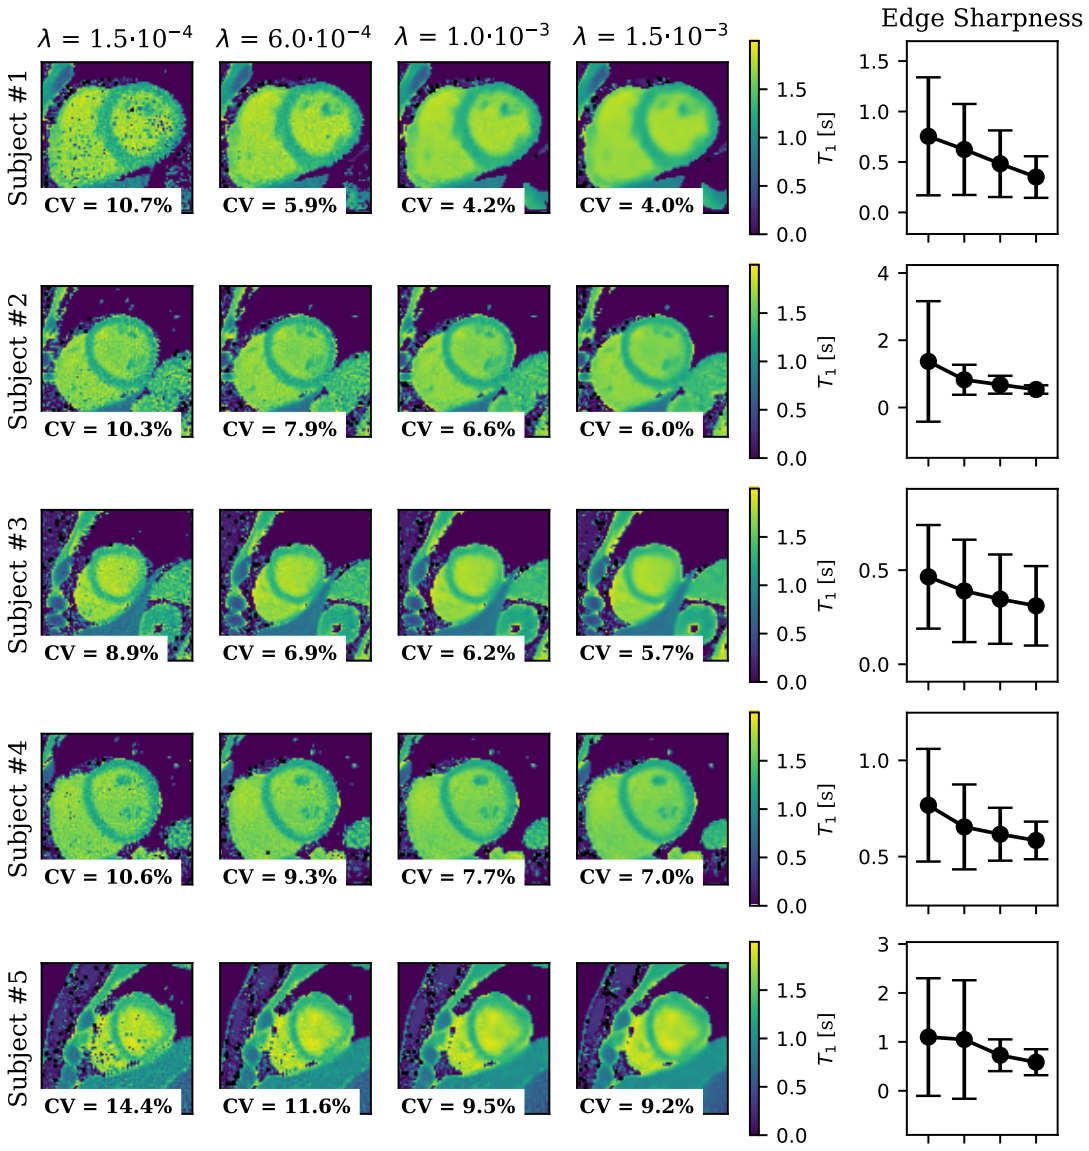

Figure 4.4: ℓ_1-Wavelet PICS regularization parameter λ tuning for test subjects from 1 to 5. Low values of λ increase noise in the myocardial T_1 maps, i.e., higher CV, while high values lead to blurring, i.e., lower edge sharpness. The CV was evaluated in the interventricular septum. Connecting lines in the edge sharpness graphs serve to guide the eye. Black pixels represent NaN values.

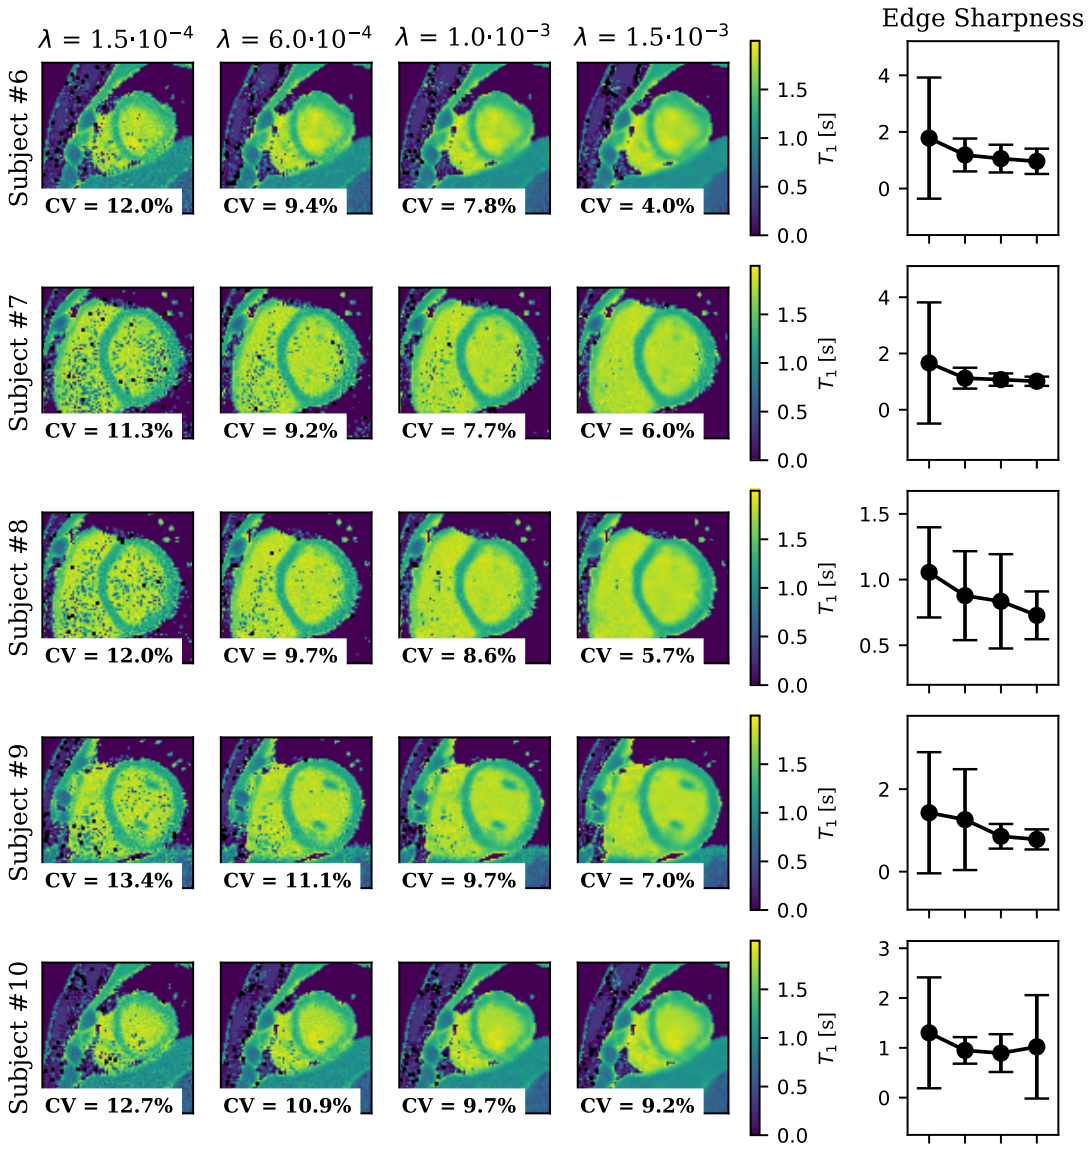

Figure 4.5: ℓ_1-Wavelet PICS regularization parameter λ tuning for test subjects from 6 to 10. Low values of λ increase noise in the myocardial T_1 maps, i.e., higher CV, while high values lead to blurring, i.e., lower edge sharpness. The CV was evaluated in the interventricular septum. Connecting lines in the edge sharpness graphs serve to guide the eye. Black pixels represent NaN values.

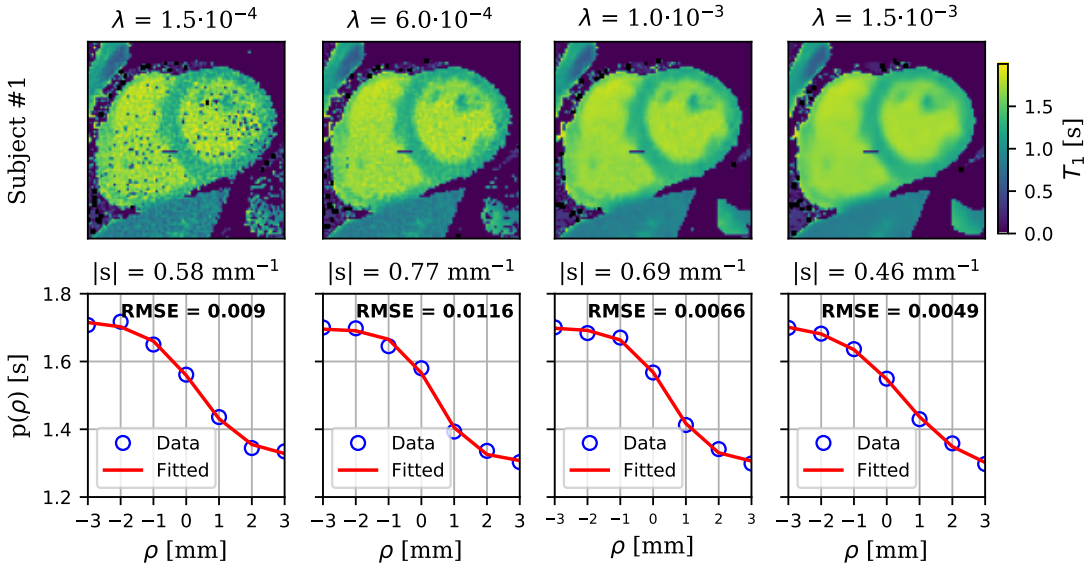

Figure 4.6: Edge sharpness as a function of the ℓ_1-Wavelet PICS regularization parameter λ for the first dataset. On the top, T_1 maps highlighting the selected T_1 intensity profile in dark blue. In the middle, the quantitative edge sharpness value |s|. On the bottom, the selected T_1 intensity profile and the fitted sigmoid function, with the goodness of fit quantified as RMSE. Higher values of λ lead to blurring, i.e., lower edge sharpness.

λ	**CV** [%]
$1.5 \cdot 10^{-4}$	11.6 ± 1.6
$6.0 \cdot 10^{-4}$	9.2 ± 1.8
$1.0 \cdot 10^{-3}$	7.8 ± 1.8
$1.5 \cdot 10^{-3}$	$\mathbf{7.1 \pm 1.6}$

Table 4.2: CV as a function of the ℓ_1-Wavelet PICS regularization parameter λ. Low values of λ increase noise in the myocardial T_1 maps, i.e., higher CV. The reported CV values are the mean $\pm$ SD of the CV values evaluated in the interventricular septum across test subjects. The best result is highlighted in bold.

4.3 Self-Supervised Subspace Learning for Myocardial T_1 Mapping

4.3.1 CV and Edge Sharpness

Figure 4.7 and Figure 4.8 illustrate the impact of extending the training duration for NLINV-Net and reports the CV values computed in the interventricular septum. A noticeable disparity is evident when comparing the T_1 maps generated after 100 epochs of training against those after 500, 1000, and 2000 epochs. However, no significant visual difference is observed among the T_1 maps produced for 500, 1000, and 2000 epochs. Table 4.3 reports the mean ± SD of the CV values computed in the interventricular septum across test subjects. The precision of the T_1 values increases for a higher number of epochs, with comparable results for 1000 and 2000 epochs. Since training the network longer increases the computational effort significantly, it is worth it only if enhanced precision in the results is achieved. Thus, NLINV-Net trained for 1000 epochs was selected as the optimal architecture. The training lasted almost 30 seconds/epoch.

Figure 4.9 and Figure 4.10 shows a comparison among NLINV-Net T_1 maps and the T_1 maps computed using the baseline methods, i.e., NLINV and ℓ_1-Wavelet PICS, and reports the CV values computed in the interventricular septum as well as the edge sharpness associated to the T_1 maps. The edge sharpness of each image is represented by the distribution of the edge sharpness of the multiple sigmoid functions fitted to the selected T_1 intensity profiles. Figure 4.11 reports the fitting results of a representative T_1 intensity profile for the first test dataset and shows a good agreement, i.e., low RMSE, between the profiles and the fitted sigmoid curves. This evidence holds to all the datasets and all the selected profiles. Statistical tests (t-tests when the edge sharpness values were normally distributed based on the Lilliefors test; Wilcoxon rank-sum tests when they were not) revealed no significant differences among NLINV-Net and PICS edge sharpness values for all the test subjects, except for the sixth. Thus, for most of the datasets NLINV-Net is not over-smoothing the images compared to PICS. NLINV-Net T_1 maps exhibit reduced level of noise compared to the T_1 maps produced by plain NLINV. Conversely, no significant visual difference in noise level is observed between NLINV-Net and PICS T_1 maps. Table 4.4 reports the

mean ± SD of the CV values computed in the interventricular septum across subjects. NLINV-Net is the method leading to the best T_1 precision, i.e., lowest CV. Once the CV values were identified as normally distributed by the Lilliefors test, t-tests revealed that (i) the NLINV-Net CV values are significantly lower than the NLINV CV values (p<0.02), (ii) no significant difference between NLINV and PICS CV values (p>0.3). Thus, the precision of NLINV-Net and PICS T_1 maps as well as the precision of NLINV and PICS T_1 maps is comparable, while the precision of NLINV-Net T_1 is higher than the precision of NLINV T_1. Figure 4.12 reports the coefficient maps reconstructed by NLINV-Net, ℓ_1-Wavelet PICS and NLINV, and shows that the NLINV-Net coefficient maps are the less noisy.

#Epochs	**CV** [%]
100	8.9 ± 1.2
500	7.5 ± 1.5
1000	**7.2 ± 1.4**
2000	**7.2 ± 1.4**

Table 4.3: CV as a function of the number of epochs. The precision of the T_1 values increases for a higher number of epochs, with comparable results for 1000 and 2000 epochs. The reported CV values are the mean ± SD of the CV values evaluated in the interventricular septum across test subjects. The best result is highlighted in bold.

Method	**CV** [%]
NLINV-Net	**7.2 ± 1.4**
PICS	7.8 ± 1.8
NLINV	9.1 ± 1.5

Table 4.4: NLINV-Net, NLINV, and ℓ_1-Wavelet PICS CV values. The reported CV values are the mean ± SD of the CV values evaluated in the interventricular septum across test subjects. The best result is highlighted in bold.

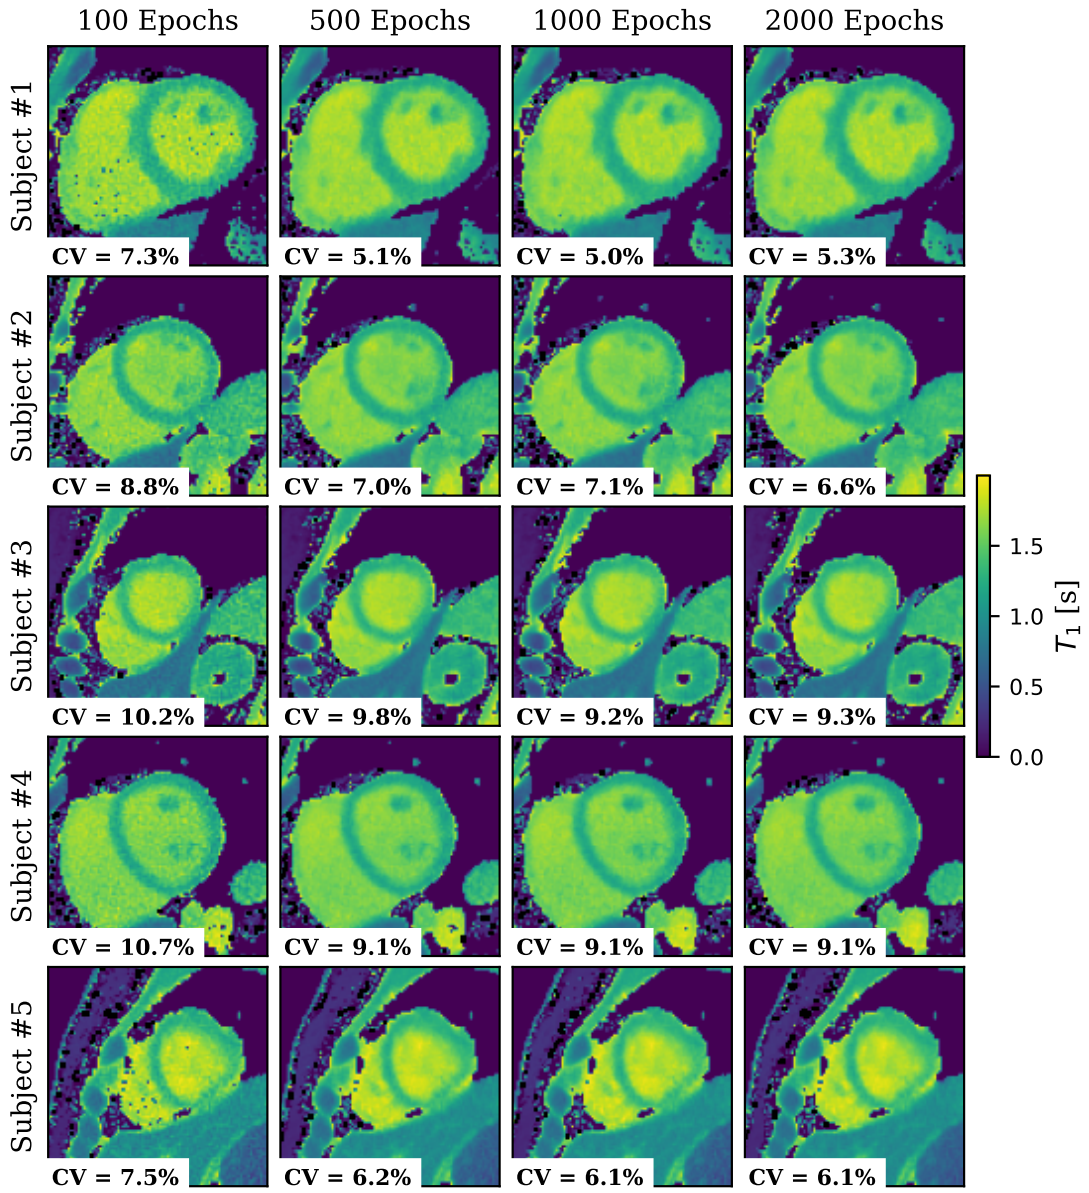

Figure 4.7: NLINV-Net T_1 maps as a function of the number of epochs for test subjects from 1 to 5. A noticeable disparity is evident when comparing the T_1 maps generated after 100 epochs of training against those after 500, 1000, and 2000 epochs. However, no significant visual difference is observed among the T_1 maps produced for 500, 1000, and 2000 epochs.

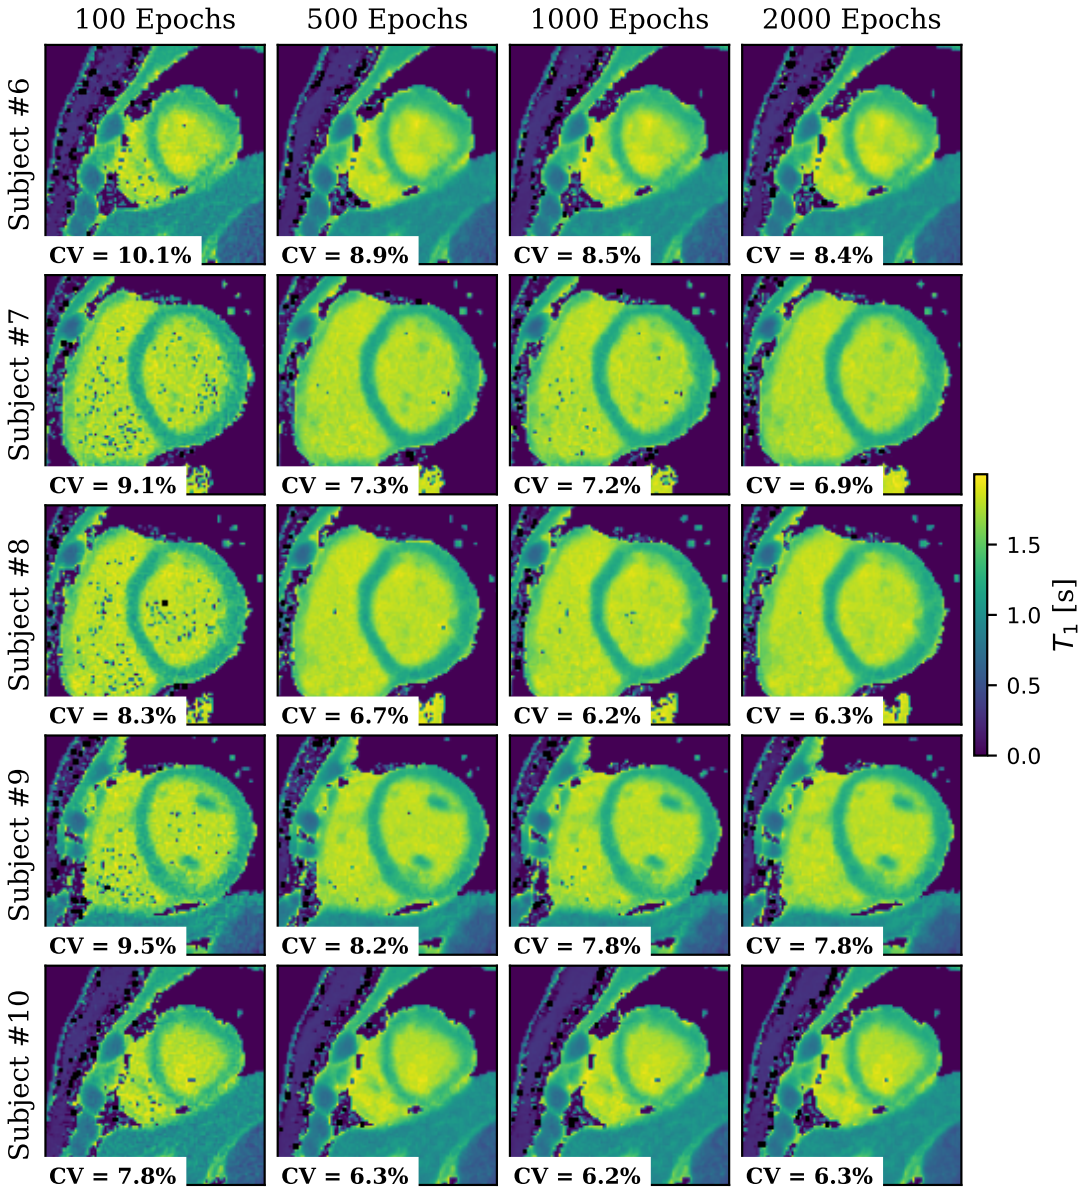

Figure 4.8: NLINV-Net T_1 maps as a function of the number of epochs for test subjects from 6 to 10. A noticeable disparity is evident when comparing the T_1 maps generated after 100 epochs of training against those after 500, 1000, and 2000 epochs. However, no significant visual difference is observed among the T_1 maps produced for 500, 1000, and 2000 epochs.

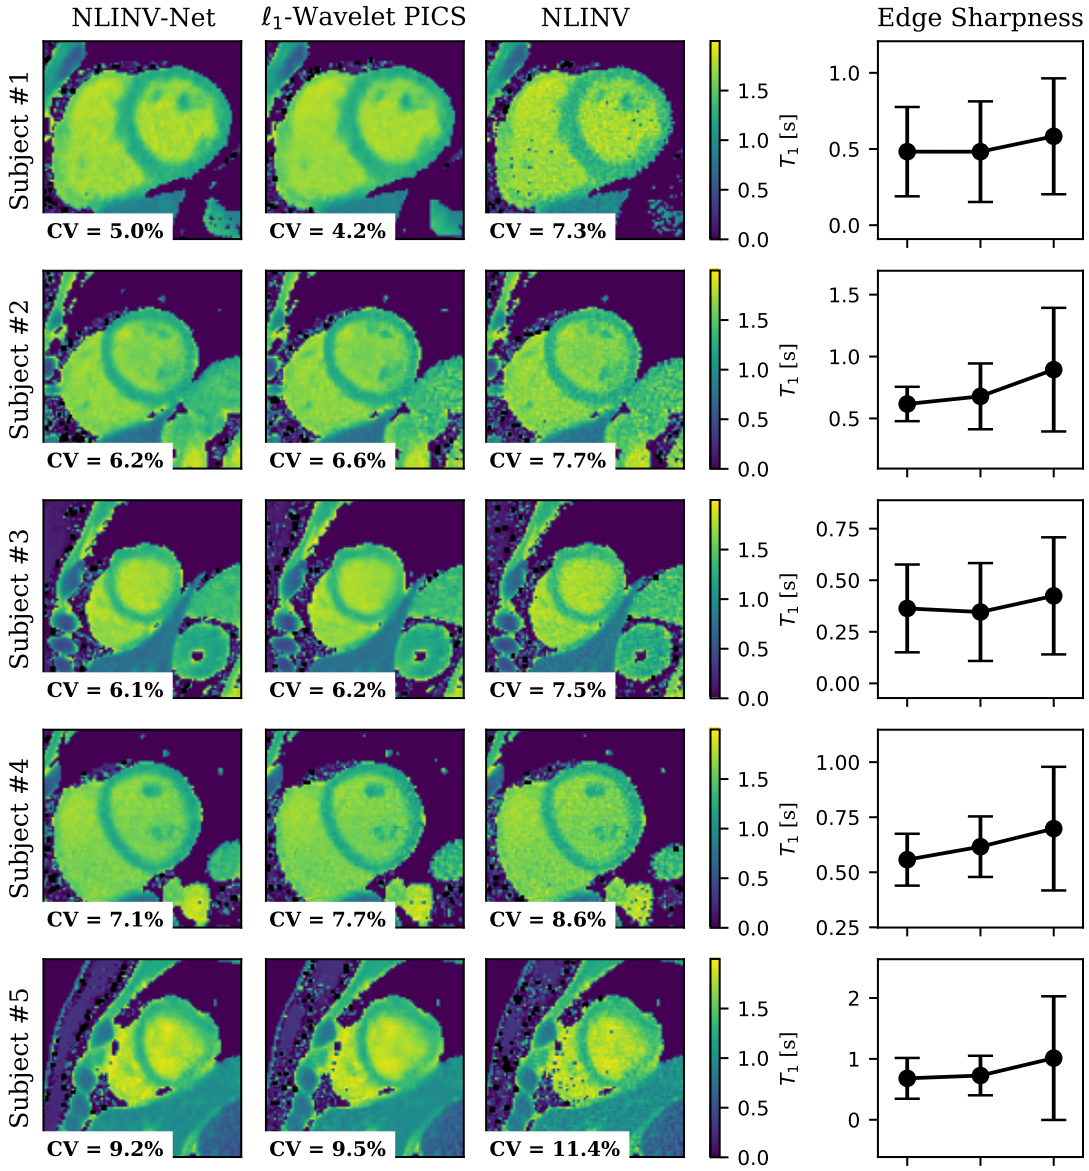

Figure 4.9: NLINV-Net, NLINV, and ℓ_1-Wavelet PICS T_1 maps for test subjects from 1 to 5. NLINV-Net T_1 maps exhibit reduced level of noise compared to the T_1 maps produced by plain NLINV. Conversely, no significant visual difference in noise level is observed between NLINV-Net and ℓ_1-Wavelet PICS T_1 maps. The T_1 maps were computed using four basis functions, and as for PICS, the regularization parameter was $\lambda = 1.0 \cdot 10^{-3}$. The learning rate was r=0.001. Black pixels represent NaN values. The images are cropped for visualization purpose.

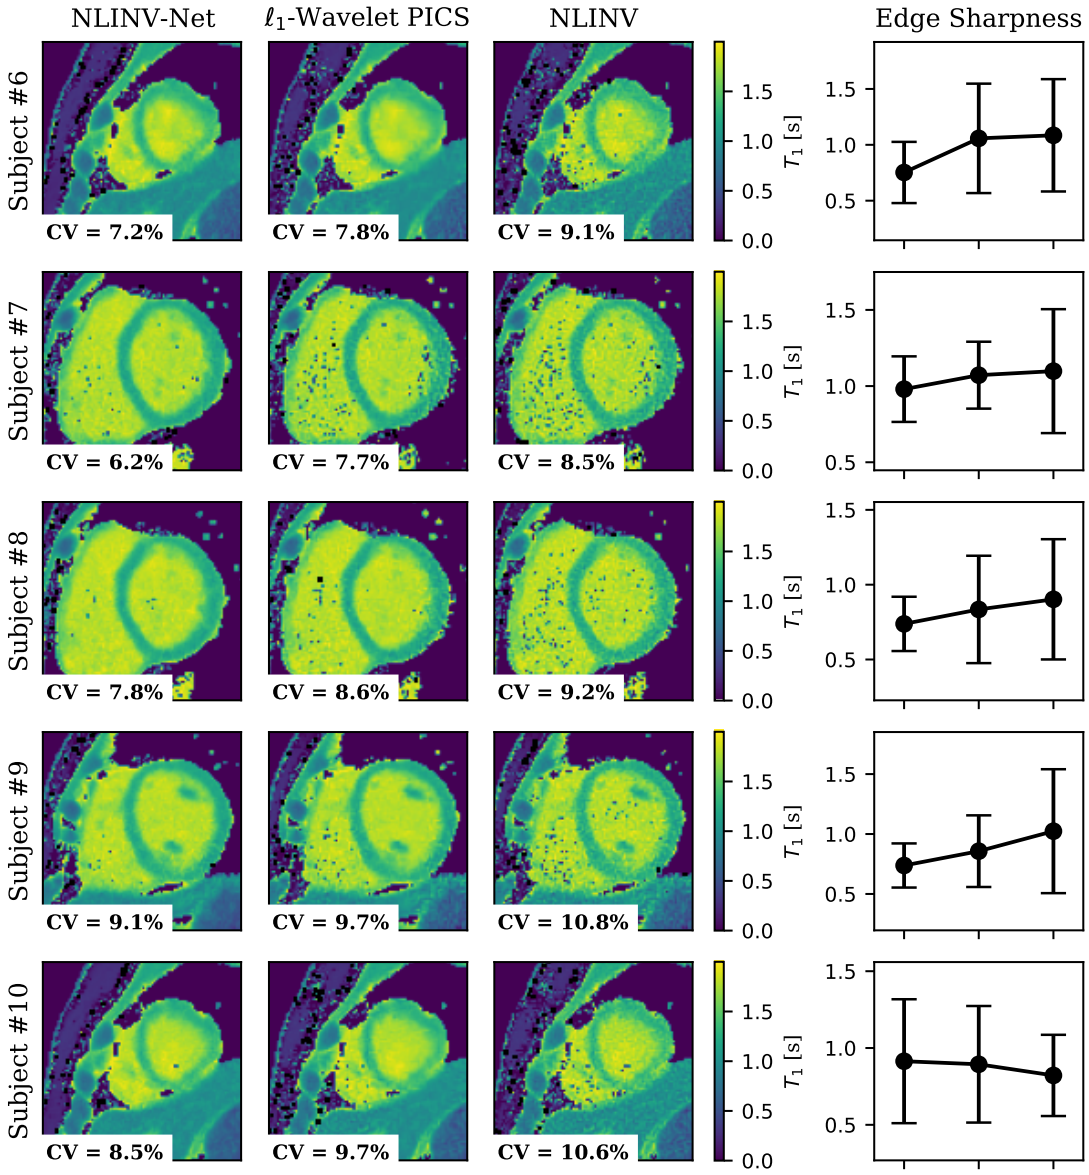

Figure 4.10: NLINV-Net, NLINV, and ℓ_1-Wavelet PICS T_1 maps for test subjects from 6 to 10. NLINV-Net T_1 maps exhibit reduced level of noise compared to the T_1 maps produced by plain NLINV. Conversely, no significant visual difference in noise level is observed between NLINV-Net and ℓ_1-Wavelet PICS T_1 maps. The T_1 maps were computed using four basis functions, and as for PICS, the regularization parameter was λ=0.01. The learning rate was r=0.001. Black pixels represent NaN values. The images are cropped for visualization purpose.

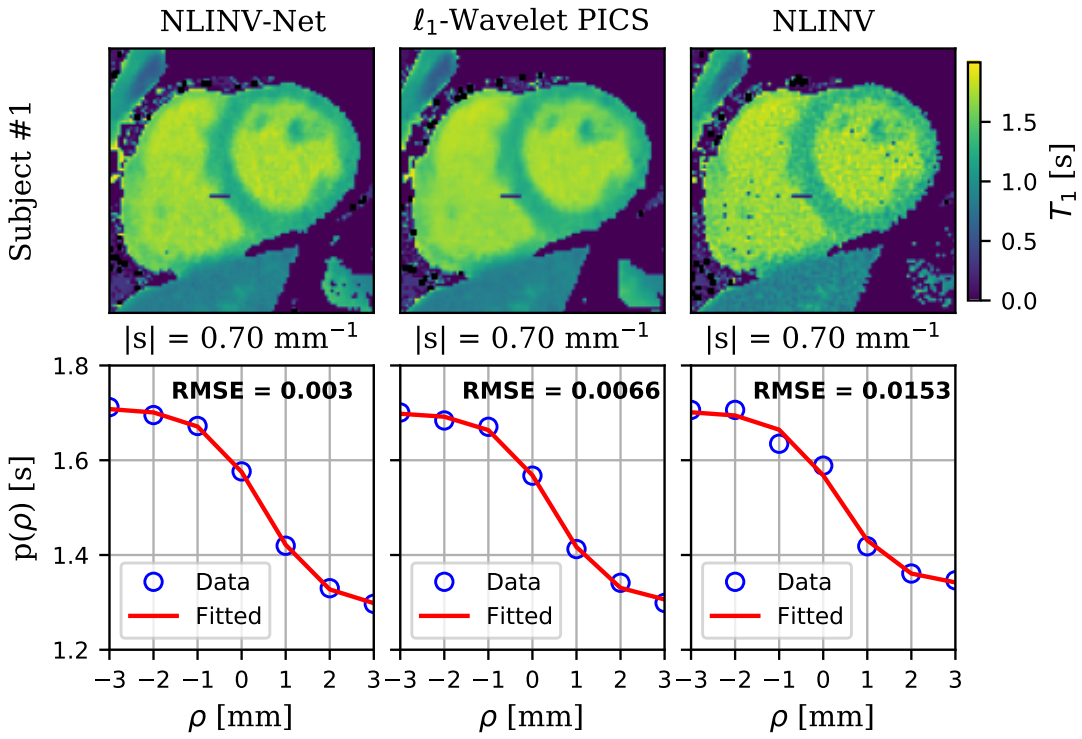

Figure 4.11: NLINV-Net, NLINV, and ℓ_1-Wavelet PICS edge sharpness for the first dataset. On the top, T_1 maps highlighting the selected T_1 intensity profile in dark blue. In the middle, the quantitative edge sharpness value |s|. On the bottom, the selected T_1 intensity profile and the fitted sigmoid function, with the goodness of fit quantified as RMSE.

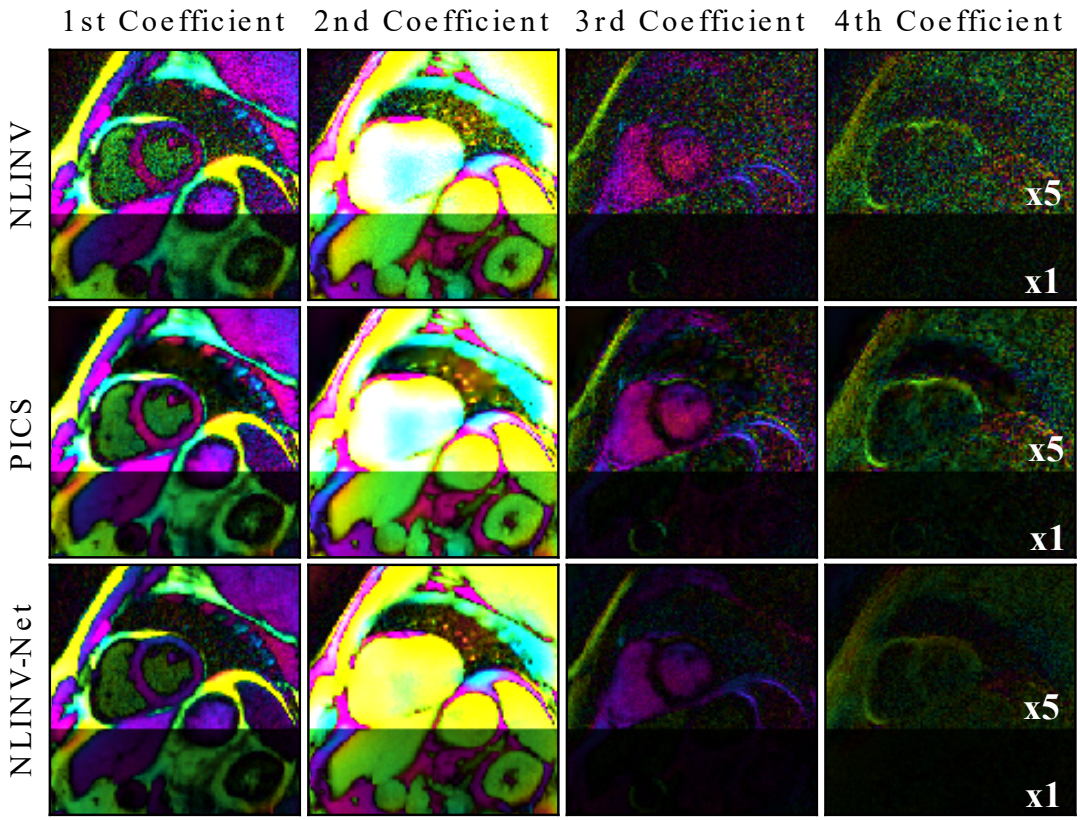

Figure 4.12: NLINV-Net, NLINV, and ℓ_1-Wavelet PICS coefficient maps. The NLINV-Net coefficient maps are the less noisy. The color encodes the phase.

4.3.2 Loss Analysis

Figure 4.13 shows the training and test losses versus number of epochs for the optimized architecture, i.e., 1000 epochs and learning rate $r = 0.001$. The loss decreases as the number of epochs increases, meaning the network is learning to reconstruct the coefficient maps even better. The randomness in the training loss curve is caused by the epoch-wise random generation of the masks selecting the input spokes for NLINV-Net.

Figure 4.14a highlights the presence of a consistent variability in the loss values across different test subjects. Upon observing Figure 4.14b, which presents the first coefficient maps of the test subjects, it becomes evident that this variability might stem from the manner in which data is collected. Notably, subjects with a horizontal display of the thorax exhibit lower loss values compared to subjects with an oblique acquisition. Figure 4.15 shows a similar pattern in the training set.

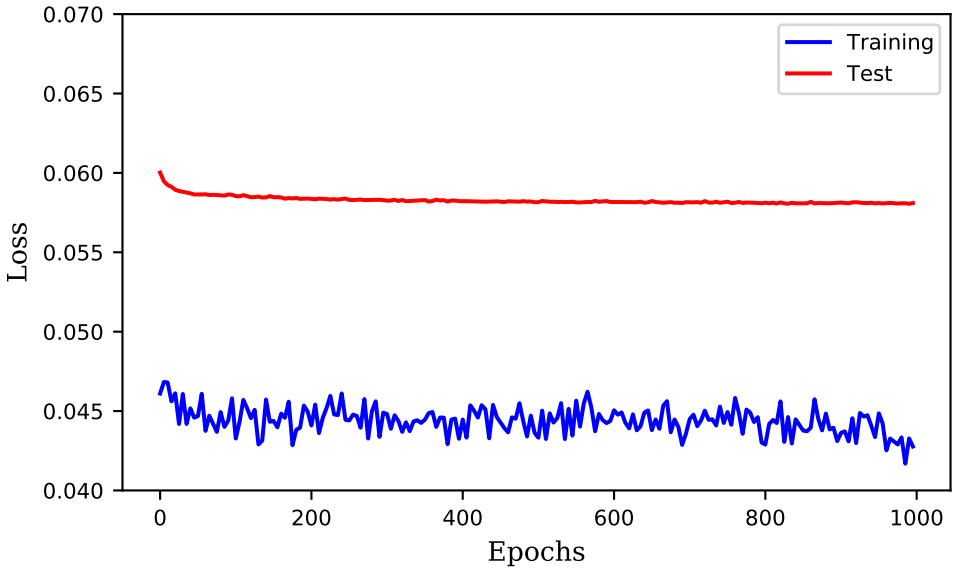

Figure 4.13: Training and test loss versus number of epochs. The loss decreases as the number of epochs increases, meaning the network is learning to reconstruct the coefficient maps even better. The randomness in the training loss curve is caused by the epoch-wise random generation of the masks selecting the input spokes for NLINV-Net.

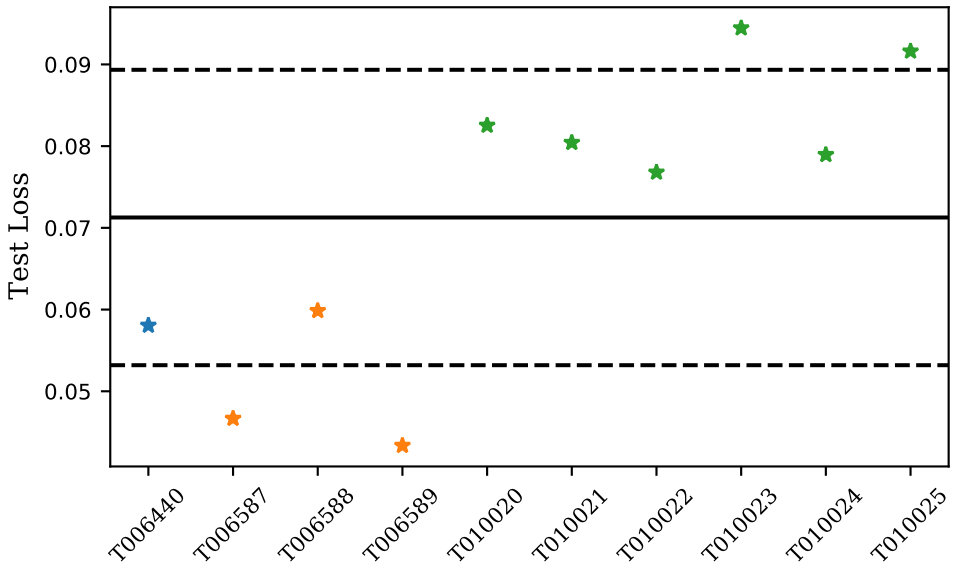

(a) Loss values at the last epoch for the test subjects.

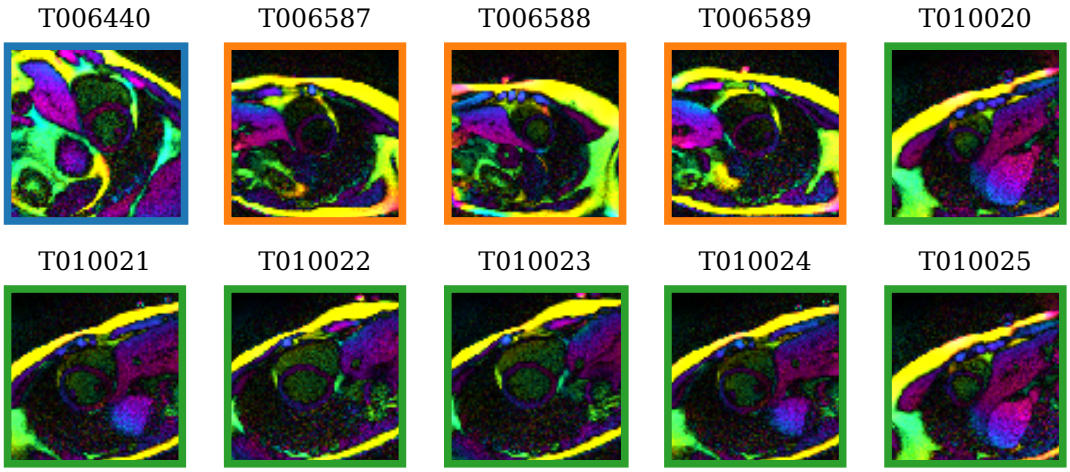

(b) First coefficient maps of the test subjects.

Figure 4.14: Variability in the loss values among test subjects.

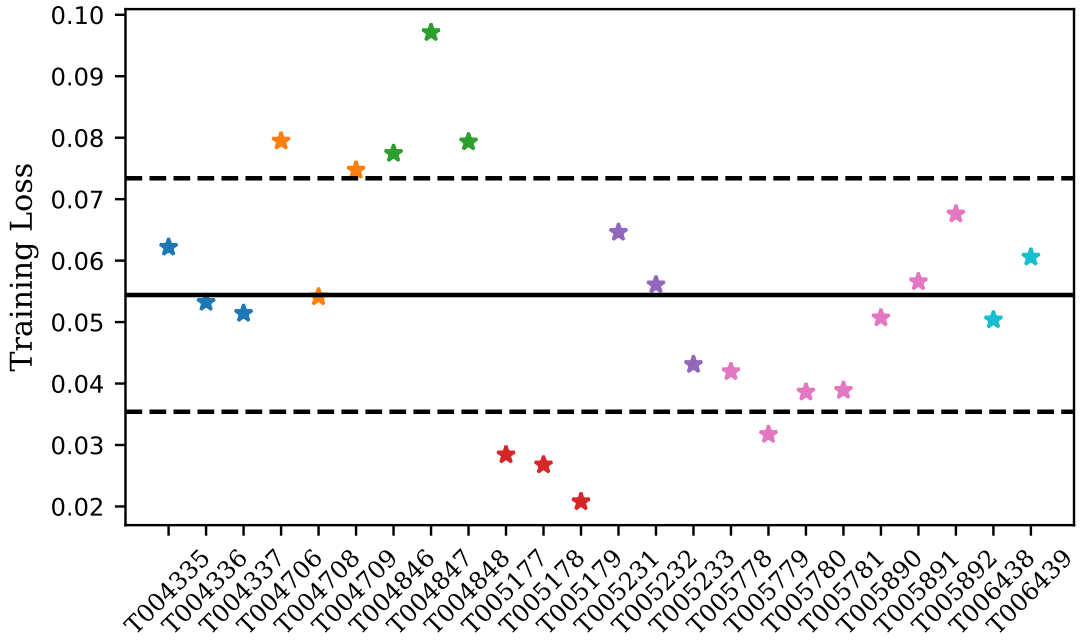

(a) Loss values at the last epoch for the training subjects.

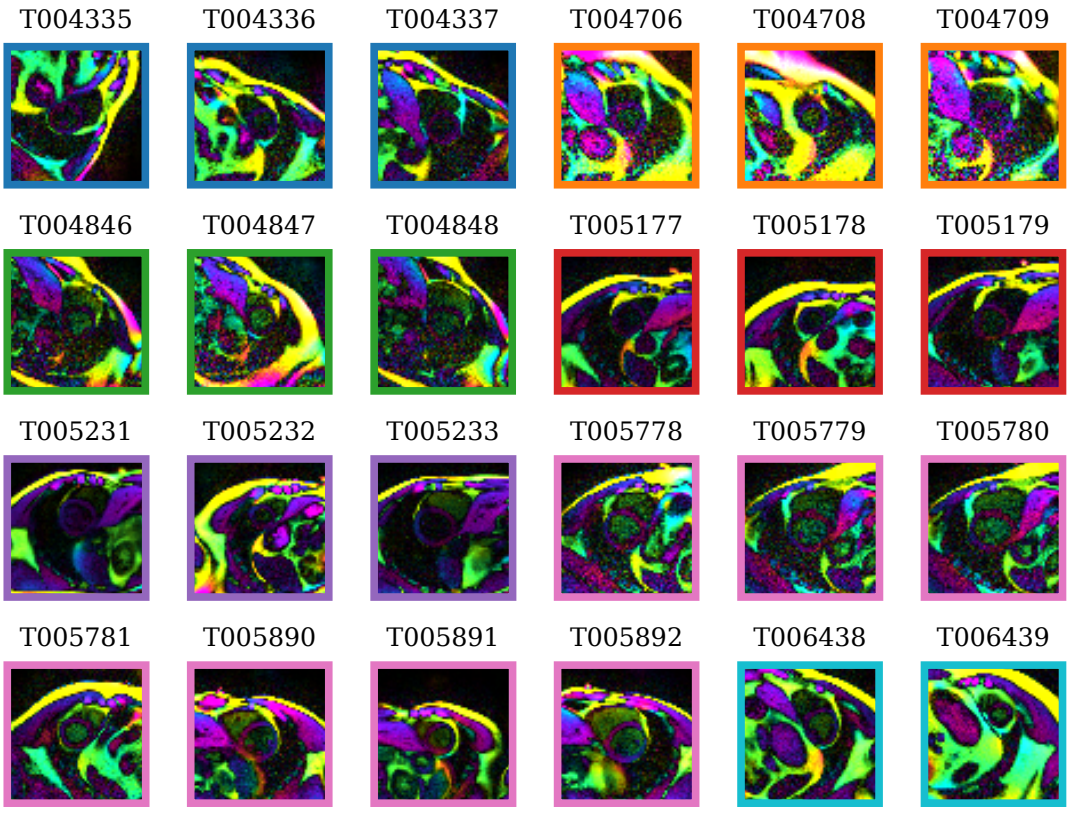

(b) First coefficient maps of the training subjects.

Figure 4.15: Variability in the loss values among training subjects.

Chapter 5

Discussion

This thesis investigates the performance of self-supervised subspace learning for myocardial T_1 mapping using single-shot inversion-recovery radial FLASH. The employed neural network is NLINV-Net [9], which is trained in a self-supervised fashion to reconstruct the coefficient maps jointly with the coil sensitivity maps from undersampled radial cardiac data. This learning strategy is crucial for cardiac data, which lack any ground truth reference. T_1 maps are estimated using pixel-wise fitting to the signal model [25]. Reconstruction and parameter mapping were performed in the subspace. The reconstruction quality of the coefficient and T_1 maps computed using NLINV-Net was evaluated considering NLINV [22] and ℓ_1-Wavelet PICS as baseline methods, where ℓ_1-Wavelet PICS combines parallel imaging and compressed sensing by using ℓ_1-Wavelet sparsity regularization when solving the inverse problem for parallel imaging.

5.1 Comparison of Different Basis Dimensions

The problem of finding the optimal basis dimension for subspace-constrained reconstruction is challenging, but several factors can help define its solution. According to [7], a lower number of principal components causes bias for myocardial T_1 mapping, while a higher number of components results in noisier T_1 maps. Thus, the chosen number of components needs to compromise between quantitative accuracy and precision. Since one of the purposes of subspace-constrained reconstruction is to reduce computational effort, another meaningful variable to keep in mind when choosing the dimension of the basis is the execution time of the reconstruction algorithm. Finally, the cumulative explained variance graph is a useful tool

for assessing how many components explain most of the variance. This metric identifies the optimal basis size where the plateau starts. Based on the suggestion found in [7] and according to which four components define an optimal trade-off between quantitative accuracy and precision, and combining the aforementioned parameters, the first four components were taken as the subspace basis for the self-supervised subspace learning performed in this thesis. Noteworthy, NLINV reconstruction fails for some datasets when performed in a twenty-dimensional subspace, providing over-smoothed coefficient maps. The coefficient maps computed by the ℓ_1-Wavelet PICS method are not as smooth as the NLINV coefficient maps and lead to at least meaningful T_1 maps, which is interesting because the coil sensitivity maps provided to ℓ_1-Wavelet PICS to compute these coefficient maps are those produced by the NLINV algorithm, which denotes that most of the effort of NLINV in the case of many components lies in the computation of the coefficient maps more than in the coil sensitivity maps. The last point about the selection of the basis dimension refers to the consideration of healthy volunteers only, as this study does. One could explore what happens when diseased subjects are involved and observe how the method should translate to produce high-quality T_1 maps. Likely, it would require selecting a higher number of basis functions to live up to the variability increase.

5.2 Comparison of Different Regularization Parameters

The inverse problem for MRI reconstruction is ill-posed, meaning that tiny errors in the acquired data can result in significant errors in the estimated image. The noise amplification decreases by adding a regularization term into the inversion. However, the regularized inverse problem requires careful tuning of the regularization parameter to balance between noise reduction and retaining fine details. In this study, to choose the optimal regularization parameter λ for the ℓ_1-Wavelet PICS method, the noise was measured as CV, while the detail level as edge sharpness. Considering these metrics, $\lambda = 1.0 \cdot 10^{-3}$ was identified as an optimal value. It is worth noting that the process of tuning the regularization parameter is influenced by user preferences, making this method partially subjective. It is one of the main driving forces of introducing deep learning into the image reconstruction process to learn automatically a regularization term. Notably, including the

ℓ_1-Wavelet regularization term for compressed sensing into the NLINV functional is a possibility. However, up to date, this adaptation has not been integrated into BART.

5.3 Self-Supervised Subspace Learning for Myocardial T_1 Mapping

The lack of any ground truth reference compounds the major challenge of assessing the performance of self-supervised subspace learning for myocardial T_1 mapping. However, the precision of NLINV-Net T_1 maps was demonstrated to be high by small CV values, ranging from 5.1 to 9.6%, without losing in image details, as proved by the edge sharpness computation. The precision of NLINV-Net T_1 maps is improved compared to the precision obtained by the baseline methods, i.e., NLINV [22] and ℓ_1-Wavelet PICS, but the difference between the network and PICS is not statistically significant. Thus, NLINV-Net learns in a self-supervised fashion, i.e., without a prior or reference, correlations between the parameters encoded with the FLASH sequence and consequently a well-tuned regularization and produces high-quality T_1 maps as granted by PICS but getting rid of the subjective regularization parameter tuning that comes with it. The first reason for the inability of NLINV-Net to significantly outperform PICS could be the small size of the sample, i.e., 24 subjects. Training the network on more data could help improve its image reconstruction performance. However, this requirement would be a possible bottleneck when projecting the application of this method in clinical practice: acquiring more training data means increasing the installation time of the technique. Thus, a fair balance between the two requirements needs to be found. The second reason could be the inclination with which data are collected. It came up that depending on the rotation of the thorax, different amount of energy in the image lies outside the reconstructed field of view, which gives a constant offset for the loss. Thus, one could improve the FOV selection for the reconstruction.

Chapter 6

Conclusion and Outlook

This thesis investigates a novel neural network-based subspace MRI reconstruction method for myocardial T_1 mapping utilizing single-shot inversion-recovery radial FLASH. NLINV-Net was able to learn in a self-supervised fashion, i.e., without a prior or reference, correlations between the parameters encoded with the FLASH sequence and consequently a well-tuned regularization, it outperformed NLINV in terms of T_1 precision and produced high-quality T_1 maps as granted by ℓ_1-Wavelet PICS but getting rid of the subjective regularization parameter tuning that comes with it, providing an excellent basis for myocardial T_1 mapping using single-shot inversion-recovery radial FLASH sequence.

The core of the proposed approach, the subspace NLINV-Net, holds significant promise beyond its application in the specific three-parameter MRI approach utilizing the FLASH sequence. For instance, it could be particularly valuable in addressing more extensive quantitative MRI challenges, such as the quantitative magnetization transfer mapping without constraints on model parameters [39] [40] using Hybrid State Free Precession (HSFP) sequences [41]. The network's capacity to learn correlations among individual parameters, now encompassing a substantial number (e.g., an eight-parameter model), could significantly enhance the reconstruction process.

References

[1] A. Haase et al. “FLASH imaging: Rapid NMR imaging using low flip-angle pulses”. In: *Journal of Magnetic Resonance* 67.2 (1986), pp. 258–266. DOI: https://doi.org/10.1016/j.jmr.2011.09.021.

[2] Daniel K. Sodickson and Warren J. Manning. “Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays”. In: *Magnetic Resonance in Medicine* 38.4 (1997), pp. 591–603. DOI: https://doi.org/10.1002/mrm.1910380414.

[3] Michael Lustig, David Donoho, and John M Pauly. “Sparse MRI: The application of compressed sensing for rapid MR imaging”. In: *Magnetic resonance in medicine* 58.6 (2007), pp. 1182–1195. DOI: https://doi.org/10.1002/mrm.21391.

[4] Block KT, Uecker M, and Frahm J. “Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint”. In: *Magnetic resonance in medicine* 57.6 (2007), pp. 1086–1098. DOI: https://doi.org/10.1002/mrm.21236.

[5] Radenkovic D et al. “T1 mapping in cardiac MRI”. In: *Heart Fail Rev.* 22 (4 2017), pp. 415–430. DOI: 10.1007/s10741-017-9627-2.

[6] Xiaoqing Wang et al. “High-resolution myocardial T1 mapping using single-shot inversion recovery fast low-angle shot MRI with radial undersampling and iterative reconstruction”. In: *The British Journal of Radiology* 89.1068 (2016), p. 20160255. DOI: 10.1259/bjr.20160255.

[7] Xiaoqing Wang et al. “Physics-based reconstruction methods for magnetic resonance imaging”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2200 (2021), p. 20200196. DOI: 10.1098/rsta.2020.0196.

[8] Blumenthal M et al. “Deep, Deep Learning with BART”. In: *Magnetic resonance in medicine* 89.2 (2023), pp. 678–693. DOI: 10.1002/mrm.29485.

[9] Moritz Blumenthal et al. “NLINV-Net: Self-Supervised End-2-End Learning for Reconstructing Undersampled Radial Cardiac Real-Time Data”. In: *Proceeding of the 31st Annual Meeting ISMRM, London, England, UK* (2022), p. 0499.

[10] Robert W. Brown et al. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. John Wiley & Sons, 2014. DOI: 10.1002/9781118633953.

[11] I. I. Rabi. “Space Quantization in a Gyrating Magnetic Field”. In: *Phys. Rev.* 51 (8 1937), pp. 652–654. DOI: 10.1103/PhysRev.51.652.

[12] F. Bloch. “Nuclear Induction”. In: *Phys. Rev.* 70 (7-8 1946), pp. 460–474. DOI: 10.1103/PhysRev.70.460.

[13] E. M. Purcell, H. C. Torrey, and R. V. Pound. “Resonance Absorption by Nuclear Magnetic Moments in a Solid”. In: *Phys. Rev.* 69 (1-2 1946), pp. 37–38. DOI: 10.1103/PhysRev.69.37.

[14] Stefanie Winkelmann et al. “An Optimal Radial Profile Order Based on the Golden Ratio for Time-Resolved MRI”. In: *IEEE Transactions on Medical Imaging* 26.1 (2007), pp. 68–76. DOI: 10.1109/TMI.2006.885337.

[15] Stefan Wundrak et al. “A Small Surrogate for the Golden Angle in Time-Resolved Radial MRI Based on Generalized Fibonacci Sequences”. In: *IEEE Transactions on Medical Imaging* 34.6 (2015), pp. 1262–1269. DOI: 10.1109/TMI.2014.2382572.

[16] J.A. Fessler and B.P. Sutton. “Nonuniform fast Fourier transforms using min-max interpolation”. In: *IEEE Transactions on Signal Processing* 51.2 (2003), pp. 560–574. DOI: 10.1109/TSP.2002.807005.

[17] Sebastian Rosenzweig, H. Christian M. Holme, and Martin Uecker. “Simple auto-calibrated gradient delay estimation from few spokes using Radial Intersections (RING)”. In: *Magnetic Resonance in Medicine* 81.3 (2019), pp. 1898–1906. DOI: https://doi.org/10.1002/mrm.27506.

[18] Li Feng. “Golden-Angle Radial MRI: Basics, Advances, and Applications”. In: *Journal of Magnetic Resonance Imaging* 56.1 (2022), pp. 45–62. DOI: https://doi.org/10.1002/jmri.28187.

[19] "Multi-Dimensional Fourier Imaging and Slice Excitation". In: *Magnetic Resonance Imaging*. John Wiley & Sons, Ltd, 2014. Chap. 10, pp. 165–206. DOI: https://doi.org/10.1002/9781118633953.ch10.

[20] Pruessmann KP et al. "SENSE: sensitivity encoding for fast MRI". In: *Magnetic resonance in medicine* 42.5 (1999), pp. 952–962. DOI: https://doi.org/10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S.

[21] Uecker M et al. "ESPIRiT–an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA". In: *Magnetic resonance in medicine* 71.3 (2014), pp. 990–1001. DOI: https://doi.org/10.1002/mrm.24751.

[22] Martin Uecker et al. "Image reconstruction by regularized nonlinear inversionJoint estimation of coil sensitivities and image content". In: *Magnetic Resonance in Medicine* 60.3 (2008), pp. 674–682. DOI: https://doi.org/10.1002/mrm.21691.

[23] M. Yu. Kokurin A. B. Bakushinsky. *Iterative Methods for Approximate Solution of Inverse Problems*. Springer Dordrecht, 2004. DOI: https://doi.org/10.1007/978-1-4020-3122-9.

[24] Weiskopf N et al. "Quantitative magnetic resonance imaging of brain anatomy and in vivo histology". In: *Nature Reviews Physics* 3 (2021), pp. 570–588. DOI: 10.1038/s42254-021-00326-1.

[25] Deichmann R. "Fast high-resolution T1 mapping of the human brain". In: *Magnetic resonance in medicine* 54 (2005), pp. 20–27. DOI: https://doi.org/10.1002/mrm.20552.

[26] J. Gordon Betts et al. *Anatomy and Physiology*. OpenStax, 2013. URL: https://openstax.org/books/anatomy-and-physiology/pages/19-introduction.

[27] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG].

[28] Sergey Ioffe and Christian Szegedy. *Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift*. 2015. arXiv: 1502.03167 [cs.LG].

[29] Kaiming He et al. *Identity Mappings in Deep Residual Networks*. 2016. arXiv: 1603.05027 [cs.CV].

[30] Kai Zhang et al. “Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising”. In: *IEEE Transactions on Image Processing* 26.7 (2017), pp. 3142–3155. DOI: 10.1109/TIP.2017.2662206.

[31] Burhaneddin Yaman et al. “Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data”. In: *Magnetic Resonance in Medicine* 84.6 (2020), pp. 3172–3191. DOI: https://doi.org/10.1002/mrm.28378.

[32] Kerstin Hammernik et al. “Learning a variational network for reconstruction of accelerated MRI data”. In: *Magnetic resonance in medicine* 79.6 (2018), pp. 3055–3071. DOI: https://doi.org/10.1002/mrm.26977.

[33] Hemant K. Aggarwal, Merry P. Mani, and Mathews Jacob. “MoDL: Model-Based Deep Learning Architecture for Inverse Problems”. In: *IEEE Transactions on Medical Imaging* 38.2 (2019), pp. 394–405. DOI: 10.1109/TMI.2018.2865356.

[34] Martin Uecker et al. “Berkeley Advanced Reconstruction Toolbox”. In: *Proceeding of the 23rd Annual Meeting ISMRM, Toronto, Ontario, Canada* (2015), p. 2486.

[35] Xiaoqing Wang et al. “Model-based myocardial T1 mapping with sparsity constraints using single-shot inversion-recovery radial FLASH cardiovascular magnetic resonance”. In: *Journal of Cardiovascular Magnetic Resonance* 21.1 (2019), p. 60. DOI: https://doi.org/10.1186/s12968-019-0570-3.

[36] Ahmad R. “Edge Sharpness Assessment by Parametric Modeling: Application to Magnetic Resonance Imaging”. In: *Concepts in Magnetic Resonance Part A* 44.3 (2015), pp. 138–149. DOI: https://doi.org/10.1002/cmr.a.21339.

[37] Yoav Benjamini and Yosef Hochberg. “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 57.1 (1995), pp. 289–300. DOI: https://doi.org/10.1111/j.2517-6161.1995.tb02031.x.

[38] Yoav Benjamini and Daniel Yekutieli. “The control of the false discovery rate in multiple testing under dependency”. In: *The Annals of Statistics* 29.4 (2001), pp. 1165–1188. DOI: 10.1214/aos/1013699998.

[39] Jakob Assländer et al. *On multi-path longitudinal spin relaxation in brain tissue*. 2023. arXiv: 2301.08394 [physics.med-ph].

[40] Jakob Assländer et al. *Rapid quantitative magnetization transfer imaging: utilizing the hybrid state and the generalized Bloch model*. 2023. arXiv: 2207.08259 [physics.med-ph].

[41] Assländer J et al. "Hybrid-state free precession in nuclear magnetic resonance." In: *Communications Physics* 2 (1 2019), p. 73. DOI: 10.1038/s42005-019-0174-0.

Acknowledgments

I am deeply grateful to the individuals whose guidance and support were instrumental in accomplishing this MSc thesis. First and foremost, I extend my sincere appreciation to Prof. Martin Uecker, head of the Biomedical Imaging Institute at the Technical University of Graz, for the opportunity to work on this project under his mentorship and for showing me the essence of a truly supportive work environment, exemplified by his institute. I am particularly grateful to Moritz Blumenthal, whose essential guidance during the project was invaluable. His humility and approachability always made me feel like a curious learner, never the latest arrival in the office. I extend my gratitude to the other members of the Biomedical Imaging Institute for welcoming me into their team and making me feel genuinely valued. I would also like to express my gratitude to Prof. Alessandra Bertoldo, head of Functional and Anatomical Imaging Research at the University of Padua, for her kind and enthusiastic support throughout the completion of this thesis. A special thank you goes to the best roommate and friend I could ever wish for, Mariaelena, who taught me how important it is to cherish even the smallest gifts that life offers. I want to thank my boyfriend, Massimiliano, for his unwavering love and support. Thank you to my parents, Lucia and Filippo, for always leaving me free to find my way in this world but never leaving me alone. Thank you to my sister, Elisa, for bringing laughter into my life. Finally, I am indebted to my friends and colleagues, Camilla and Nicol, for making our last two years unforgettable.