POLITECNICO DI TORINO

Master of Science in Biomedical Engineering
Biomedical Instrumentation
Department of Mechanical and Aerospace Engineering (DIMEAS)

Master Thesis

Self-Supervised, Deep Learning Denoising of GRAPPATINI T2-Maps in Magnetic Resonance Imaging

---

Supervisors
Prof. Filippo Molinari
Dr. Tom Hilbert
Dr. Thomas Yu

Candidate
Paolo Garelli

March 2025

Abstract

Magnetic Resonance Imaging (MRI) is a widely used medical imaging modality, offering non-invasive and non-ionizing diagnostic capabilities. However, conventional MRI relies on contrast-weighted images influenced by various physical parameters and scanner characteristics, making their interpretation still highly qualitative and subjective. This approach limits analysis only to relative changes in intensities or morphological features and no absolute comparisons can be performed. Quantitative MRI (qMRI) is shifting this perspective by enabling the generation of parameter-specific maps, such as relaxation times T_1 and T_2 maps, capable of providing absolute and quantitative information, enabling longitudinal studies, inter-subject comparisons, and the establishment of normative ranges. Specifically, T_2 relaxation times have been shown to be highly sensitive to physiological and pathological changes, such as intra- and extra-cellular water accumulation and myelin loss. Despite these advantages, gold standard T_2 mapping techniques are still constrained by long acquisition times, impacting patient comfort and increasing the susceptibility to motion artifacts. GRAPPATINI, an accelerated technique, address these limitations by leveraging the orthogonality of two previously established methods in accelerating MRI, GRAPPA and MARTINI. However, this acceleration comes at the cost of increased noise levels and reduced signal-to-noise ratio (SNR) in the reconstructed T_2 maps. Consequently, there is an urgent need to improve the final quality of T_2 maps reconstructed using GRAPPATINI. This thesis proposes two novel strategies to denoise GRAPPATINI T_2 maps using a self-supervised machine learning framework to train a deep learning model, with the aim of bypassing the need for large datasets typically required for supervised learning approaches. The first strategy, operating in the k-space domain, was found to be ineffective, while the second, implemented in the image-space domain, demonstrated exceptional performance on the test set. Furthermore, the image-space strategy showed remarkable generalizability considering 7T brain and knee datasets acquired with different resolutions, field strengths, anatomies, and orientations compared to those used during the training process. Statistical validation was performed through scan-rescan analyses, with results confirming its ability to preserve unbiased and reproducible T_2 values. Reproducibility and generalizability are the key milestones of this work, enabling potential integration of the strategy into scanner reconstruction pipelines without retraining or modifications, independently from the anatomy or the acquisition settings. Eventually, achieving superior denoising performance compared to traditional methods, this work represents a significant advancement toward the potential clinical adoption of GRAPPATINI, both at 3T and 7T, where the combination of high-resolution and high-quality imaging can drive innovations in both research and clinical practice.

Sommario

La Risonanza Magnetica per Immagini (MRI) è una modalità di imaging medicale non invasiva e priva di radiazioni ampiamente utilizzata e diffusa. Tuttavia, l'MRI convenzionale si basa su immagini il cui contrasto è influenzato sia da parametri fisici sia dalle caratteristiche dello scanner, rendendo di conseguenza la loro interpretazione altamente qualitativa e soggettiva. Questo approccio limita l'analisi alla valutazione di soli cambiamenti relativi di intensità o all'estrazione di *features* di carattere morfologico, impedendo comparazioni di tipo assoluto. La Risonanza Magnetica Quantitativa (qMRI) permette di cambiare prospettiva, consentendo la generazione di mappe quantitative specifiche per ciascun parametro, come, ad esempio, le mappe dei tempi di rilassamento cerebrali longitudinale e trasversale T_1 e T_2, in grado di fornire informazioni assolute e quantitative, abilitando di conseguenza studi longitudinali, comparazioni inter-soggetto e l'estrazione di range normativi. In particolare, i tempi di rilassamento trasversali T_2 si sono dimostrati altamente sensibili a cambiamenti sia fisiologici sia patologici, come l'accumulo di acqua intra- ed extracellulare e la perdita di mielina a livello assonale. Nonostante questi vantaggi, le tecniche *gold standard* per la generazione di mappe T_2 sono ancora limitate a causa dei lunghi tempi di acquisizione, i quali influiscono sul comfort del paziente e aumentano la suscettibilità agli artefatti da movimento. GRAPPATINI è un metodo sviluppato per superare tali limitazioni unendo due approcci precedentemente sviluppati nell'ambito dell'accelerazione del MRI e sfruttandone l'ortogonalità: GRAPPA e MARTINI. Tuttavia, questa accelerazione comporta una riduzione del rapporto segnale-rumore (SNR) nelle mappe T_2 ricostruite; di conseguenza, risulta rilevante sviluppare un metodo in grado di migliorarne la qualità finale. Questo lavoro di tesi propone dunque due strategie innovative con lo scopo di ridurre il rumore presente nelle mappe T_2 GRAPPATINI utilizzando un approccio *self-supervised* di *Machine Learning* per l'allenamento di un modello di *Deep Learning*, con l'obiettivo di superare la necessità di grandi dataset tipicamente richiesta dagli approcci di tipo *supervised*. La prima strategia, operante nel k-spazio, si è rivelata inefficace, mentre la seconda, implementata nello spazio dell'immagine, ha dimostrato prestazioni eccezionali sul *test set*. Inoltre, tale strategia ha mostrato una notevole capacità di generalizzare su dataset cerebrali e del ginocchio acquisiti a 7T con diverse risoluzioni, intensità di campo magnetico esterno, anatomie e orientazioni rispetto a quelli utilizzati durante il *training* del modello. La validazione statistica è stata effettuata attraverso analisi *scan-rescan*, con risultati che confermano la capacità di preservare valori T_2 non alterati e riproducibili. Riproducibilità e generalizzabilità rappresentano quindi i traguardi chiave di questo lavoro, permettendo una rapida e potenziale integrazione della strategia nella *pipeline* di ricostruzione all'interno degli scanner,

senza necessità di ulteriori *retrainings* o modifiche, indipendentemente dall'anatomia o dai parametri di acquisizione. Con prestazioni di riduzione del rumore superiori rispetto ai metodi tradizionali, questo lavoro rappresenta un importante progresso nella direzione di un'adozione clinica di GRAPPATINI, sia a 3T che a 7T, dove la combinazione di imaging ad alta risoluzione e alta qualità può portare a innovazioni sia nella ricerca sia nella pratica clinica.

Acknowledgements

At the conclusion of this work, I would like to acknowledge those who supported me throughout this journey. Firstly, my deepest gratitude goes to Prof. Filippo Molinari for supervising this thesis from the Politecnico di Torino's side. Thank you for your valuable suggestions and unwavering belief in this project from the very beginning.

I extend my heartfelt thanks to everyone at Siemens Healthineers Swiss Innovation Hub for warmly welcoming me from day one. A special thank you to Dr.Tom Hilbert and Dr.Thomas Yu for supervising this master's thesis.

Thank you for the opportunity, for believing in me from day one, perhaps even more than I did, and for sharing your expertise. Your support through both the highs and lows has meant a lot to me. I truly look up to you.

A special acknowledgement goes to Gian and Tommi. Thank you for taking me under your wing from day one. I truly appreciate your advices, support, and guidance, both in the project and in life.

A heartfelt acknowledgement goes to all the PhD students, master's students, and interns who shared these past six months with me in the Lausanne office. Thank you, Jocelyn, Natalia, Lina, Lucas, Mathieu, and Chloè , for making me feel part of a special group.

I would also like to acknowledge all the people I met throughout these years at Politecnico di Torino and Alta Scuola Politecnica. A special thank you goes to Giulio, best teammate in every project, and to Marco, for his incredible expertise in Swiss bureaucracy and hikings, but most importantly, for his genuine friendship.

Thank you Chri, Gio, Gian, Luchi for your support during these years and your never-ending friendship. Thanks Chri for your loyalty, Gio for your passion, Gian for your sincerity and Luchi for your "bisognerebbe...".

Most importantly, thank you, Chiara, for your love and patience. You encouraged me every day, supported me through both fears and hopes, and, even from distance, made me always feel as if you were right there beside me.

Merci à Olivier et Lisa de m'avoir accueillie dans leur appartement pendant cette pèriode à Lausanne et pour tous les souvenirs que nous avons crèès ensemble.

Per concludere però, vorrei ringraziare e dedicare tutto questo percorso alla mia famiglia. Grazie mamma e papi per l'amore incondizionato, per essere sempre dalla mia parte e avermi supportato in tutte le mie scelte. Grazie Friski e Ste, siete e sarete sempre per me l'esempio in tutto. Dove sono arrivato e chi sono oggi è tutto merito vostro.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) has its roots in the phenomenon of Nuclear Magnetic Resonance (NMR), first discovered and conceptualized by Stern and Gerlach in 1922 [1] and then reported and experimentally demonstrated by Bloch and Purcell in 1946 [2, 3]. Their work earned them the joint Nobel Prize in Physics in 1952. Subsequently, Lauterbur opened the path to acquire the first NMR-based image introducing the concept of spatial encoding in the NMR signal using magnetic field gradients [4]. This innovation laid the foundation for MRI as we know it today. In the 1980s, the first clinical MRI scanners were installed [5], and since then the technology has seen remarkable advancements, establishing MRI as one of the most widely used imaging modalities for the human body. In particular, MRI has the capability of acquiring images of the brain, in a non-ionizing and non-invasive way. Its field of applications spans from anatomical imaging, up to spectroscopy [6], diffusion imaging [7], and neuroscience, where, for instance, the functional activity of the brain is studied (fMRI)[8].

In a classical MRI experiment, protons in a medium or in a body are exposed to a strong magnetic field, and their nuclear moment tends to align and precess around the direction of the external magnetic field, causing a net magnetization to appear in the body or object. With the aim of perturbing this condition, radio frequency (RF) pulses are then applied at the precession frequency of the protons (known as Larmor frequency ω_0.), and due to this transmission of energy, the magnetization is tilted in the transverse plane. This is the condition that is usually known as resonance. The magnetization in the transverse plane can be measured as a fluctuation in voltage in nearby coils, explained by the law of induction, or Faraday's Law. This effect forms the basis of the MR signal.

After the application of the RF pulse, the spins undergo a relaxation process back to the equilibrium direction along the external magnetic field resulting in a decay of the MR signal. As protons relax, two key processes (T_1 and T_2 relaxation) come into play. In T_1 relaxation, the magnetization realigns with the external magnetic field, while in T_2 relaxation, spins lose phase coherence in the transverse plane, causing loss of magnetization in the transverse plane. In order to form images, it is essential to spatially encode the MR signal, during the acquisition. This is typically done using magnetic field gradients to sample the frequency domain of the object in the scanner, the so-called k-space. In the k-space, different points correspond to various spatial frequencies within the image. The final image is reconstructed by applying a mathematical transformation, the inverse Fourier transformation, to convert this frequency information into the spatial domain.

Conventional MRI images are often described as "weighted" toward one of these relaxation parameters, or diffusion parameters (ADC, FA, among others) as a function of their contribution to the overall contrast of the image. By adjusting the parameters of an MRI sequence, the effect of one specific property can be emphasized over the other, producing, considering the relaxation parameters, either T_1-weighted or T_2-weighted images. T_1-weighted images excel at highlighting anatomical details and tissues with a high fat content, while T_2-weighted images make fluid-filled regions, e.g. cerebrospinal fluid or edema, appear brighter, facilitating the identification of inflammation and lesions [9]. However, completely isolating specific contrast mechanisms in MRI is inherently challenging due to the intricate relation between these tissue or system properties. Even considering the choice of MRI sequence, whether spin echo (SE), gradient echo (GE), or inversion recovery (IR), as basis for a particular contrast inevitably allows for some residual contributions from other factors. Additionally, hardware characteristics such as magnetic field strength (B_0), gradient performance, system imperfections, and coil sensitivities introduce additional variability, even when imaging the same subject under similar conditions. Eventually, in biological tissues, physiological variables such as blood flow, temperature, and hydration could further modulate relaxation times, adding an additional layer of complexity to the resulting image contrast.

1.2 Quantitative MRI

Since the early days of MRI, pioneers in the field were convinced that the future of the technology would lean more toward quantitative rather than qualitative analysis. Two key questions arise: why is the term *qualitative* always associated with the interpretation of conventional MRI images? And what exactly does *quantitative* mean in the context of MRI? As mentioned earlier, conventional MRI images are typically weighted toward a specific parameter. This means that the imaging sequence and acquisition are designed to emphasize one particular aspect, such as T_1 or T_2 relaxation times, making it the dominant feature in the resulting image. Moreover, the pixel intensity values in MRI images are proportional to the MR signal. However, their interpretation is valid only from a relative perspective. This means that only the relative differences in image intensities carry meaning, and absolute comparisons between images from different individuals or imaging sequences are not reliable. Absolute pixel intensities cannot be compared directly.

On the other hand, the term "quantitative" in the MRI field is defined as the processes necessary to obtain a map, which is a spatial representation of the distribution of absolute values for a specific physical or chemical parameter. In fact, every intensity is represented by a value with a specific physical or chemical unit of measurement [11]. For example, a T_1 map represents a spatial representation of the object in the scanner as an image where each voxel provides the longitudinal relaxation time in milliseconds. An overview on potential applications of qMRI is shown in **Figure 1.1**.

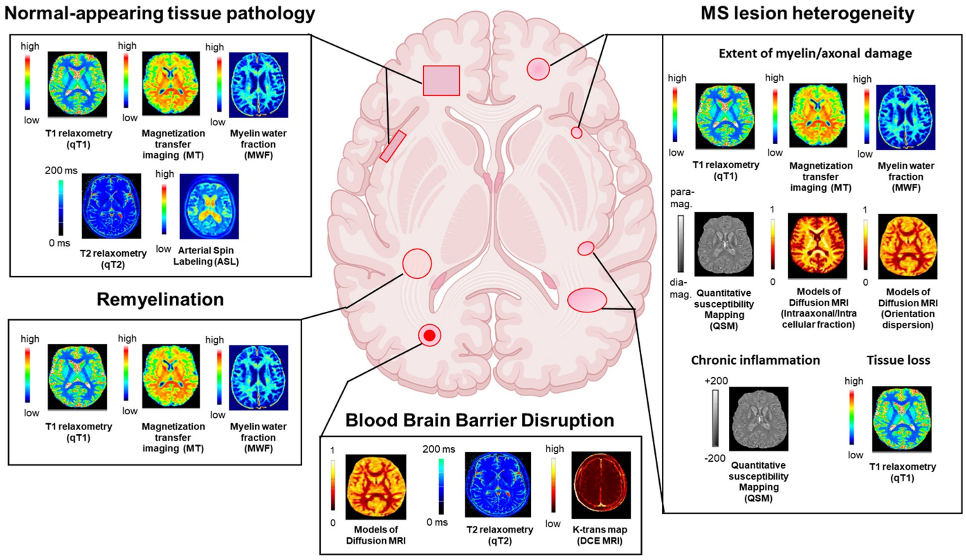

Figure 1.1: General overview of qMRI techniques applied in the context of Multiple Sclerosis. Figure entirely reproduced from [10].

Currently, the MRI images interpretation is still highly qualitative, and the need to rely on robust and reproducible techniques and methodologies capable of extracting quantitative information from the raw MRI signal is relevant. In conventional MRI, only tissue volumes and spatial characteristics can be effectively measured. Differences in tissue, based on biological or chemical properties cannot be quantified. For example, a variation in contrast on a T_2-weighted image, may be correlated with multiple underlying contrasts mechanisms such as relaxation, diffusion or other physiological effects making it harder to differentiate phenomena such as edema, necrosis, demyelination etc,[11] leading to difficulties in the diagnosis, or limiting it only to cases where there is clear evidence of the presence of an abnormal structure. This is the reason why a more quantitative, robust and reproducible approach may be ideal to open frontiers also in longitudinal studies, comparing inter-subject with the use of normative values and intra-subject by monitoring the temporal evolution of a pathological tissue. However, two primary challenges have historically constrained the widespread clinical adoption of quantitative tissue properties measurements: the long acquisition times required and, at least initially, and to some extent still today, a lack of reproducibility for some of the quantitative measurements available [12]. Currently, for qMRI to be translated for clinical use, there is a need to address both the above concerns for reproducible and interpretable images which can be acquired feasibly in the clinical routine. (see **Figure 1.2**). However, high reproducibility may come at the cost of reduced accuracy, while high accuracy can result in lower reproducibility [10]. Reducing acquisition times often leads to worse Signal to Noise Ratio (SNR) and image quality. As a result, numerous methods and techniques have been developed throughout the parametric map acquisition pipeline with the goal of reducing

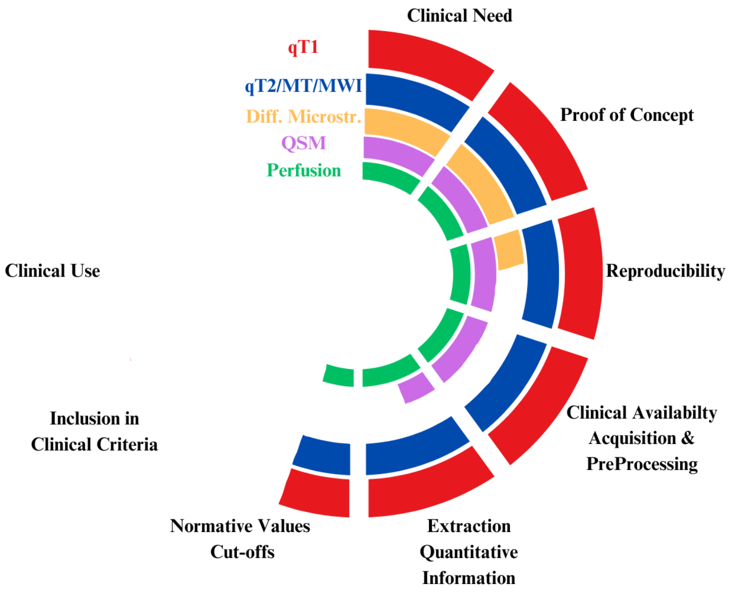

Figure 1.2: A schematic overview on the clinical maturity stages of key approaches in quantitative MRI and their advancements from 2021. Schematic representation inspired by the work in [10]

scanning times while improving the quality of the final maps.

1.3 Thesis Rationale and Outline

This work aims to enhance the quality of T_2 maps generated using the GRAPPATINI method by means of a self-supervised ML approach (using Deep Learning) with the aim of getting better results than traditional methods. The focus of this work is to apply this approach to the two main domains of MRI: the k-space domain (or Fourier domain) and the image domain, which in this context is better defined as the parameter domain. A denoising step is needed since the GRAPPATINI method reduces the acquisition time of T_2 maps, and, consequently, a reduction in the SNR and an amplification of the noise level is expected. The objective is to develop a robust and unbiased denoising strategy that could potentially improve image quality and address common complaints about standard GRAPPATINI. This approach ensures that the denoising process does not disrupt or statistically alter the values obtained using the standard technique, preserving the diagnostic integrity as a consequence. In this context, a self-supervised approach was selected due to the lengthy acquisition times required to obtain a fully sampled and extensive dataset, which would be necessary for use as ground truth in a supervised approach.

The structure of the thesis is as follows. In **Chapter 2**, a general background on MRI physics is provided to lay the foundation for the concepts discussed throughout the thesis. The basic principles of image acquisition are also introduced, along with relevant techniques and technologies, explaining how data is typically collected during an MRI session. **Chapter 3** then addresses the techniques employed in accelerated MRI, focusing in particular on parallel imaging and model-based reconstruction methods such as the GRAPPATINI method. These techniques form the basis for the methodologies explored in this project. In **Chapter 4**, the methods implemented throughout the project are explained, starting from the datasets used, the self-supervised framework and in the end the two strategies implemented in the two different domains. In **Chapter 5** the results are shown and then discussed in **Chapter 6**. Finally, **Chapter 7** concludes the thesis with a summary of the key findings and a discussion of the results, along with possibilities for future work.

Chapter 2

Background

The physics and technology behind MRI are crucial for understanding how a natural phenomenon, like Nuclear Magnetic Resonance (NMR), can be leveraged to generate signals that are transformed into detailed images. A solid grasp of these principles is necessary for both working with and interpreting MRI images in clinical and research settings. In this chapter, the essential physical concepts underlying MRI are covered, including how magnetic fields interact with protons in the body, the role of radio-frequency (RF) pulses in generating signals, the importance of T_1 and T_2 relaxation processes in determining the contrast of the image, and how spatial encoding and gradients enable the creation of images. The key references for the theory content of this chapter are from textbooks [13, 14] and in research articles cited below.

2.1 MRI Physics

The basic concept of MRI involves placing the body in an external magnetic field, after which radio-frequency (RF) coils transmit and receive radio waves within the scanned area. The key protagonists in this process are the magnetic moments of the protons in the body's tissues, particularly in water molecules, which possess the ability to align with the magnetic field and interact with the RF waves. These interactions generate signals that are detected as voltages by the RF coils, digitized by an analog-to-digital converter, and then processed by a computer to create the final images.

2.1.1 Spin and Magnetic Moment

In classical physics, any rotating object possesses a property known as angular momentum, which is analogous to linear momentum in linear motion. Angular momentum reflects how much rotational inertia a body has, and it depends on extensive properties of an object, e.g. mass, dimensions, shape, and rotational velocity. This concept can also be extended to subatomic particles, such as protons and electrons.

These particles exhibit a property called spin, similar as a concept to classical angular momentum. Spin can only be described with discretized values. For particles like protons, the spin is non-zero, which means that the particle has an intrinsic magnetic moment,

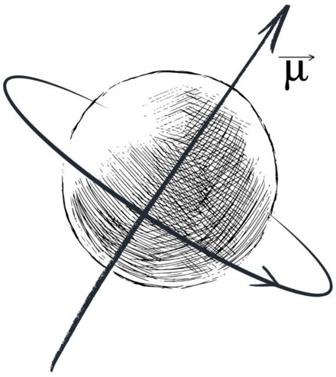

Figure 2.1: Graphical representation of a spinning nucleus around its axis.

often denoted by $\vec{\mu}$. The magnetic moment can be thought of as a vector that describes both the direction and strength of the magnetic dipole that arises from the "spinning" motion of the particle (see **Figure 2.1**). Moreover, there is a relation between the magnetic moment and the spin angular momentum $\vec{J}$ of a particle, such that:

$$\vec{\mu} = \gamma \vec{J} \tag{2.1}$$

, where γ is the so-called gyromagnetic ratio. In MRI, the hydrogen atom, H_1^1, plays a crucial role because the single proton in the nucleus has a spin of $s = +\frac{1}{2}$.

2.1.2 Protons in a Magnetic Field

When protons, and thus magnetic moments $\vec{\mu}$, are placed in a region where the external magnetic field, usually referred to as $\vec{B_0}$, is different from zero, in a similar way as the needle of a compass, they tend to align along the direction of the magnetic field. This is the effect of a torque, that the external magnetic field $\vec{B_0}$ exerts on the magnetic moments (see **Figure 2.2**). The protons are not capable of aligning exactly with the direction of the external magnetic field, and the consequence is that they precess around the direction of the external magnetic field like a gyroscope under the effect of the gravitational force. This precession motion occurs at a specific frequency, known as the Larmor Frequency:

$$\omega = \gamma B_0 \tag{2.2}$$

Of relevant interest is the net magnetization, $\vec{M}$, defined as the vector sum of the magnetic moments within a sample, e.g. an ensemble of hydrogen atoms or water molecules:

$$\vec{M} = \sum \vec{\mu} \tag{2.3}$$

When the external magnetic field is equal to zero, the different magnetic moments are randomly oriented, such that the vector sum is equal to zero. If an external magnetic field is present, the ensemble of spins will produce a net magnetization which also precesses

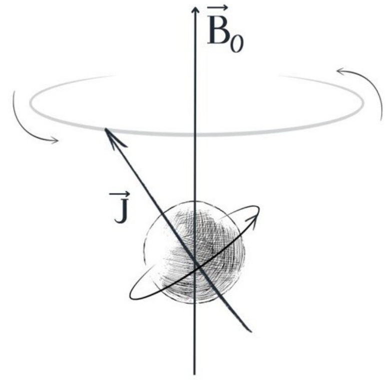

Figure 2.2: Magnetic Moment precession.

around the external magnetic field at the Larmor Frequency. One can show that the alignment of the magnetic moments with the external magnetic field corresponds to the lowest energy configuration, which is preferred at equilibrium. These correspond to the quantized spins being either up (parallel) or down (anti-parallel). However, thermal fluctuations result in magnetic moments also oriented anti-parallel to the B_0 field (see **Figure 2.3**). The thermal effects result in a distribution of magnetic moments which has a relatively small bias toward spins parallel to the external magnetic field, i.e. spin up.

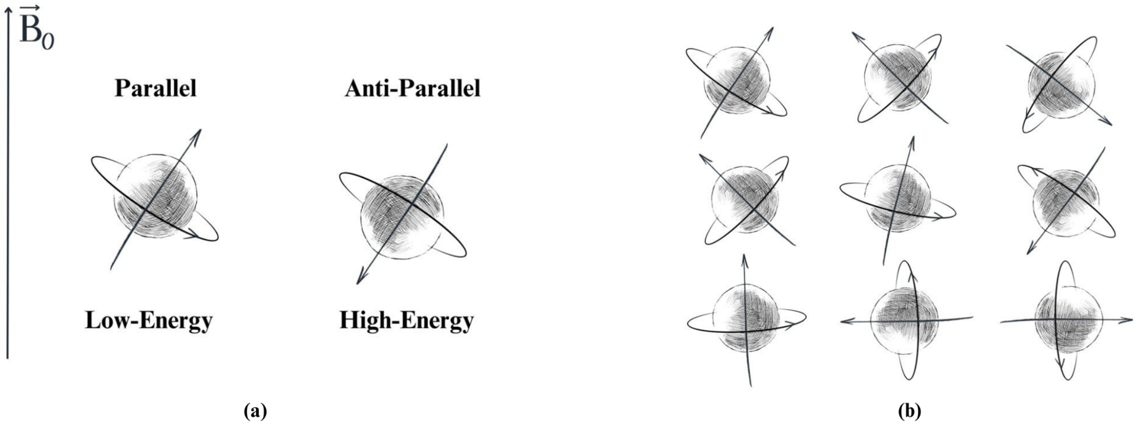

Figure 2.3: (a) Spin-up and spin-down orientations and the associated energy level. (b) Distributions of random oriented spins in a region where the external magnetic field $B_0 = 0$.

The proportion of spin up to spin down moments can be derived from the Boltzmann distribution:

$$\frac{N^+}{N^-} = e^{\frac{-\Delta E}{kT}} \tag{2.4}$$

, where N^+ and N^- are the total number of spins in the parallel and anti-parallel orientations, k is the Boltzmann constant, ΔE is the energy difference between the two states and T is the temperature in Kelvin. Given this consideration, studying the system from a macroscopic perspective, on average there will be more spins aligned with the external magnetic field than the ones anti-parallel to the external magnetic field. The net magnetization, in this case, will be aligned with the magnetic field direction and the value can be referred to as $\vec{M_0}$:

$$M_0 = \frac{\gamma^2 h^2 N B_0}{4kT} \tag{2.5}$$

In **Equation 2.5**, γ is the gyromagnetic ratio, h is Planck's constant, N is the number of spins, B_0 is the magnitude of the external magnetic field, k is Boltzmann constant, and T is the temperature in Kelvin. Eventually, the receiver coils measure voltages coming from change in the magnetic flux driven by changes in the net magnetization M_0. This is why the net magnetization M_0 is one of the key factors for the SNR of the final acquisition. Note that the net magnetization is proportional to the strength of the external magnetic field.

2.1.3 Resonance

To summarize the previous section, an external magnetic field induces a net magnetization aligned with the direction of the field since the energetically favorable orientation for the magnetic moment is aligned with the field, modulo the effects of thermal fluctuations. The question then arises, whether the net magnetization can be manipulated by e.g. rotating it into another plane. From the quantum mechanical perspective, the energy difference between the spin up and spin down in which a proton can be found is equal to:

$$\Delta E = hf \tag{2.6}$$

where f is the Larmor Frequency of the precession motion. In order to induce protons to change their spin state and consequently the components of the net magnetization, we need to give the system discrete amounts of energy which are integer multiples of ΔE. **Equation 2.6** tells us that we can give energy to the system, and the system will enter a resonance condition if we introduce a perturbation exactly at the Larmor frequency. In our case, this perturbation is another magnetic field, typically referred to as the B_1 field.

The B_1 field is a time-varying magnetic field oscillating at the Larmor frequency. Considering **Equation 2.2**, that $\gamma = 42.57 MHz/T$, and that the external magnetic field magnitude is typically on the order of 1.5T or 3T, it is straightforward to compute that the Larmor frequency is in the RF spectrum. The RF pulse is introduced into the system through a transmit coil and the net magnetization is no longer parallel to the direction of the external magnetic field, but most of the protons will align their precession motion with the direction of the combined magnetic field.

2.1.4 Bloch Equations

To reach a deeper understanding of what is happening mathematically in the process of tilting the magnetization vector the Block Equations are presented. The Bloch equations are phenomenological equations that describe the motion of the net magnetization vector when the relaxation times T_1 and T_2 are known. Moreover, to better align with the conventions in terms of the common nomenclature in MRI, it will be assumed that the external magnetic field B_0 is aligned with the z direction of a Cartesian frame of reference, so that $\vec{B_0} = B_0\hat{z}$. The z direction will be referred to as *longitudinal*, while the xy plane as *transverse plane.* As for the magnetization vector $\vec{M}$, the longitudinal component is:

$$\vec{M}_{\|} = \vec{M_z} = M_z\hat{z} \tag{2.7}$$

while the transversal component, in the xy plane is:

$$\vec{M}_{\perp} = \vec{M_{xy}} = M_x\hat{x} + M_y\hat{y} \tag{2.8}$$

Given these definitions, the Bloch Equations can be derived considering that:

- γ , the gyromagnetic ratio, which establishes a relationship between the angular momentum and the magnetic moment.
- $\vec{\tau}$, the torque on a magnetic moment $\vec{\mu}$, given by $\vec{\tau} = \vec{\mu} \times \vec{B}$
- The torque $\vec{\tau}$ can be written also as the time derivative of the angular momentum, $\vec{\tau} = \frac{d\vec{J}}{dt}$

When no relaxation mechanisms are considered, the time evolution of $\vec{M}$ is described by the following torque (or precession) equation:

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B} = \gamma \begin{bmatrix} (M_yB_z - M_zB_y)\hat{x} \\ +(M_zB_x - M_xB_z)\hat{y} \\ +(M_xB_y - M_yB_x)\hat{z} \end{bmatrix} \tag{2.9}$$

The total magnetic field $\mathbf{B}$, can be considered as including a z component, which is the typical B_0, and two other components, forming a second magnetic field in the transverse plane (*xy plane*), the so-called B_1 field:

$$\vec{B_x} = B_1\cos(\omega t)\hat{x} \quad \vec{B_y} = B_1\sin(\omega t)\hat{y} \quad \vec{B_z} = B_0\hat{z} \tag{2.10}$$

If the RF pulse is designed so that the energy it is carrying is exactly equal to **Equation 2.6**, and such that it can be defined by:

$$\vec{B_1} = B_{1x}\hat{x} \tag{2.11}$$

the net magnetization will be tilted in the transverse plane. In the case where only a single RF pulse is present, the angle by which $\vec{M}$ moves, referred to as *flip angle* is equal to :

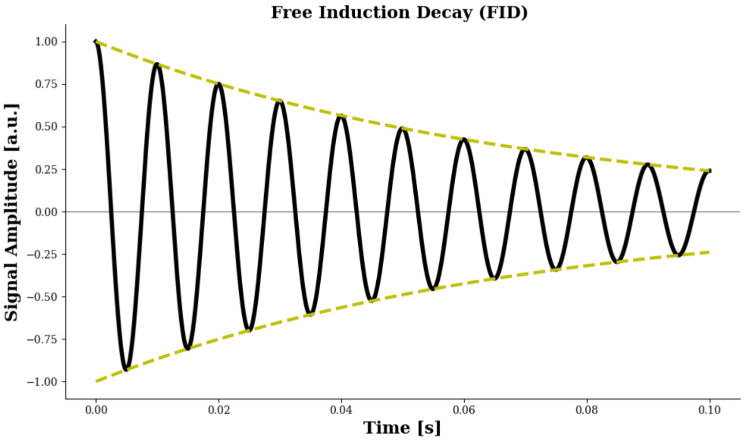

Figure 2.4: Free Induction Decay (FID) signal.

$$\alpha = \gamma B_1 t_p \tag{2.12}$$

, which is the angle by which the net magnetization vector is tilted away from the z-axis. If the magnetization is tilted exactly in the transverse plane, the RF pulse is defined as a 90° pulse.

After this perturbation, a receiver coil can measure the transverse magnetization M_{xy} and the signal produced is typically referred to as *Free Induction Decay (FID)*, which is characterized by a rapid and exponential decrease due to relaxation phenomena (see **Figure 2.4**).

2.1.5 Longitudinal and Transversal Relaxation

The resonance condition is not something that lasts forever. As a natural tendency, every system will tend spontaneously toward equilibrium in the absence of an external force. For this reason, whenever the perturbation is removed from the system, the magnetization vector will again tend to align with the external magnetic field. This recovery phenomenon is also called *relaxation*. This relaxation phenomena happens in both the longitudinal and transversal components of the net magnetization, arising from different mechanisms.

Recovery of longitudinal magnetization arises from a mechanism that is called *spin-lattice relaxation*. In particular, protons will lose energy by interacting with the environment in which they are located, typically referred to as the lattice. Transversal relaxation, where the transversal component of the net magnetization goes to zero, non-homogeneity and the small magnetic fields that arise due to precession of the spins. This second mechanism is also known as *spin-spin relaxation*. Considering a complete recovery after the application of an RF pulse, the longitudinal and the transverse components of the magnetization follow a relaxation process that can be modeled by the following expressions:

$$\vec{M}_z(t) = M_z(0)e^{-\frac{t}{T1}} + M_0(1 - e^{-\frac{t}{T1}}) \tag{2.13}$$

$$\vec{M}_{xy}(t) = M_z(0)(1 - e^{-\frac{t}{T1}})e^{-\frac{t}{T2}} \tag{2.14}$$

The relaxation time constants T_1 and T_2 are tissue-specific properties that play a crucial role in determining the contrast observed in MR images, and relaxation curves of the main brain tissues are reported in the figure below (see **Figure 2.5**).

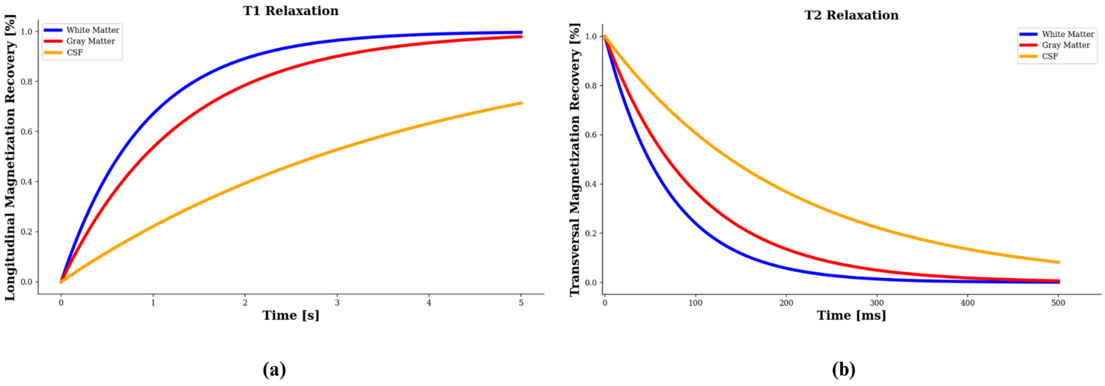

Figure 2.5: T_1 (a) and T_2 (b) relaxations for the main tissues in the brain: white matter, gray matter and CSF. Differences in the curves are the consequence of different relaxation parameters, due to their tissue-specificity.

2.2 Image Acquisition

Building on the fundamental physics of MRI introduced in the previous chapter, this following chapter focuses on the core tools essential for translating raw MRI signals into images. The concepts covered form the backbone of MRI image formation and reconstruction.

2.2.1 Gradients

The first step to form MR images is to impose a spatial variation to the z-component of the main magnetic field $\vec{B_0}$ using magnetic field gradients. Generally, magnetic fields gradients can be defined by the following relations:

$$G_x = \frac{\partial B_Z}{\partial x} \qquad G_y = \frac{\partial B_Z}{\partial y} \qquad G_z = \frac{\partial B_Z}{\partial z} \tag{2.15}$$

Most commonly, a linear variation is imposed, such that, for instance, the z-component along the x axis changes linearly with respect to x. A linear change in the total magnetic field experienced by the protons will cause the Larmor frequencies to shift proportionally along the direction of the gradient.

Considering for instance a gradient applied in the x direction:

$$f = \bar{\gamma}(B_0 + G_x x) \tag{2.16}$$

In regions where the total magnetic field is lower, the protons precess at a lower frequency, whereas in regions with a higher magnetic field, their precession frequency increases. This principle of adjusting Larmor frequencies in protons forms the basis of spatial encoding in MRI. Spatial Encoding is the process that allows the formation of an image in MR. A spatial variation of the magnetic field eventually imposes a spatial variation in the measured signals. Through some modeling and signal processing, one can show that these spatially varying signals can be transformed into images thanks to the inverse Fourier transformation. Typically, magnetic field gradients are applied as short pulses, altering the local magnetic field for a brief period.

2.2.2 K-space

The fundamental concept for image reconstruction in magnetic resonance imaging is the Fourier representation of an image, commonly referred to as the k-space. The Fourier Transform (FT) is a mathematical tool that is used to analyze signals by decomposing them into their constituent frequencies (see **Figure 2.6**). In essence, it transforms a signal from a specific domain(time, space,..), into the frequency domain, where it is represented as a sum of sinusoidal components, each characterized by a certain frequency and amplitude.

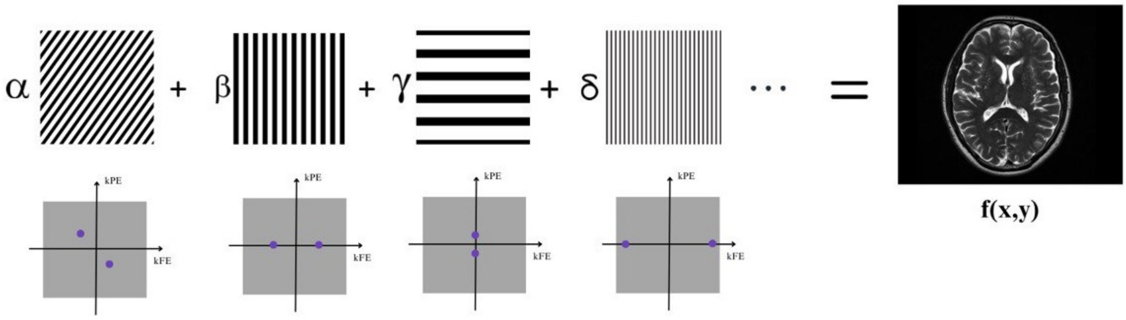

Figure 2.6: The image f(x,y) is represented as a weighted sum of 2D orthogonal basis functions, each with specific frequencies and orientations. Each basis function (such as the components weighted by α, β, and γ) contributes uniquely to the overall image, similar to decomposing a finite dimensional vector along basis vectors. This illustrates how an image can be constructed from various frequency components through Fourier coefficients.

The raw data measured by the MR scanner are points in k-space, representing specific spatial frequency components of the underlying image. The relationship between the k-space domain and the spatial domain is governed by the Fourier Transform (FT) and its inverse, known as the Inverse Fourier Transform (IFT).

Typically, a centered representation of k-space is used, where low frequency components are concentrated at the origin. In this representation, the central portion of the k-space corresponds to the low-frequency area, which contains the majority of information related to the contrast and structure of the underlying image. In contrast, the peripheral area represents the high frequencies, which store information about edges, sharpness, and image resolution.

2.2.3 Spatial Encoding

The entire encoding process can be summarized by the action of three main gradients: the *slice selective gradient*, the *frequency encoding*, and the *phase encoding*. By properly modulating their activations and combinations, the entire k-space can be acquired, i.e. every spatial frequency component, for a specific slice, can be probed.

Slice selection is the process that allows spin excitation to be confined to a specific two-dimensional plane or slice within the patient's body. It is realized by exciting the spins with a RF-pulse while, at the same time, a linear gradient is used to have spatially dependent Larmor frequencies. The relationship between the selected slice, the RF bandwidth, and the gradient applied to accurately excite the spins in that particular region is illustrated in the figure below. (see **Figure 2.7**)

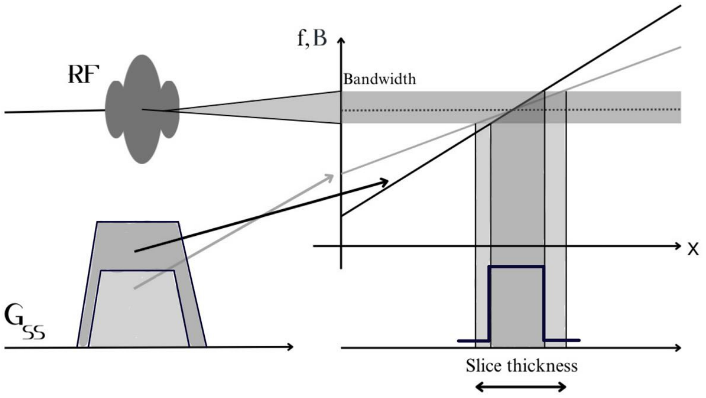

Figure 2.7: Slice selection process in MRI. The thickness of the selected slice is determined by the bandwidth of the RF pulse and the amplitude of the magnetic field gradient applied.

Once a slice is selected, the next step is *in-plane localization*, which involves spatially encoding information along the two remaining dimensions. This is achieved through the application of *frequency* and *phase encoding* gradients. The *frequency encoding* gradient is applied during signal acquisition, causing spins to precess at varying frequencies based on their position along one axis. Previously, the *phase-encoding* gradient was applied to assign a specific phase to the spins along that direction.

Each application of the frequency encoding and phase encoding gradients results in the acquisition of one line in k-space, aligned with the read-out direction. To sample different lines in k-space, a different phase encoding gradient is applied with larger gradients resulting in larger phase differences due to the spatial location and hence lines further away from the k-space center. This process is summarized in the diagram below (see **Figure 2.8**).

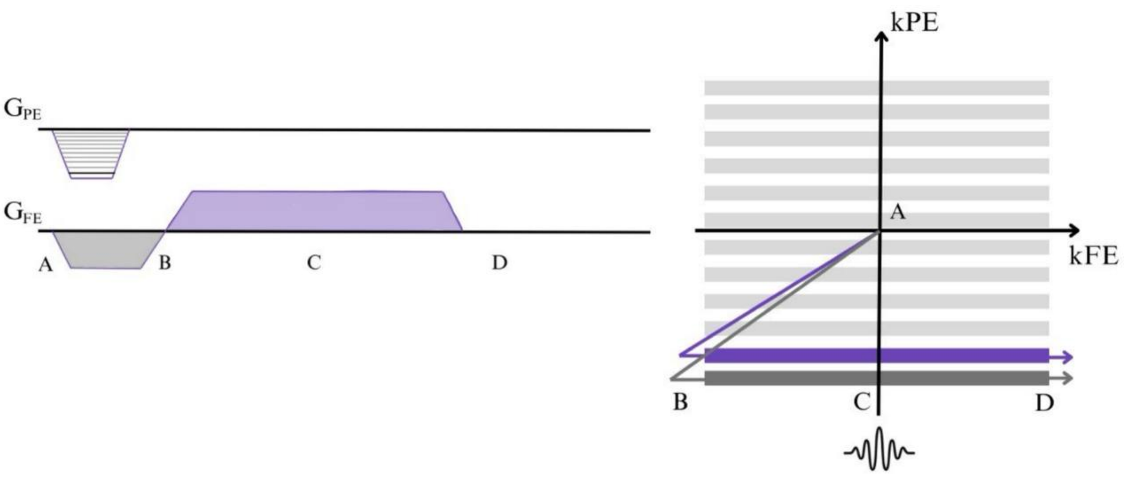

Figure 2.8: Schematic representation of a k-space trajectory and the acquisition of two lines considering a Cartesian sampling.

2.3 Quantitative Mapping

Quantitative MRI, or parameter mapping in general, aims to generate accurate and reliable maps of physical and biophysical parameters, offering significant potential for clinical applications. The general approach involves reconstructing a series of images acquired with a specific contrast tailored to the parameter being mapped, followed by a pixel-wise fit to a signal evolution model to estimate the parameter. This section will specifically focus on T_2 mapping, examining its current uses, technical characteristics, and current challenges.

2.3.1 T_2 Mapping

T_2 mapping methods are methods which generate spatial maps of the relaxation time T_2. These T_2 values are typically measured in milliseconds (ms). These techniques can provide crucial insights into the tissue composition and chemical state by reflecting variations mainly in water content, tissue structure, and pathological changes, making it a potential biomarker of interest in both research and clinical settings.

The gold standard for T_2 mapping is the Carr-Purcell-Meiboom-Gill (CPMG) sequence [15], also known as the Multi-Echo Spin-Echo (MESE) sequence. The sequence diagram is illustrated in the figure below (see **Figure 2.9**). The process begins with a 90° pulse which tilts the magnetization into the transverse plane, and then a 180° RF pulse is applied after some time t to inverse the spins in the transverse plane, causing the spins to rephase back to their original position at time $2t$. This rephasing produces an echo signal at a specific echo time referred to as TE. Without the 180° RF pulse, this would generate a free induction decay (FID) signal. Due to T_2* effects, the FID signal decays rapidly. As T_2 relaxation progresses, the amplitude of successive echoes decreases, following the relation described in **Equation 2.17**.

$$M_{xy}(TE) = M_0 e^{-\frac{TE}{T_2}} \tag{2.17}$$

T_2 values can be determined by acquiring images at different echo times and performing a voxel-wise fitting of the signal decay model. In a typical multi-echo protocol, multiple slices can be acquired within a single repetition time (TR) using interleaved slice acquisition. This approach enables the acquisition of several slices during the extended TR period, thereby accelerating the imaging process. However, there is a trade-off between the number of slices acquired simultaneously and the total number of echoes in a single TR, also known as the Echo Train Length (ETL). ETL is a critical parameter since it determines the number of data points available per voxel for the fitting process. A higher ETL generally improves the precision of T_2 estimation. Even when interleaved slice acquisition is considered, the overall process remains time-consuming, requiring the application of optimized strategies to further enhance its efficiency.

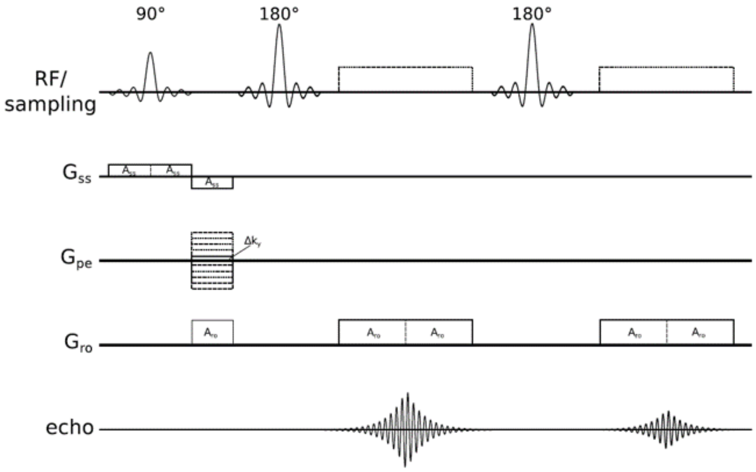

Figure 2.9: Sequence Diagram for the Carr-Purcell-Meiboom-Gill sequence, also known as Multi-Echo Spin-Echo sequence.

2.3.2 Clinical Applications of T_2 Measurements

One of the main fields of application for T_2 measurements is brain imaging. In particular, both relaxation times, T_1 and T_2, are correlated with relaxation processes that dissipate energy due to the interaction of protons with the surrounding lattice or with other protons. This is why they are closely linked to the biophysical properties of the tissues. The amount of water, the iron content, and the myelin content (i.e. lipid-based sheaths surrounding axons of nerves) are all factors that influence the relaxation times T_1 and T_2. Differences in white matter (WM) and gray matter (GM) relaxation times have been attributed to variations in water and myelin content around the nerves [16]. Moreover, as demonstrated in [17, 18], increases in myelin content are associated with shortening of relaxation times

T_1 and T_2. Similarly, [19] showed that iron accumulation also leads to a reduction in these relaxation times.

Moreover, both T_1 and T_2 have the potential to serve as biomarkers in the detection of various pathologies, such as Multiple Sclerosis (MS), Parkinson's Disease (PD), and Alzheimer's Disease (AD). In MS, increases in T_2 have been linked to one of the typical signs of the pathology, namely inflammation of normal-appearing white matter (NAWM) [20]. In PD, myelin degradation, particularly in the brainstem region, the origin of the well known motor disorders associated with PD, has been strongly correlated with shortening of T_2 values [21]. For patients with AD and memory loss, increases in T_2 have been observed in NAWM [22].

Lastly, T_2 values, in particular, have shown potential in the study and diagnosis of temporal lobe epilepsy (TLE). In TLE, the primary driver of epilepsy is hippocampal sclerosis, and increases in T_2 values have proven useful as potential biomarkers for identifying seizure foci within the hippocampus or temporal lobe[23]. The fields of application of T_2 mapping are not restricted to brain imaging only, but span from cardiac to musculoskeletal and prostate imaging, where a quantitative approach can potentially carry more information about biological properties, structural organization, and composition than traditional qualitative or weighted approaches. In cardiac magnetic resonance imaging, high values of T_2 are associated with an increase in free water, considered an indicator of the presence of edema [24]. Differences in T_2 values are also considered indicators of acute myocardial infarction, myocarditis, rejection of heart transplant [25] and potential dilated cardiomyopathy or general myocardial injury[26]. Moreover, several works are present in the T_2 mapping of knee cartilage, where with these techniques it is possible to probe and study the composition of cartilage disks with a non-invasive approach, to better study the short and long-term effects of pathologies [27] or the effects of physical training[28].

Chapter 3

Accelerated MRI

As mentioned in the previous chapter, the principal disadvantage of acquiring T_2 maps with the gold standard CPMG sequence is the long acquisition time. This is a problem intrinsic to the nature of MR itself, in contrast to other imaging modalities like Computed Tomography (CT). This limitation has driven significant research into acceleration methods for k-space acquisition to make MRI more efficient. One natural approach to speeding up the process would be sampling fewer points in k-space (i.e. undersampling), but this must be balanced against the well-established Shannon-Nyquist theorem [29] , which establishes a theoretical condition for how much of the k-space must be sampled for perfect reconstruction. According to this principle, inadequate sampling leads to aliasing artifacts, which compromise image quality. To address these challenges, two primary categories of techniques have emerged: parallel imaging, which takes advantage of the possibility of relying on multiple receiver coils, and model-based reconstruction methods, which incorporate prior knowledge directly into the reconstruction process [30].

In this section, the two primary techniques underlying the GRAPPATINI are introduced, both in the parallel imaging and in the model-based family of methods.

3.1 Parallel Imaging

Parallel imaging represents one of the key approaches developed to accelerate MRI acquisition. These methods operate by undersampling the k-space and multiple receiver coils acquisitions can be used to compensate for artifacts. Although this significantly reduces acquisition time on the one hand, undersampling introduces aliasing effects that complicate the overall reconstruction process. A clear relationship exists between the sampling density, resolution and field of view (FOV) of the k-space: if acquiring more k-space lines leads to a higher resolution, acquiring k-space lines with higher inter-space in between reduces the FOV, leading to folding artifacts[31] (see **Figure 3.2**). These artifacts are particularly detrimental in the T_2 mapping, where voxel-wise exponential fitting is required, making accurate reconstruction critical. To address this, techniques have been developed to recover the missing k-space lines. Methods such as GRAPPA achieve this in the k-space domain, while approaches like SENSE [32] unfold the undersampling artifacts using optimization in image space.

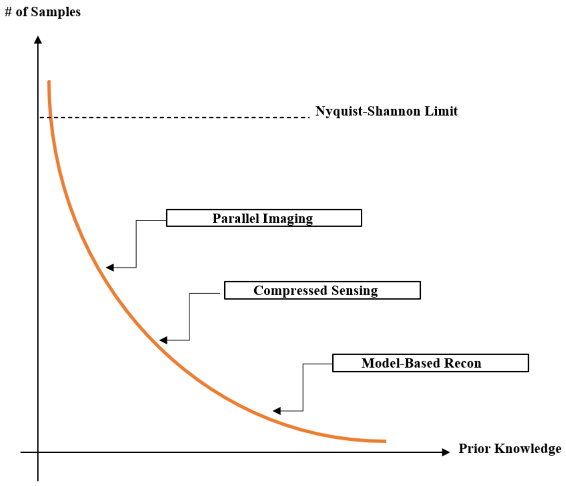

Figure 3.1: Schematic showing the relationship between the number of MRI samples acquired and prior knowledge used in reconstruction, highlighting how the more prior knowledge is introduced, the lower the total number of samples acquired and necessary.

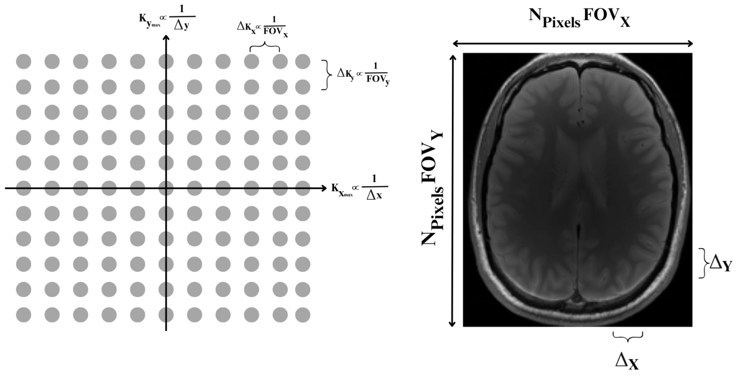

Figure 3.2: A schematic representation illustrating the relationship between the field of view in the x and y directions and the spacing between lines in the k-space shown on the left. The total number of lines in one of the two directions of the 2D k-space is proportional to the resolution of the underlying image.

3.1.1 GRAPPA

The GRAPPA (GeneRalized Autocalibrating Partially Parallel Acquisitions) method is designed to recover missing k-space lines after MRI data acquisition, allowing the generation of an artifact-free image through Fourier transform reconstruction [33]. Each coil could, in theory, independently acquire a k-space and GRAPPA addresses the problem of undersampling by synthesizing missing data combining information from all of them. In a multi-receiver acquisition, the final image intensities are weighted by the coils' sensitivity profiles. When considering this effect in terms of image intensity as a multiplication, it corresponds to a convolution in the frequency domain (k-space), based on the known relationship between the image and frequency domains. This convolution means that information about specific k-space locations can be inferred from neighboring k-space points. In **Figure 3.3**, a flow chart is presented illustrating the main steps of GRAPPA.

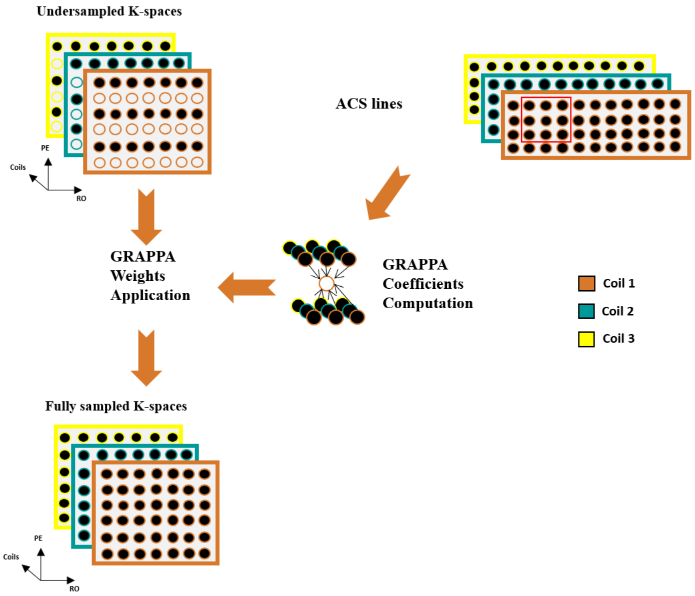

Figure 3.3: Schematic representation of the GRAPPA pipeline. In the scheme, every black dot represents an acquired location in the k-space (sources), while every white dot is not acquired (targets). The three different colors of the boxes, are to highlight k-spaces corresponding from different coils.

A key component of GRAPPA is the Auto Calibrating Signal (ACS) region, a smaller, fully sampled patch of k-space. The ACS region has the same Field of View (FOV) as the desired image but is acquired at a lower resolution, i.e. fewer k-space lines with the same spacing as the final target image. The ACS is fully sampled across the coil dimension and can be acquired during the accelerated scan. This ACS data is used to compute the weights for the GRAPPA method, which are then applied to the undersampled k-space to recover the missing lines. In accelerated MRI methods, a key parameter commonly defined is the *Acceleration Factor (AF)*. Since phase encoding is typically the most time-consuming part of the acquisition process, the AF can be described as the ratio between the total number of phase encoding lines required for a specific resolution and the number of lines actually acquired. This parameter effectively quantifies the degree of undersampling in the scan.

Mathematically, the GRAPPA reconstruction can be defined by the following expression:

$$S_j(k_y - m\Delta k_y) = \sum_{c=1}^{C}\sum_{b=0}^{N_b-1} n(j,b,l,m)S_l(k_y - bA\Delta k_y), \tag{3.1}$$

where S_j is a missing line at coil j, n(j,b,l,m) the weights obtained after the calibration phase, A the acceleration factor and Nb the number of blocks. The GRAPPA reconstruction results in one image per channel; these are typically combined into a single image by taking the sum of squares across the channels or performing a weighted sum using the coil sensitivities.

3.2 Model-Based Reconstruction

Unlike methods such as GRAPPA, which first recover k-space locations and then compute the image using the Inverse Fourier Transform, model-based reconstruction involves an iterative optimization process to directly recover the parameter maps or the underlying images [34]. In these methods, the reconstruction is obtained directly by solving an optimization problem. This problem minimizes a cost function, typically using the square of the L_2 norm of the residual between the acquired k-space and the simulated k-space from applying the forward model to the optimized image. The final result is a parametric map or image that directly reflects the desired solution. The iterative reconstruction starting from an undersampled k-space is essentially an ill-posed inverse problem.Given a forward model which maps parameters of interest to measurable signals, an inverse problem refers to the process of estimating parameters of interest based on measured signals. In the context of MRI, this involves reconstructing an image from the observed data, which corresponds to k-space samples, with the forward model being the application of the Fourier transform and masking in k-space. The difficulty of an inverse problem depends on how well-posed it is. For an inverse problem to be well-posed, three conditions must be satisfied: the solution must exist, it must be unique, and it must be stable, with respect to input data [35]. However, when k-space undersampling is introduced to accelerate acquisition, the problem becomes increasingly ill-posed, deviating further from these conditions. This highlights the importance of accurately defining the forward model, which represents the mapping from image space to k-space, to achieve a high-quality reconstruction despite the challenges posed by undersampling.

In the literature there are different examples for the application of iterative reconstruction problems, many of them specialized in a particular sequence.

3.2.1 MARTINI

One example of a model based reconstruction in T_2 mapping is the MARTINI (Model Based Iterative Accelerated RelaxomeTry by Iterative Non-Linear Inversion) [36]. A classical reconstruction problem can be implemented starting from the definition of the underlying inverse problem. Let y be the acquired k-space data and x the underlying image. It is possible to define an operator, E, called *forward model operator*, with the function of mapping x from the image space to y in k-space.

$$y = Ex + n, \tag{3.2}$$

where n is the noise typically present in the acquired data.

In a conventional image reconstruction, the forward operator E consists of multiplication by the coils sensitivities C (if the acquisition includes multiple receiver coils), application of the Fourier Transform F, and potentially multiplication by an undersampling pattern P, when not all the k-space locations are acquired.

$$y_{i,j} = PFC_j x_i + n \qquad i = 1....T, j = 1.....C \tag{3.3}$$

, where T is the total number of echo times and C the number of coils used in the acquisition.

In a parametric mapping method, another step is included in the forward operator E, typically including a model of the signal evolution considering the specific relaxation time that is weighting the acquired data, e.g. the model M included in **Equation 2.17**.

$$y_{i,j} = PFC_j M_0 e^{-\frac{TE_i}{T_2}} + n, \qquad i = 1....T, j = 1.....C \tag{3.4}$$

The problem of recovering M_0 and T_2 from undersampled k-space data, is solved using an iterative optimization algorithm using the following cost function:

$$\Phi(M_0, T_2) = argmin_{M_0,T_2} \frac{1}{2} \sum_{j=1}^{C} \sum_{i=1}^{T} \|PFC_j M_0 e^{-\frac{TE_i}{T_2}} - y_{i,j}\|_2^2 \tag{3.5}$$

In the MARTINI method, the inverse problem is solved using a non-linear optimization approach, specifically a Non-Linear Conjugate Gradient algorithm with the Hager-Zhang line-search method [37] (see **Figure 3.4a** for the pseudocode). The MARTINI k-space undersampling pattern is different from those shown previously, known as block sampling. A block consists of a collection of adjacent phase encoding (PE) lines, with its width defined as the total number of PE lines divided by the acceleration factor (AF). In a T_2 mapping sequence, multiple k-spaces are acquired across the echo dimension. Consequently, the block in the MARTINI method shifts along the PE dimension with each echo, as illustrated in **Figure 3.5**.

3.3 GRAPPATINI

The GRAPPATINI method, the core T_2 mapping approach in this project, is a natural extension of the MARTINI and GRAPPA methods. GRAPPATINI combines MARTINI's iterative model-based reconstruction with GRAPPA's concepts, leveraging their orthogonality to achieve higher acceleration factors, shortening acquisition times, and potentially achieving higher resolutions. The overall undersampling pattern uses a block-sampled MARTINI pattern but with interleaved sampling within each block, as is done in GRAPPA. The overall acceleration factor is the product of the individual factors from both methods (see **Figure 3.5**). The reconstruction proceeds in two steps: first, a GRAPPA reconstruction fills the missing PE lines within each block using coefficients derived from a small ACS region in the first echo. Subsequently, the resulting k-space is treated as a standard MARTINI set, enabling iterative model-based reconstruction using the same optimization algorithm as the one present in MARTINI. This approach requires only the cost function and gradient formulations to operate effectively.

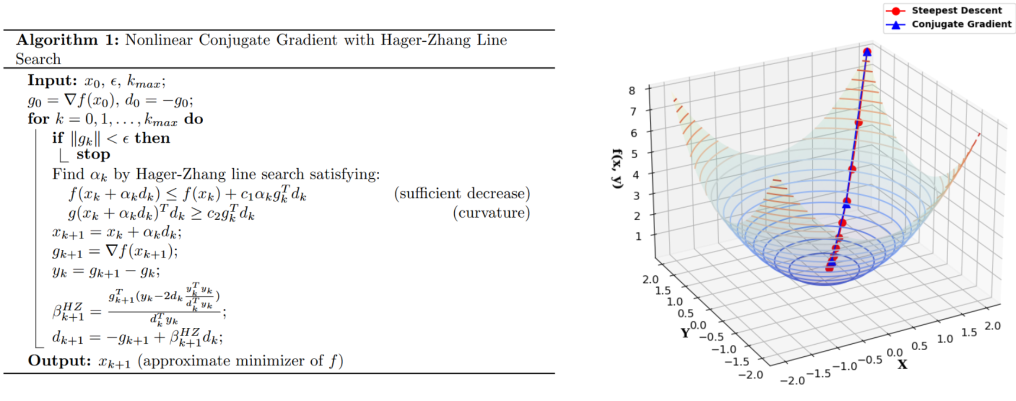

```
Algorithm 1: Nonlinear Conjugate Gradient with Hager-Zhang Line Search
Input: $x_0$, $\epsilon$, $k_{max}$;
$g_0 = \nabla f(x_0)$, $d_0 = -g_0$;
for $k = 0, 1, \ldots, k_{max}$ do
    if $\|g_k\| < \epsilon$ then
        stop
    Find $\alpha_k$ by Hager-Zhang line search satisfying:
        $f(x_k + \alpha_k d_k) \leq f(x_k) + c_1 \alpha_k g_k^T d_k$        (sufficient decrease)
        $g(x_k + \alpha_k d_k)^T d_k \geq c_2 g_k^T d_k$                    (curvature)
    $x_{k+1} = x_k + \alpha_k d_k$;
    $g_{k+1} = \nabla f(x_{k+1})$;
    $y_k = g_{k+1} - g_k$;
    $\beta_{k+1}^{HZ} = \frac{g_{k+1}^T (y_k - 2d_k \frac{y_k^T y_k}{d_k^T y_k})}{d_k^T y_k}$;
    $d_{k+1} = -g_{k+1} + \beta_{k+1}^{HZ} d_k$;
Output: $x_{k+1}$ (approximate minimizer of $f$)
```

(a)

(b)

Figure 3.4: (a) Pseudocode for the non linear conjugate gradient used in the MARTINI method. (b) Diagram showing the advantages in terms of convergence of a conjugate gradient method with respect to a classical steepest descent.

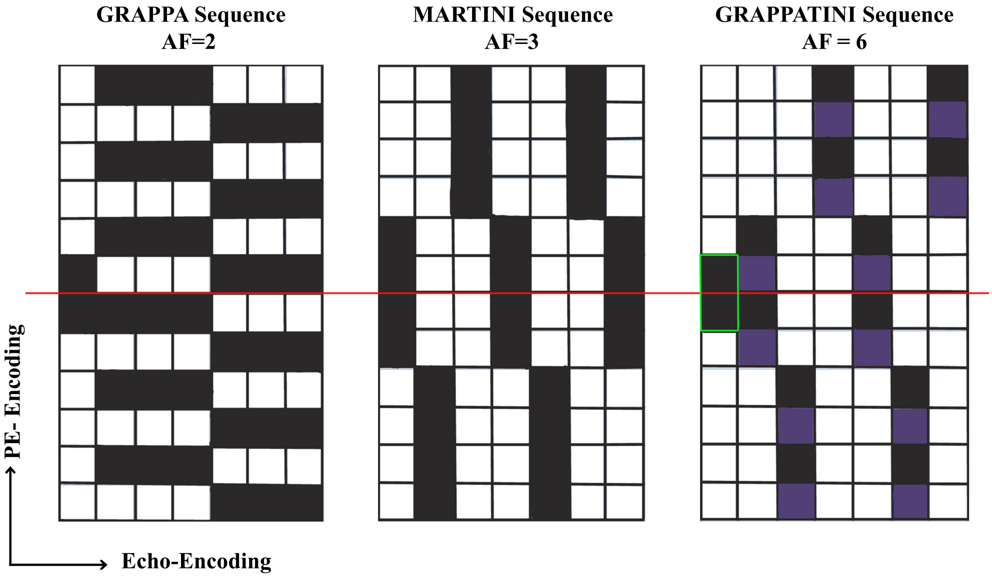

Figure 3.5: Schematic representation of GRAPPA, MARTINI and GRAPPATINI sampling patterns. Every column is a different echo, while every line is a different phase encoding step. The gain in acceleration factor of the GRAPPATINI method is highlighted. The red line indicates the k-space center, while the green box indicates the small ACS region acquired during the first echo in the GRAPPATINI sequence.

The GRAPPATINI method has the capability to accelerate acquisition times while maintaining at the same time acceptable image quality in comparison to the standard CPMG. The reduction in time, with respect to the other state-of-the-art method, is highlighted by the following table and in **Figure 3.6**:

Parameter		**Brain**	
	a. CPMG	**b. GRAPPATINI**	**c. Fast SE**
TA min	17:23	3:46	2:37
Acc.	Fully sampled	10-fold	Factor 9, GRAPPA 2
TR ms	6520	4000	4000
TE ms	9	10.9	77
Echoes	17	16	1
Matrix	160 x 192	512 x 270	512 x 270
Resol. mm^3	1.1 x 1.0 x 1.3	0.4 x 0.8 x 3	0.4 x 0.8 x 3
Slices	30	43	43

Table 3.1: Comparison of three T2 mapping brain scanning methods (CPMG (a) , GRAPPATINI (b), and Fast SE(c)) highlighting differences in acquisition time, resolution, and number of slices acquired. The GRAPPATINI method significantly reduces acquisition time while providing higher resolution and more slices than the CPMG fully sampled method.

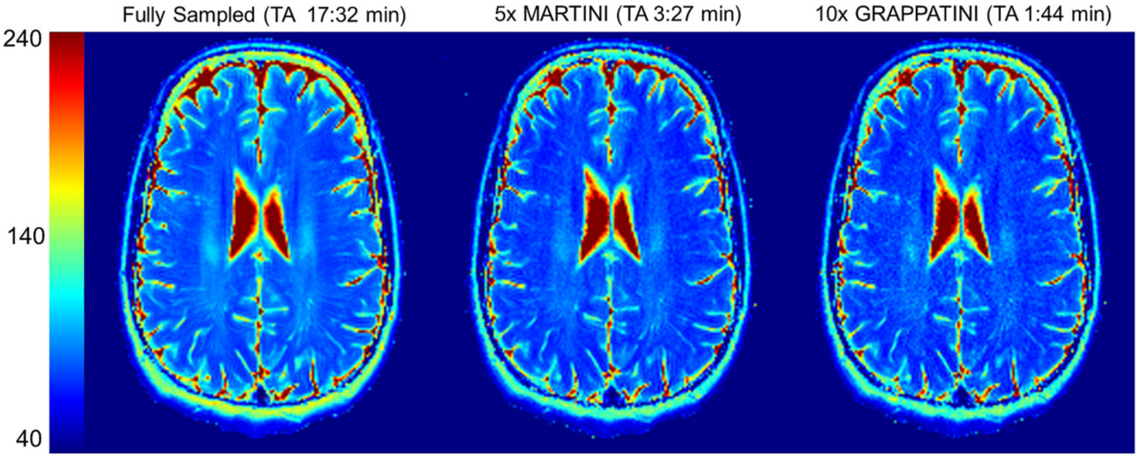

Figure 3.6: Comparison of T2 maps computed with the GRAPPATINI method (AF = 10), MARTINI (AF = 5) and a fully sampled acquisition. The different time of acquisition are also highlighted. The image is reproduced entirely from the paper [38].

3.4 Machine Learning and Deep Learning in MRI

Machine Learning (ML) and Deep Learning (DL) have revolutionized medical imaging by enabling data-driven methods that learn complex features from large-scale datasets.

Among DL approaches, Convolutional Neural Networks (CNNs) have established themselves as a milestone architecture over the last decade, considering their ability to automatically learn hierarchical feature representations from raw input data, capturing both local patterns and more global contextual information. In the MRI domain, CNN-based models have demonstrated remarkable success in a variety of tasks, such as image reconstruction starting from undersampled k-space data to, in the end, accelerate acquisitions[39], denoising noisy scans to improve the signal-to-noise ratio (SNR) [40], and segmenting anatomical structures with high accuracy [41]. Nevertheless, many existing approaches continue to rely on a fully supervised training paradigm, where the model learns to map, for instance, noisy (or undersampled) inputs to high-quality ground-truth data. In the context of T_2 mapping, as previously discussed in **Section 2.3.1**, acquiring fully sampled k-space data that would yield high-SNR T_2 maps is prohibitively time-consuming. This challenge is amplified when trying to build large datasets required under a classic supervised approach, as lengthy acquisition sessions not only require considerable resources but also pose limitations to patient comfort. Consequently, self-supervised learning has emerged as an appealing alternative and, in the case of T_2 mapping, is often the only feasible option because it can remove the need for clean or fully sampled ground-truth data. As a consequence, several methods have already been proposed to tackle various MRI-related tasks using self-supervised paradigms [42, 43, 44]. Under a self-supervised framework, a model learns directly from, for instance, noisy or undersampled measurements, by leveraging inherent data redundancies and internal data structures.

The previous chapters provided an overview of MR physics, MRI acquisition procedures, and the fundamental concepts behind quantitative T_2 mapping, illustrating why denoising is crucial for accelerating T_2 mapping techniques such as GRAPPATINI. The following section introduces the specific self-supervised strategies developed and evaluated in this work, with the aim of effectively denoising GRAPPATINI T_2 maps.

Chapter 4

Materials and Methods

In this chapter, the methods developed and tested in this project for training a neural network with a self-supervised approach to denoise T_2 maps acquired and reconstructed using the GRAPPATINI method are presented. First, the datasets used throughout this study are introduced, and their key characteristics are highlighted. Subsequently, the self-supervised framework, the core of both strategies, is described. Since it was applied in different domains, each of the implemented strategies and the corresponding methods are then explained. Finally, the statistical tools and quantitative metrics employed to evaluate the performance of these strategies are outlined.

4.1 Datasets

4.1.1 Fully Sampled 3T Dataset

A fully sampled MRI dataset was acquired at 3T (MAGNETOM Skyra, Siemens Healthineers AG, Erlangen, Germany). In this work, the term 'fully sampled dataset' refers to acquisitions in which all k-space locations necessary for achieving the desired resolution and field of view were collected at the scanner, thus leading to longer acquisition times. The dataset was obtained using a standard Carr-Purcell-Meiboom-Gill (CPMG) sequence, with the acquisition parameters summarized in **Table 4.1**. This dataset was primarily used during the implementation phases of the methods and served as a reference for the final refinement step of one of the two denoising strategies evaluated in this study.

Parameter	**Value**
TA (Total Acquisition Time)	17:23 min
TR (Repetition Time)	6520 ms
TE (Echo Time)	9 ms
Number of Echoes	17
Resolution (mm^3)	1.1 x 1.1 x 3
Slices	30

Table 4.1: Acquisition parameters of the fully sampled 3T Dataset used in this work.

4.1.2 3T LPM Dataset

The principal dataset used in this project was acquired at 3T (MAGNETOM Skyra, Siemens Healthineers AG, Erlangen, Germany) and is referred to as the Lausanne Parametric Mapping (LPM) dataset. The scans were performed using a commercial 20−channel head/neck coil. All experimental sessions were approved by the local Hospital Review Board, and written informed consent was obtained from each subject. All participants were healthy controls, resulting in a total of 52 subjects (28 females, 24 males; mean age ± standard deviation: 28.58±5.94, age range: 21−48 years). The sequence parameters used to acquire the data are those typically employed in a GRAPPATINI sequence (see **Table 4.2**). To enable a more comprehensive analysis of the denoising methods studied and proposed in this work, each subject underwent a second scanning session immediately after the first. Consequently, both scan and rescan data were available, allowing the assessment of the reproducibility of T_2 values in terms of intra−subject and inter−subject variability following the application of denoising methods. The demographics of the study population are summarized in **Table 4.3**.

Parameter	**Value**
TA (Total Acquisition Time)	3:22 min
Acceleration (AF)	10
TR (Repetition Time)	4000 ms
TE (Echo Time)	10 ms
Number of Echoes	16
Resolution (mm^3)	0.7 x 0.7 x 3.3
Slices	44

Table 4.2: Acquisition parameteres for the 3T LPM dataset used in the training processes of the models throughout the project.

Subjects	**Gender**		**Age (years)**	
	M	**F**	**mean ± SD**	**range**
Healthy Controls	24	28	28.59 ± 5.84	[21 - 48]

Table 4.3: Demographics of the population in the 3T LPM Dataset.

For the two strategies implemented in this work and presented in **Section 4.3** and **Section 4.4** a partition of the dataset into two independent sets was performed such that 32 subjects out of the 52 available were included in a training set, while the remaining subjects were included in a test set. The criteria selected for the division is a *stratified sampling*. Although not really necessary considering the age distribution of the subjects in the dataset, differences in ages could potentially be a bias for the training of a model in the T_2 mapping field. With a stratified sampling strategy, at least one subject was included in the training set for each bin created from the histogram of the ages.

To assess the model's performance and generalizability during the training phase, a

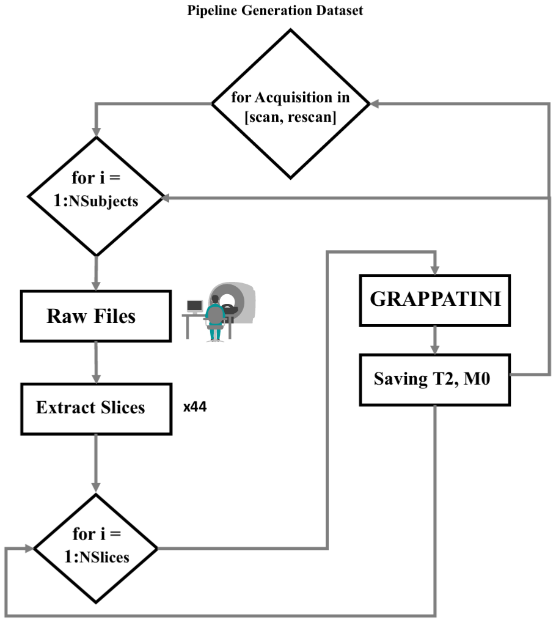

Figure 4.1: Flowchart illustrating the generation process of the 3T LPM T_2 map dataset: starting with acquisitions obtained from the scanner, the data is processed to reconstruct T_2 maps for each subject, slice, and session, i.e. scan and rescan.

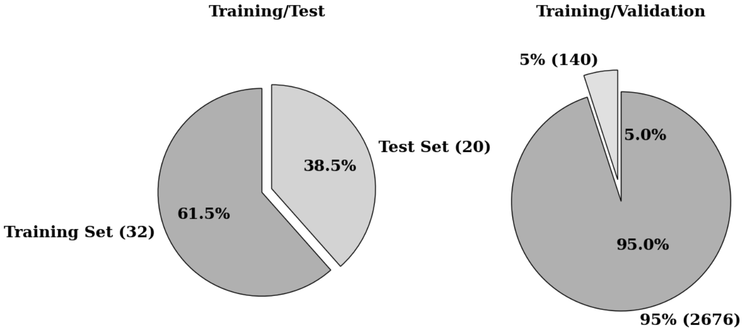

Figure 4.2: A pie chart illustrating the division of subjects into training and validation sets based on the number of subjects. The training set includes 32 subjects, while the test set consists of 20 subjects. Additionally, the chart shows the division of the training set into training and validation subsets based on slices, where 95% of the total slices are used for training and 5% for validation.

second and random partition was applied to create a validation set. The partition ratio was set to 0.95, leaving 5% of the training samples for validation. During the final testing phase, the scan-rescan variability was evaluated. To avoid biasing the model's learning process, only the scan and rescan data from subjects included in the training set were used during training. However, at inference time, the model's performance was evaluated under both scan and rescan conditions. In summary, the training dataset consisted of 32 subjects, each with scan and rescan acquisitions, and each acquisition included 44 slices. This results in a total of $44 \times 2 \times 32 = 2816$ samples. Of these, 140 samples (5% of the total) were set aside in the validation set (see **Figure 4.2**).

4.1.3 7T Datasets

To evaluate the generalizability of the model, four separate datasets were acquired at 7T. Each dataset differed from those used during the training phase in terms of field strength, resolution, orientation, anatomy, and consequently the range of T_2 values. Each dataset consisted of a single acquisition from one subject, providing diverse testing scenarios to assess the performance of the model on data that substantially deviates from the training conditions.

Parameter	**Brain Low Res.**	**Brain High Res.**	**Sagitt. Knee**
TA (Time)	2:55 min	6:11 min	7:50 min
AF	10	10	6
TR (ms)	4680	4680	4880
TE (ms)	10.7	11.40	10.9
Echoes	16	16	10
Res. (mm^3)	0.5 x 0.5 x 2.0	0.3 x 0.3 x 1.5	0.4 x 0.4 x 3.0
Slices	15	15	24

Table 4.4: Acquisition parameters of the 7T datasets acquired and used in the project to probe the generalizability of the model.

All scans were acquired at 7T (MAGNETOM Terra.X, Siemens Healthineers, Forchheim, Germany). All experimental sessions were approved by the local Hospital Review Board, and both informed and written consent were obtained from each subject. The acquisition parameters are summarized in **Table 4.4**. Three brain datasets were acquired at two different resolutions, resulting in T_2 maps with higher in-plane resolution in the coronal plane. For the fourth dataset, a knee scan was performed to obtain sagittal T_2 maps.

4.2 Self-Supervised Blind Denoising

The concepts, notation, and background of the self-supervised method used throughout the project are thoroughly presented in the article [45]. This method belongs to the family of approaches known as *Blind Denoising* [46], as it does not rely on prior assumptions about the noise model, the signal itself, or the availability of clean training data. The need for a self-supervised approach arises due to the lengthy acquisition times required to generate a fully sampled dataset of T_2 maps, as highlighted in **Section 3**, which could potentially be used as a reference in supervised learning for model training. In general, it is possible to model a signal y, high-dimensional with a total of m dimensions, and the corresponding noisy measurements as two random variables in $\mathbb{R}^m$, as:

$$x = y + n \tag{4.1}$$

, where n is the additive noise present in the measurement.

A partition of the set containing all dimensions of y can be defined and denoted by J, where the values of the noisy measurement x restricted to this partition are denoted by x_J. Furthermore, let $\mathcal{J}$ be the subspace of all possible partitions of the space formed by the dimensions of y, such that $J \subseteq \mathcal{J}$. A function $f : \mathbb{R}^m \to \mathbb{R}^m$ is called J-invariant if the value of f, restricted to x within the partition J, does not depend on the specific value of x_J. Thus, if f is J-invariant across all possible partitions of the space of the dimensions of y, it can also be considered $\mathcal{J}$-invariant. Conversely, if a function is $\mathcal{J}$-invariant, then it is automatically J-invariant for all partitions $J \subseteq \mathcal{J}$.

The method then defines the self-supervised loss function, considering $\mathcal{J}$-invariant functions as follows:

$$\mathcal{L}(f) = \mathbb{E}\|f(x) - x\|^2 \tag{4.2}$$

, where the expectation is taken over the partitions. It is possible to demonstrate, considering the noise in a subset J independent from the noise in the complementary set J_c, that this loss function is equal to:

$$\mathbb{E}\|f(x) - x\|^2 = \mathbb{E}\|f(x) - y\|^2 + \mathbb{E}\|x - y\|^2 \tag{4.3}$$

As shown in **Equation 4.3**, the self-supervised loss is equal to the sum of the usual supervised loss and the noise variance in the initial measurement. Consequently, minimizing the self-supervised loss in expectation leads to minimizing the error with respect to the ground truth, noise free signal. This self-supervised loss can be employed as the loss function for e.g. training a deep convolutional network to denoise a high−dimensional measurement x. In order for a neural network to be considered $\mathcal{J}$-invariant, an additional operation must be applied to its input. Specifically, a masking procedure can ensure that the network output remains independent of the original noisy values at specific points within a partition of the input dimensions.

In this project, the high-dimensional and noisy measurements are GRAPPATINI k-spaces and GRAPPATINI T_2 maps. The described approach of training a model in a self-supervised manner by masking the input to ensure the model behaves as a $\mathcal{J}$-invariant function was implemented in the two primary domains where the information is stored and represented in MRI: the k-space and the image space.

4.3 K-Space Strategy (SSDU)

The first strategy implemented in this work is based on the concepts introduced in **Section 4.2**, but adapting them to consider the k-spaces, acquired with a standard GRAPPATINI sequence, as the starting high-dimensional and noisy measurement. The objective was to train a model using a self-supervised loss, as defined in **Equation 4.2**, with the specific aim of formulating this loss directly in the k-space domain. To achieve this, as a first step, the network needed to be transformed into a $\mathcal{J}$-invariant function, which

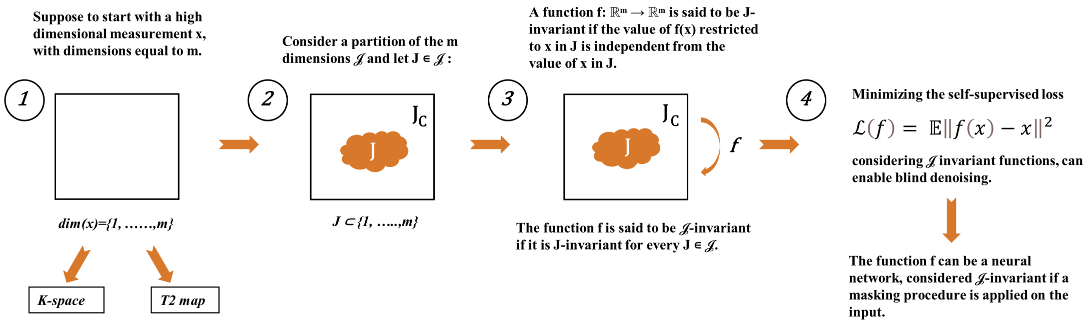

Figure 4.3: Diagram summarizing the pipeline of the self-supervised method implemented throughout the project. The method involves training a model with a self-supervised approach by modifying the input to ensure the model behaves as a $\mathcal{J}$-invariant function.

was accomplished through a masking operation applied directly to the k-space locations acquired during the GRAPPATINI sequence. In this way, the output of the network, was no longer dependent on the original values in those k-space locations, and the weights of the model could be trained by looking at the discrepancy between those originally acquired k-space locations and the one "simulated" (using the forward MRI model) from the output of the model itself.

To provide a clearer understanding of the nomenclature and concepts used in this method, the definition of the inverse problem underlying the T_2 mapping is here reported again. The theory and notation presented in this section are largely based on the work of [47].

Let y_Ω represent the acquired locations in the k-space during a standard GRAPPATINI acquisition. The forward model of the inverse problem is then given by:

$$y_\Omega = E_\Omega x + n \tag{4.4}$$

For T_2 mapping reconstruction, the forward model operator E_Ω is made up of several components: the relaxation model M, which is a function of parameters M_0 and T_2; multiplication by the coils sensitivities C; application of the Fourier transform F; and the sampling pattern P_Ω, which selects the acquired locations of the k-space Ω. The recovery of the maps is then formulated in the context of the GRAPPATINI reconstruction as:

$$\underset{M_0,T_2}{\operatorname{argmin}} \|P_\Omega F C M(M_0, T_2) - y_\Omega\|_2^2 \tag{4.5}$$

In the self-supervised strategy implemented in this work, the first step was to consider the acquired k-space locations Ω and divide them into two distinct sets, Θ and Λ. The division was made such that:

$$\begin{cases} \Omega = \Theta \cup \Lambda, \\ \Lambda \cap \Theta = \emptyset \end{cases} \tag{4.6}$$

Equation 4.6 demonstrates that the two subsets, Θ and Λ, are disjoint, i.e. there are no k-space locations shared between the two sets. Furthermore, their union is equal to the original set of acquired k-space locations, ensuring that all locations are taken into account in the strategy.

The set Θ was used to perform the GRAPPATINI reconstruction, yielding a pair of M_0 and T_2 maps for each scan and subject. In contrast, the set Λ of k-space locations was used to define the loss function for the model, which was trained with the ultimate goal of denoising the T_2 maps.

The general form of the loss function can be expressed by the following equation:

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^{N} \mathcal{L}\left(y_{\Lambda}^{i}, E_{\Lambda}^{i}\left(f\left(y_{\Theta}^{i}, E_{\Theta}^{i}; \theta\right)\right)\right) \tag{4.7}$$

Breaking down this expression, the network output $f\left(y_{\Theta}^{i}, E_{\Theta}^{i}; \theta\right)$ is a function of the k-space locations in the set Θ and the network parameters θ. This output is transformed back to the k-space by the operator E_{Λ}^{i}, which selects only the k-space locations specified by the indices in the set Λ. The key distinction between the operators E_{Θ}^{i} and E_{Λ}^{i} lies in the definition of their sampling patterns, P_{Θ} and P_{Λ}, respectively.

The loss is then computed as the discrepancy between the original k-space values at the locations in Λ (y_{Λ}) and the transformed network output, which has been mapped back to k-space and evaluated at the same locations using E_{Λ}^{i}. This discrepancy is quantified by the loss function $\mathcal{L}$ and averaged over all N training samples.

The division of the acquired k-space locations Ω into the disjoint sets Θ and Λ is a critical step to ensure that the function f, in this case a neural network, can be considered J-invariant. By using only a subset of k-space locations (Θ) to compute the M_0 and T_2 maps, this approach effectively performs a masking procedure on the original k-space data. The model output is then evaluated at the remaining locations (Λ), ensuring that the network does not directly rely on information from the k-space values it aims to predict.

The originally acquired k-space locations, which serve as the starting point, correspond to standard GRAPPATINI datasets. For each echo in the sequence and each coil in the system, a distinct k-space is acquired. The data are inherently undersampled, as is typical in the GRAPPATINI sequence. Excluding the first echo, starting from the second, a block of k-space lines is sampled and shifted along the phase encoding (PE) direction in subsequent k-spaces at different echo times. Additionally, within each block, interleaved undersampling is applied, enabling GRAPPA reconstruction to recover the missing lines.

The division of the acquired k-space locations into two distinct sets, Θ and Λ, was performed after the GRAPPA reconstruction. For a given dataset, an example of the masks used to select the Θ and Λ locations is shown in **Figure 4.4**.

The masks used are binary, selecting specific PE lines to exclude during the reconstruction of the T_2 maps. The selection of the lines was performed randomly. The ratio ρ between the k-space locations removed from the reconstruction process and the total acquired locations was set as:

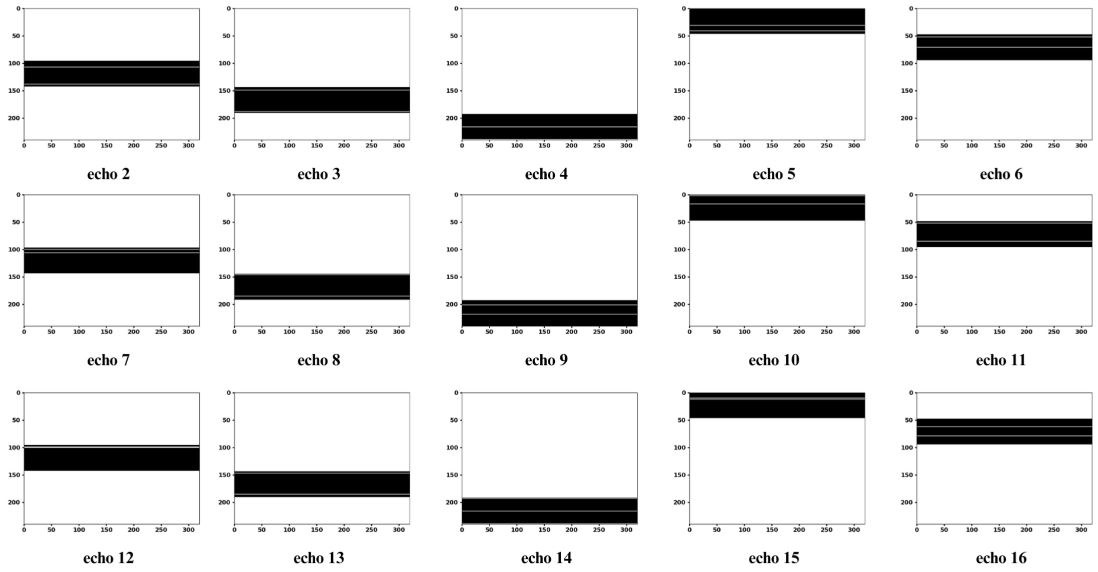

echo 2 echo 3 echo 4 echo 5 echo 6

echo 7 echo 8 echo 9 echo 10 echo 11

echo 12 echo 13 echo 14 echo 15 echo 16

Figure 4.4: Set of masks along the echo dimension for a given GRAPPATINI dataset. Black lines represent sampled lines, while white lines indicate lines that are not sampled and included in the reconstruction process. Within each block, a certain number of phase encoding (PE) lines is removed, forming the set of k-space locations Λ.

$$\rho = \frac{\|\Lambda\|}{\|\Omega\|} = 0.05 \tag{4.8}$$

, where $\|\cdot\|$ is the cardinality of the set of k-space locations. Moreover, in each echo k-space, a specific number of lines were consistently kept sampled if the center of the k-space had already been acquired according to the GRAPPATINI sequence and the acceleration factor (AF). The number of lines always sampled and excluded from the random masking procedure, in the center of the k-space, was fixed at 24.

The set of masks was randomly varied across subjects and slices within the entire volume. However, the masks remained consistent across the coil dimension, ensuring that, for a given slice, all k-spaces along the coil dimension were masked identically for each echo. **Figure 4.5** provides a graphical summary of the acquisition process, the selection of random lines, and the division of the acquired k-space locations into the two disjoint sets, Θ and Λ. The k-space is represented in a 4D space, highlighting the different dimensions of the original data: Phase Encoding (PE), Readout (RO), echo, and coil.

4.3.1 Training

The initial goal of the training process using this self-supervised strategy was to deliberately overfit the proposed method to a single slice from the 3T LPM dataset. This

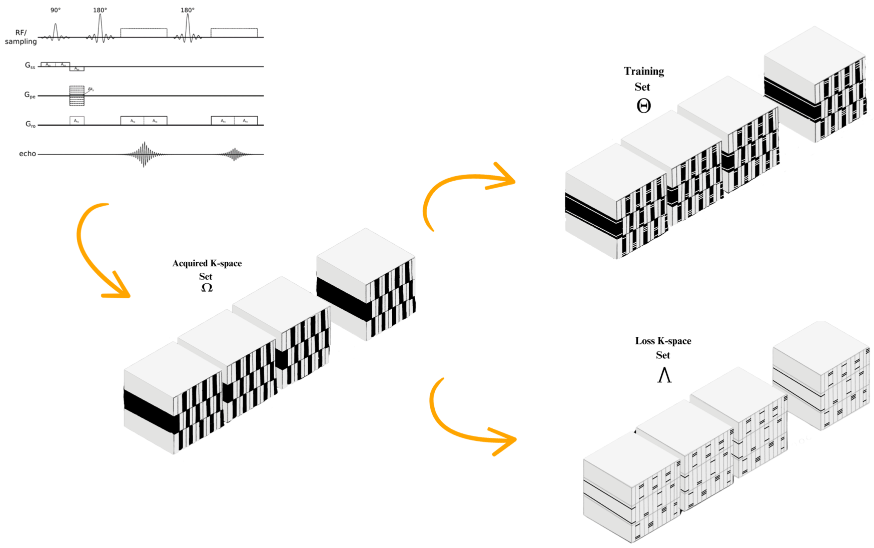

Figure 4.5: A graphical representation of the division of the acquired k-space locations, Ω, into two disjoint sets, Θ and Λ, is provided. The k-space is shown in a 4D space, where the different dimensions of the original data—Phase Encoding (PE), Readout (RO), echo, and coil—are highlighted.

was implemented by applying 10 distinct random masking patterns (as shown in **Figure 4.4**) to the k-space data of that slice. As detailed in the previous section, random masking was implemented such that, for each pattern, the masks were consistent along the coil dimension but varied across the echo dimension. This approach exposed the model to multiple masked versions of the same input, potentially accelerating the overfitting process. Overfitting is a critical initial step in the development of a machine learning model. If overfitting is not achieved on a single slice, it may indicate potential limitations in the model's architecture or its ability to generalize effectively across a larger dataset. The first step in this process involved reconstructing T_2 and M_0 maps from the k-space data of a single slice using the different random masking patterns. The reconstruction was limited to the k-space locations defined by the set Θ, which varied for each random pattern. As a result, each reconstruction produced a unique set of T_2 and M_0 maps, all derived from the same slice but using different randomly masked k-space data. These maps were then used as input to train the model.

Key decisions were made regarding the model training process. Instead of using both the T_2 and M_0 maps as input, only the R_2 map was used. The R_2 map is defined as the inverse of the T_2 map, i.e., $R_2 = \frac{1}{T_2}$. This choice was made to enhance the stability of pixel values during training and simplify the task by focusing the model on denoising a single map at a time, rather than addressing both maps simultaneously. Before feeding the R_2 map into the model, a standardization process was applied. The mean value, μ, was subtracted from the R_2 map, and the result was divided by the standard deviation, σ, to normalize the data. The final output was obtained by adding the original normalized R_2 map to the normalized residual output of the model, weighted by a λ coefficient. This coefficient was used to control the influence of the model's output on the input to the loss function.

Next, a denormalization step was applied to restore the R_2 map to its original scale. This was done by reintroducing the mean and standard deviation values. After denormalization, the model's output remained in the parametric space. As part of the self-supervised strategy, the loss was computed based on the discrepancy between two elements: the k-space locations Λ (which were acquired but not used during the reconstruction) and the model's output. The model's output was transformed back into the k-space domain and evaluated at these same locations, using the forward model of the inverse problem. Specifically, the self-supervised k-space loss used in this work was defined as follows:

$$\mathcal{L} = \frac{1}{2}\left(\sum_{t=2}^{TE} \frac{\left\|E_{\Lambda,t}\left(f\left(y_{\Theta}, E_{\Theta}; \theta\right)\right) - y_{\Lambda,t}\right\|_2}{\left\|y_{\Lambda,t}\right\|_2} + \sum_{t=2}^{TE} \frac{\left\|E_{\Lambda,t}\left(f\left(y_{\Theta}, E_{\Theta}; \theta\right)\right) - y_{\Lambda,t}\right\|_1}{\left\|y_{\Lambda,t}\right\|_1}\right), \tag{4.9}$$

where TE represents the total number of echo times. The total loss is computed as the average of the normalized L_1 and L_2-norms, calculated from the residuals (discrepancies) between the Λ-locations of the model's output, transformed back into the k-space, and the originally acquired data. To correctly apply the forward model to the output of the model, the M_0 map was also required. The chosen approach was to use the same M_0 map reconstructed from the k-space locations in Θ.

The complete pipeline of the proposed strategy is summarized in the flow chart below (see **Figure 4.6**).

The model and training scripts were implemented in Python using the PyTorch framework [48]. In the initial attempt to overfit the model to a single slice of the 3T LPM dataset, training was conducted over 80 epochs, with different learning rates explored. The λ coefficient was varied within the range [0,1], while the ratio between the k-space locations Λ and Ω was maintained constant at 0.05. The selected model was a U-Net, a fully convolutional neural network originally designed originally for image segmentation tasks [49]. The architecture consists of three main components: the encoder, the bottleneck, and the decoder, all interconnected by skip connections. The encoder progressively reduces spatial resolution while increasing feature depth through convolutional blocks followed by average pooling. The bottleneck processes the features at the lowest resolution, and the decoder reconstructs the output by applying transpose convolutions and skip connections, gradually increasing the resolution. Each convolutional block included both instance normalization and LeakyReLU activation functions. The hyperparameters used during the training process are summarized in the table below (see **Table 4.5**).

4.4 Image Space Strategy (N2S)

In the second denoising strategy implemented in this work, in this work referred to as N2S, the perspective on the problem of denoising GRAPPATINI T_2 maps was shifted. Instead of considering the noisy k-space measurements as the initial high-dimensional data (as done in the k-space strategy), the focus was placed on the T_2 maps obtained

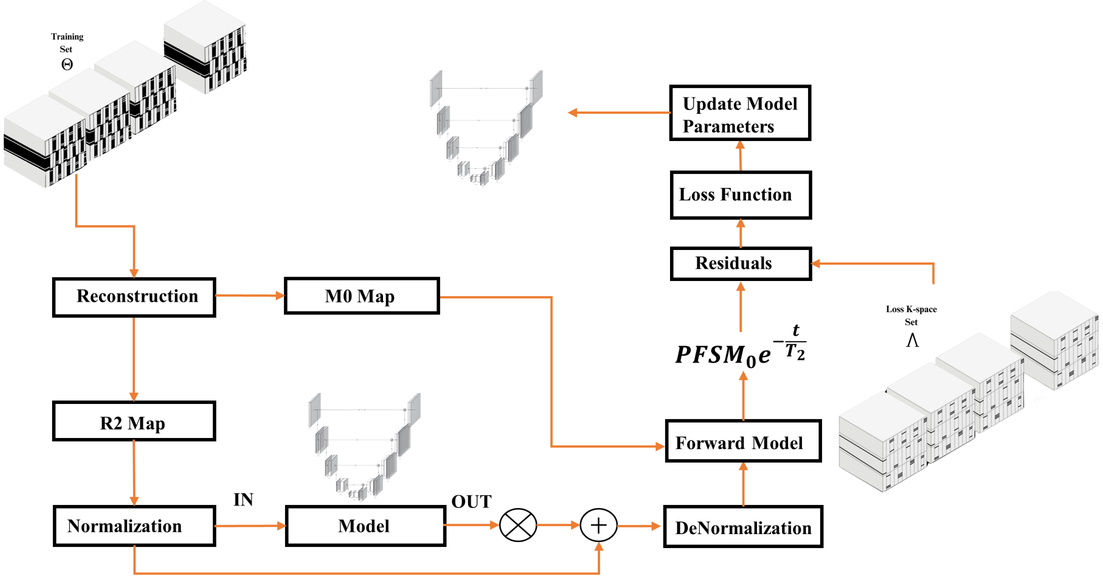

Figure 4.6: Flowchart for the k-space strategy implemented in this work.

Hyperparameter	Value
Epochs	80
Batch Size	1
Biases in Convolutions	False
LeakyReLU Negative Slope	0.2

Table 4.5: Hyperparameters selected during the training of the model in the k-space strategy.

after the reconstruction considering the 3T LPM dataset. In this case, the dimensions of the starting data could no longer be defined in the k-space domain but rather directly in the image-space, consisting in the pixels forming a single T_2 map.

As a result, the masking procedure, necessary to ensure that the neural network could be considered $\mathcal{J}$-invariant, was applied directly in the image space. This shift in approach also implied that the loss function was defined in a different domain compared to the k-space strategy. Specifically, the loss function was defined in the image space as the discrepancy between the original values of specific pixels selected in the input map and the predictions made by the neural network at those locations.

In each T_2 map used as input to the model, a grid of pixels was replaced by a weighted average of the surrounding neighborhood after a kernel was applied to each pixel in the 2D map (see **Equation 4.10**). The surrounding neighborhood was defined as the 4-connected set of pixels with respect to the pixel where the kernel was centered. The network could then be considered $\mathcal{J}$-invariant, as the weighted average was computed using only the surrounding pixel values rather than the original noisy value selected by

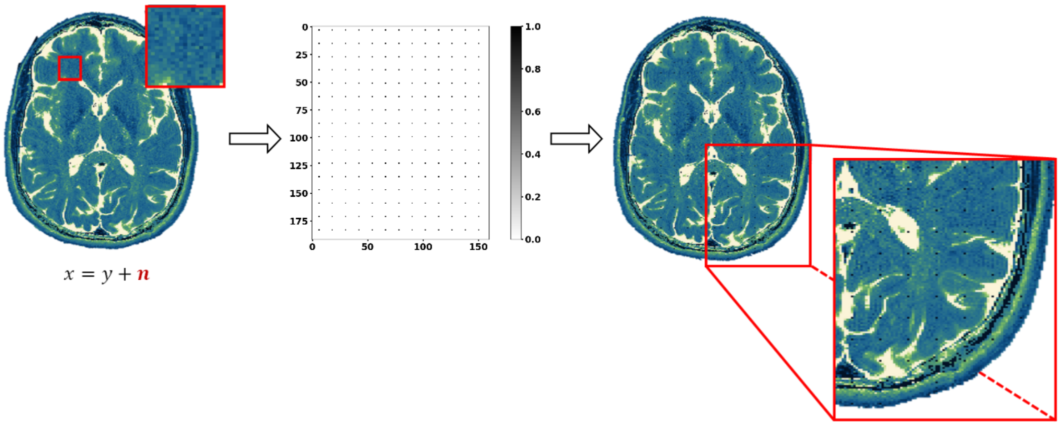

Figure 4.7: Flow chart masking procedure on input and noisy T_2 map.

the center of the kernel. The use of the average kernel provided an initial estimate of the pixel value for the masked positions, but the approach would still be valid if all values in these masked positions had been set to zero, as was done in the k-space strategy, where k-space locations were directly masked and set to zero following the application of a binary mask to the k-spaces acquired at different coils and echo times.

The masks used in this strategy were binary (see **Figure 4.7**), selecting a grid of pixels from each map provided as input to the model. The crucial hyperparameter defining the masks was the grid size, which was determined by the number of pixels in both directions of a Cartesian grid between one selected pixel and another.

$$kernel = \begin{bmatrix} 0 & 0.25 & 0 \\ 0.25 & 0 & 0.25 \\ 0 & 0.25 & 0 \end{bmatrix} \tag{4.10}$$

The model's loss function was then designed to minimize the discrepancy between the predictions made by the model at the initially masked pixel locations and their original values. The model selected for this strategy was a deep convolutional neural network specifically designed for image denoising (DnCNN) [50]. It consisted of a series of convolutional layers with ReLU activations and batch normalization. The network had an input layer that mapped the input image channels to 64 feature maps using a 3 × 3 convolutional filter. This was followed by multiple hidden layers, each employing 3 × 3 convolutions, batch normalization, and ReLU activation to progressively refine the learned features. The final layer was designed to map the 64 feature maps back to the original number of channels, ideally producing a denoised version of the input image (see **Figure 4.8**).

4.4.1 Training

The model and training scripts were implemented in Python using the PyTorch framework [48]. Training was conducted over 500 epochs, with a learning rate of 0.01. The chosen

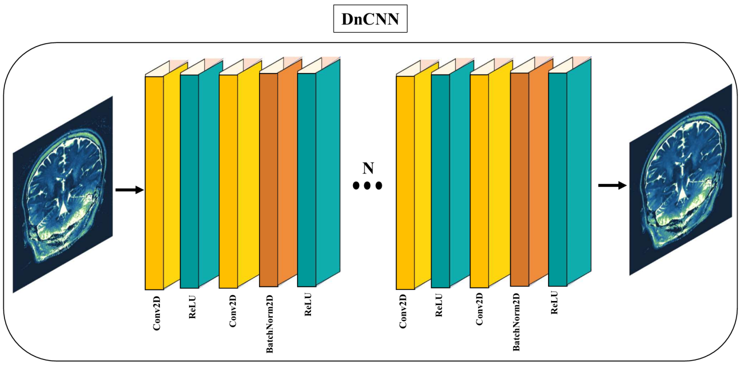

Figure 4.8: Schematic representation of the DnCNN model used in the Image Space strategy.

loss function was the L_1 loss, and the ADAM optimizer was employed. The batch size was fixed at 5, and, as mentioned in **Section 4.1**, the available data were split to create a validation set for tracking the model's performance during training. In each iteration of training, a different masking operation was applied to the entire batch, relative to the previous iteration. This operation depended on the current iteration number and the grid size of the mask, which introduced an offset in the grid of selected pixels. The mask was changed dynamically during each iteration. Specifically, the grid of pixels was shifted according to the iteration number and the grid size, creating a new selection of pixels at every training step. The shifting pattern was influenced by the current iteration, ensuring that the batch was masked differently compared to the prior iteration. This resulted in varied input configurations, which helped the model generalize better by learning from different masked perspectives throughout each training iteration.

As pre-processing steps, an additional masking operation and a normalization were applied to the T_2 maps. The masking operation consisted of a threshold-based selection using the values of the M_0 map, where the threshold was set to 10% of the mean M_0 value. This procedure was performed to isolate and discard points with near-zero signal (e.g., background regions). For normalization, a Min-Max scaling technique was employed to ensure all input values to the model fell within the range [0,1]. Specifically, the Min-Max scaling method was defined by:

$$y_{\text{norm}} = \frac{y - y_{\text{min}}}{y_{\text{max}} - y_{\text{min}}} \tag{4.11}$$

where y is the input T_2 map to the model.

The hyperparameters selected during the training process of the model are summarized in the table below (see **Table 4.6**).

The complete pipeline of the steps involved in the training process for the image space strategy is shown in the **Figure 4.9**.

Hyperparameter	**Value**
Epochs	500
Learning Rate	1e−2
Batch Size	5
Loss	L_1
Optimizer	Adam
Mask Grid Size	12
Num. hidden layers	10

Table 4.6: Hyperparameters selected during the training of the model in the image-space strategy.

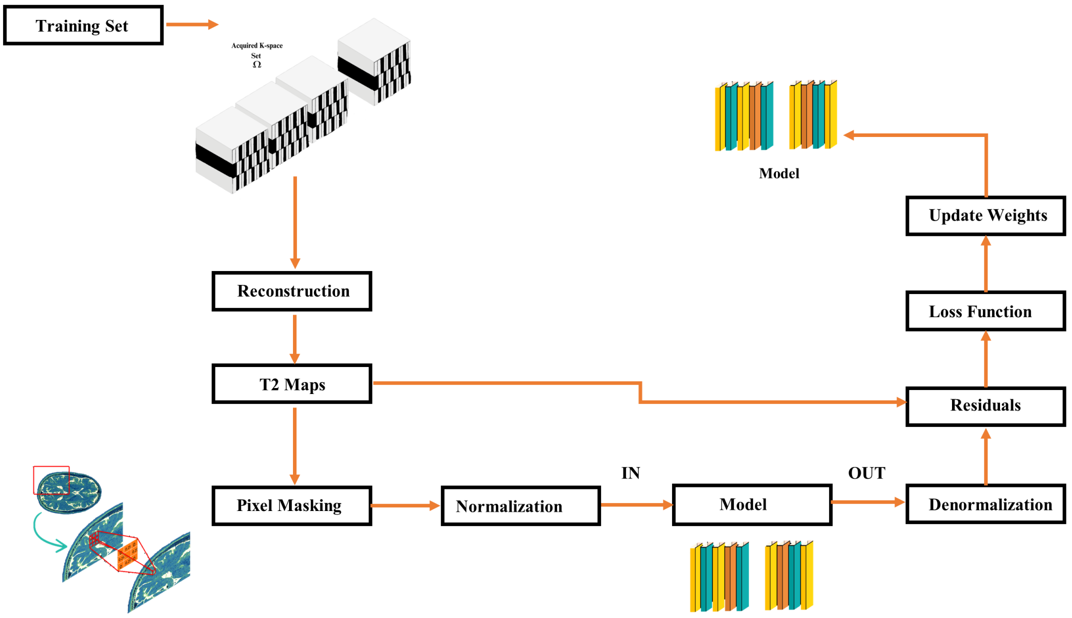

Figure 4.9: Pipeline of the training process in the Image Space strategy.

4.4.2 Inference

At inference, the performances of the model were evaluated on the entire T_2 map test set. The generation of the test set was performed as explained in **Section 4.1.2**. The complete pipeline for the inference process is shown in **Figure 4.10**.

For the final output of this strategy, a decision was made to consider the linear combination of the model's output and the original T_2 map provided as input. This approach was implemented to better control the blurring effect introduced by the denoising method. As highlighted in the original paper discussed in the previous section (see **Section 4.2**), this linear combination yielded the best results in terms of signal-to-noise ratio (SNR) improvement. To justify the choice of the interpolation coefficient, α, a fully sampled

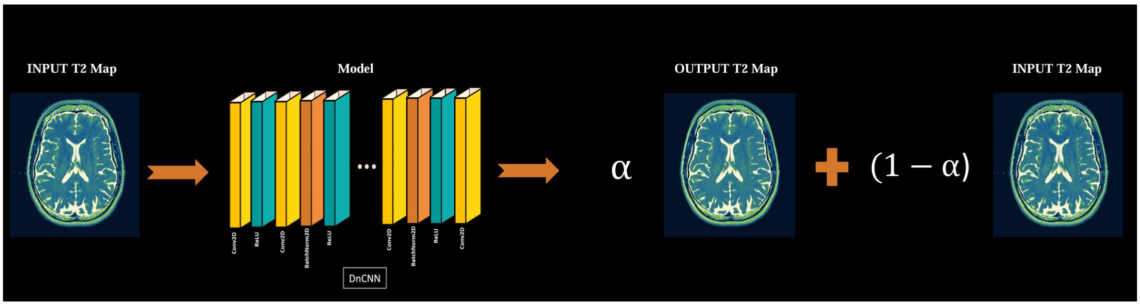

Figure 4.10: Schematic representation of the inference step in the Image Space strategy.

dataset was utilized for fine-tuning, specifically the 3T fully sampled dataset, detailed in **Section 4.1**. The tuning process involved performing a fully sampled reconstruction to serve as a reference, followed by the introduction of Gaussian noise ($\mu = 0$, $\sigma = 3 \times 10^{-5}$) to the fully sampled k-space data. Then, T_2 maps were reconstructed from the corrupted k-space data, and the mean squared error (MSE) was computed relative to the fully sampled reference, where MSE was defined by the following relation:

$$MSE = \frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2 \tag{4.12}$$

The interpolation coefficient α was varied between 0 and 1, with the final output computed as the weighted sum:

$$\text{Denoised } T_2 \text{ map} = \alpha \cdot \text{Model Output} + (1 - \alpha) \cdot \text{Input } T_2 \text{ Map}. \tag{4.13}$$

4.5 Statistical Analyses

4.5.1 Scan-Rescan Reproducibility

Scan-rescan reproducibility is a critical aspect when evaluating quantitative MRI measurements. Ensuring the robustness and accuracy of parameter estimation is essential for a potential clinical adoption of the technique. The reproducibility of T_2 values using the GRAPPATINI sequence has already been demonstrated in the study by Gruenebach et al. [51]. With the aim to perform the same evaluation, two different tools were used to ensure that the denoising methods did not alter the T_2 values originally estimated with the GRAPPATINI method. First, a definition of regions of interest (ROIs) was necessary. To accurately and automatically segment specific tissues across the entire volume probed by the GRAPPATINI sequence, an internal tool for brain segmentation, Morphobox [52], was employed. The segmentation tool was applied across all slices of both scan and rescan acquisitions for all subjects included in the test set. The selected ROIs for this study were:

- Frontal White Matter (FWM)

- Frontal Gray Matter (FGM)
- Thalamus (TH)
- Globus Pallidus (GP),

considering both the right and the left hemisphere.

The output of the segmentation tool consisted of binary masks for each of the regions across the entire volume, as demonstrated in the figure below for a single slice (see **Figure 4.11**). Once the ROIs were defined, statistical analysis were performed on the selected regions and the extracted T_2 values. As a pre-processing step, morphological erosion was applied to the binary masks using a 3×3 structural element (kernel) for a single iteration (see **Figure 4.11** for a visual representation).

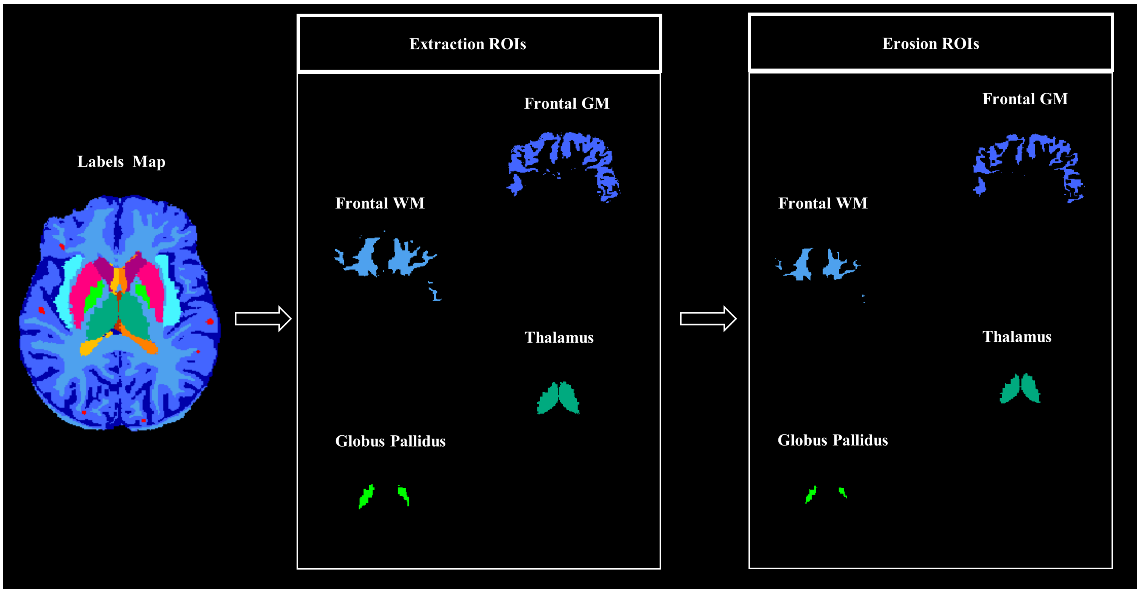

Figure 4.11: ROIs selected in the scan-rescan reproducibility analysis. The effect of the erosion process is highlighted.

This operation was performed to reduce the impact of partial volume effects when calculating the median T_2 values within the selected regions of interest (ROIs). Afterward, the median T_2 values were computed for each of the four different ROIs and to further minimize the influence of any remaining partial volume effects, outliers were removed from the distributions using the Median Absolute Deviation (MAD) method. Specifically, T_2 values with a modified Z-score exceeding a threshold τ were excluded:

$$\mathbf{T_2}_{\text{, filtered}} = \{T_{2,i} \in \mathbf{T_2} \mid Z_i = \frac{|T_{2,i} - \tilde{T_2}|}{\text{MAD}} \leq \tau\},$$

where $\tilde{T_2}$ is the median T_2 in one of the four ROIs, and MAD is the median absolute deviation between a T_2 value in the ROI and the median value of that ROI.

Consequently, a linear regression was conducted between the median T_2 values in the four different regions considering the scan and rescan sessions. Ideally, if the T_2 values were fully reproducible across sessions, the linear regression would yield a correlation coefficient of 1 and an intercept of 0. To assess the quality of the linear fit, the R^2 coefficient, known as the coefficient of determination or the squared value of Pearson's correlation coefficient, was also computed using the relation:

$$R^2 = \left(\frac{\mathrm{Cov}(X, Y)}{\sigma_X \sigma_Y} \right)^2 \tag{4.14}$$

Additionally, scan-rescan reproducibility was evaluated using a Bland-Altman analysis [53]. The Bland-Altman plot was computed for both scans and rescans within the same method, to assess any potential intra-method bias and to establish intra-method limits of agreement. Furthermore, the Bland-Altman analysis was also applied between the two methods, identifying any potential bias between them and highlighting the variability in the measurements. In the Bland-Altman analysis, the median T_2 values were computed for the four different ROIs selected for each subject in the test set. The difference between the median values of scan and rescan (when assessing intra-method reproducibility) and the difference between the median values in the two methods (when assessing inter-method reproducibility) was plotted on the y-axis, while the mean of the medians was plotted on the x-axis. To calculate the limits of agreement, a factor of 1.96 times the standard deviation of all the differences in median values across the four regions was applied.

4.5.2 Quantitative Metrics

Coefficient of Variation (COV)

With the final aim of assessing the quality and effect of the denoising, different metrics were used considering different datasets. Considering the 3T LPM Dataset, in addition to the evaluation of the scan and rescan reproducibility in four different regions, the Coefficient of Variation (COV) was also evaluated considering the same regions, between the original and noisy T_2 maps and the denoised version for every single subject included in the test set. The COV was computed considering the following relation:

$$\mathrm{COV}\ [\%] = \frac{\sigma}{\mu} \times 100, \tag{4.15}$$

where μ is the mean T_2 value and σ the standard deviation of the T_2 values for the specific region considered. In addition, a statistical test was implemented to evaluate whether a significant statistical difference existed between the distribution of COV values in the regions before and after denoising using the implemented methods. The statistical test used is the Wilcoxon test. The Wilcoxon test is a non-parametric statistical test specifically designed to compare two distributions of data, determining whether there is or not a significant difference in the means [54].

The p-value selected for this study was 0.05, which was used as the threshold for determining statistical significance.

Contrast to Noise Ratio (CNR) and Contrast Ratio (CR)

Considering the 4 different datasets at 7T, two other metrics were implemented to assess the quality and performances of the denoising method. Let M_i be the median value of T_2 in the selected ROIs i and IQR the Interquantile Range considering the 25^{th} and 75^{th} percentile of one of the two ROIs being considered, it is possible to define the Contrast to Noise Ratio (CNR) and the Contrast Ratio (CR) as following:

$$CNR = \frac{M_1 - M_2}{IQR} \tag{4.16}$$

$$CR = \frac{M_1}{M_2} \tag{4.17}$$

In this case, the selection of regions of interest (ROIs) was performed manually rather than automatically. Specifically, the segmentation tool implemented in ITK-SNAP [55] was used to manually delineate the ROIs. For the brain datasets, one ROI was drawn in the thalamus, and another in the hippocampal tissue. In the knee dataset, one ROI was drawn in the cartilage, and another one in the muscle. The selected ROIs used in the metric evaluation are shown in the figure below (see **Figure 4.12**).

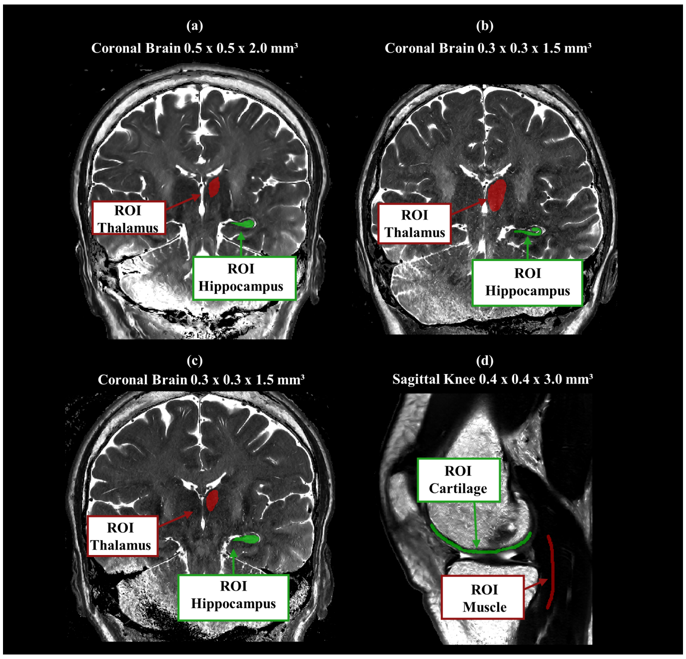

Figure 4.12: ROIs selected for the computation of the metrics of Contrast to Noise Ratio (CNR) and Contrast Ratio (CR) in the four available dataset at 7T, different in resolution and in anatomy.

Finally, the same metrics were used to evaluate the effect of the denoising achieved through the self-supervised strategies, in comparison to traditional methods already implemented inside the pipeline of the online reconstruction, i.e. the direct output obtained at the scanner. The scanner's reconstruction pipeline includes built-in post-processing and filtering using standard denoising methods, and comparing the performance of the deep learning-based denoising methods to these traditional techniques was essential for assessing the potential advantages of the proposed approaches.

Chapter 5

Results

In this chapter, the results of the experiments outlined in **Chapter 4** are presented. First, the outcomes of the k-space strategy are discussed, followed by the results of the image-space strategy. Next, the findings from the scan-rescan reproducibility analysis and the quantitative metrics used to evaluate the denoising performance are detailed. Finally, four distinct proofs-of-concepts on 7T datasets are introduced, with the final objective of demonstrating the generalizability of the model to datasets entirely different from those used during its training.

5.1 K-Space Strategy

As outlined in **Section 4.3**, the initial goal of the k-space strategy was to deliberately overfit the model to a single slice from a single subject in the 3T LPM dataset. The network was trained for a total of 80 epochs, and the results are summarized in **Figure 5.1**. Three learning rates were tested (1e−3, 1e−4, and 1e−5) along with three different values (0.2, 0.5, and 1) for the λ coefficient, which weights the model output in the final T_2 map. Despite exploring these parameter combinations, none successfully denoised the T_2 map. While certain relative difference maps appeared promising for noise reduction, they did not translate into effective denoising when evaluated visually on the T_2 maps themselves, where high relative error persisted. Moreover, the strategy consistently exhibited a negative bias in the T_2 values, resulting, under the same color scale, in hypointense regions that were particularly evident in the white matter.

5.2 Image Space Strategy

Training Results

The model's performances were initially evaluated by analyzing the training and validation loss curves. As shown in **Figure 5.2** , the training loss, in blue, consistently decreased over the course of the training process, indicating that the model was effectively learning from the training data. The validation loss, in red, also exhibited a downward trend, though it was observed to slightly reach a plateau after 300-400 epochs, suggesting that

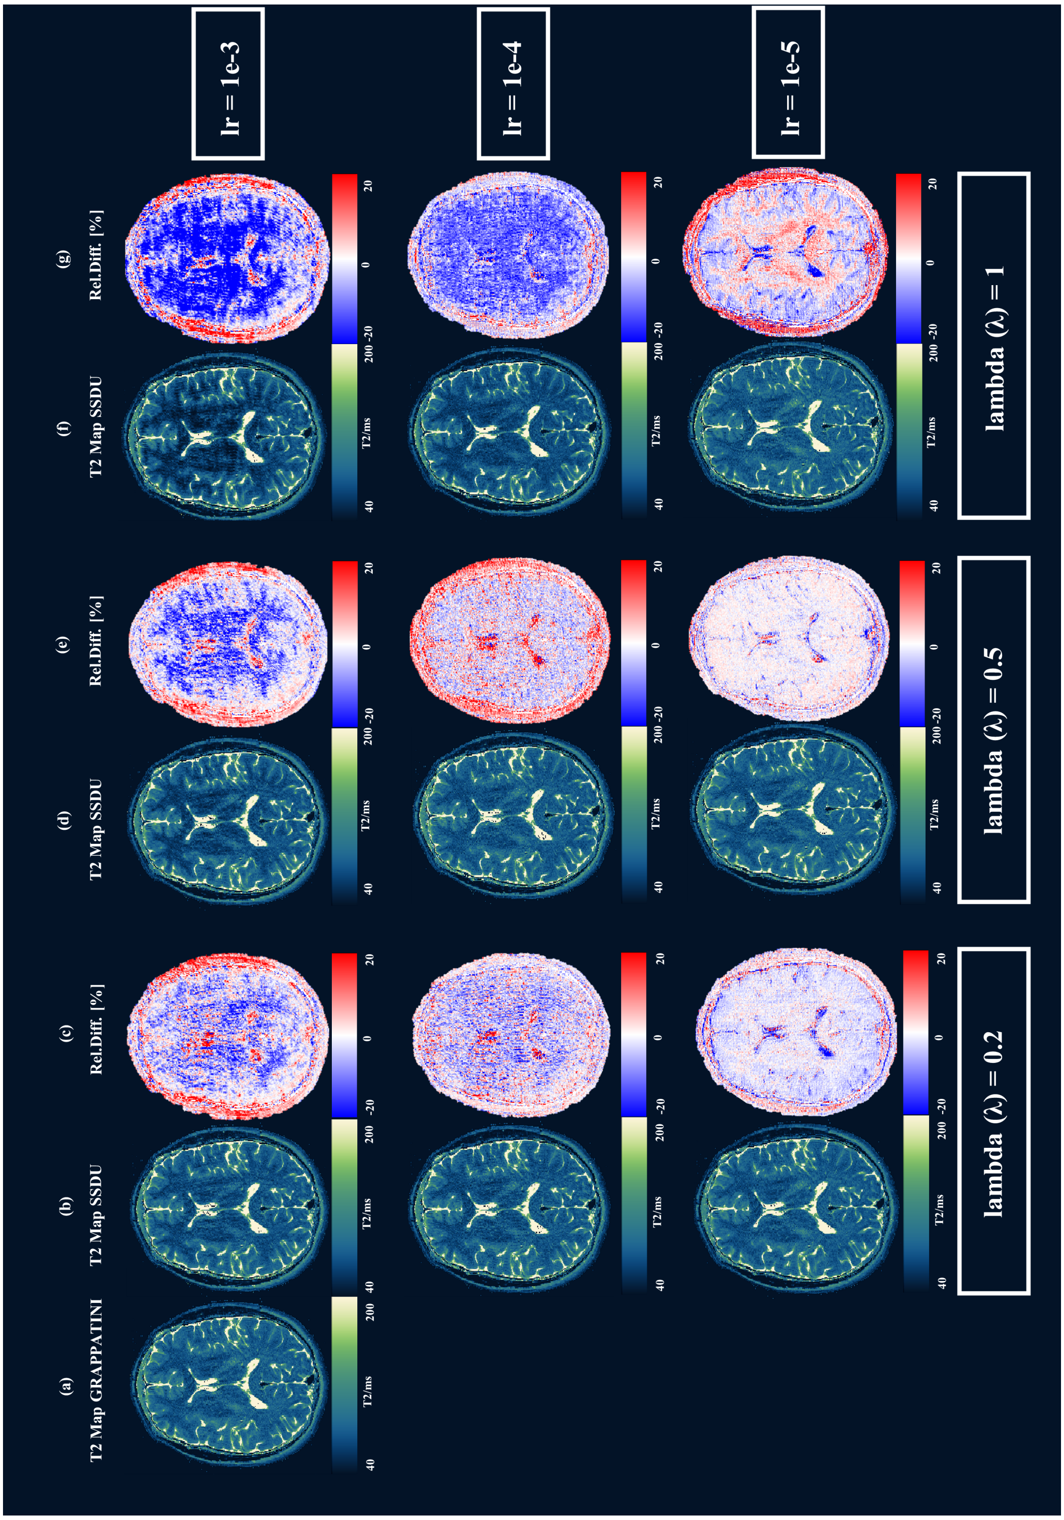

Figure 5.1: Results of the overfitting process on a single slice obtained from a subject of the 3T LPM Dataset. (a) shows the original and noisy GRAPPATINI T_2 map. In the subsequent columns, different combinations of the λ parameter and learning rate of the model were tested.

the model had reached a point of diminishing returns on the validation set. The gap between the training and validation losses remained narrow across the epochs, implying minimal overfitting, and further supporting the model's ability to generalize well to unseen data.

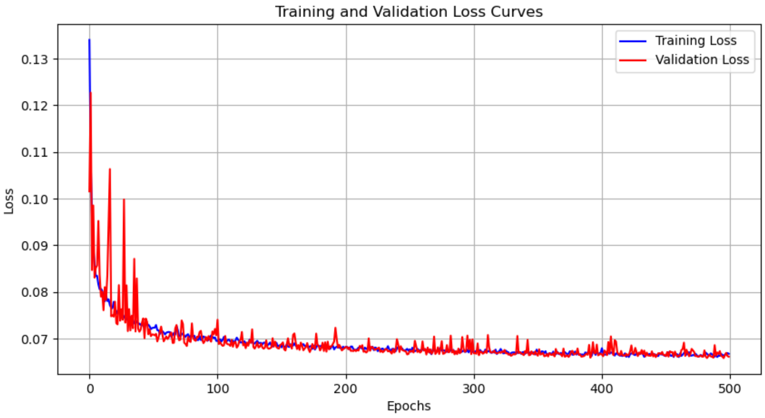

Figure 5.2: Training and validation loss curves in the training process of the model in the image-space strategy.

Interpolation Coefficient α

In the image-space strategy, the final model output at inference was linearly combined with the input T_2 map (see **Equation 4.13**), necessitating the selection of an interpolation coefficient, α. The results of this selection process, applied to a fully sampled dataset acquired at 3T, artificially corrupted with Gaussian noise, and subsequently denoised using the same strategy adapted for the larger dataset, are presented below (see **Figure 5.3**). Panel (a) illustrates the T_2 map obtained from fully sampled k-spaces. Panel (b) shows the impact of Gaussian noise corruption on the T_2 map, derived from artificially corrupted k-spaces. For this analysis, the DnCNN model was trained to overfit a single T_2 map derived from the corrupted fully sampled k-spaces, as described in **Section 4.4.2**. Panel (c) in **Figure 5.3** shows the denoised T_2 map resulting from this overfitting process. Finally, the MSE was evaluated between the denoised version of the T_2 map and the original fully sampled GRAPPATINI considering values of α in the range [0,1]. The minimum point of the MSE curve as a function of the α coefficient is highlighted with a red dot (see **Figure 5.3**). For the results presented in this chapter, the interpolation coefficient α was set to 0.7. The choice to interpolate the original input and the denoised output was made to improve reconstruction accuracy and with future clinical applications in mind, envisioning a user-friendly implementation in which α could be adjusted directly on the scanner immediately after the reconstruction, to the preference of the radiologist.

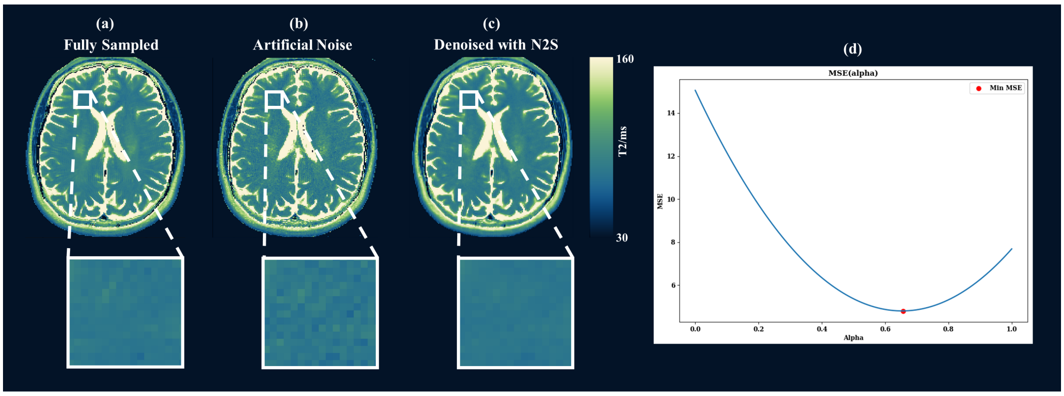

Figure 5.3: (a) GRAPPATINI reconstruction of the fully sampled 3T Dataset. (b) GRAPPATINI reconstruction starting with fully-sampled k-spaces but with added artificial Gaussian noise ($\mu = 0$, $\sigma = 3e - 5$). (c) shows the result of the image-space strategy applied to the corrupted T_2 map. (d) MSE as a function of the linear interpolation coefficient α.

3T LPM Dataset Results

As detailed in **Section 4.1**, the 3T LPM Dataset was divided into two sets for the model's training process. Specifically, 32 subjects were included in the training set, while the remaining 20 formed the test set. The model's performances on the test set were evaluated and are presented in this section. **Figure 5.4** and **Figure 5.5** provide a representative, though non-exhaustive, illustration of the denoising results achieved with the image-space strategy. A subset of subjects from the test set was selected, focusing on a specific subset of slices to demonstrate the denoising effect across the full range of slices. This selection spans from the lowest slices, emphasizing structures like the cerebellum, to the uppermost slices, where different anatomical structures are visible simultaneously. Columns (a) and (b) display T_2 maps reconstructed using the standard GRAPPATINI method. Columns (c) and (d) present the results of the image-space strategy implemented in this work, applied to the same slices shown in Columns (a) and (b). Specifically, Columns (b) and (d) provide zoomed-in views of anatomical structures, highlighting the impact of the denoising method and the level of noise originally present in the T_2 maps immediately after the GRAPPATINI reconstruction. Finally, column (e) shows the relative difference maps between the original and noisy T_2 map and the denoised version.

Scan-Rescan Reproducibility

In this section, the results of the scan-rescan reproducibility analysis performed on the 3T LPM dataset are presented. Linear regression was performed on the median T_2 values obtained from four different ROIs, as described in **Section 4.5.1**. The analysis was conducted for both the original GRAPPATINI method and the T_2 maps denoised using the image-space strategy. This study included all subjects from the test set, as detailed in **Section 4.1**. The results, shown in **Figure 5.6**, demonstrate that for the GRAPPATINI

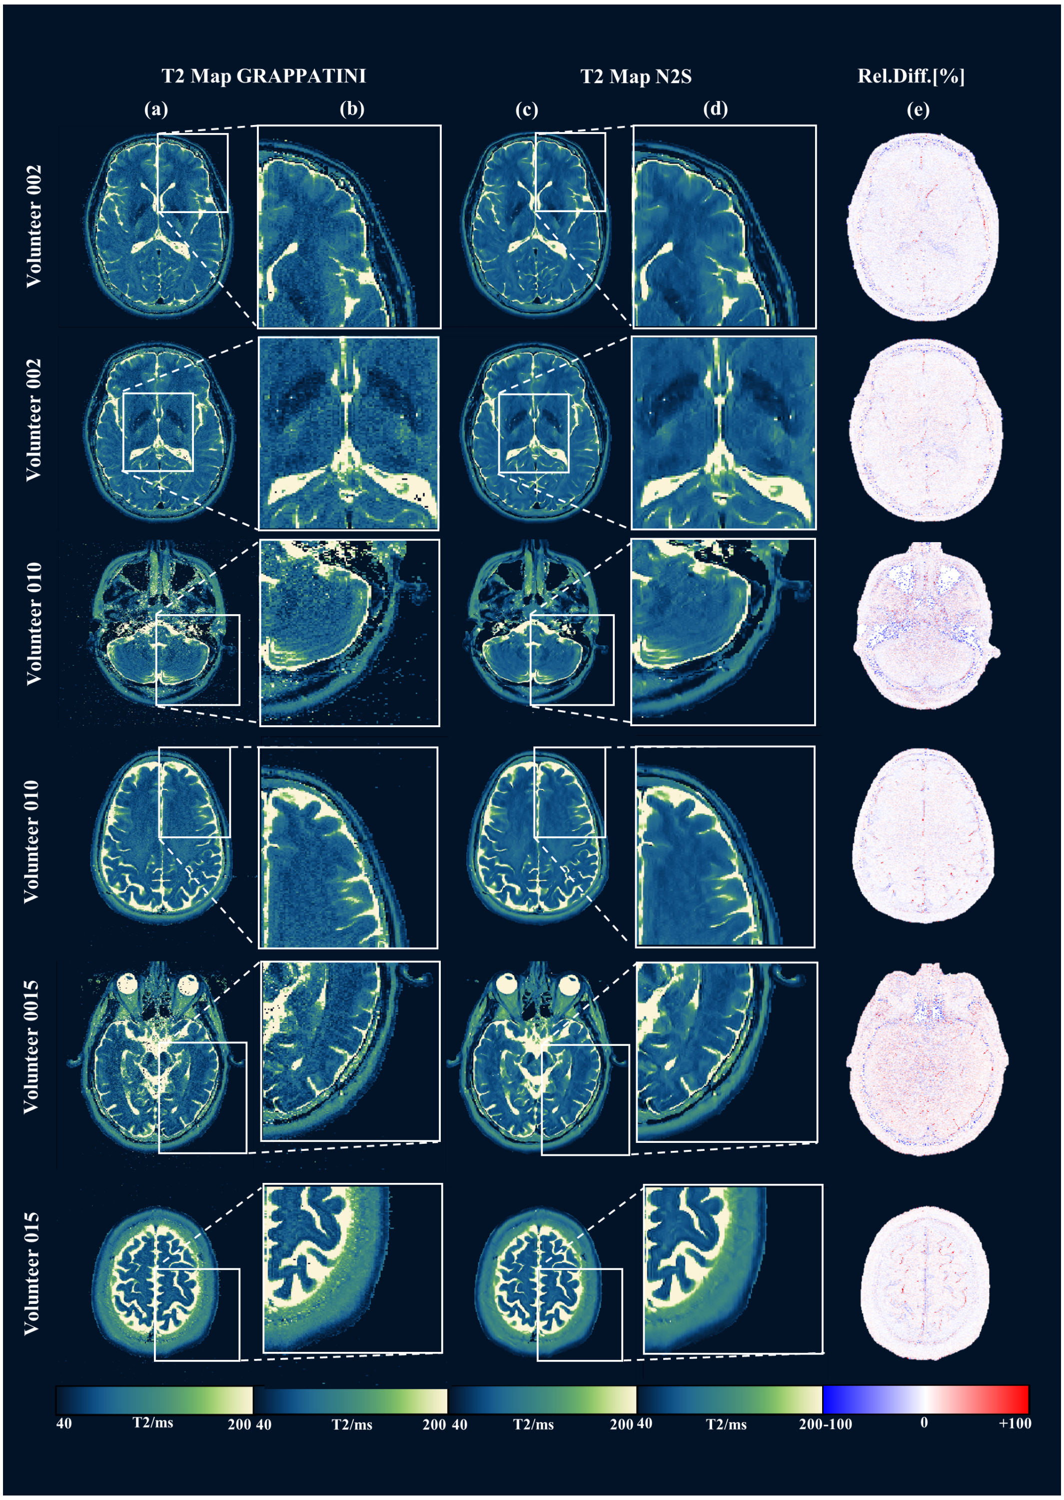

Figure 5.4: T_2 maps reconstructed using (a) standard GRAPPATINI, (c) applying the image-space strategy (N2S), with detailed views on anatomical regions shown in (b) and (d) respectively.

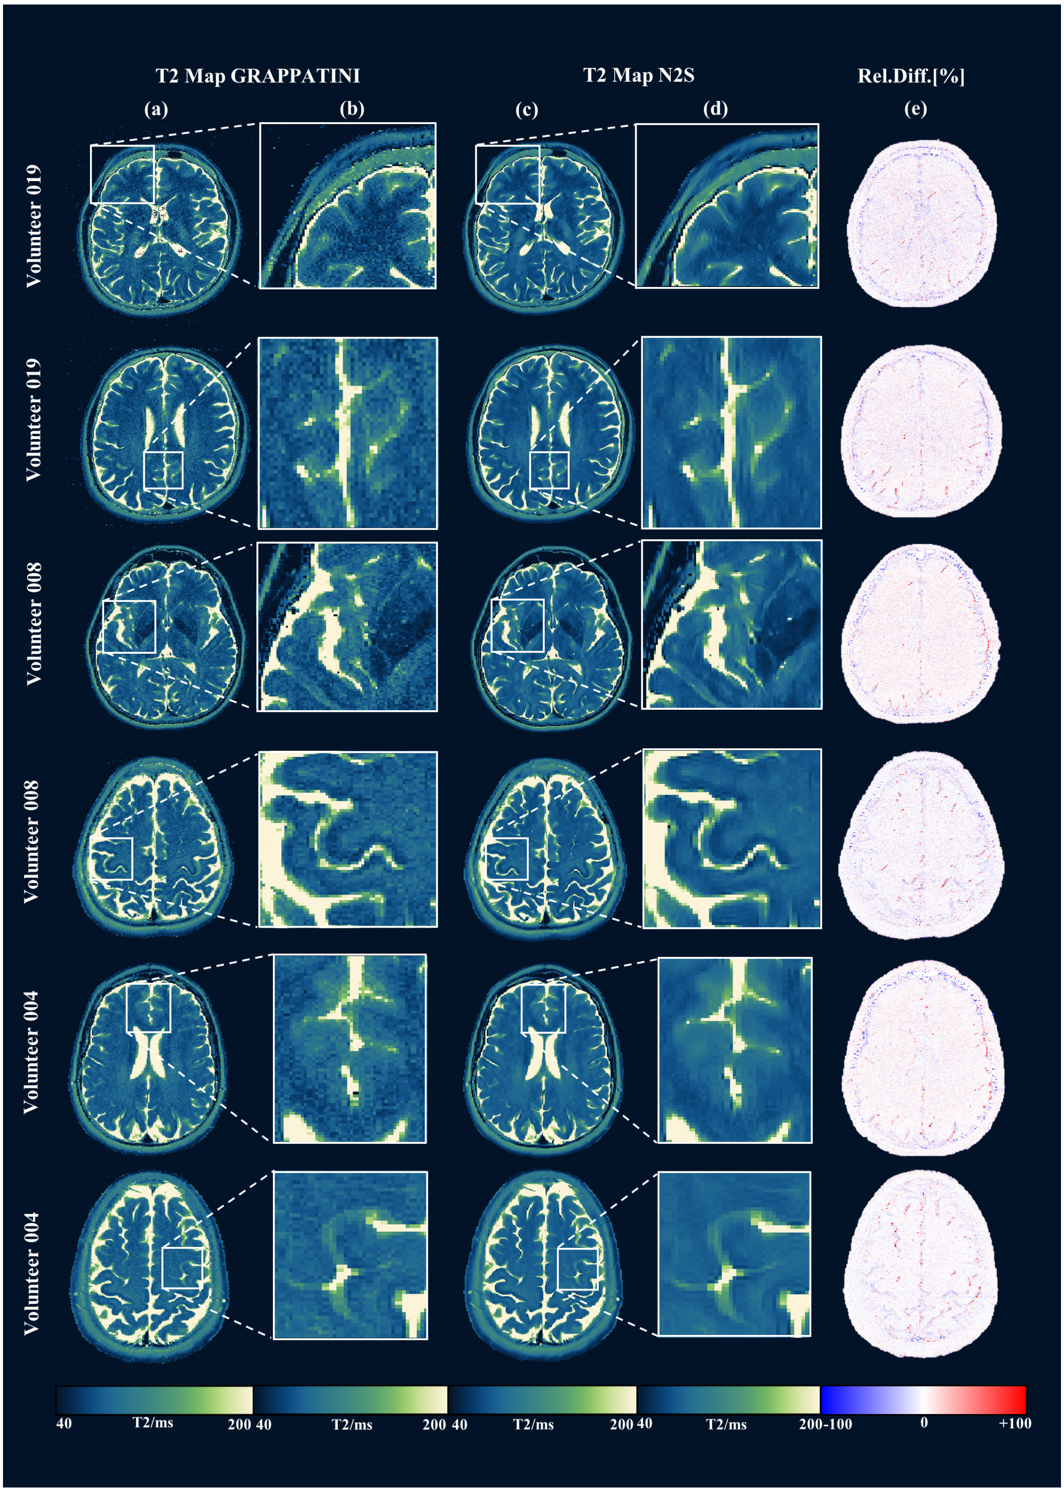

Figure 5.5: T_2 maps reconstructed using (a) standard GRAPPATINI, (c) applying the image-space strategy (N2S), with detailed views on anatomical regions shown in (b) and (d) respectively.

method, the linear regression yielded a slope of 0.99, an intercept (bias) of 0.94 ms, and an R^2 coefficient of 0.98. Similarly, for the image-space strategy (referred to as N2S in the figure), the regression produced a slope of 0.98, an intercept of 1.78 ms, and an R^2 coefficient of 0.98.

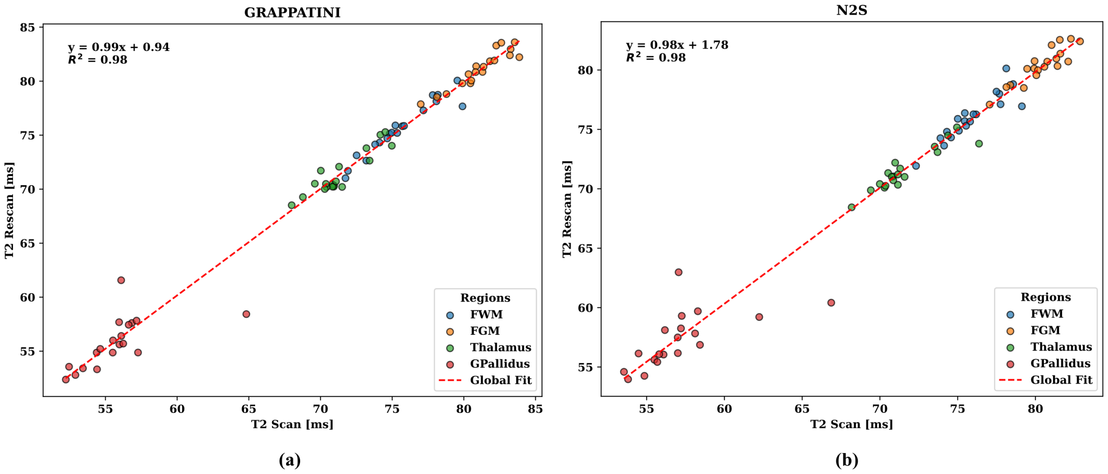

Figure 5.6: Scatter plot diagrams comparing median T_2 values obtained using GRAPPATINI (a) and N2S (b) across the selected ROIs during scan and rescan acquisitions (FWM: Frontal White Matter, FGM: Frontal Gray Matter, Thalamus, and GPallidus: Globus Pallidus). A linear regression is performed, with the resulting linear model shown as a red line. The equation for the linear model and the corresponding R^2 coefficient are displayed in the top left corner of each plot.

The second analysis performed was a Bland-Altman analysis to evaluate potential biases or systematic differences between the original GRAPPATINI method and the proposed image-space denoising strategy. This analysis demonstrates that the new method does not statistically alter the T_2 values, given that the reproducibility of the GRAPPATINI T_2 maps had already been established in previous studies. The Bland-Altman analysis was conducted to evaluate the agreement between GRAPPATINI and N2S measurements considering two consecutive scanning sessions and across four brain regions of interest: Frontal White Matter (FWM), Frontal Gray Matter (FGM), Thalamus, and Globus Pallidus (GPallidus). For every subject in the test set the median T_2 was extracted from the four ROIs and then represented with a dot in the plot. As highlighted by the results shown in **Figure 5.7** the mean bias between the two methods was -0.34 ms, as indicated by the red dashed line, suggesting a small average offset between the two methods. The limits of agreement in this case ranged from -2.05 ms to +1.38 ms, as represented by the gray dashed lines and the blue area. Approximately 95% of the data points fall within these limits, as expected under the assumption of the normality of the data. The plot also reveals the distribution of differences across the range of mean measurements (55 ms to 85 ms). Differences do not appear evenly distributed; for instance, a slight clustering of data points near the upper limit of agreement is observed in certain regions, particularly in the FGM. Additionally, there is one clear outlier with a difference exceeding -6 ms in the GPallidus region. Regional variations are evident, as indicated by the color-coded points: differences in the FWM and Thalamus regions cluster near the mean bias, while differences in the GPallidus and FGM regions show larger variability.

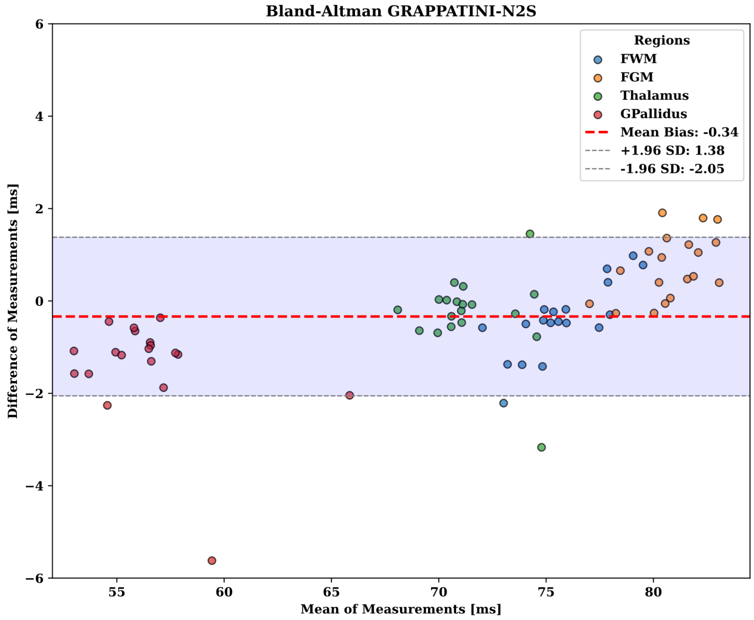

Figure 5.7: Bland-Altman diagram illustrates the comparison of GRAPPATINI median T_2 values across the selected ROIs (FWM: Frontal White Matter, FGM: Frontal Gray Matter, Thalamus, and GPallidus: Globus Pallidus). The gray dotted lines represent the 95% confidence interval, while the red dotted line indicates the bias between the two methods. The y-axis shows the differences between the median T_2 values and the mean of the two measurements, with values expressed in milliseconds (ms). Each dot represents a subject in the test set.

Subsequently, a second Bland-Altman analysis was conducted to assess the scan-rescan reproducibility of the two methods, this time evaluated separately. The primary objective was to identify any potential intra-method biases or systematic differences. **Figure 5.8a** shows the result of the analysis considering the original GRAPPATINI method. The mean bias evaluated was -0.02 ms, indicating a small average difference between the two measurements. The limits of agreement ranged from -2.13 ms to +2.09 ms, encompassing the expected 95% of the differences between repeated scans. The differences were plotted across the range of mean T_2 values (55 ms to 85 ms). The distribution of differences appeared consistent, with no apparent trend or proportional bias observed. Most data points lie within the limits of agreement, with only a single outlier detected at a difference below -6 ms in the GPallidus region. Regional differences show similar clustering behavior, with FWM and Thalamus exhibiting tighter distributions around the mean bias. The remaining two regions displayed slightly greater variability in differences but remained within the limits of agreement. **Figure 5.8b** presents the results from the analysis using the denoising strategy in the image-space (N2S). The mean bias calculated was -0.05 ms, indicating a negligible average difference between scan and rescan measurements. The limits of agreement delimited an area from -2.35 ms to +2.25 ms, including approximately 95% of the differences between the two scans. The differences are plotted against the mean T_2 values, which range from 55 ms to 85 ms. The distribution of differences appears consistent, with no evident proportional bias across the measurement range. One outlier, falling below -6 ms in the GPallidus region, was observed, consistent with findings from the other analyses. As seen in the GRAPPATINI method, the FWM and Thalamus regions show a tight distribution of points near the mean bias, indicating low variability. However, the other two regions exhibit higher variability. Despite this, all subjects remain within the established limits of agreement.

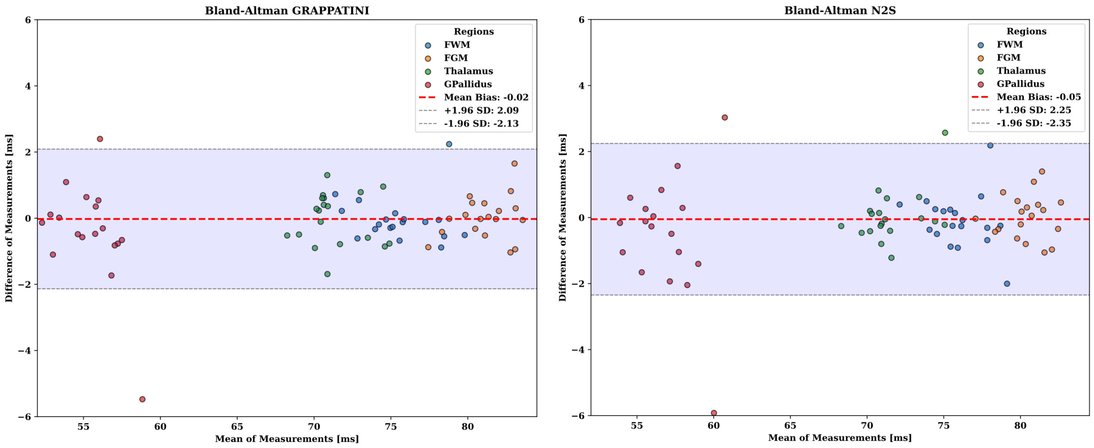

(a) **(b)**

Figure 5.8: Bland-Altman diagrams comparing GRAPPATINI (a) and N2S (b) median T_2 values across the selected ROIs considering scan and rescan acquisitions (FWM: Frontal White Matter, FGM: Frontal Gray Matter, Thalamus, and GPallidus: Globus Pallidus). The gray dotted lines represent the 95% confidence intervals, and the red dotted lines indicate the bias between the two methods. The y-axis shows the differences between the median T_2 values and the mean of the two measurements, with values expressed in milliseconds (ms). Each dot corresponds to a subject in the test set.

An additional analysis was performed on the same four extracted ROIs, considering all subjects in the test set, to evaluate the coefficient of variation (COV), as defined in **Section 4.5.2**. For the FWM, the GRAPPATINI method yielded similar median COV values of 11.99% and 11.91% for the scan and rescan sessions, respectively. In contrast, the N2S method produced notably lower median COV values of 7.75% in the scan session and 7.77% in the rescan session. However, a few outliers were observed for GRAPPATINI in the rescan session, with values exceeding 15%, suggesting higher variability under these conditions. In the FGM, the GRAPPATINI method demonstrated similar inter-scan median COV values of 14.71% and 14.20%, which were significantly higher than the corresponding values obtained with the N2S method, at 10.52% and 10.28% for scan and rescan sessions, respectively. For the Thalamus, the GRAPPATINI method reached the highest median COV values among all regions, with 16.22% and 15.88% for the scan and rescan sessions. By comparison, the N2S method resulted in much lower median values of 8.62% and 8.66% for the same conditions. Finally, in the Globus Pallidus, the N2S reported an higher COV with respect to GRAPPATINI, with median values of 11.08% and 11.40% for scan and rescan sessions, respectively, compared to 15.88% and 16.07% for GRAPPATINI.

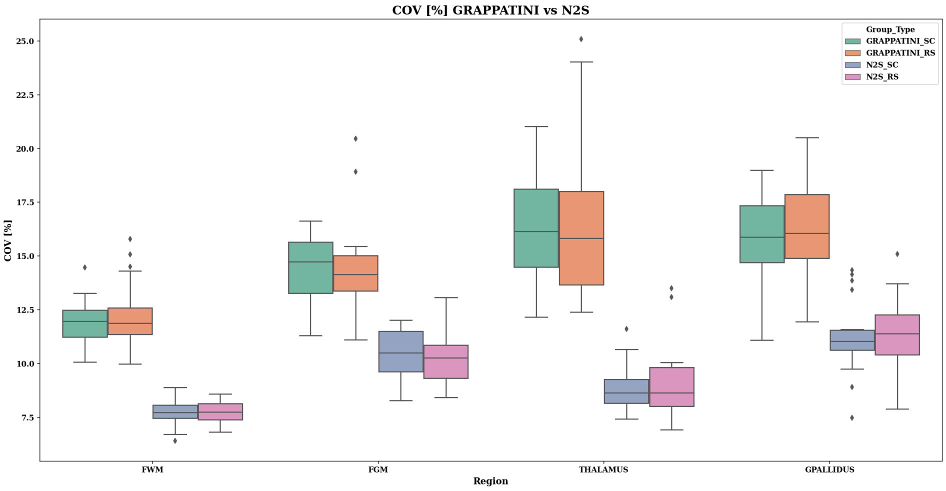

Figure 5.9: Box plot diagrams comparing the coefficients of variation for GRAPPATINI and N2S across the selected ROIs (FWM: Frontal White Matter, FGM: Frontal Gray Matter, Thalamus, and GPallidus: Globus Pallidus) and considering scan and rescan sessions (SC-RS).

Figure 5.10 represents a subset of the box plots shown in the previous figure (see **Figure 5.9**), focusing on the COV values for the GRAPPATINI and N2S methods across the same four ROIs in the first of the two sessions of scanning. In this plot, a Wilcoxon statistical test was performed to assess the significance of the differences between the two methods. The regions marked with an asterisk (*) indicate a statistically significant difference, with p-values below 0.05. Statistically significant differences were found across all the four ROIs.

5.2.1 Proofs of Concept : 7T Datasets

As a final step in the evaluation of the performances of the model trained using the image-space strategy, an analysis of the capabilities of the model to generalize to T_2 maps completely different from the ones used during the training and testing phases was performed. Four distinct datasets were acquired at 7T and the results are shown in this section.

Coronal Brain

Three datasets with the same anatomical structure (brain) as the training dataset were acquired, but with higher resolutions and a different orientation as well as acquired at

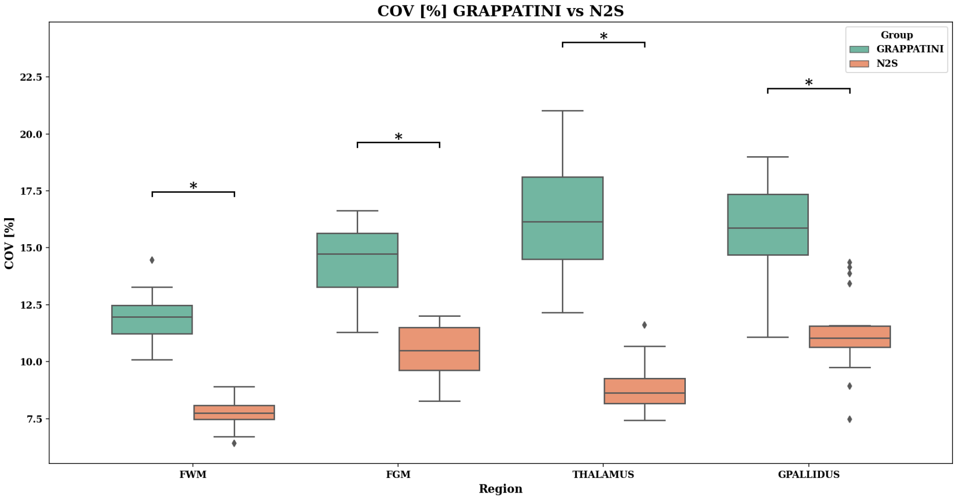

Figure 5.10: Box plot diagrams comparing the coefficients of variation for GRAPPATINI and N2S across the selected ROIs (FWM: Frontal White Matter, FGM: Frontal Gray Matter, Thalamus, and GPallidus: Globus Pallidus). A Wilcoxon statistical test was performed, and statistically significant differences are highlighted.

different field strength. Specifically, the lower in-plane resolution was set in a different plane compared to the one used for the training samples. For all these datasets, T_2 map reconstruction was performed using the standard GRAPPATINI method, followed by denoising using the model trained on the 3T LPM dataset. As a comparison with traditional methods of denoising, a comparison with the reconstruction directly performed at the scanner was evaluated.

The first dataset had a resolution of $0.5 \times 0.5 \times 2.0\text{mm}^3$, in contrast to the $0.7 \times 0.7 \times 3.3\text{mm}^3$ resolution of the 3T LPM dataset, and the T_2 maps were coronal rather than axial slices. **Figure 5.11** presents a comparison of T_2 maps generated with the original and standard GRAPPATINI method, denoised with the image-space strategy (N2S), and the online reconstruction, directly obtained at the scanner. Panel (a) displays the original T_2 map obtained with the GRAPPATINI method, while panel (b) shows the T_2 map after denoising with the image-space strategy (N2S). Panel (c) represents the T_2 map directly obtained from the scanner, referred to as "online reconstruction". Panels (d), (e), and (f) provide detailed views of the hippocampus region corresponding to panels (a), (b), and (c), respectively, while panels (h), (i), and (j) present detailed views of the thalamus region. The consistent colormap scale across all panels allows for direct comparison of signal intensities and visual differences in noise level between the methods. Moreover, panels (d), (e) and (f) highlights notable differences in the noise levels and structural clarity among the three methods. Panel (d), corresponding to the GRAPPATINI method, shows higher levels of noise, which are obscuring finer anatomical details in the hippocampus tissue. In contrast, panel (e), obtained after applying the image-space denoising strategy (N2S), demonstrates a significant reduction in noise, resulting in improved clarity and better preservation and recovery of the hippocampal structure and edges. Panel (f), representing the online reconstruction directly obtained at the scanner, exhibits higher noise levels compared to panel (e), but the noise pattern appears more uniform than in panel (d). Overall, the N2S method provides the cleanest view of the hippocampus and thalamus, with both reduced noise and sharper visualization of anatomical boundaries.

Method	**CNR** (↑)	**CR** (≈)	$\mathbf{COV}_{Thal}$(%)(↓)	$\mathbf{COV}_{Hippo}$(%)(↓)
GRAPPATINI	2.697	1.370	9.925	16.590
N2S	**5.566**	1.382	**6.259**	**11.208**
Online Reconstruction	3.863	1.407	9.148	15.355

Table 5.1: CNR, CR and COV values for the coronal brain dataset at 7T with $0.5 \times 0.5 \times 2.0$ mm^3 resolution.

In **Table 5.1**, the Contrast-to-Noise Ratio (CNR) and Contrast Ratio (CR) for the three methods are reported. Additionally, the COV, as defined in **Section 4.5.2**, was computed for the two ROIs presented in **Section 4.5.2**, and the results are presented in the same table. The N2S method achieved the highest CNR value (5.566), with almost double the values from both GRAPPATINI (2.697) and the online reconstruction (3.863). The CR values, representing the ratio between the median T_2 values across the ROIs selected, were relatively consistent across all methods, with GRAPPATINI at 1.370, N2S at 1.382, and the online reconstruction at 1.407. Regarding the coefficient of variation (COV), the lowest COV value was observed with the N2S method. In the thalamus ROI, the COV was 6.259%, representing a reduction of approximately three percentage points compared to both GRAPPATINI (9.925%) and the online reconstruction (9.148%). Similarly, in the hippocampus ROI, the COV was 11.208%, showing a decrease of nearly five percentage points compared to GRAPPATINI (16.590%) and the online reconstruction (15.355%).

The second and the third brain datasets acquired at 7T were acquired with a higher spatial resolution compared to the first one. As a consequence, the initial level of noise in the T_2 maps is significantly higher than in the previous case. In **Figure 5.12** and **Figure 5.13**, the comparison between the three methods are presented, considering the two datasets obtained from two different subject, here referred to as "001" and "002". As expected, the higher resolution in these datasets further amplifies the initial noise level in the T_2 maps, which is particularly evident in the GRAPPATINI reconstruction (a). The denoising effect of the N2S method (b) remains consistent even at higher resolution, effectively reducing the noise and allowing for improved visualization of the hippocampus compared to both GRAPPATINI and the online reconstruction (c). The zoomed-in views of the hippocampus (d), (e), and (f) highlight the capability of N2S to preserve anatomical structures and enhance the edges and sharpness of the map more effectively while maintaining a much lower noise level. The level of noise in the GRAPPATINI map is quite high, and even the traditional denoising methods included in the online reconstruction are not able to maintain their performance at this high resolution.

Table 5.2 reports the Contrast-to-Noise Ratio (CNR) and Contrast Ratio (CR)

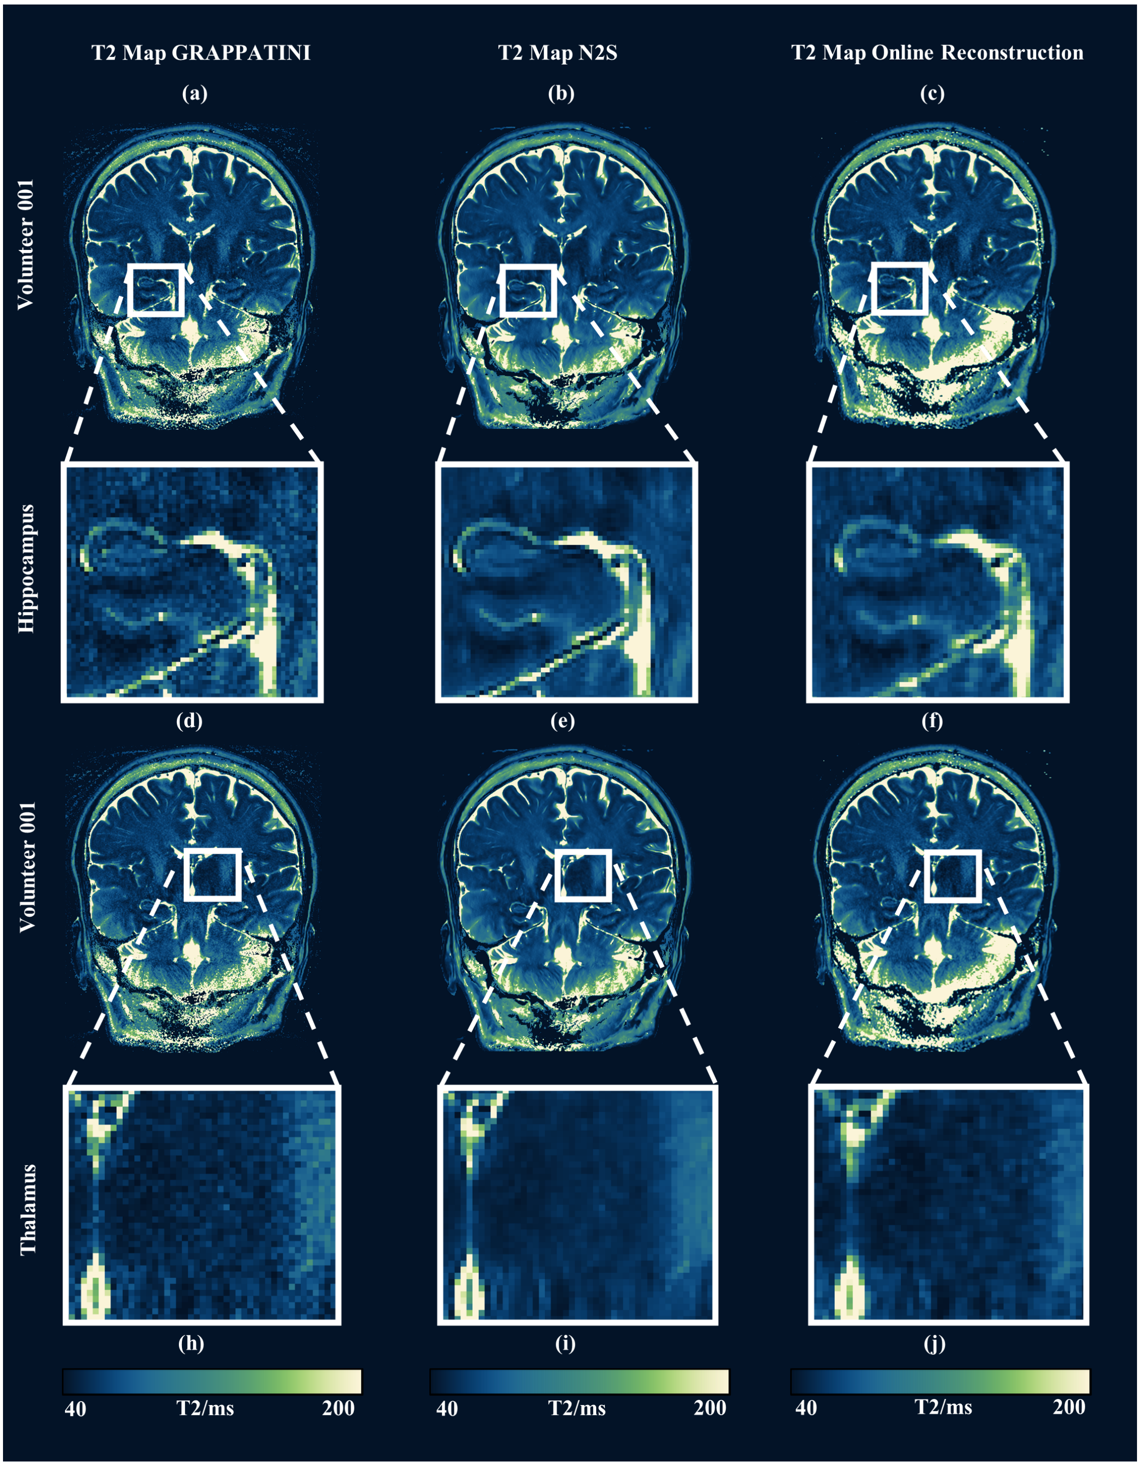

Figure 5.11: Comparison between T_2 maps using 3 methods at $0.5 \times 0.5 \times 2.0$ mm^3 resolution. (a) shows the GRAPPATINI T_2 map used as input to the model, (b) displays the output of the denoising method, and (c) the T_2 map obtained directly at the scanner. Figures (d), (e) and (f) provide more detailed views of the hippocampus regions, while (h), (i) and (j) present detailed views of the thalamus.

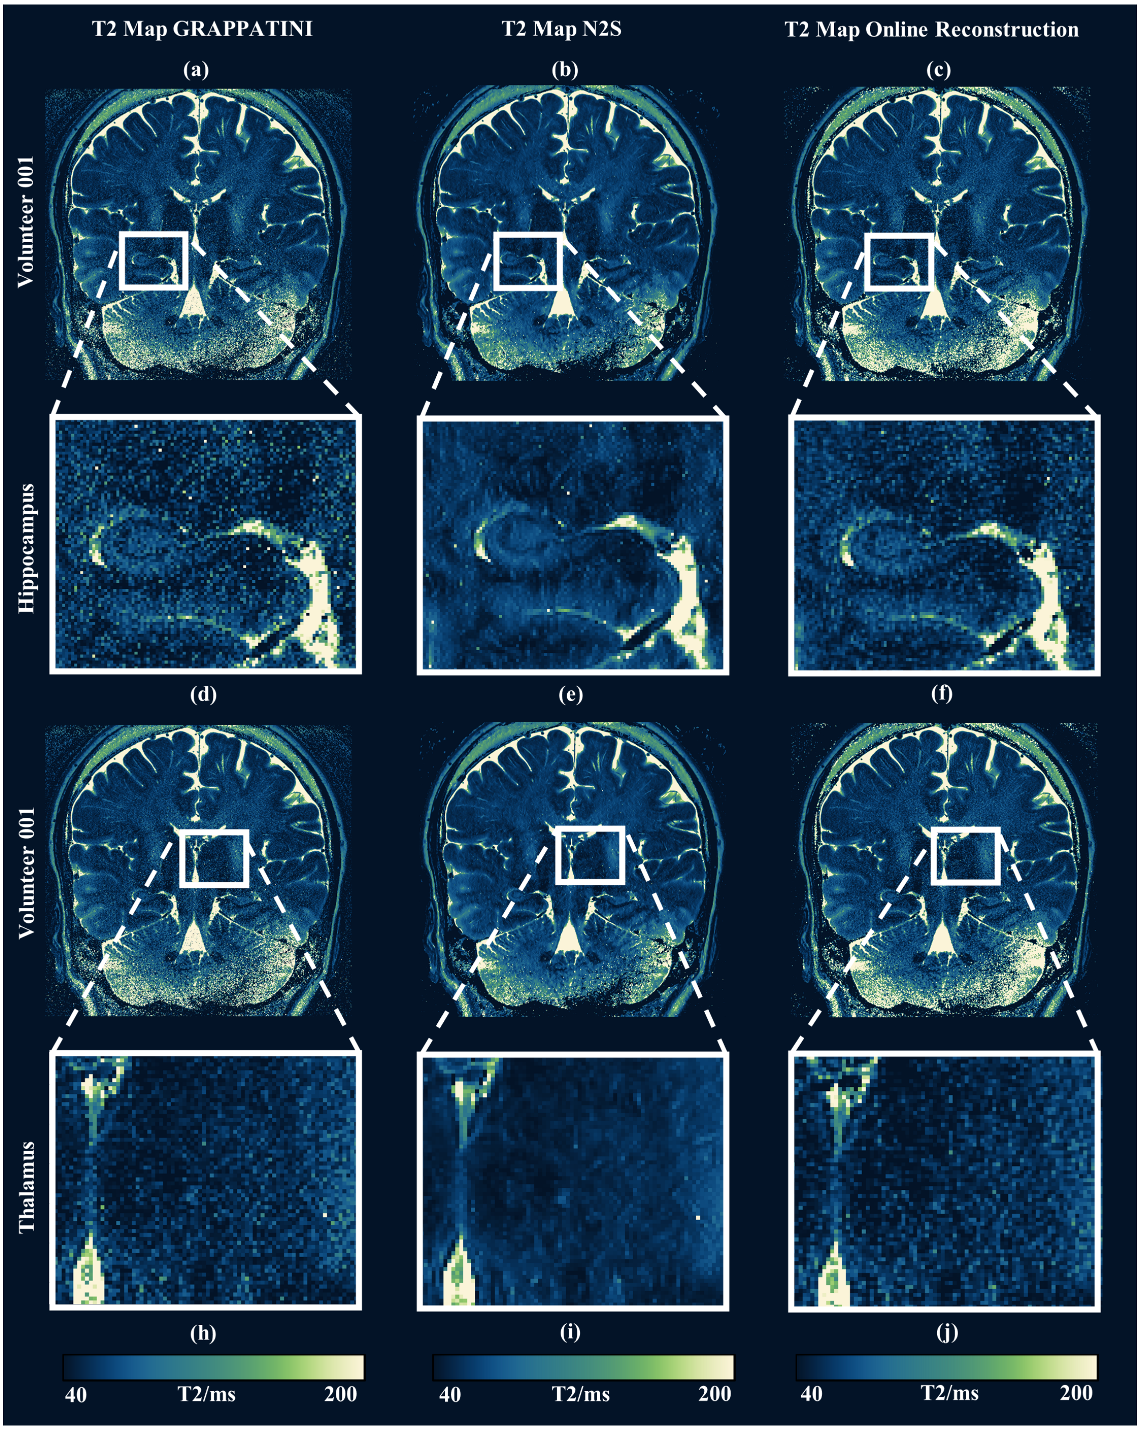

Figure 5.12: Comparison between T_2 maps using 3 methods at $0.3 \times 0.3 \times 1.5$ mm^3 resolution, considering the subject 001. (a) shows the GRAPPATINI T_2 map used as input to the model, (b) displays the output of the denoising method, and (c) the T_2 map obtained directly at the scanner. Figures (d), (e) and (f) provide more detailed views of the hippocampus regions, while (h), (i) and (j) present detailed views of the thalamus.

Method	CNR (↑)	CR (≈)	COV$_{Thal}$(%)(↓)	COV$_{Hippo}$(%)(↓)
GRAPPATINI	1.845	1.397	17.961	22.640
N2S	**3.240**	1.385	**9.616**	**15.817**
Online Reconstruction	1.485	1.412	22.350	18.575

Table 5.2: CNR, CR and COV values for the coronal brain dataset at 7T with 0.3 × 0.3 × 1.5 mm^3 resolution for the Subject 001.

values for the coronal brain dataset at 7T with a resolution of $0.3 \times 0.3 \times 1.5\text{mm}^3$. Additionally, the COV was computed for the two ROIs, and the results are presented in the same table. The N2S method achieved a CNR value of 3.240, which is an increase of 75% with respect to GRAPPATINI (1.845) and 95% with respect to the online reconstruction (1.661). The CR values remain consistent across methods with GRAPPATINI at 1.397, N2S at 1.385, and the online reconstruction slightly higher at 1.443. Considering the COV values, the lowest values were observed with the N2S method for both ROIs, at 9.616% and 15.817%. This represents a reduction of approximately eight percentage points compared to the GRAPPATINI method (17.961% and 22.640%) for both ROIs. Furthermore, with respect to the online reconstruction, the COV value in the thalamus ROI was reduced by more than ten percentage points (22.350%), and by approximately three percentage points in the hippocampus ROI (18.575%).

Method	CNR (↑)	CR (≈)	COV$_{Thal}$(%)(↓)	COV$_{Hippo}$(%)(↓)
GRAPPATINI	0.846	1.238	26.419	24.878
N2S	**1.513**	1.216	**12.180**	**16.446**
Online Reconstruction	0.556	1.225	27.453	31.642

Table 5.3: CNR, CR and COV values for the coronal brain dataset at 7T with 0.3 × 0.3 × 1.5 mm^3 resolution for the Subject 002.

Eventually, **Table 5.3** reports the CNR and CR for the third and last brain dataset acquired at 7T. The COV evaluated on the two ROIs are also reported. The N2S method achieved a value of 1.513, higher than both GRAPPATINI (0.8462) and the online reconstruction (0.673). The CR values remain consistent across methods with GRAPPATINI at 1.238, N2S at 1.216, and the online reconstruction slightly higher at 1.225. The COV was higher in both the GRAPPATINI and the online reconstructions compared to the N2S method for both ROIs.

Sagittal Knee

The final dataset was a knee scan acquired at 7T, with a resolution of 0.4 × 0.4 × 3.0 mm^3. **Figure 5.14** presents the resulting T_2 maps: Panel (a) depicts the reconstruction obtained using the GRAPPATINI method, column (b) shows the outcome of the image-space strategy applied to the same slice, and Panel (c) illustrates the online reconstruction. Panels (d), (e), and (f) provide a magnified view of the cartilage region to highlight the denoising effect, while panels (h), (i) and (j) present a view on a muscle region. The CNR and CR values for these methods are summarized in **Table 5.4**. In this

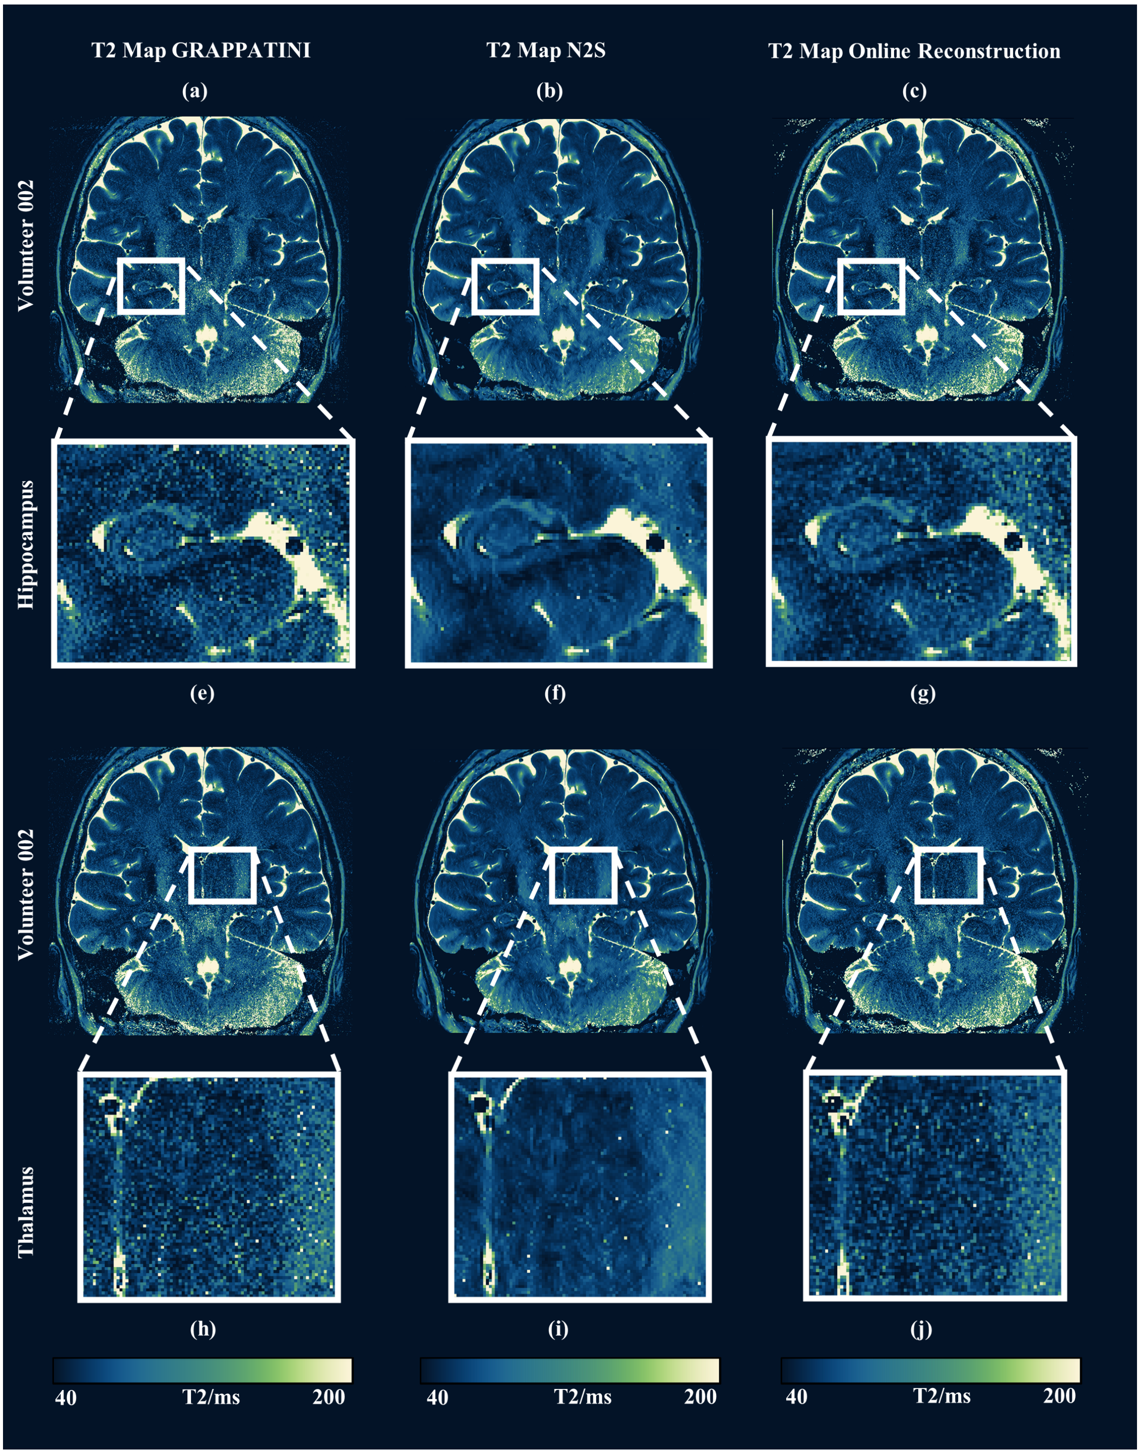

Figure 5.13: Comparison between T_2 maps using three methods at $0.3\times0.3\times1.5$ mm^3 resolution, considering the subject 002. (a) shows the GRAPPATINI T_2 map used as input to the model, (b) displays the output of the denoising method, and (c) the T_2 map obtained directly at the scanner. Figures (d), (e) and (f) provide more detailed views of the hippocampus regions, while (h), (i) and (j) present detailed views of the thalamus.

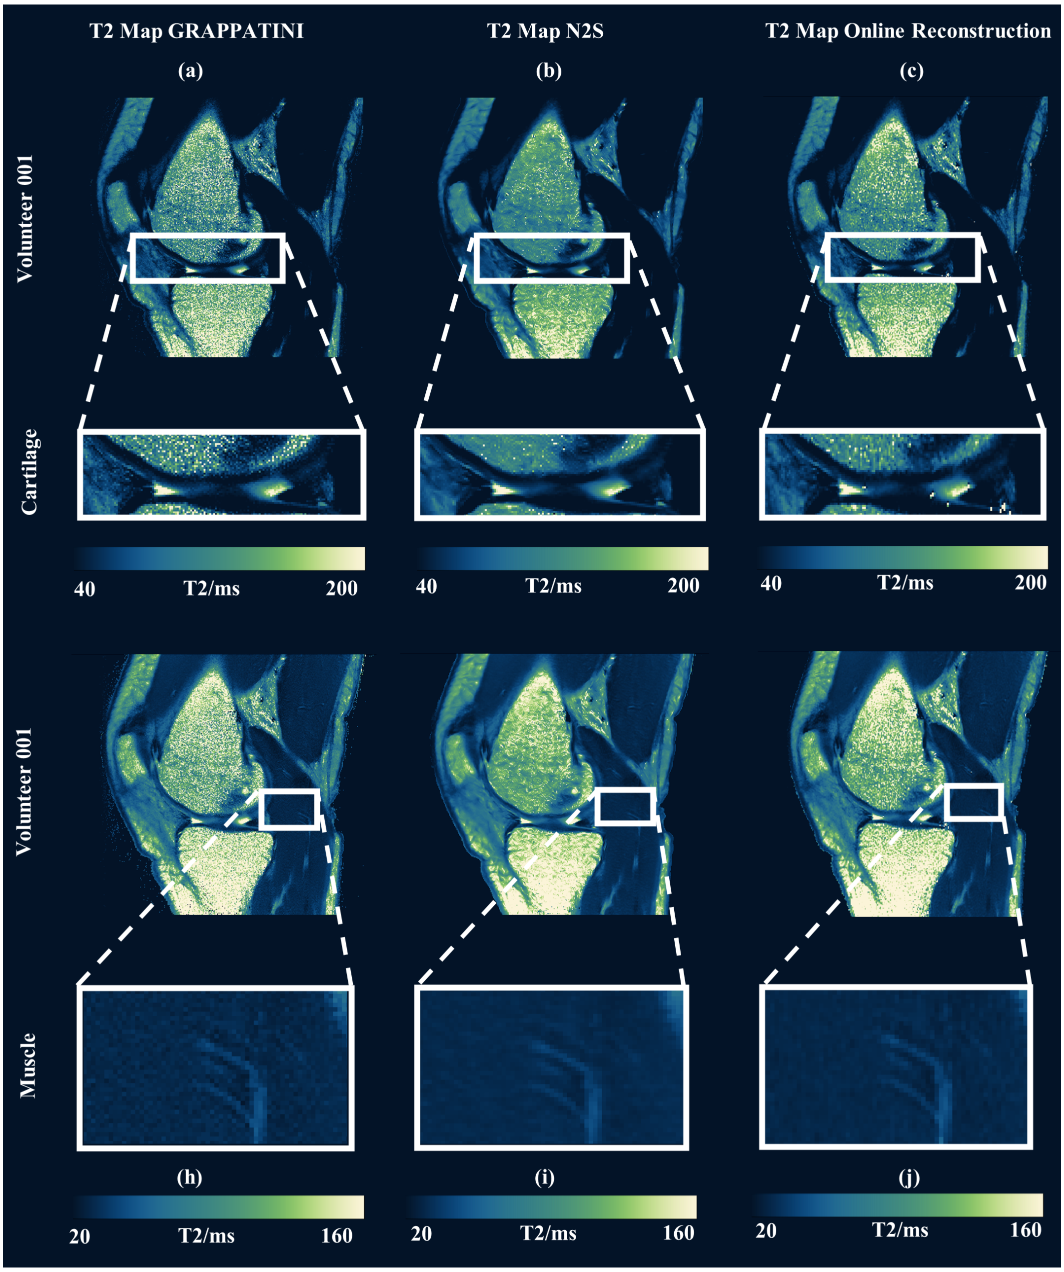

Figure 5.14: Comparison between three T_2 maps of a sagittal slice of the 7T Knee dataset acquired at $0.4 \times 0.4 \times 3.0$ mm^3 resolution. (a) shows the GRAPPATINI T_2 map used as input to the model, (b) displays the output of the denoising method, and (c) the T_2 map obtained directly at the scanner. Figures (d), (e) and (f) provide more detailed views of the cartilage region, while (h), (i) and (j) present detailed views of a region in the muscle.

dataset, the N2S method achieved a higher CNR value (7.667) compared to the original GRAPPATINI reconstruction (3.763). Interestingly, the CNR of the online reconstruction was nearly identical to that of the N2S method (7.666). The contrast ratio (CR) remained consistent across all three methods. Regarding the COV, better performance was observed with the N2S method (3.616% and 13.278%) compared to the other two.

Method	**CNR** (↑)	**CR** (≈)	**COV**$_{Muscle}$(%)(↓)	**COV**$_{Cartil}$(%)(↓)
GRAPPATINI	3.763	1.393	6.981	16.620
N2S	**7.667**	1.378	**3.616**	**13.278**
Online Recon.	**7.666**	1.379	3.645	16.896

Table 5.4: CNR, CR and COV values for the sagittal knee dataset at 7T with $0.4 \times 0.4 \times 3.0$ mm^3 resolution.

Chapter 6

Discussions

In the previous chapter, the results of the two denoising strategies were presented. The initial goal of this work, as described in the introduction, was to:

> *"...enhance the quality of T_2 maps generated using the GRAPPATINI method by means of a self-supervised ML approach (using Deep Learning), aiming for better results than traditional methods."*

Overall, the findings aligned well with this objective, particularly for the image-space strategy. In contrast, the k-space strategy did not meet the initial expectations for its potential. As an additional result, the performance of the model on the 7T datasets is promising, highlighting the capability of the model to generalize. This chapter delves into both the positive and negative outcomes of the work, as well as the main limitations and insights that emerged over the course of the project.

6.1 K-Space Strategy

At least theoretically, this strategy appeared to be the most natural and logical approach to follow in order to achieve a denoising strategy capable of "respecting" the original T_2 values. By considering a loss function directly defined in k-space, a better smoothing of the T_2 values was expected, as the physics underlying the imaging process would be inherently accounted for within the self-supervised framework itself. The concept was to incorporate a data-consistency step into the denoising strategy, as one of the most important aspects of the approach was to ensure that the statistical values remained unaltered, given their intrinsic physical significance. Undoubtedly, the knowledge of the forward model was, and still is, a significant advantage to leverage. However, despite these theoretical strengths, the k-space strategy did not perform as intended, revealing substantial challenges that limited its practical success. Different combinations of learning rate and the λ coefficient failed to produce the expected results. These two parameters were initially examined to manage the influence of the model itself on the input to the loss function. This was done to ensure better control over the computation of the loss and the subsequent updates to the model's weights. Moreover, the model was designed to be residual, meaning the input was added back to the output. This approach aimed to focus the learning process specifically on accurately modeling the behavior of the noise itself. Additionally, the initial choice of the ratio between the acquired k-space locations Ω and those used in the loss computation Λ was set at 0.05. This ratio was already constrained by the initially available k-space locations in the dataset. For the 3T LPM dataset, acquired with a GRAPPATINI acceleration factor of 10, a resolution of $0.7 \times 0.7 \times 3.3$ mm^3, and matrix sizes of 240×320 (considering 80% phase resolution), only 48 phase encoding lines per echo were available after performing the GRAPPA reconstruction. With a ratio ρ of 0.05, just two lines per echo were removed from the reconstruction and subsequently used to evaluate the loss during each training iteration. On one hand, this limitation constrains the amount of information the model can use during its learning process. On the other hand, increasing the number of phase encoding (PE) lines removed from the reconstruction, i.e. decreasing the cardinality of the Θ set of k-space locations, would likely degrade the final quality of the T_2 maps used as input to the model itself. Additional undersampling along the PE dimension would further amplify the ill-posed nature of the inverse problem, leading to a degradation in reconstruction quality. Moreover, in the GRAPPATINI sequence, block sampling ensures that the center of the k-space, where the highest concentration of information is stored, is sampled only every AF/2 echoes (where AF is the total GRAPPATINI acceleration factor). To preserve critical information for the reconstruction process, the central 24 lines of these echo k-spaces were always included across all the random patterns generated, for instance, in the attempt to reach overfitting on a single slice. This is why one of the initial conjectures regarding the method's failure was that attempting to denoise only two PE lines in the k-spaces, most of which were far from the center of the k-space, may not have provided the model with sufficient information to effectively learn the necessary features for denoising the T_2 maps. A second conjecture as to why the approach did not succeed may lie in the structure of the loss function. The application of the forward model, which includes the classical monoexponential decay, transforms the pair of M_0 and T_2 maps into a single image for each echo in the sequence. The signal-to-noise ratio (SNR) of these images might already be sufficiently high that, when transformed into k-space, the model does not receive enough information to effectively learn the noise pattern. Finally, this strategy is implemented, as described in **Section 4.3** with the k-spaces after the GRAPPA reconstruction. Technically, this violates the original assumption of independence between the input and output of the model. However, this was the only viable solution, as dividing the k-space locations into the two sets Θ and Λ prior to the GRAPPA reconstruction would have resulted in an unfeasible reconstruction. The limited number of PE lines would not provide sufficient data to achieve a minimum of the cost function during the non-linear optimization process defined inside the MARTINI method.

6.2 Image Space Strategy

As an initial note on the developed strategy in the image-space, it is important to acknowledge that the method technically violates one of the underlying assumptions of the general self-supervised framework proposed in [45]. Specifically, in the image-space, the noise is spatially varying and, more importantly, correlated. However, the results presented in **Section 5.2** demonstrate that, empirically, the method works effectively, with a visible denoising effect.

Interpolation Coefficient

As described in **Section 5.2**, the final output of the image-space denoising strategy involves performing a linear interpolation between the original GRAPPATINI T_2 map (the input to the model) and the model's output. This interpolation is controlled by a coefficient α, whose selection, for the results of this work, was guided by the fully sampled 3T dataset (see **Section 4.1.1**). This approach was selected as a first attempt to justify the selection of the coefficient. Starting from fully sampled k-spaces, the reconstructed T_2 map is characterized by an higher SNR. By introducing Gaussian noise to the k-spaces, a noisy version of the same T_2 map can be generated. Training the same DnCNN model with identical hyperparameters to denoise this noisy T_2 map would provide an initial baseline method for evaluating the effectiveness of the α coefficient. Specifically, selecting α considering the minimum of the mean squared error (MSE) curve with respect to the original GRAPPATINI T_2 map could serve as an initial proof of its effectiveness in achieving optimal denoising. However, this method introduces a potential limitation: the lack of an automated and optimal process for selecting the interpolation coefficient. A future improvement could involve developing a method to automatically determine the optimal value of α, tailored to each specific case, perhaps including it directly into the learning process of the model. Additionally, as discussed in **Section 5.2**, an alternative solution might involve incorporating the interpolation coefficient as a configurable parameter within the reconstruction pipeline. This parameter could then be adjusted by the radiologist analyzing the T_2 map directly at the scanner, based on their assessment of the overall blurriness and sharpness of the T_2 map itself.

3T LPM Dataset

The 3T LPM dataset was partitioned into two sets: a training set comprising 32 out of the 52 subjects and a test set with the remaining 20 subjects. The model's performance was evaluated on the test set, with the results detailed in **Section 5.2**. The denoising effect is visually evident in the zoomed views of selected anatomical structures, where it highlights tissue boundaries more clearly, enhancing the visualization of specific regions. Relative difference maps further illustrate the effectiveness of noise removal with the typical noise pattern and no general bias. Most relative errors were confined around the 0% level, indicating, potentially, an effective noise reduction. From a quantitative perspective, the coefficient of variation (COV) was computed across four ROIs: Frontal White Matter (FWM), Frontal Gray Matter (FGM), Thalamus, and Globus Pallidus (GPallidus). The average reduction in the median COV across these ROIs was 5.20% and 4.98% for the scan and rescan sessions, respectively, demonstrating the strategy's capability to smooth and reduce the variability in T_2 values present in the original GRAPPATINI T_2 maps and being consistent in doing this considering two consecutive scanning sessions of the same subjects. Additionally, statistically significant differences ($p<0.05$) were observed in the

COV distributions for all four regions when considering the test set subjects between the standard GRAPPATINI method and the image-space strategy. However, it is important to note that a reduction in COVs may also indicate the potential suppression of natural and physiological variability. For instance, the COV of a fully homogeneous image would be 0. Nonetheless, this is not the case here, as evidenced by the relative difference maps. As mentioned above, these maps present the typical pattern of noise, confirming that the natural variability is in the end preserved.

6.2.1 Scan-Rescan Reproducibility Analyses

The scan-rescan reproducibility analyses aimed to assess the statistical robustness of the T_2 values after applying the denoising strategy developed. Calculating and evaluating the COV alone was not sufficient, as the goal of this project goes beyond a typical computer vision problem; the pixel intensities in the final map hold a physical meaning and potential diagnostic value. Consequently, a more robust statistical analysis of the method's effectiveness was required, as a single metric assessing variability within specific ROIs was not enough. While a reduction in variability within a specific ROI are important factors, they are not the only effects to consider when evaluating the impact of denoising.

Specifically, the goal of the analyses was to identify any potential biases in T_2 values between two consecutive scanning sessions. The analyses were performed on the subjects included in the test set, evaluating the performance of the model on subjects not seen during the training process. The first step was to evaluate whether the original and standard GRAPPATINI reconstruction provided reliable and reproducible T_2 values. This objective was addressed through the application of a linear regression, as presented in **Section 5.2** and **Figure 5.6**.

This analysis was conducted using median T_2 values across four selected ROIs: Frontal White Matter (FWM), Frontal Gray Matter (FGM), Thalamus, and Globus Pallidus (GPallidus). The mean bias between scan and rescan median T_2 values was 0.94 ms, and the slope of the linear regression was 0.99. These results confirm that the GRAPPATINI reconstruction is highly reproducible across two consecutive scanning sessions. However, two outliers were identified in the GPallidus region and a potential explanation for these outliers could lie in how the labels were extracted using the internal tool Morphobox or in how T_2 outlier values were excluded prior to calculating the final median, considering the MAD distance. Notably, the same outliers observed in the original GRAPPATINI analysis were also present in the linear regression performed on the median T_2 values obtained after applying the denoising model. Importantly, in addition to yielding a comparable bias and linear regression coefficient relative to the original GRAPPATINI, the distribution of median T_2 values appeared more clustered with the denoising method, i.e. the median T_2 values evaluated for different subjects within the same ROI appeared more tightly grouped, with a distribution closer to the expected T_2 values for that region. This increased clustering may indicate an improvement resulting from the application of the denoising model and a more precise computation of the median T_2 values as a consequence.

Considering the Bland-Altman analyses, the first analysis aimed to evaluate potential inter-method biases, while the subsequent two focused separately on the performance of the method during scan and rescan sessions. The results confirmed the effectiveness of the image-space denoising strategy in reducing noise from the T_2 maps. Notably, the Bland-Altman analyses demonstrated that the area between the limits of agreement, that can be considered as a measure of variability between two measurements, was smaller when comparing the two methods than when evaluating the same method across scan and rescan sessions. This finding indicates that inter-method variability is lower than intra-method variability, which is in alignment with initial expectations. One point for discussion could be the higher discrepancy between the two methods in the T_2 values within the GPallidus. A possible conjecture is that the standard GRAPPATINI, being more noisy in hypointense regions like the GPallidus, may incorporate higher and erroneous T_2 values in the evaluation of the median T_2, not effectively eliminated by the metrics used to remove potential outliers, and eventually contributing to the observed differences.

6.2.2 Proofs of Concept

In **Section 5.2.1**, the results of the image-space strategy applied to the 7T datasets are presented. Four different datasets were acquired, considering various anatomies, resolutions, orientations and field strengths with respect to the T_2 maps included in the training set of the model. Specifically, one brain dataset consisting of T_2 maps in the coronal plane was acquired with a resolution of $0.5 \times 0.5 \times 2.0$ mm^3. Additionally, two more brain datasets with a spatial resolution of $0.3 \times 0.3 \times 1.5$ mm^3 were analyzed and lastly, a knee dataset, which resulted in sagittal T_2 maps with spatial resolution of $0.4 \times 0.4 \times 3.0$ mm^3, was also included in the study. To the best of the author's knowledge, the GRAPPATINI T_2 maps acquired at 7T in this work represent the highest resolution ever achieved and attempted for such data. The denoising results on the reconstructed T_2 maps are promising, especially considering the model's training conditions. The model was trained exclusively on brain T_2 maps with a resolution of $0.7 \times 0.7 \times 3.3$ mm^3 and an axial orientation. Despite these significant differences between the training set and the 7T datasets, the model exhibited an exceptional ability to generalize. Starting from the brain datasets, two specific anatomical structures were highlighted. Detailed views of the hippocampus region were provided with the final aim of assessing the effect of the denoising strategy as this is one of the more relevant structures that can be studied and seen in a coronal T_2 map. T_2 relaxometry of the hippocampus is indeed one of the most promising tools to assess and study the so-called temporal lobe epilepsy (TLE), particularly by detecting subtle changes like gliosis and neuronal loss. It is of crucial importance in diagnosis, seizure focus localization, and surgical planning for TLE patients [56]. As illustrated in **Figures 5.11, 5.12, 5.13**, the image-space strategy produces hippocampal views in which tissue areas, boundaries, and anatomical details are more prominently visible compared to the original GRAPPATINI reconstruction. In the 7T dataset with lower resolution (see **Figure 5.11**), the inherent noise level is already relatively low, which explains why the "baseline" GRAPPATINI reconstruction appears less noisy than the corresponding GRAPPATINI reconstructions from the other two datasets.

From a quantitative point of view, considering the brain datasets, the image-space strategy outperformed the original GRAPPATINI, yielding higher contrast-to-noise ratio (CNR) values in all three brain datasets with an average improvement of 87% (see

Tables 5.1, 5.2, 5.3). Furthermore, this approach outperformed also the conventional filtering methods integrated into the reconstruction directly performed at the scanner, showing an average increase in the CNR value of 50.5%. Nearly doubling the performance of traditional denoising methods currently used on the scanner represents a significant success and an optimal starting point for future integration of the model trained in this work into the scanner reconstruction pipeline. Additionally, the CR (Contrast Ratio) value remained consistent and stable across all three datasets, indicating that the denoising strategy was not introducing significant alterations to the T_2 values. This is evidence that the original level of contrast in the selected ROIs was preserved. Furthermore, across the three brain datasets, the N2S method outperformed both reconstruction methods in terms of COV for both ROIs (thalamus and hippocampus).

On the other hand, considering the knee datasets, the results are in line with what happens inside the reconstruction at the scanner. In particular, **Figure 5.14** provides zoomed-in views of the most critical structure for T_2 mapping of the knee: the cartilage. Numerous studies have demonstrated that alterations or variations in cartilage maturity can be tracked and analyzed through T_2 values. This is especially relevant in the context of pathologies, injuries, or the effects of training programs [57]. Regarding the CNR value and considering the ROIs in the muscle and in the cartilage (see **Table 5.4**), the online reconstruction is similar to that of the proposed image-space strategy (7.666 with respect to 7.667). However, after a visual analysis, the blurring effect of traditional methods included in the online reconstruction appears worse than that observed with the denoising strategy proposed in this work. Notably, the effect of the proposed denoising strategy on bone regions is significant. In the original GRAPPATINI protocol used for knee dataset acquisition, a fat-saturation step is applied, resulting in highly noisy T_2 values in the bones. After applying the model, potential bone patterns seemed to emerge. However, it remains unclear whether this reflects a beneficial effect of denoising, revealing previously hidden structures or patterns, or whether it represents model hallucinations or regions where the T_2 values lack meaningful interpretation due to the fat-saturation step inherent in the sequence design. The COVs in the two ROIs (cartilage and muscle) were lower considering the proposed image-space denoising strategy, with respect to both GRAPPATINI and the online reconstruction.

Overall, the successful adaptation and generalization to different anatomical regions, alongside the capabilities of handling varying T_2 value ranges, is particularly significant because it demonstrates that the ability of the model to denoise is not limited to the training conditions but can be applied across diverse datasets. This is crucial for potential and future clinical applications, as the model could potentially be used in a wide range of patients and protocols without needing to retrain or adapt the model for each new case.

Demonstrating these generalization capabilities can be considered one of the key outcomes of this work. The prospect of relying on a potentially universal model, capable of denoising any T_2 map regardless of its resolution, orientation, field strength or anatomical structure being mapped, is highly appealing and attractive. Moreover, this could reduce the amount of data required to still achieve high-quality results, potentially addressing the challenge of building large datasets in the T_2 mapping field. This strategy could generate impact in research environments but, more critically, in clinical settings.

In such contexts, the adoption of new methodologies and technologies is often delayed or constrained by the specific conditions required for their proper functionality, which are typically quite restrictive or limited to few subcases. An additional key strength of the proposed method lies in its inherent simplicity. The proposed strategy and trained model serve as an effective, straightforward, and easily integrable addition to the standard GRAPPATINI reconstruction. By functioning, in the end, as a classical post-processing strategy, it seamlessly enhances and improves the existing pipeline without requiring substantial additional modifications.

6.3 Future Developments and Limitations

The performance of the image-space denoising strategy, in terms of robustness and preservation of T_2 values without statistical alterations, was evaluated using both the test set and the datasets acquired at 7T. The latter served as evidence of the model's capability to handle data not seen during the training process. As mentioned in this section, the dataset comprised only healthy controls, meaning that, theoretically, no pathologies were present that could have influenced the model's performance. Therefore, a natural next step would be to validate the model in cases with pathology. Such a study could help in assessing whether the denoising method introduces masking effects or artifacts on T_2 maps. However, it is worth emphasizing that, even though the subjects were all healthy controls, the T_2 value distributions between 3T and 7T datasets, as well as across different anatomies, are already inherently distinct and different. This demonstrates the model's ability to generalize to values outside the training set. Consequently, it is reasonable to expect that the model would still remain robust when applied to datasets with T_2 values typically associated with lesions in some regions. As a potential future improvement, leveraging synthetic data could address the challenge of limited sample sizes during the training process. Given the performance of the model on 7T datasets with only axial brain T_2 maps from 52 subjects, training including synthetic data could address the issue of prolonged acquisition times in T_2 mapping, which restricts the creation of large datasets. As a final remark, the primary limitation of the image-space denoising strategy implemented in this work is considering the T_2 maps as conventional images. This perspective is not entirely accurate, as every pixel intensity in a T_2 map represents a physical parameter with a specific meaning behind and unit of measurement. By not incorporating the underlying MRI physics, with for instance the knowledge of the forward model, this approach risks altering the values in a manner that may not fully align with the originally acquired data. For this reason, a potential future improvement to the strategy could involve incorporating, in the model or in the pipeline, a data consistency step. This addition would eventually bring the approach closer to the underlying concepts of the k-space strategy, and such a hybrid method could potentially yield better final results than both the strategies. Nonetheless, this work demonstrated, through statistical analyses, the reproducibility of T_2 values before and after denoising. These results were presented to ensure that the denoising process did not significantly alter the GRAPPATINI T_2 values, in a consistent way as reported in previous studies. However, further in-depth analyses are encouraged, particularly using a larger dataset and a broader group of subjects, to strengthen the statistical validation and enhance the reliability of the findings.

Chapter 7

Conclusions

This thesis work introduced two distinct strategies aimed at denoising GRAPPATINI T_2 maps using a self-supervised approach to train a deep learning model. While one of the initially conceptualized strategies was found to be ineffective, the second strategy, in the image-space, delivered promising results on the test set. More importantly, it demonstrated remarkable generalization capabilities, achieving results beyond the original expectations on the 7T datasets acquired during the project, improving the performances of traditional denoising methods already implemented in the online reconstruction performed at the scanner.

Additionally, the scan-rescan analysis supported the quality of the denoising achieved, confirming the reproducibility of T_2 values after the application of the method without introducing significant biases or statistical alterations. This study holds great promise as a tailored solution directly integrated into the standard GRAPPATINI pipeline, with potential applicability in scanner reconstruction workflows. The model's generalizability could be a key advantage, enabling a ready-to-use solution for T_2 maps across various anatomies, resolutions, orientations and field strengths without requiring further training or adjustments.

In conclusion, this work represents a significant step forward toward the clinical application of quantitative MRI (qMRI) techniques, such as T_2 mapping, particularly considering the GRAPPATINI method. Eventually, it paves the way for applying GRAPPATINI T_2 maps at 7T, where the combination of high-quality and high-resolution imaging can generate a huge impact across a range of clinical and research applications.

Bibliography

[1] Walther Gerlach and Otto Stern. "Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld". German. In: *Zeitschrift fÃ¼r Physik* 9 (1922), pp. 349–352.

[2] F. Bloch. "Nuclear Induction". In: *Phys. Rev.* 70 (7-8 Oct. 1946), pp. 460–474. DOI: `10.1103/PhysRev.70.460`. URL: `https://link.aps.org/doi/10.1103/PhysRev.70.460`.

[3] 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 Jan. 1946), pp. 37–38. DOI: `10.1103/PhysRev.69.37`. URL: `https://link.aps.org/doi/10.1103/PhysRev.69.37`.

[4] P. Lauterbur. "Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance". In: *Nature* 242 (1973), pp. 190–191. DOI: `10.1038/242190a0`. URL: `https://doi.org/10.1038/242190a0`.

[5] H. Kabasawa. "MR Imaging in the 21st Century: Technical Innovation over the First Two Decades". In: *Magnetic Resonance in Medical Sciences* 21.1 (Mar. 2022). Epub 2021 Apr 16, pp. 71–82. DOI: `10.2463/mrms.rev.2021-0011`.

[6] Eloy Martinez-Heras et al. "Diffusion-Weighted Imaging: Recent Advances and Applications". In: *Seminars in Ultrasound, CT and MRI* 42.5 (2021), pp. 490–506. ISSN: 0887-2171. DOI: `10.1053/j.sult.2021.07.006`. URL: `https://www.sciencedirect.com/science/article/pii/S0887217121000846`.

[7] D. P. Soares and M. Law. "Magnetic resonance spectroscopy of the brain: review of metabolites and clinical applications". In: *Clinical Radiology* 64.1 (Jan. 2009). Epub 2008 Aug 30, pp. 12–21. DOI: `10.1016/j.crad.2008.07.002`.

[8] M. S. Chow et al. "Functional magnetic resonance imaging and the brain: A brief review". In: *World Journal of Radiology* 9.1 (Jan. 2017), pp. 5–9. DOI: `10.4329/wjr.v9.i1.5`.

[9] V. P. Grover et al. "Magnetic Resonance Imaging: Principles and Techniques: Lessons for Clinicians". In: *Journal of Clinical and Experimental Hepatology* 5.3 (Sept. 2015). Epub 2015 Aug 20, pp. 246–255. DOI: `10.1016/j.jceh.2015.08.001`.

[10] Cristina Granziera et al. "Quantitative magnetic resonance imaging towards clinical application in multiple sclerosis". In: *Brain* 144.5 (2021), pp. 1296–1311. ISSN: 0006-8950. DOI: `10.1093/brain/awab029`.

[11] Carlo Pierpaoli. “Quantitative brain MRI”. In: *Topics in Magnetic Resonance Imaging: TMRI* 21 (2010), p. 63. DOI: `10.1097/RMR.0b013e31821e56f8`.

[12] Vikas Gulani and Nicole Seiberlich. “Quantitative MRI: Rationale and Challenges”. In: *Advances in Magnetic Resonance Technology and Applications*. Ed. by Nicole Seiberlich et al. Vol. 1. Academic Press, 2020, pp. xxxvii–li. ISBN: 9780128170571. DOI: `10.1016/B978-0-12-817057-1.00001-9`. URL: `https://www.sciencedirect.com/science/article/pii/B97801%2028170571000019`.

[13] Mara Cercignani, Nicholas G. Dowell, and Paul S. Tofts. *Quantitative MRI of the Brain: Principles of Physical Measurement*. Boca Raton, Florida: CRC Press, 2018. ISBN: 9781138197255.

[14] David W. McRobbie et al. *MRI From Picture to Proton*. 3rd. Cambridge, UK: Cambridge University Press, 2017. ISBN: 9781107643239.

[15] G. S. Pell, S. R. Thomas, and M. A. Lones. “Optimized clinical T2 relaxometry with a standard CPMG sequence”. In: *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* 23.2 (2006), pp. 248–252.

[16] Tomás Paus et al. “Maturation of white matter in the human brain: a review of magnetic resonance studies”. In: *Brain Research Bulletin* 54.3 (2001), pp. 255–266. DOI: `10.1016/s0361-9230(00)00434-2`.

[17] K. Schmierer et al. “Magnetization transfer ratio and myelin in postmortem multiple sclerosis brain”. In: *Annals of Neurology* 56.3 (2004), pp. 407–415. DOI: `10.1002/ana.20202`.

[18] C. Laule et al. “Myelin water imaging in multiple sclerosis: quantitative correlations with histopathology”. In: *Multiple Sclerosis* 12.6 (2006), pp. 747–753. DOI: `10.1177/1352458506070928`.

[19] R. J. Ogg and R. G. Steen. “Age-related changes in brain T1 are correlated with iron concentration”. In: *Magnetic Resonance in Medicine* 40.5 (1998), pp. 749–753. DOI: `10.1002/mrm.1910400516`.

[20] G. Bonnier et al. “Advanced MRI unravels the nature of tissue alterations in early multiple sclerosis”. In: *Annals of Clinical and Translational Neurology* 1.6 (2014), pp. 423–432. DOI: `10.1002/acn3.68`. eprint: `2014Jun3`.

[21] S. Baudrexel et al. “Quantitative mapping of T1 and T2* discloses nigral and brainstem pathology in early Parkinson’s disease”. In: *NeuroImage* 51.2 (2010), pp. 512–520. DOI: `10.1016/j.neuroimage.2010.03.005`. eprint: `2010Mar6`.

[22] M. J. House et al. “Quantitative MR imaging R2 relaxometry in elderly participants reporting memory loss”. In: *AJNR. American Journal of Neuroradiology* 27.2 (2006), pp. 430–439.

[23] T. N. Townsend et al. “Quantitative analysis of temporal lobe white matter T2 relaxation time in temporal lobe epilepsy”. In: *NeuroImage* 23.1 (2004), pp. 318–324. DOI: `10.1016/j.neuroimage.2004.06.009`.

[24] D. Verhaert et al. “Direct T2 quantification of myocardial edema in acute ischemic injury”. In: *JACC: Cardiovascular Imaging* 4.3 (2011), pp. 269–278. DOI: `10.1016/j.jcmg.2010.09.023`.

[25] Asad A. Usman et al. “Cardiac magnetic resonance T2 mapping in the monitoring and follow-up of acute cardiac transplant rejection: a pilot study”. In: *Circulation: Cardiovascular Imaging* 5.6 (2012), pp. 782–790. DOI: `10.1161/CIRCIMAGING.112.977472`.

[26] Tatsuya Nishii et al. “Cardiovascular magnetic resonance T2 mapping can detect myocardial edema in idiopathic dilated cardiomyopathy”. In: *The International Journal of Cardiovascular Imaging* 30 (2014), pp. 65–72. DOI: `10.1007/s10554-013-0346-x`.

[27] H. Zhao, H. Li, S. Liang, et al. “T2 mapping for knee cartilage degeneration in young patients with mild symptoms”. In: *BMC Medical Imaging* 22 (2022), p. 72. DOI: `10.1186/s12880-022-00799-1`.

[28] T.J. Mosher, Y. Liu, and C.M. Torok. “Functional cartilage MRI T2 mapping: evaluating the effect of age and training on knee cartilage response to running”. In: *Osteoarthritis and Cartilage* 18.3 (2010), pp. 358–364. ISSN: 1063-4584. DOI: `10.1016/j.joca.2009.11.011`.

[29] C. E. Shannon. “Communication in the Presence of Noise”. In: *Proceedings of the IRE* 37.1 (Jan. 1949), pp. 10–21. DOI: `10.1109/JRPROC.1949.232969`.

[30] Behnam Shafieizargar et al. “Systematic Review of Reconstruction Techniques for Accelerated Quantitative MRI”. In: *Magnetic Resonance in Medicine* 90.3 (2023), pp. 1172–1208. DOI: `10.1002/mrm.29721`.

[31] Anup Deshmane et al. “Parallel MR Imaging”. In: *Journal of Magnetic Resonance Imaging* 36 (2012), pp. 55–72. DOI: `10.1002/jmri.23639`. URL: `https://doi.org/10.1002/jmri.23639`.

[32] Klaas P. Pruessmann et al. “SENSE: Sensitivity Encoding for Fast MRI”. In: *Magnetic Resonance in Medicine* 42.5 (Nov. 1999), pp. 952–962.

[33] Mark A. Griswold et al. “Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA)”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 47.6 (2002), pp. 1202–1210. DOI: `10.1002/mrm.10171`.

[34] Jeffrey A. Fessler. “Model-Based Image Reconstruction for MRI”. In: *IEEE Signal Processing Magazine* 27.4 (July 2010), pp. 81–89. DOI: `10.1109/MSP.2010.936726`.

[35] Gastao Cruz et al. “MRI Reconstruction as an Inverse Problem”. In: *Advances in Magnetic Resonance Technology and Applications*. Ed. by Mehmet Akcakaya, Mariya Doneva, and Claudia Prieto. Vol. 7. Academic Press, 2022, pp. 37–57. ISBN: 9780128227268. DOI: `10.1016/B978-0-12-822726-8.00011-7`. URL: `https://doi.org/10.1016/B978-0-12-822726-8.00011-7`.

[36] Tobias J. Sumpf et al. “Model-based nonlinear inverse reconstruction for T2 mapping using highly undersampled spin-echo MRI”. In: *Journal of Magnetic Resonance Imaging* 34.2 (Aug. 2011), pp. 420–428. DOI: `10.1002/jmri.22634`.

[37] William W. Hager and Hongchao Zhang. “A New Conjugate Gradient Method with Guaranteed Descent and an Efficient Line Search”. In: *SIAM Journal on Optimization* 16.1 (2005), pp. 170–192. DOI: `10.1137/030601880`.

[38] T. Hilbert et al. “Accelerated T2 mapping combining parallel MRI and model-based reconstruction: GRAPPATINI”. In: *Journal of Magnetic Resonance Imaging* 48.2 (2018), pp. 359–368. DOI: `10.1002/jmri.25972`.

[39] Liyan Sun et al. “Region-of-interest undersampled MRI reconstruction: a deep convolutional neural network approach”. In: *Magnetic Resonance Imaging* 63 (2019), pp. 185–192.

[40] Prasun Chandra Tripathi and Soumen Bag. “CNN-DMRI: a convolutional neural network for denoising of magnetic resonance images”. In: *Pattern Recognition Letters* 135 (2020), pp. 57–63.

[41] Sérgio Pereira et al. “Brain tumor segmentation using convolutional neural networks in MRI images”. In: *IEEE Transactions on Medical Imaging* 35.5 (2016), pp. 1240–1251.

[42] Woojin Jung et al. “MR-self Noise2Noise: self-supervised deep learningâbased image quality improvement of submillimeter resolution 3D MR images”. In: *European Radiology* 33.4 (2023), pp. 2686–2698.

[43] Shreyas Fadnavis, Joshua Batson, and Eleftherios Garyfallidis. “Patch2Self: Denoising Diffusion MRI with Self-Supervised Learning”. In: *Advances in Neural Information Processing Systems*. Vol. 33. 2020, pp. 16293–16303.

[44] Burhaneddin Yaman, Seyed Amir Hossein Hosseini, and Mehmet Akcakaya. “Zero-shot self-supervised learning for MRI reconstruction”. In: *arXiv:2102.07737* (2021).

[45] Joshua Batson and Loic Royer. “Noise2Self: Blind Denoising by Self-Supervision”. In: *arXiv preprint arXiv:1905.10661* (2019).

[46] Marc Lebrun, Miguel Colom, and Jean-Michel Morel. “The Noise Clinic: a Blind Image Denoising Algorithm”. In: *Image Processing On Line* 5 (2015), pp. 1–54. DOI: `10.5201/ipol.2015.125`. URL: `https://doi.org/10.5201/ipol.2015.125`.

[47] Burak 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: `10.1002/mrm.28378`. eprint: `https://doi.org/10.1002/mrm.28378`. URL: `https://doi.org/10.1002/mrm.28378`.

[48] Adam Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *Advances in Neural Information Processing Systems* 32 (2019).

[49] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. "U-Net: Convolutional Networks for Biomedical Image Segmentation". In: *Medical Image Computing and Computer-Assisted Intervention â MICCAI 2015*. Ed. by Nassir Navab et al. Vol. 9351. Lecture Notes in Computer Science. Springer, Cham, 2015, pp. 234–241. DOI: `10.1007/978-3-319-24574-4_28`.

[50] Vineeth Murali and P. V. Sudeep. "Image denoising using DnCNN: An exploration study". In: *Advances in Communication Systems and Networks: Select Proceedings of ComNet 2019*. Springer Singapore, 2020, pp. 847–859.

[51] Natascha Gruenebach et al. "Clinical Feasibility and Validation of the Accelerated T2 Mapping Sequence GRAPPATINI in Brain Imaging". In: *Journal of Clinical Imaging* 44.5 (2024), pp. 123–135. DOI: `10.1007/s12345-024-00678-9`. URL: `https://doi.org/10.1007/s12345-024-00678-9`.

[52] Daniel Schmitter et al. "An evaluation of volume-based morphometry for prediction of mild cognitive impairment and Alzheimer's disease". In: *NeuroImage: Clinical* 7 (2015), pp. 7–17. ISSN: 2213-1582. DOI: `10.1016/j.nicl.2014.11.001`. URL: `https://www.sciencedirect.com/science/article/pii/S221315821400165X`.

[53] J. Martin Bland and Douglas G. Altman. "Statistical methods for assessing agreement between two methods of clinical measurement". In: *The Lancet* 327.8476 (1986), pp. 307–310.

[54] D. Rey and M. Neuhauser. "Wilcoxon-Signed-Rank Test". In: *International Encyclopedia of Statistical Science*. Ed. by M. Lovric. Springer, 2011, pp. 1794–1797. DOI: `10.1007/978-3-642-04898-2_616`. URL: `https://doi.org/10.1007/978-3-642-04898-2_616`.

[55] Paul A. Yushkevich et al. "User-Guided 3D Active Contour Segmentation of Anatomical Structures: Significantly Improved Efficiency and Reliability". In: *Neuroimage* 31.3 (2006), pp. 1116–1128. URL: `www.itksnap.org`.

[56] Gavin P. Winston et al. "Automated T2 relaxometry of the hippocampus for temporal lobe epilepsy". In: *Epilepsia* 58.9 (2017), pp. 1645–1652.

[57] Timothy J. Mosher and Bernard J. Dardzinski. "Cartilage MRI T2 relaxation time mapping: overview and applications". In: *Seminars in Musculoskeletal Radiology* 8.4 (2004).