

Deep Learning-Based Image Optimization for Low-Field MRI

Dissertation

zur Erlangung des akademischen Grades

**Doktoringenieur
(Dr.-Ing.)**

von M.Sc. David Schote

geb. am 30.07.1995 in Braunschweig

genehmigt durch die Fakultät für Elektrotechnik und Informationstechnik
der Otto-von-Guericke-Universität Magdeburg

Gutachter:

Prof. Dr. rer. nat. Georg Rose

Prof. Dr. rer. nat. habil. Oliver Speck

Prof. Dr. Mikko Nissi

Promotionskolloquium am 11. September 2025

Abstract

Point-of-care low-field MRI systems become increasingly interesting in modern medicine because of their small footprint and cost-effective profile. The portability of such devices poses new challenges, as they are often operated under uncontrolled and changing environmental conditions. The low magnetic field (50 mT) and the high B_0 field inhomogeneities, in combination with temperature fluctuations, and electromagnetic interferences culminate in a poor signal-to-noise ratio and severe image artifacts. The current state of research lacks suitable and robust correction methods that can be implemented directly into point-of-care devices using high-performance computing systems.

In this work, the combination of a novel, highly adaptable MR console and deep-learning based image reconstructions is presented and evaluated to boost the imaging performance of point-of-care low-field MRI devices. The console shows comparable in vivo imaging results to state-of-the-art proprietary systems with additional features such as a fully transparent acquisition process, integration of auxiliary sensors and direct application of deep-learning based image reconstruction techniques.

By combining data driven image denoising and B_0 field prediction with model-based image reconstruction, a physics-informed end-to-end reconstruction network is introduced and evaluated for low-field MRI. The proposed method outperforms conventional reconstruction techniques based on the training and evaluation with simulated low-field MRI data. However, when the model is applied to real in vivo low-field MRI data, considerable challenges are identifiable. To address these challenges, a conventional regularization approach is combined with the estimation of regularization parameter maps by a neural network in an iterative approach in order to optimize image denoising. Thereby, the applicability of the neural network to real low-field MRI data could be improved. Finally, it could be shown that a model trained with supervision can be adapted to other data by applying a subsequent self-supervised refinement step. By integrating the methods and strategies presented in this work, low-field MRI image quality is successfully improved, enhancing its diagnostic value for future clinical applications.

Zusammenfassung

Niedrigfeld MRT Systeme für die Point-of-Care Anwendung erfreuen sich aufgrund ihres geringen Platzbedarfs und ihrer kostengünstigen Nutzung an zunehmendem Interesse in der modernen Medizin. Die Mobilität derartiger Systeme birgt jedoch neue Herausforderungen, da sie unter anderem bei veränderlichen und unkontrollierten Umgebungsbedingungen betrieben werden. Das niedrige Magnetfeld (50 mT) und die hohen B_0 -Feldinhomogenitäten in Kombination mit Temperaturschwankungen und elektromagnetischen Interferenzen resultieren in einem schlechten Signal-zu-Rausch-Verhältnis und starken Bildartefakte. Im aktuellen Stand der Forschung fehlen passende und robuste Korrekturverfahren, die mit leistungsfähigen Computersystemen direkt in Point-of-Care-Geräte implementiert werden können.

In dieser Arbeit, wird eine neuartige, anpassbare MR Konsole in Kombination mit Deep Learning-basierter Bildrekonstruktion vorgestellt und evaluiert, um die Bildgebungsleistung von Point-of-Care-Niedrigfeldsystemen zu optimieren. Die Konsole erlaubt In-vivo Bildgebung, die mit proprietären Systemen auf dem neuesten Stand der Technik vergleichbar ist, und ermöglicht darüber hinaus zusätzliche Funktionen wie einen vollständig transparenten Akquisitionsprozess, die Integration zusätzlicher Sensoren sowie die direkte Anwendung Deep-Learning-basierter Bildrekonstruktionstechniken.

Durch die Kombination von datengetriebener Rauschenunterdrückung und B_0 -Feldschätzung mit modellbasierter Bildrekonstruktion wird ein physikalisch-informiertes End-to-End-Rekonstruktionsnetzwerk für die Niedrigfeld-MRT vorgestellt und evaluiert. Die vorgeschlagene Methode übertrifft konventionelle Rekonstruktionsverfahren basierend auf Training und Evaluation mit simulierten Niedrigfeld-MRT-Daten. Bei der Anwendung auf reale in vivo Niedrigfeld-MRT-Daten, wurden jedoch erhebliche Herausforderungen identifiziert. Um diese Herausforderungen zu adressieren, wird ein alternativer Ansatz untersucht, der konventionelle Regularisierung mit der Schätzung von Regularisierungsparameterkarten durch ein neuronales Netzwerk in einem iterativen Ansatz kombiniert, um die Rauschenunterdrückung zu optimieren. Dadurch konnte die Übertragbarkeit des neuronalen Netzes auf reale Niedrigfeld-MRT-Daten verbessert werden. Schließlich konnte gezeigt werden, dass ein mit Supervision trainiertes Modell durch einen nachfolgenden selbstüberwachten Verfeinerungsschritt an andere Daten angepasst werden kann. Durch die Integration der in dieser Arbeit vorgestellten Methoden und Strategien wird die Bildqualität der Low-Field-MRT erfolgreich verbessert, wodurch ihr diagnostischer Wert für zukünftige klinische Anwendungen gesteigert wird.

Contents

List of Acronyms	3
1 Introduction	13
1.1 Affordable low-field MRI	13
1.2 Overcoming limitations by deep learning	14
1.3 Scope of this thesis	16
2 Theoretical background	19
2.1 Magnetic resonance imaging	19
2.2 Challenges in low-field MRI	22
2.3 Conventional image reconstruction for low-field MRI	24
2.4 AI-based image reconstruction and optimization	27
3 Development of the Nexus console	33
3.1 Architecture description of the console	34
3.2 Experimental setup	40
3.3 Results	47
3.4 Discussion	53
3.5 Summary	55
4 Image and B_0 field reconstruction	57
4.1 Formulation of the joint reconstruction problem	58
4.2 Neural network architecture	62
4.3 Computational experiments	66
4.4 Results	71
4.5 Discussion	78
4.6 Summary	80
5 Application of image and B_0 field reconstruction in vivo	81
5.1 Method adaption	81
5.2 Computational experiments	82
5.3 Model training	82
5.4 Model testing	83
5.5 Results	84

5.6	Discussion	90
5.7	Summary	92
6	Learning regularization parameter maps	93
6.1	Denoising network	94
6.2	Computational experiments	96
6.3	Results	99
6.4	Discussion	106
6.5	Summary	106
7	Model adaptation by test time training	109
7.1	Two-stage training approach	110
7.2	Computational experiments	112
7.3	Results	114
7.4	Discussion	119
7.5	Summary	121
8	Discussion and conclusion	123
	Bibliography	129
A	Appendix	147
A.1	Sequence calculation and execution with the Nexus console	148

List of Acronyms

ADC	Analog-to-Digital Converter.
AWG	Arbitrary Waveform Generator.
CG	Conjugate Gradient.
CIC	Cascaded-Integrator-Comb.
CNN	Convolutional Neural Network.
DAC	Digital-to-Analog Converter.
EDITER	External Dynamic InTerference Estimation and Removal.
EMI	Electromagnetic Interference.
ESP	Echo Spacing.
ETL	Echo-train length.
FFT	Fast Fourier Transform.
FID	Free Induction Decay.
FIR	Finite Impulse Response.
FOV	Field-of-View.
GPA	Gradient Power Amplifier.
GPIO	General Purpose Input/Output.
GPU	Graphics Processing Unit.
IQR	Interquartile Range.
ISMRM	International Society for Magnetic Resonance in Medicine.
ISMRMRD	ISMRM Raw Data.
LUMC	Leiden University Medical Center.
MAE	Mean Absolute Error.

MR	Magnetic Resonance.
MRI	Magnetic Resonance Imaging.
MSE	Mean Squared Error.
NEMA	National Electrical Manufacturers Association.
NMR	Nuclear Magnetic Resonance.
NN	Neural Network.
NRMSE	Normalized Root Mean Squared Error.
ORN	Off-Resonance Network.
OSI ²	Open-Source Imaging Initiative.
PD	Proton Density.
PDHG	Primal Dual Hybrid Gradient.
PE	Phase Encoding.
POC	Point-of-Care.
PSNR	Peak Signal-to-Noise Ratio.
PTB	Physikalisch-Technische Bundesanstalt.
RARE	Rapid acquisition with relaxation enhancement.
RF	Radio Frequency.
RFPA	RF Power Amplifier.
RMSE	Root Mean Squared Error.
RO	Readout.
ROI	Region of Interest.
SDR	Software Defined Radio.
SH	Spherical Harmonics.
SNR	Signal-to-Noise Ratio.
SSIM	Structural Similarity.
T/R	Transmit/Receive.
TE	Echo Time.
TR	Repetition Time.
TSE	Turbo Spin-Echo.
TV	Total Variation.

List of Figures

2.1	Simulation of factors contributing to differences between 3 T and 50 mT low-field MRI: (a) 3 T image, (b) downsampled to typical 50 mT resolution, (c) geometric distortions due to the B_0 field, (d) low SNR, (e) a narrow receive profile of the RF coil.	24
3.1	Halbach-based MRI scanner (~ 50 mT) at PTB with installed gradient coils, the main MR coil (top left) and additional noise sensing coils (left). GPA with power supply, RFPA, T/R switch with pre-amplifier and the Nexus console are installed below the magnet. A detailed view of the Nexus console (right) shows the 8-channel receive and the 4-channel transmit module with the GPIO expander. (Adapted from [96] CC BY 4.0, 2025). .	35
3.2	Schematic of the Nexus console showing the flow of information from the sequences defined by <i>PyPulseq</i> to the acquisition raw data in ISMRMRD format. The Python-based console software interpretes the MRI sequences and replays gradient and RF waveforms by the AWG module (transmit device). A connection from AWG to digitizer module (receive device), enables precise sampling of the pre-amplified MR signal. (Adapted from [96] CC BY 4.0, 2025).	37
3.3	Comparison of the calculated and replayed sequence waveforms at the Nexus console (Reproduced from [96] CC BY 4.0, 2025).	42
3.4	Different phase encoding techniques within an echo train of the RARE sequence to obtain PD and T_2 weighted image contrasts (adopted from [54]). The sub-figures plot the signal intensity over the PE location in k -space.	43
3.5	In vivo brain images of a healthy volunteer acquired with the Nexus console on a 50 mT Halbach-based system. For each contrast, 15 image slices of a 3D TSE sequence with $240 \times 200 \times 210$ mm (RO \times PE1 \times PE2) FOV and $2 \times 2 \times 5$ mm resolution are displayed. (Reproduced from [96] CC BY 4.0, 2025).	48

- 3.6 Comparison of the image quality obtainable by a three-dimensional TSE sequence across consoles. The proposed Nexus console is compared to a reference console which is commercially available. SNR is calculated from signal (S) and noise (N) regions as denoted in the images, and reported at the bottom of each image. The profiles below compare the intensity values along the diagonals in the images. (Adaptated from [96] CC BY 4.0, 2025). 49
- 3.7 Comparison of the image reconstruction (top row) with full (left) and a truncated (right) k -space (bottom row). The truncation and reconstruction has been done in 3D along the readout dimension. Each readout is truncated by 12 samples in the beginning and 12 samples at the end. For illustration, the magnitude data of a single slice is shown. 50
- 3.8 Monitoring results of RF and gradient waveforms, demonstrating system characterization with the console. Left: Calculated RF waveform (black) compared to measurements using a pickup coil (purple), with corresponding forward (blue) and reflected (yellow) power at the RFPA. Right: Comparison of calculated (black) and GPA-monitored (purple) gradient waveforms. (Adapted from [96] CC BY 4.0, 2025). 51
- 3.9 Experimental results for EMI mitigation using three additional noise sensing coils with the Nexus console. Compared to the image obtainable without any interferences (a), single frequency (b) and broad band (c) EMI could be mitigated successfully using the EDITER algorithm [88]. (Adapted from [96] CC BY 4.0, 2025). 52
- 3.10 Positioning of the fiber Bragg grating temperature sensors (T1-T7) and the NMR probe (left). Measurement of the room temperature increase and the correlated decrease of the B_0 field strength at different stages: System on (Sys. on), subject in the scanner (Sub. in), execution of a TSE sequence (TSE), and subject removed (Sub. out). (Adapted from [96] CC BY 4.0, 2025). 53
- 4.1 SH-Net to estimate a field map prior ω_{NN} from a complex-valued B_0 -uninformed FFT reconstruction. The network is built of an encoding path with two convolutional layers per stage and a constant number of 16 filters. The last layer is a fully connected layer that maps the values to 25 SH coefficients, which are used to calculate a spatially smooth field map by (4.24). (Reproduced from [121] CC BY 4.0, 2024). 63

4.2	Overview of the end-to-end reconstruction network. The initial NNs estimates of the B_0 field map and the denoised image are obtained from B_0 -uninformed FFT reconstruction. Using $(\mathbf{E}_{\omega_{\text{NN}}}^L)^{\mathbf{H}}$ the noisy data is reconstructed under consideration of the estimate ω prior to denoising. Unrolling the end-to-end reconstruction network yields the final results \mathbf{x}^T and ω^T . Although the CNNs SH_{Ω} and u_{Θ} are applied only once, their parameters as well as the regularization parameters λ_{ω} and $\lambda_{\mathbf{x}}$ are optimized during training. (Reproduced from [121] CC BY 4.0, 2024).	65
4.3	Comparison of SH-Net and ORN on the test dataset after pretraining. (Reproduced from [121] CC BY 4.0, 2024).	72
4.4	Comparison of different metrics during further fine-tuning of the pretrained u_{Θ} and SH_{Ω} by unrolling the proposed scheme in Algorithm 1 with SP-A, $T = 1$ and $T = 3$ alternations. For SP-A, the field map is entirely defined by the output of the SH-Net and only the image is ensured to be data-consistent by solving problem (4.12). (Reproduced from [121] CC BY 4.0, 2024). . .	73
4.5	Boxplot comparison of image RMSE, PSNR and SSIM for B_0 -uninformed FFT reconstruction, joint-model, bi-model, and end-to-end trained unrolled optimization with $T = 3$. (Reproduced from [121] CC BY 4.0, 2024). . . .	75
4.6	Comparison of reconstruction results from B_0 -uninformed FFT or phase difference map, joint-model, bi-model and the proposed end-to-end trained unrolled optimization. The first and the third row show the magnitude images or the field maps respectively for noise levels $\sigma_y = 0.3$ and $\sigma_y = 0.6$. Rows two and four depict the absolute error map with respect to the ground truth image or the field map respectively. (Reproduced from [121] CC BY 4.0, 2024).	77
5.1	Photograph of the ACR phantom with a diameter of 203 mm and a height of 173 mm, mounted in a solenoid RF coil inside the OSI ² One low-field MRI scanner. The phantom contains a ring and a grid structure as can be seen on the photograph.	83
5.2	Image and field map results of the joint optimization with intermediate NN results for $T = 3$ alternations of the two sub-problems.	86
5.3	Evaluation of input, NN and optimized complex-valued images from the test dataset with respect to the ground truth.	87
5.4	Image magnitude and B_0 field map reconstruction of the joint-model reconstruction with $T = 3$ compared to B_0 -uninformed FFT reconstruction, when applied to the PD-weighted data from an ACR phantom measurement at PTB. The closeup of the circular structure illustrates the correction capabilities of geometric distortion.	87

5.5	Image magnitude and B_0 field map from the joint-model reconstruction with $T = 3$ compared to FFT of the unprocessed k -space data for two different image slices of the T_2 -weighted in vivo brain data, acquired at LUMC. The cutout below each image sample shows a zoomed in section for a more detailed comparison.	89
5.6	Image magnitude and B_0 field map from the joint-model reconstruction with $T = 3$ compared to FFT of the unprocessed k -space data for image slice 16 of the PD-weighted in vivo brain data, acquired at LUMC. The cutout below each image sample shows a zoomed in section for a more detailed comparison.	90
6.1	Overview of the network architecture facilitating regularization parameter map estimation by CNN and unrolling of iterative denoising scheme with TV regularization.	96
6.2	Simulation process to obtain pairs of noise-free ground truth and noisy low-field MRI samples: Noise standard deviation and coil profile are estimated from real data. A smooth coil profile is fitted to the readout profiles derived from pure noise measurements along the PE dimension, using a 4th-degree polynomial before applying it to the simulated low-field MRI image. . . .	98
6.3	Convergence of training and validation ℓ_1 -loss over epochs (left) and convergence of λ^{xyz} , λ^{xy} and λ^z over epochs during training (right).	100
6.4	Evaluation of NRMSE and PSNR for the denoising performance on the test dataset comparing λ^{xyz} , λ^{xy} and λ^z , and the proposed Λ_Θ regularization.	101
6.5	Comparison of reconstruction results for a representative sample from the test data set of simulated low-field images with T_2 -weighted contrast. The first row compares the magnitude images of the input, i.e. $\mathbf{x}_{\text{Noisy}}$, regularization with λ^{xyz} , the combination of λ^z and λ^{xy} , and the map of regularization parameters against the ground-truth. The second row shows the regularization parameter maps Λ_Θ^{xy} and Λ_Θ^z , which allow interpreting the impact of the model.	102
6.6	Reconstruction of the T_2 weighted in vivo results obtained in Chapter 3 using the unrolled network $f_\Theta^T(\mathbf{x}_{\text{Noisy}})$. For each slice, the original noisy input image, and the denoised model output is shown along with the regularization parameter maps Λ_Θ^{xy} and Λ_Θ^z	104
6.7	Reconstruction of the PD-weighted in vivo results obtained in Chapter 3 using the unrolled network $f_\Theta^T(\mathbf{x}_{\text{Noisy}})$. For each slice, the original noisy input image, and the denoised model output is shown along with the regularization parameter maps Λ_Θ^{xy} and Λ_Θ^z . The training was performed on simulated low-field data with T_2 weighted contrast.	105

7.1	The two-stage training approach integrates an initial phase of supervised learning (top) with a subsequent phase of unsupervised refinement (bottom). In this framework, the pretrained model weights undergo optimization for individual samples utilizing the unsupervised <i>Noise2Self</i> training strategy. The objective is to exploit supervised learning during the testing phase through adaptations of the denoising model tailored to individual samples.	110
7.2	Denoising example for simulated low-field MRI data obtained from the U-Net architecture, when trained by supervised learning. The example shows the center slice of a three-dimensional sample from the test dataset which originates from the same distribution as the training and validation data. The model output in the second column demonstrates the performance of the supervised training.	115
7.3	Comparison of supervised training (ST) with distributional discrepancy, self-supervised test-time training (TTT) and supervised pretraining with self-supervised test-time training (SPT + TTT) across the test data by NRMSE and PSNR. The results correspond to the observations in Figure 7.4.	116
7.4	Comparison of supervised training (ST) with distributional discrepancy, self-supervised test-time training (TTT) and supervised pretraining with self-supervised test-time training (SPT + TTT) on a set of simulated T ₁ weighted test data.	117
7.5	Denoising results for T2-weighted 3D in vivo brain images obtained by TSE sequence with TE/TR of 20/2000 ms. The FOV of the data is $240 \times 200 \times 210$ mm with a resolution of $2 \times 2 \times 5$ mm for RO \times PE1 \times PE2. Comparison of supervised training (ST), self-supervised test-time training (TTT) and self-supervised test-time training with supervised pretraining (SPT + TTT).	118
7.6	Denoising results for PD-weighted 3D in vivo brain images obtained by TSE sequence with TE/TR of 20/500 ms. The FOV of the data is $240 \times 200 \times 210$ mm with a resolution of $2 \times 2 \times 5$ mm for RO \times PE1 \times PE2. Comparison of supervised training (ST), self-supervised test-time training (TTT) and self-supervised test-time training with supervised pretraining (SPT + TTT).	120
8.1	Comparison of magnitude images from in vivo brain data acquired with the Nexus console: Images are obtained through FFT reconstruction, joint image and B ₀ field reconstruction, and FFT reconstruction with denoising using estimated regularization parameter maps and the two-stage training strategy. The reference-free blur metric and SNR is calculated for each image.	125

A.1	Diagram which illustrates the sequence calculation by the implemented sequence provider within the Nexus console. The gradient and RF waveforms of a <i>Pulseq</i> sequence are calculated dependent on adjustable acquisition parameters and a static device configuration. Each sample point of the gradient waveform encodes a digital signal, replayed synchronously with the analog waveform. The computed sequence is stored as a 16-bit <i>Numpy</i> array in <i>Fortran</i> order.	148
A.2	Schematic of the sequence execution in the Nexus console after computation, utilizing the implemented transmit and receive device classes. The sequence is streamed to the AWG card's memory and subsequently replayed. Whenever new timestamp data becomes available, the digitizer card transfers it from the card to the buffer. Digital post-processing includes demodulation, decimation, and phase correction.	149

List of Tables

4.1	Median and IQR of MSE, PSNR and SSIM for the reconstructed image \mathbf{x} with joint-model, bi-model, and unrolled optimization $T = 3$. The second table shows the MAE for the field map of the corresponding methods. The data is shown for the noise levels given by $\sigma_y \in \{0.2, 0.4, 0.6\}$. (Reproduced from [121] CC BY 4.0, 2024).	75
-----	--	----

1 Introduction

1.1 Affordable low-field MRI

Magnetic Resonance Imaging (MRI) is widely recognized as one of the most powerful diagnostic imaging modalities available in modern medicine. Its key benefits stem from its capacity to generate high-contrast cross-sectional images of soft tissues, without the use of ionizing radiation. This non-invasive property makes MRI an ideal choice for diagnosing and monitoring a wide range of conditions, ranging from neurological disorders to musculoskeletal injuries and cardiovascular diseases [1]. Beyond anatomical details, MRI also enables functional assessments, such as diffusion imaging, providing insights into physiological processes [2]. This versatility has established MRI as the gold standard in numerous clinical and research settings.

Since its early days, MRI has been dominated by high-field systems and applications [3]. However, high-field systems lack accessibility, require a lot of infrastructure, and are expensive [4], [5]. These key challenges, combined with technical advances made in recent decades, contribute to the growing interest in low-field MRI systems. The definition of a low-field MRI system is ambiguous, because the terminology is used by a wide range of field strength between 10 mT and 1 T [3]. In recent years, a notable shift has occurred towards 0.55 T systems for clinical applications, while current research focuses on the more compact Point-of-Care (POC) systems with field strengths less than 100 mT [3], [6]–[10], for which the terminology *low-field* is used in the context of this thesis. Instead of superconductors, which require extensive infrastructure and maintenance, permanent magnets can be used to generate the static magnet field of a low-field MRI scanner. A major advantage of such systems is their portability: smaller magnets and a compact footprint allow these systems to be deployed in decentralized settings such as intensive care units [11], [12], at bedside [13]–[16], in ambulances [17], or even in patients’ homes [18]. Due to reduced field strength, systems are inherently safer. This facilitates access for implants patients and makes the system interesting for intraoperative use [19]. Open-source low-field systems have emerged, whose design encourages innovation, customization, and cost effectiveness, with the ability to reduce the material costs of a fully functional MRI scanner to under 50 000 € [20]–[22]. This development has the potential to truly democratize MRI, particularly

in resource-limited environments, regardless of infrastructural or economic constraints [23]–[25].

The clinical questions that are targeted by most of the low-field MRI systems include neurological issues, stroke assessment, subdural or epidural hematoma, hydrocephalus, extra ventricular drain placement/shunt assessment, or elevated intracranial pressure [25], [26]. In addition to the brain, the extremities are also a major point of interest [18], [27], [28]. Most recently, a whole body low-field scanner has been presented, which extends the field of application to cardiac and lung imaging, among others, which broadens the clinical scope of low-field MRI [29].

Although low-field MRI presents numerous advantages, its renaissance has also introduced significant technical challenges that need to be addressed to achieve clinical viability and reliability [3], [30], [31]:

- Inherently lower Signal-to-Noise Ratio (SNR) poses challenges in detecting subtle anatomical details.
- Higher relative magnetic B_0 field inhomogeneities.
- Operation in uncontrolled environments introduces temperature fluctuations and various sources of Electromagnetic Interference (EMI) that can severely compromise images.
- Differences in relaxivity at lower field strengths generate different image contrasts compared to conventional systems.

Overcoming these issues is an active field of research, as they can have an impact on diagnostic findings [32]. Considerable advances in artificial intelligence and its application to image reconstruction and optimization provide promising solutions to these challenges.

1.2 Overcoming limitations by deep learning

In recent years, deep learning-based approaches have emerged as promising solutions to enhance image quality in MRI. More specifically, most methods address artifact and bias correction, noise reduction, or resolution enhancement across all different field strengths [33]. With regard to low-field MRI, various deep learning-based techniques have been developed, targeting the challenges mentioned above.

Most of the methods presented for low-field MRI employ deep Neural Networks (NNs) trained with supervision to enhance image quality by denoising. Training such models

usually requires pairs of noisy and clean images for effective noise removal. However, noise-free, undistorted low-field reference data is technically unfeasible to obtain. Consequently, low-field MRI images are often retrospectively simulated from high-field MRI data, such as the 3 T *FastMRI* dataset [34], serving as a noise-free reference [35]. In [36], a k -space to image deep learning model is proposed which reconstructs noise-free images from noisy low-field k -space data at 6.5 mT and 47 mT. Supervised training is conducted with pairs of clean and noisy samples simulated with a dataset acquired at 3 T. Similarly, the authors of [37] employ a denoising autoencoder architecture which is trained on 0.35 T images using 1.5 T data which is retrospectively deteriorated by additive Rician noise. The simulation of training data based on high-field images can be extended by involving quantitative parameter maps, as shown by the authors of [38] which propose a denoising network for a 47 mT Halbach system. However, for supervised training, this requires access to an extensive quantitative dataset. Beyond pure denoising applications, extensions of the NNs for undersampled data at 0.1 T and 55 mT [39], [40] and increased resolution [35], [41] at 50 mT have been proposed. After training with simulated data, methods are often directly transferred to the application on actual low-field MRI data, showing promising results. However, due to the lack of a ground truth image, a quantitative evaluation is challenging.

Low-field MRI denoising techniques based on supervised learning assume that the data during training and the data at model inference is similar. That is, the domain shift between the simulated training data and actual low-field MRI data is sufficiently small. Although many approaches show promising results, this assumption is never entirely true [42], [43]. As an alternative to supervised learning, self-supervised learning methods have been proposed, which exploit inherent structures and redundancies in the data to learn the separation of signal and noise. In the domain of low-field MRI at around 50 mT however, such techniques are not yet very common. A dual-domain self-supervised learning approach for non-cartesian low-field MRI at 64 mT is proposed in [44]. The authors trained a reconstruction network with a loss function combining loss components in image and k -space domain, showing promising results for denoised low-field MRI data.

In addition to poor SNR, high B_0 field inhomogeneous can pose a major issue with low-field MRI systems, as it degrades the image quality by introducing geometric distortions. However, if the B_0 field map is known, the distortions can be corrected in the image reconstruction. For some specific applications at 3 T MRI, supervised methods have been proposed to estimate and correct for B_0 field inhomogeneities using NNs [45], [46]. In [45], the authors proposed a NN that estimates an undistorted image from a complex-valued distorted image. The authors of [46] employ a comparable NN, but estimate the B_0 field map instead of a corrected image, which is used in a subsequent model-based reconstruction step to correct for distortions. Both methods yield promising results, but

they target conventional systems with a field strength of 3 T and have not yet been applied to low-field MRI. At low field strength, the relative B_0 inhomogeneities are usually >1000 ppm and can cause stronger geometric distortions. Therefore, the authors of [47] proposed a method specifically targeting low-field MRI systems at ~ 60 mT. Their approach employs two U-Net models to subsequently denoise and correct for B_0 induced artifacts in the image domain, after reconstruction. Each U-Net is individually trained with supervision using paired images, where the distorted images are retrospectively simulated. However, interpreting the role of the individual U-Net layers, particularly the one for B_0 field distortion correction, remains challenging. The use of such black-box methods can introduce instabilities in image reconstruction, making it sensitive to small perturbations in the input [48]. To enhance the interpretability and stability of NNs, the authors of [49] propose to integrate them in an iterative algorithm and train the network in an end-to-end fashion. Thereby the acquisition model of the data can be incorporated into the training process to constrain the model output and improve the interpretability, robustness, and reliability of the solution. Such an end-to-end-trainable network architecture has been successfully applied to accelerated cine MRI reconstruction at 3 T [50]. Applying a similar approach to low-field MRI, where the acquisition model explicitly accounts for B_0 field inhomogeneities, could help to improve the interpretability and quality of the reconstruction. However, in the current state of research, no such methods have been proposed for low-field MRI.

1.3 Scope of this thesis

Within the scope of this work, deep learning-based techniques are investigated for Halbach-based low-field MRI systems with a B_0 field strength of 50 mT, and particularly for the system proposed by the Open-Source Imaging Initiative (OSI²) called OSI²One [22]. In general, the system has proven its imaging capabilities [7], [51], [52], however, the image quality is limited due to the various technical challenges introduced in Section 1.1. The goal of this thesis is to investigate deep learning techniques which improve image quality and imaging performance of such systems and thereby enhance the clinical value of affordable POC low-field MRI system. Implementing such techniques, necessitates a versatile console with a transparent, adaptable and interpretable image acquisition pipeline. The consoles which have been used so far are based on proprietary hardware and software, limiting their functionality. By governance of the data acquisition pipeline, important prerequisites for the development of deep learning-based reconstruction and optimization methods are met. Thus, the development of a versatile and open console for low-field MRI which enables the application of deep learning-based methods builds the foundation of this work.

Deep learning methods have proven their value for image quality enhancement in low-field

MRI. However, the majority of methods are purely data-driven black box approaches what raises questions about their reliability and robustness, particularly in clinical contexts. Inspired by the results of algorithm unrolling techniques presented for 3 T MRI, the combination of model-based reconstruction with deep learning for image quality enhancement in low-field MRI is investigated for the first time. More specifically, the method is targeting both, image and B_0 field reconstruction which have been identified as the main issues in low-field MRI. The combination of model-based reconstruction with deep learning is expected to yield more robust reconstruction results and constraint impact of NNs.

To this end, the main objectives of the work are defined as follows:

- Development of a versatile and transparent acquisition system for the OSI²One low-field scanner, capable for in vivo measurements and image reconstruction.
- Investigation of combined model-based reconstruction with deep learning for low-field MRI, targeting image and B_0 field reconstruction.
- Evaluation of different training concepts, including supervised learning, end-to-end training with algorithm unrolling, unsupervised learning, and multi-stage training.

After establishing the theoretical foundations in Chapter 2, Chapter 3 presents the development of the Nexus console for advanced low-field MRI acquisitions. It introduces the hardware and software architecture of the console, demonstrates its versatility, and compares acquired in vivo images with those obtainable from a state-of-the-art console at a 50 mT low-field MRI system. The goal is to establish a transparent acquisition and reconstruction framework, in the sense that the acquired data and its processing pipeline from Analog-to-Digital Converter (ADC) samples to raw data are interpretable. Building on these imaging results, Chapter 4 introduces a signal model for low-field MRI and integrates it with NNs in an iterative reconstruction algorithm. By unrolling this algorithm, end-to-end network training is achieved. The network is first evaluated on simulated low-field MRI data before being applied to in vivo scans acquired with the Nexus console in Chapter 5. Chapter 6 explores an alternative regularization approach, while Chapter 7 investigates a multi-stage training strategy to improve model adaptability across different data distributions. Finally, Chapter 8 compares the performance of all proposed methods and concludes the work with an integral discussion.

2 Theoretical background

2.1 Magnetic resonance imaging

MRI is based on the principles of Nuclear Magnetic Resonance (NMR), where the magnetic properties of atomic nuclei are used to generate detailed images of the human body. This section explains the core concepts of MRI, including spin, Larmor frequency, Larmor precession, net magnetization, the Bloch equations, spatial encoding, and the signal equation.

2.1.1 Magnetization

Atomic nuclei with an odd number of protons or neutrons possess a property known as *nuclear spin*, denoted by I . When a static magnetic field B_0 is applied to an ensemble of nuclei, they occupy discrete energy levels depending on their spin value. For hydrogen nuclei ^1H , the spin value is $\frac{1}{2}$ which allows two discrete spin manifestations $\pm\frac{1}{2}$ and two possible energy levels [53], parallel or antiparallel to B_0 .

The difference between the two energy levels defines the resonance condition [1]. In an external magnetic B_0 field, the nuclear magnetic moments experience a torque that causes them to precess around the axis of the magnetic field. The frequency of precession is the resonance frequency ω_0 , named after Joseph Larmor (1857-1942) [54]:

$$\omega_0 = \gamma B_0. \quad (2.1)$$

The gyromagnetic ratio γ is a constant specific to the type of nucleus. For the hydrogen nuclei it is 42.58 MHz/T. Due to its spin properties and abundance in water and biological tissue, hydrogen is the most important nucleus for MRI.

The spin property of nuclei gives rise to a magnetic moment $\boldsymbol{\mu}$, contributing to the magnetization vector of the object under investigation. Within the scope of this work, vectors are denoted by a bold symbol. In a small volume element or voxel V , in which the

magnetic field is approximately constant, the set of spins in V forms an ensemble with a (net) magnetization [53]

$$\mathbf{M} = \frac{1}{V} \sum_{\mu \in V} \boldsymbol{\mu}. \quad (2.2)$$

2.1.2 Bloch equations

The time evolution of the magnetization is described by the *Bloch equation*, which account for both the precession of the magnetization and relaxation processes. After interaction with an external magnetic field $\mathbf{B}_{\text{ext.}}$ it is described by the differential equation [53]

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B}_{\text{ext.}} - \frac{M_x}{T_2} \hat{x} - \frac{M_y}{T_2} \hat{y} - \frac{M_z - M_0}{T_1} \hat{z}. \quad (2.3)$$

Proton interactions are modeled by different time constants for the different components of the magnetization vector \mathbf{M} . M_x and M_y account for the *transverse* components and M_z for the *longitudinal* component, parallel to B_0 . M_0 is the magnetization at equilibrium. In (2.3), the unit vectors \hat{x} and \hat{y} span the transversal plane while \hat{z} points in the direction of the static magnetic B_0 field. The terms T_1 and T_2 are referred to as spin-lattice relaxation time and spin-spin relaxation time, respectively. These terms describe how the magnetization, i.e. the transverse and longitudinal components, returns to equilibrium after being disturbed, e.g. by a temporally applied, external magnetic field $\mathbf{B}_{\text{ext.}}$, generated by an Radio Frequency (RF) pulse. The process is governed by two relaxation mechanisms:

- *Longitudinal relaxation* governed by T_1 , describing the recovery of the z -component of the magnetization.
- *Transverse relaxation* by T_2 which is the the spin-spin relaxation decay. However, effectively the transverse relaxation is governed by T_2^* , which is the superposition T_2 and the B_0 field inhomogeneity induced dephasing by T_2' in the xy -plane. It is defined by $\frac{1}{T_2^*} = \frac{1}{T_2'} + \frac{1}{T_2}$.

These relaxation times influence image contrast and are fundamental in MRI pulse sequences. It should be noted, that this is a simplified yet effective model, which is sufficient for the requirements of this work.

2.1.3 Signal equation

To create an image, spatial information must be encoded in the MRI signal. This is achieved by superimposing linear magnetic field gradients $\mathbf{G} = [G_x, G_y, G_z]$ which spatially vary B_z . These gradients can be defined by $G_x = \frac{\partial B_z}{\partial x}$, $G_y = \frac{\partial B_z}{\partial y}$, and $G_z = \frac{\partial B_z}{\partial z}$ respectively.

The gradients cause the Larmor frequency to vary linearly with position, allowing different spatial locations to be distinguished on the basis of frequency and phase. According to Maxwell's equations, named after James Clerk Maxwell (1831-1879), the transversal field components B_x and B_y , known as concomitant fields, are also unavoidable included in the gradient field. They cause an additional phase accumulation of the transverse magnetization when a gradient is active [54], [55]. However, for the purpose of this work, their effect is neglected. The frequency domain, containing the trajectories $k(t)$ is called k -space and is defined by:

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(t') dt', \quad (2.4)$$

where $\mathbf{G}(t)$ is the time-varying gradient field.

By Faraday's law, the time-varying (transversal) magnetization in Equation (2.3) induces an electromotive force. It is proportional to the volume integral of the time derivative of the magnetization and results in an electrical potential across the receiving coil. In a basic configuration, a single RF coil is used to transmit the excitation pulse and subsequently receives the Magnetic Resonance (MR) signal. The MR signal is amplified and demodulated at the center frequency ω_0 and is found to be the Fourier transform of the spin density $\rho(\mathbf{r})$ over the imaged volume V :

$$s(t) = \int_V \rho(\mathbf{r}) e^{-i2\pi\mathbf{k}(t) \cdot \mathbf{r}} d\mathbf{r}. \quad (2.5)$$

For the sake of simplicity, the relaxation effects introduced in (2.3) are absorbed into the spin density.

2.1.4 Acquisition process

The core operational component of a MRI system is the console. This unit is responsible for computing the gradient waveforms and RF pulses as delineated in the pulse sequence. The gradient waveforms are subjected to amplification by the Gradient Power Amplifier (GPA), resulting in the generation of transient electromagnetic fields at the gradient coils with frequencies up the kHz range, which serve to spatially encode the MRI signal along three orthogonal dimensions. In order to alter the orientation of the magnetization vector, RF pulses near the resonance frequency ω_0 are employed which depends on the B_0 field strength according to (2.1) and ranges from ~ 2 MHz at low-field system with approximately 50 mT to ~ 128 MHz for conventional clinical 3 T system. They are amplified by a RF Power Amplifier (RFPA) to establish a transverse magnetic field \mathbf{B}_{ext} . Various pulse shapes can be employed to modify the excitation profile of the transverse field, generating transverse magnetization. Immediately after switching off the transmitted RF pulse, the transversal magnetization begins to decay according to Equation 2.3. To facilitate the transition

between two operational modes, a Transmit/Receive (T/R) switch is utilized, which safeguards the pre-amplifier from high peak amplitude RF during transmission mode and ensures seamless amplification of low amplitude signals in the receiving mode.

Processing involves demodulation, quadrature detection, band limiting, sampling, and decimation [54]. Given the phase sensitivity of the MRI signal, it is imperative to maintain phase-consistent modulation and demodulation. Typically, quadrature demodulation is employed to disaggregate the signal into its in-phase (I) and quadrature (Q) components. In an analog receiver design, this procedure entails mixing the received signal with a reference signal at the Larmor frequency, followed by low-pass or band-pass filtering to extract the I and Q components. Hereby, the Larmor frequency, which does not carry any meaningful information, is removed from the signal, which is centered at zero frequency. With the unwanted sidebands removed the analog signal is digitized by the ADC. The digitized signal may be further processed by digital filters and resampling operations, yielding the final raw data.

Alternatively, the entire signal processing pipeline can be implemented using a fully digital design. Demodulation, filtering, and resampling are achieved solely by digital processing of the MR signal [56]. While this approach requires much higher sampling rates and more computational resources, due to the digital signal demodulation, it offers enhanced flexibility and control over the processing stages. At a field strength of 50 mT ($\omega_0 \sim 2$ MHz), oversampling by a factor of 10 to 20 is technically feasible. When designing the digital decimation filter, a key consideration is the prevention of signal loss or distortion. In this context, Finite Impulse Response (FIR) filters have demonstrated great efficacy, particularly in low-field MRI applications [57].

2.2 Challenges in low-field MRI

Conventional clinical MRI scanners are usually operated with field strength at 1.5 T or 3 T. In principle, the fundamental technology does not differ between conventional systems and low-field MRI systems. However, there are some distinctions that have an impact on the image quality and thus are covered within this Section.

2.2.1 Signal-to-noise ratio

In MRI, the static magnetic B_0 field is responsible for aligning the nuclear spins, creating a net magnetization in the tissue. The strength of M_0 is proportional to B_0 , which in turn directly influences the strength of the detectable signal, measured as SNR. SNR is a critical parameter in MRI, as it determines the clarity and quality of the images produced

and directly affects the diagnostic utility of the image. In low-field MRI, the weaker B_0 field significantly reduces net magnetization, leading to a smaller SNR. This reduction in magnetization is problematic because it limits the overall signal available for imaging, contributing to poorer image quality [5].

Noise in a MRI system mainly arises from two sources: the resistive losses of the imaging hardware and the thermal noise of the subject being imaged. The contribution of noise is frequency dependent and at lower frequencies dominated by coil noise. In comparison to high-field systems, the overall magnetization is much lower in low-field MRI leading to less signal and a reduced SNR. Dependent on the B_0 field strength, the SNR scales approximately by $\text{SNR} \sim B_0^{3/2}$ [31]. In addition, POC low-field MRI systems are operated in environments that are less controlled than those used for high-field systems, further compounding the SNR problem. EMI from external sources can introduce additional noise, which is particularly problematic in portable low-field MRI systems. These systems are typically more exposed to environmental noise sources, which degrade the already low SNR even further.

In conclusion, the diminished SNR observed in low-field MRI systems is primarily attributed to the weakened net magnetization resulting from the reduced B_0 field, further compounded by environmental influences and field inhomogeneities, as expounded in the subsequent section. This combination of factors represents a significant obstacle to achieving high quality imaging, especially in clinical contexts where precise anatomical delineation is essential.

2.2.2 Static magnetic B_0 field

As discussed in the previous section, the intrinsically lower B_0 field leads to less SNR. Another critical issue in low-field MRI is the inhomogeneity of the B_0 field. Ideally, the B_0 field should be uniform across the imaging volume to ensure that all spins experience the same magnetic environment. Inhomogeneities in the B_0 field – where the field strength varies spatially – cause spins to precess at slightly different frequencies, resulting in phase dispersion according to equation (2.1). This phase dispersion leads to signal dephasing described by T'_2 , degrading the coherence of the spin system and reducing the measured signal.

The roots of the inhomogeneities in the field B_0 are intrinsic and extrinsic. Intrinsic sources are related to the fundamental properties of the magnet used in the MRI system. In low-field systems, permanent magnets are often employed. Variations in the manufacturing of these magnets, material imperfections, and the geometry of the magnet all contribute to inhomogeneities in the B_0 field. Extrinsic factors, such as temperature fluctuations

of the surrounding environment, can also exacerbate B_0 field inhomogeneities. These external influences can alter the magnetic field distribution, further compromising the homogeneity of the B_0 field. Considering a standard clinical 3 T system, the relative B_0 field inhomogeneities observable in low-field MRI are about 1000 times larger [58]. The absolute field inhomogeneity which effects the $\frac{1}{T_2^*}$ relaxation is therefore increased by a factor of ~ 15 . Figure 2.1 illustrates the combination of image degradation factors, including low SNR as discussed in Section 2.2.1, and the strong B_0 field inhomogeneities.




Figure 2.1: Simulation of factors contributing to differences between 3 T and 50 mT low-field MRI: (a) 3 T image, (b) downsampled to typical 50 mT resolution, (c) geometric distortions due to the B_0 field, (d) low SNR, (e) a narrow receive profile of the RF coil.

2.2.3 Relaxation times

Although relaxation times are affected by the field strength, they introduce secondary effects related to SNR and image contrast. The T_1 relaxation time decreases with the field strength, leading to faster signal recovery at low field strengths. This can be explained by the magnetic coupling between hydrogen and the solid tissue components [51], [59]. The enhanced T_1 relaxation can be beneficial in terms of scan efficiency and image contrast [5]. Due to the larger T_1 value differences in the tissues, the tissue contrast is increased. The T_2 relaxation, which governs the decay of the transverse magnetization, is conversely less affected by the field strength and is often offset by rapid dephasing due to inhomogeneities of the B_0 field, accounted by T_2' .

2.3 Conventional image reconstruction for low-field MRI

As described earlier, the MRI data are acquired in the k -space. To obtain an interpretable representation, the k -space data need to be transformed from the frequency domain to the image domain. This section derives an operator that is suitable for performing the transformation and that considers the impact of the B_0 field.

2.3.1 Problem formulation

Based on the signal equation (2.5) a linear image reconstruction problem can be formulated in matrix-vector notation by

$$\mathbf{y} = \mathbf{A}\mathbf{x}_{\text{True}} + \boldsymbol{\varepsilon}, \quad (2.6)$$

where the system matrix \mathbf{A} can be defined by the Fourier transform $\mathbf{F} : \mathbb{C}^N \rightarrow \mathbb{C}^N$ in the simplest case, that maps the complex-valued clean image $\mathbf{x}_{\text{True}} \in \mathbb{C}^N$ showing the true object to the measured noisy data $\mathbf{y} \in \mathbb{C}^N$. In this context, N equals the number of image voxels or acquired k -space sample points respectively. The measurement error that is modeled by complex-valued, zero mean, additive white Gaussian noise [60], [61], is denoted by $\boldsymbol{\varepsilon}$. Within the scope of this work, only single-coil Cartesian acquisitions are considered as these are usually used in low-field MRI. Instead of \mathbf{F} more sophisticated system matrices can be used, which consider undersampling or partial sampling techniques, where the number of sample points may deviate from the number of image voxels. As introduced in Section 2.2.2, the B_0 field plays a crucial role in low-field MRI and can lead to geometric distortions in the image. The forward problem under consideration of the B_0 field can be extended to

$$\mathbf{y} = \mathbf{E}(\mathbf{x}, \boldsymbol{\omega}) + \boldsymbol{\varepsilon}. \quad (2.7)$$

Here, the forward operator $\mathbf{E} : \mathbb{C}^N \times \mathbb{R}^N \rightarrow \mathbb{C}^N$ maps the (unknown) complex-valued clean image $\mathbf{x} \in \mathbb{C}^N$ and the (unknown) real-valued static magnetic field deviation $\boldsymbol{\omega} \in \mathbb{R}^N$ to a set of measurements $\mathbf{y} \in \mathbb{C}^N$. To simplify the problem, the extended operator is restricted to equispaced, two-dimensional fully sampled single-coil Cartesian acquisitions. The values of the discrete B_0 field vector $\boldsymbol{\omega}$ are given in radians, which equals the frequency deviation in Hertz multiplied by 2π .

The elements of \mathbf{E} can be defined according to the signal equation (2.5) taking into account the magnetic field inhomogeneities. Thus, the time-dependent, discrete signal equation considering field inhomogeneities $\boldsymbol{\omega}$ and image \mathbf{x} for a discrete-time point $\mathbf{t} = [t_1, t_2, \dots, t_N]^T \in \mathbb{R}^N$ is given by

$$y_m = \sum_{n=1}^N x_n e^{-i\omega_n t_m} e^{-ik_m r_n} + \varepsilon_m, \quad m = 1, \dots, N. \quad (2.8)$$

Hereby, $x_n = x(r_n)$ and $\omega_n = \omega(r_n)$ are dependent on the spatial voxel location $\mathbf{r} \in \mathbb{R}^N$, while the acquired signal $\mathbf{y} \in \mathbb{C}^N$, with $y_m := y(t_m)$ and the k -space trajectory $k_m := k(t_m)$ are dependent on the sample time points vector \mathbf{t} . Within the discretization step, the center of a voxel approximation is used implicitly, which is common practice in MRI [54].

2.3.2 Model-based image reconstruction

Assuming the general expression of the forward problem (2.6), \mathbf{x} can be estimated by minimizing the regularized least-squares misfit

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \mathcal{F}(\mathbf{x}), \quad \mathcal{F}(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}(\mathbf{x}) - \mathbf{y}\|_2^2 + \alpha \mathcal{R}(\mathbf{x}). \quad (2.9)$$

Since $\boldsymbol{\varepsilon}$ is assumed to be uncorrelated white noise with uniform variance, the Euclidean norm is appropriate [62]. For fully sampled Cartesian single-coil imaging with uniform coil sensitivity and ignored relaxation and B_0 inhomogeneity, the system matrix \mathbf{A} is orthogonal and its inverse is given by $\mathbf{A}^{-1} = \mathbf{A}^\mathbf{T}$. If the acquired data would be noise-free, or if $\boldsymbol{\varepsilon}$ is sufficiently small, regularization would not be required and the solution is given by $\mathbf{A}^\mathbf{T}\mathbf{y}$. However, because of the poor SNR in low-field MRI, the problem is ill-posed and regularization is required to obtain a denoised solution close to \mathbf{x}_{True} . In general, the regularization term $\mathcal{R}(\mathbf{x})$ introduces prior knowledge about \mathbf{x} that can be used to improve the quality of the estimate $\hat{\mathbf{x}}$. Usually, the regularizer is weighted by a regularization parameter α that balances the data fidelity term and the regularization term. The optimal choice of regularizer $\mathcal{R}(\mathbf{x})$ does not exist and remains an active research field. Different approaches exist, where two prominent examples that are also relevant for this work, are Tikhonov [63] and Total Variation (TV) [64], [65] regularization. Tikhonov regularization is defined by $\mathcal{R}(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}'\|$, where \mathbf{x}' is a prior or reference of \mathbf{x} that can potentially also be defined as zero. The disadvantage of this regularization method is the bias towards the prior or the reduced image contrast in $\hat{\mathbf{x}}$ if the prior is zero. TV regularization in turn enforces piecewise smoothness in the image and is defined by $\mathcal{R}(\mathbf{x}) = \|\nabla \mathbf{x}\|_1$, where ∇ is the gradient operator. It is particularly useful for preserving edges and fine details in the image.

To find a solution $\hat{\mathbf{x}}$ for which the functional $\mathcal{F}(\mathbf{x})$ in (2.9) is a minimum, the derivative of the functional $\nabla \mathcal{F}(\mathbf{x})$ is enforced to be zero, i.e. solving $\nabla \mathcal{F}(\mathbf{x}) \stackrel{!}{=} 0$. Therefore, often iterative methods are employed in practice, dependent on the forward problem. One approach is the Conjugate Gradient (CG) method, which iterates from an initial guess and updates the solution in directions that are mutually conjugate with respect to $\mathbf{A}^\mathbf{H}\mathbf{A}$ [66]. This process continues until a convergence criterion is met, yielding a solution to the problem, that is, the reconstructed image. Another approach is the Primal Dual Hybrid Gradient (PDHG) algorithm, also known as the Chambolle–Pock algorithm [67]. It aims to solve convex problems, particularly those with a saddle-point structure, using a primal-dual construct. The primal problem aims to solve for the unknowns in an optimization task, while the dual problem is derived from it, optimizing the trade-offs between the objective and constraints of the primal. The algorithm updates the primal and dual components in

an alternating pattern and terminates when a desired accuracy or stopping criterion is reached.

2.4 AI-based image reconstruction and optimization

The conventional model-based approaches for image reconstruction and denoising introduced in the last section rely on iterative methods to optimize data-consistency terms with regularization. These methods can be computationally complex or suboptimal in capturing complex patterns in the acquired data. For many different applications, such as for low-field MRI, deep NNs have shown superior performance in contrast to conventional methods [44], [68], [69]. As a cornerstone of this work, the following section introduces (deep) NNs in general and the two main categories supervised and unsupervised learning. The fundamentals are supplemented by current state-of-the-art applications and by an outline of the objectives of this thesis.

2.4.1 Deep neural networks

Deep NNs can be understood as functions, mapping elements of an input space to elements of an output space. In the case of image denoising, the input space contains the noisy images, while the output space contains noise-free reference images. These functions depend on a set of trainable parameters, $\Theta \in \mathbb{R}^\ell$, which are optimized throughout the training process. A function $f_\Theta : V \rightarrow W$, where $v \mapsto w := f_\Theta(v)$, is made up of several contributing layers. Each layer is represented by a sub-function f_{Θ_i} , for $i = 1, \dots, N_{\text{layers}}$, which takes the form $f_{\Theta_i}(\cdot) = \phi_i(\mathbf{W}_i \cdot + \mathbf{b}_i)$. Here, \mathbf{W}_i and \mathbf{b}_i represent the weights and biases of a layer applied to the input \cdot , while ϕ_i is a nonlinear activation function applied element-wise to the vector. The entire NN, i.e. the function f_Θ is a composition of these layers, forming a deep structure [70]:

$$f_\theta := \bigcirc_{i=1}^{N_{\text{layers}}} f_{\theta_i}, \quad \text{with} \quad \theta := \bigcup_{i=1}^{N_{\text{layers}}} \theta_i, \quad (2.10)$$

In the context of image denoising, the NN learns the complex relationships between the noise and the underlying image features, estimating a denoised version of a noisy input. This approach is particularly interesting for low-field MRI where the image quality particularly suffers from poor SNR.

There are many different approaches to construct such a deep NN. A widely used architecture in modern deep learning is the U-Net model [71], which is widely applied in various image denoising tasks [72]. It has a characteristic shape composed of two main

parts: a contracting encoder path and an expanding decoder path. The encoder gradually reduces the spatial resolution of the input by applying convolutional operations followed by downsampling. This process captures the context information, extracting high-level features helping the network to understand the structure of noise and information. Each layer in the encoder decreases the resolution but increases the depth of the feature maps, thereby encoding increasingly abstract information. The decoder, on the other hand, progressively up-samples the feature maps, reconstructing spatial details, and focuses on restoring the fine details, producing a high-resolution output. Conventionally, the U-Net architecture incorporates so-called skip connections between corresponding layers in the encoder and decoder. They allow high-resolution features from the encoder to be directly transferred to the decoder, preserving crucial details that might otherwise be lost during the downsampling process. The symmetric structure, combined with skip connections, makes it particularly effective for tasks like image denoising, where both global context and fine details need to be preserved in the output. More recent advances of the U-Net extend the encoder-decoder architecture by additional techniques. Without any claim to completeness, some examples of these techniques are listed below.

- **Residual groups**, that is, residual blocks stacked to preserve identity mapping and improve gradient flow [73].
- **Channel attention** to adaptively reweighting channel-wise features to highlight the most informative channels [74].
- **Multi-head self-attention**, splitting the representation into multiple attention heads to capture various dependencies of features in parallel [75].
- **Upsampling** to increase the spatial resolution of feature maps for high-resolution output [76].
- **Additional skip-connections** to bypass intermediate layers, retaining rich spatial information from earlier stages [71].

2.4.2 Minimization of the objective function

Giving a set of elements from the input and the output space, a dataset can be defined by

$$\mathcal{D} := \{(v_i, w_i)_{i=1}^{N_{\text{samples}}}\}, \quad (2.11)$$

where N_{samples} counts the pairs of inputs and outputs. For the general formulation, the meaning of inputs and outputs is not specified further. Usually, the data set \mathcal{D} is divided into fractions of training, validation, and test data. The goal is to optimize Θ for any mapping $v \mapsto w := f_{\Theta}(v)$ only by using a limited number of samples drawn from the

training data set. The validation data is used to validate f_{Θ} at different stages during the training while the test data remains unknown for the model until the end of the training and is usually employed to evaluate the performance of the model.

Finding the optimal parameters Θ of the NN f_{Θ} , which was defined in the previous section, is achieved by minimizing an objective function, or loss function:

$$\min_{\Theta} \mathcal{L}(\Theta), \quad \mathcal{L}(\Theta) := \frac{1}{|\mathcal{D}|} \sum_{(v,w) \in \mathcal{D}} \ell(f_{\Theta}(v), w). \quad (2.12)$$

The performance measure ℓ is chosen depending on the task fulfilled by f_{Θ} , while there are many options [77]. For image denoising, Mean Squared Error (MSE) is often used [78]. In general, it measures the accuracy or similarity between the networks prediction given the input v with respect to the output or target w . Potentially, ℓ can also combine different measures using, by a convex combination for instance. Equation (2.12) can also be extended by an optional regularization term to improve the learning process. With the backpropagation algorithm [79], the gradient of the loss with respect to the parameters Θ can be calculated. Based on the gradient, the learning step can be performed using techniques such as stochastic gradient descent [80] or a variation of it like Adam [81].

2.4.3 Supervised learning

In supervised learning, the training dataset usually contains pairs of inputs and an associated labels or targets. For low-field MRI this translates to pairs of noisy inputs and noise-free targets. Supervised learning involves observing images from the input and target space and learning to predict a noise-free output from the input [80], in the case of low-field MRI.

Obtaining the dataset of inputs and targets is not trivial. For the low-field MRI denoising problem, it means obtaining a noise-free representation of the scanned object. This can be achieved by paired low-field and high-field acquisitions using co-registration [12]. However, acquiring such data is very costly and introduces errors due to the registration process. Instead, training data for supervision can be retrospectively simulated from high field MRI data, which is considered as the noise-free reference [39], [40]. The simulated data usually contains errors caused by approximations, e.g. due to contrast differences at low and high field, but it is inexpensive to collect as different datasets are publicly available [34], [82], [83]. If the forward problem is known, it can be used to simulated input data retrospectively for a given target, which originates from an acquisition at higher field strength, e.g. at 3 T. However, due to differences in hardware, the transferability from

simulated to real low-field MRI data is limited. A dataset for supervised training can be described by

$$\mathcal{D}^{\text{Supervised}} = \{(\mathbf{x}_{\text{Noisy}}^i, \mathbf{x}_{\text{True}}^i) \mid \mathbf{x}_{\text{Noisy}}^i = \mathbf{x}_{\text{True}}^i + \boldsymbol{\varepsilon}^i\}, \quad (2.13)$$

with $\boldsymbol{\varepsilon}^i \sim \mathcal{N}(0, \sigma^2)$. Supervised training is carried out by performance measures between prediction and the target which serves as a ground truth. The objective function to minimize can be defined using an MSE loss function:

$$\mathcal{L}^{\text{Supervised}}(\Theta) = \frac{1}{|\mathcal{D}^{\text{Supervised}}|} \sum_{\substack{(\mathbf{x}_{\text{Noisy}}, \mathbf{x}_{\text{True}}) \\ \in \mathcal{D}^{\text{Supervised}}}} \|f_{\theta}(\mathbf{x}_{\text{Noisy}}) - \mathbf{x}_{\text{True}}\|_2^2. \quad (2.14)$$

To be able to transfer the trained model and apply it to real low-field MRI data, the choice of the input space is particularly important, as it determines the generalizability of the trained model. For example, if a model is trained only on image data with a T_2 -weighted contrast, its performance is probably worse when applied to image data with a T_1 -weighted contrast.

2.4.4 Unsupervised learning

Unsupervised learning usually aims to implicitly or explicitly learn the probability distribution of input [80]. Here, the data set does not provide a target that is associated with the input. Instead, different techniques have emerged, exploiting only noisy data in pairs or in single instances. The approaches are introduced in the following.

Noise2Noise

Noise2Noise [84] is an unsupervised technique that addresses the central challenge of training a denoising model without requiring noise-free ground-truth images. Instead of clean targets, the method only requires two (or more) noisy observations of the same underlying signal of interest. However, the noise must originate from the same statistical distribution. Since the signal of interest is highly correlated in both observations, whereas noise is not, the model can be trained to predict the signal of interest from one of the observations given another noise observation with the same underlying signal. In low-field MRI this corresponds to having two measurements of the same anatomy under identical parameters, differently corrupted by noise.

To illustrate this, let $\mathcal{D}^{\text{Noise2Noise}}$ be a dataset with pairs of separate noisy measurements of the object \mathbf{x}_{True} . The dataset can be defined as

$$\mathcal{D}^{\text{Noise2Noise}} = \{(\mathbf{x}_{\text{Noisy},1}^i, \mathbf{x}_{\text{Noisy},2}^i) \mid \mathbf{x}_{\text{Noisy},1}^i = \mathbf{x}_{\text{True}}^i + \boldsymbol{\varepsilon}_1^i, \mathbf{x}_{\text{Noisy},2}^i = \mathbf{x}_{\text{True}}^i + \boldsymbol{\varepsilon}_2^i\}, \quad (2.15)$$

with $\epsilon_1^i, \epsilon_2^i \sim \mathcal{N}(0, \sigma^2)$. Although \mathbf{x}_{True} itself is never observed, it is assumed that the noise in $\mathbf{x}_{\text{Noisy},1}$ and $\mathbf{x}_{\text{Noisy},2}$ is zero mean and independent across the two acquisitions. Because each noisy observation is an unbiased estimator of \mathbf{x}_{True} , one can train a denoising network $f_\theta(\mathbf{x})$ to map from one noisy observation $\mathbf{x}_{\text{Noisy},1}$ to the other $\mathbf{x}_{\text{Noisy},2}$ while relying on the following mean squared error objective

$$\mathcal{L}_{\text{Noise2Noise}}(\Theta) = \frac{1}{|\mathcal{D}^{\text{Noise2Noise}}|} \sum_{\substack{(\mathbf{x}_{\text{Noisy},1}, \mathbf{x}_{\text{Noisy},2}) \\ \in \mathcal{D}^{\text{Noise2Noise}}}} \|f_\theta(\mathbf{x}_{\text{Noisy},1}) - \mathbf{x}_{\text{Noisy},2}\|_2^2. \quad (2.16)$$

By minimizing the objective function, the prediction $f_\theta(\mathbf{x}_{\text{Noisy},1})$ converges to \mathbf{x}_{True} , as the expectation value $\mathbf{x}_{\text{Noisy},2}$ given \mathbf{x}_{True} remains \mathbf{x}_{True} . It is also possible to jointly denoise multicontrast low-field MRI data using the *Noise2Noise* strategy, which has the advantage of preserving more of the detailed information in the images [85].

Noisier2Noise

In contrast to the *Noise2Noise* approach, where two noisy measurements of the same anatomical structure form a paired training set, *Noisier2Noise* relaxes this constraint by learning a denoising function from unpaired noisy data without requiring ground truth [86]. Data collection for the *Noise2Noise* approach can be challenging, especially in low-field MRI because two acquisitions of the same subject are required for training. The key idea of *Noisier2Noise* is to artificially introduce additional noise into one set of real noisy images, thereby creating a noisier domain that approximates the distribution of the original noisy data. Although this noisier set is derived directly from the real noisy images, the added noise is independent of the underlying clean signal. This preserves the unbiasedness required to train a denoising network.

Let $\mathbf{x}_{\text{Noisy}}$ denote a noisy image drawn from a real distribution of noisy low-field MRI images. By adding more corruption to $\mathbf{x}_{\text{Noisy}}$, a second image $\mathbf{x}_{\text{Noisier}}$ is produced. The corresponding dataset can be written as

$$\mathcal{D}^{\text{Noisier2Noise}} = \left\{ (\mathbf{x}_{\text{Noisy}}^i, \mathbf{x}_{\text{Noisier}}^i) \mid \mathbf{x}_{\text{Noisy}}^i = \mathbf{x}_{\text{True}}^i + \epsilon_1^i, \mathbf{x}_{\text{Noisier}}^i = \mathbf{x}_{\text{Noisy}}^i + \epsilon_2^i \right\}, \quad (2.17)$$

with $\epsilon_1^i, \epsilon_2^i \sim \mathcal{N}(0, \sigma^2)$. Even though $\mathbf{x}_{\text{Noisy}}$ itself is noisy, the added noise can be controlled to be zero-mean and signal-independent, so that $\mathbf{x}_{\text{Noisier}}$ represents a noisier version of the same scene. Then, a denoising network f_θ is trained to map from $\mathbf{x}_{\text{Noisy}}$ to $\mathbf{x}_{\text{Noisier}}$. Because $\mathbf{x}_{\text{Noisier}}$ still carries the same signal information of interest as $\mathbf{x}_{\text{Noisy}}$, it can act as

a target in a MSE objective function:

$$\mathcal{L}_{\text{Noiser2Noise}}(\Theta) = \frac{1}{|\mathcal{D}^{\text{Noiser2Noise}}|} \sum_{\substack{(\mathbf{x}^{\text{Noiser}}, \mathbf{x}^{\text{Noisy}}) \\ \in \mathcal{D}^{\text{Noiser2Noise}}}} \|f_{\theta}(\mathbf{x}^{\text{Noiser}}) - \mathbf{x}^{\text{Noisy}}\|_2^2 \quad (2.18)$$

Noise2Self

In contrast to previously introduced approaches, the *Noise2Self* method [87] relies solely on a single noisy observation. The idea is that each pixel can be predicted from its surrounding context without directly observing its own noisy value, thereby preventing the network from trivially learning an identity mapping. In the context of low-field MRI, this allows to train a denoising model on an individual image or volume, by exploiting the correlations between neighboring pixels or voxels. The method introduces a blind-spot masking strategy that manipulates pixels or voxels on a predefined grid. Usually, the masked pixel is replaced with a local mean value, because setting the masked pixels to zero introduces intensity discontinuities and can reveal the location of the masked pixel, which impairs the network of learning the contextural information in the data. During training, a mask is chosen which manipulates the input of the model, such that the model prediction depends on the remaining contextural information. The selection of masked positions, is defined by the masking operator \mathbf{M} . The objective function is only evaluated on positions restricted to indices defined by the complement of the mask \mathbf{M}^c .

Given a noisy image $\mathbf{x}_{\text{Noisy}}$, the denoising model f_{θ} and a masking operator \mathbf{M}_i which manipulates certain positions by local mean values at locations i and its complement \mathbf{M}_i^c , the objective function, to train the denoising network can be defined by

$$\mathcal{L}_{\text{Noise2Self}}(\Theta) = \|\mathbf{M}_i^c f_{\theta}(\mathbf{M}_i \mathbf{x}_{\text{Noisy}}) - \mathbf{M}_i^c \mathbf{x}_{\text{Noisy}}\|_2^2. \quad (2.19)$$

Because training with *Noise2Self* is usually performed per sample, the definition of a dataset is not necessary. The noisy input $\mathbf{x}_{\text{Noisy}}$ is manipulated at locations i by the masking operator \mathbf{M}_i while the prediction error is only evaluated at the unseen locations i using the complement operator \mathbf{M}_i^c . During training, the locations i are changed over the epochs. Because the masked noisy target is an unbiased estimate of the underlying clean image \mathbf{x} , matching the networks prediction to it still drives it toward a denoised output.

Within this thesis, the *Noise2Self* self-supervised training strategy is considered, as it neither requires a dedicated training dataset nor assumptions about the actual noise distribution. However, the review of the various self-supervised training strategies should provide a broader context and a better understanding of the chosen approach.

3 Development of the Nexus console

One of the most critical components of an MRI system is the console which controls the acquisition process by interacting with the different system components. In essence, the console transmits gradient and RF waveforms to the different amplifiers, i.e. GPA and RFPA, and receives the amplified MR signal from the pre-amplifier. As introduced in Section 2.1.4, different console designs exist which are either fully digital or incorporate analog components for specific processing stages. Regardless of whether the MRI system is a conventional clinical system and POC low-field research systems, the console carries out the data acquisition, serving as the interface between the analog and digital worlds. The acquired MR data may then undergo subsequent signal processing steps and image reconstruction.

To develop effective deep learning-based image reconstruction and optimization techniques, targeting POC low-field MRI, a transparent and performant, yet compact MRI console is desired. For the integration of data-driven deep learning techniques with model-based reconstruction, as explored in the subsequent chapters, a detailed traceability of the acquisition process is advantageous. A transparent system enables an accurate implementation of the signal model and ensures a close match between simulated training data and actual low-field MRI data. As this work is particularly targeting the compact OSI²One POC low-field MRI system, image optimization and reconstruction is intended to be performed directly on the portable device. With respect to the limited space, the acquisition and deep learning-based processing should be managed by a single system, which provides enough computational resources for these purposes. Another important consideration for portable systems is their susceptibility to changing environmental conditions. Temperature fluctuations or EMI can affect the acquired data and negatively impact image quality. Advanced techniques have demonstrated that EMI can be mitigated by additional sensing coils positioned around the system [88]–[90]. Consequently, a console is required that supports the integration of additional sensors which provide environmental information to be integrated into the data processing pipeline or into the image reconstruction model. In addition, the console should be cost-effective to align with the overarching goal of the OSI²One system, which is to democratize MRI.

Existing consoles are constrained by limitations in technical specifications, accessibility, and cost, which restrict the implementation of advanced low-field techniques. Consoles with

proprietary software pose constraints in customizability, interpretability and compatibility to existing pulse sequence standards [91], [92], data formats [93] and image reconstruction frameworks [94]. The limited access impedes technical innovations for this critical component, resulting in high costs and redundant engineering efforts [22]. In recent years, open source software has been developed for the use with proprietary Software Defined Radios (SDRs) [95]. The device integrates a field-programmable gate array (FPGA) and a microcontroller within a compact design. Although the software can be customized in principle, it requires specialized knowledge in FPGA programming. Additionally, the current memory limitations of the SDR board restrict the maximum duration of *Pulseseq* sequences that can be executed by the console. With its current configuration, the console supports only two transmit and receive channels, which is insufficient for accommodating additional sensors. Moreover, the lack of on-board computational power necessitates an external system for processing the acquired data.

To address the increasing technical requirements of state-of-the-art POC low-field imaging, this chapter introduces the Nexus console, which is a versatile and high-performance console targeting low-field MRI systems. The console is capable of integrating various sensor information from additional receive channels. Throughout the subsequent section the hardware and software architecture of the system is introduced, which is used to conduct in vivo brain measurements with a healthy volunteer. Additional experiments emphasize the versatility of the proposed system. The results presented in this chapter have been published in Magnetic Resonance in Medicine [96] and were previously presented at the annual meeting of the International Society for Magnetic Resonance in Medicine (ISMRM) in 2024 [97]. They serve as the basis for image reconstruction and optimization techniques, which are presented in Chapters 5, 6 and 7.

3.1 Architecture description of the console

The aim of this work is to propose a low-field MRI console that not only enables MR data acquisition with additional sensor information but also functions as a high-performance reconstruction system, capable of integrating advanced deep learning-based image reconstruction techniques. Leveraging Python-based software and adhering to established open-source MRI sequence standards [91], [92] and data formats [94], the console ensures seamless integration into existing MRI workflows, while providing exceptional customization options and enabling rapid prototyping. Within the scope of this work, the console is targeting the portable POC low-field MRI scanner OSI²One [22] with a B_0 field strength of ~ 50 mT. The following section introduces the hardware architecture and the software framework of the Nexus console [97] and summarizes the capabilities of the system.




Figure 3.1: Halbach-based MRI scanner (~ 50 mT) at Physikalisch-Technische Bundesanstalt (PTB) with installed gradient coils, the main MR coil (top left) and additional noise sensing coils (left). GPA with power supply, RFPA, T/R switch with pre-amplifier and the Nexus console are installed below the magnet. A detailed view of the Nexus console (right) shows the 8-channel receive and the 4-channel transmit module with the General Purpose Input/Output (GPIO) expander. (Adapted from [96] CC BY 4.0, 2025).

3.1.1 Hardware setup

The console design follows the framework outlined by Winter et al. [98] and is built around a Supermicro X12SPL-F motherboard (Supermicro Computer Inc., San Jose, USA) equipped with a 24-core Intel 6312U CPU (Intel Corporation, Santa Clara, USA) and 256 GB of RAM. The setup is mounted in a 19-inch rack case as shown in Figure 3.1 and includes seven PCI express slots, accommodating the measurement cards (Spectrum Instrumentation GmbH, Großhansdorf, Germany).

The Digital-to-Analog Converter (DAC) is the 16-bit Arbitrary Waveform Generator (AWG) module M2p.6546-x4, offering four analog transmit channels at a maximum sampling rate of 40 MS/s and a ± 12 V output. This module is supported by 512 MS of on-board memory and has four ports, expandable to 20. The ADC is the 16-bit M2p.5933-x4 digitizer module, capable of 40 MS/s sampling on up to eight single-ended receive channels, with 512 MS of on-board memory and a streaming rate of 700 MB/s. These cards can directly stream data to *CUDA* Graphics Processing Units (GPUs) and synchronize with up to 15 additional cards, allowing for expanded channels if needed. The installed cards support MR systems with operating frequencies of up to ~ 20 MHz, with the option to upgrade to higher-performance cards for systems at even higher field strengths.

The architecture of the OSI²One low-field MRI system with a Larmor frequency of ~ 2 MHz is depicted in 3.2. The RF transmit signal (TX1) is routed to an RFPA, which subsequently sends the amplified pulses through a passive T/R switch to a single-channel RF coil. The AWG module generates the gradient waveforms (G_x , G_y , G_z), which are directly connected to the GPA. The received MR signal is pre-amplified before being digitized by the RX1 channel of the digitizer module. Additionally, two GPIO ports are utilized for RFPA unblanking and digitizer gating. Precise synchronization and timing are ensured by transmitting clock and phase reference signals through two other GPIO ports. In this configuration seven analog receive channels remain which can be used for additional sensors, e.g. for EMI sensing coils, as well as the additional 16 digital GPIO lines.

3.1.2 Software architecture

To operate the hardware setup described in the previous section, an open source software framework is implemented. It handles all the parts of an MRI acquisition including the sequence calculation, transmission of RF and gradient waveforms, sampling and processing of the MR signal. All the steps are performed digitally, giving full control over the transmitted and received signals. This Section describes the console's system environment, the sequence execution, signal acquisition and processing steps, which are implemented within the Nexus console software architecture.

System environment

At PTB, the console operates on AlmaLinux version 9.3, with the spectrum cards running on kernel version M2p 2.17 build 17916 and library version 6.05 build 21120. At Leiden University Medical Center (LUMC), Microsoft Windows 10 is used as the operating system for the console.

The console software, which is completely written in Python, integrates several modules as shown in Figure 3.2, including acquisition control, sequence provider, transmit device and receive devices. Upon startup, the corresponding instances are created using the device configuration file. The acquisition control instance administrates the sequence provider, the transmit and the receive device. It facilitates script-based execution of MR experiments what is particularly suitable for a research environment. In combination with a user interface, the presented framework can be deployed as the backend of the console software [99]. The Python-based implementation allows for easy customization up to the low level control of the measurement cards, i.e. the DAC and ADC modules.




Figure 3.2: Schematic of the Nexus console showing the flow of information from the sequences defined by *PyPulseq* to the acquisition raw data in ISMRMR Raw Data (ISMRMRD) format. The Python-based console software interpretes the MRI sequences and replays gradient and RF waveforms by the AWG module (transmit device). A connection from AWG to digitizer module (receive device), enables precise sampling of the pre-amplified MR signal. (Adapted from [96] CC BY 4.0, 2025).

Sequence execution

The console software seamlessly builds upon *Pulseq* sequences [91] using the *PyPulseq* [92] implementation. As indicated in Figure 3.2, the sequence provider class directly inherits from the *PyPulseq* sequence definition, denoted as `pp.Sequence`, and adds functionality to interpret the sequence for the AWG module. In *Pulseq* sequences are composed of blocks, where each block may contain up to three gradient events, an RF event, and an ADC event. When calculating the sequence, gradient events are analyzed and interpolated on the raster time of the AWG module. *Pulseq* defines the gradient amplitudes in kHz/m which need to be converted to corresponding voltage values in mV. This conversion is determined by the gradient coil efficiency, the GPA gain, the gyromagnetic ratio, and a scaling factor for the Field-of-View (FOV). The resulting waveform is then stored as 16-bit integer values in a pre-allocated *Numpy* array. The maximum channel output can be individually set for each channel and determines the resolution of the waveform. This allows to configure different resolutions for gradient and RF outputs. A sanity check ensures, that the calculated waveforms do not exceed the maximum which would result in a truncation. RF events in *Pulseq* are defined by their envelope waveform which is interpolated similarly to the gradient waveforms. Additionally, the RF waveforms needs to be modulated by the Larmor frequency, which is done digitally. Unlike the gradient conversions, the transformation of RF amplitudes expressed in Hz into 16-bit voltage levels in mV is non-trivial due to its dependence on the coil load. To address this, an arbitrary B_1 scaling factor is introduced which allows fine-tuning the conversion by dedicated flip

angles calibration protocol. All the waveforms of the four output channels are stored in a single *Numpy* array using *Fortran* order.

Beside the analog waveforms some digital control signals are required to enable the RFPA or trigger the digitizer module, i.e. the ADC, by a gating signal. In addition to the ADC gating signal a phase reference signal is transferred between transmit and receive device to maintain phase stability of the system. These digital control signals are transferred synchronously with the analog waveforms by utilizing the most significant bits of the three gradient waveforms. This effectively reduces the available gradient precision to 15-bit which is a trade-off that is also commonly adopted in commercial MRI systems. Despite this reduction in nominal resolution, techniques such as oversampling and dithering can be employed to enhance the effective resolution of the waveforms.

Each sequence calculation is tied to a specific set of acquisition parameters, which include the Larmor frequency, B_1 scaling, FOV scaling, and gradient offsets. These parameters are encapsulated in a unique, hashable dataclass, facilitating validation of the computed sequence before execution.

Signal acquisition

The signal acquisition is control by an instance of the TX device class and an instance of the RX device class which are both managed by the acquisition control instance and inherited from a common abstract device class. This abstract base device implements common functionality to establish or close the connection to the measurement card or to read and report the current device status. Since the transmit and receive device need to operate in parallel, the acquisition control maintains multiple operation loops which are running in concurrent python threads.

Prior to the sequence execution, the buffer of the AWG card is populated with the pre-calculated sequence samples. During sequence execution, the pre-filled buffer is replayed by the AWG module. The buffer works as a ring buffer that is constantly monitored within the operation loop of the TX device and is refilled with new sequence data points, as soon as memory chunks are replayed by the AWG card. This approach enables the streaming of longer sequences in sections, effectively overcoming the on-board memory limitations of the measurement card. The operation loop persists until the sequence is fully completed, or the acquisition control interrupts the loop.

In parallel to the TX device, the RX device operates the digitizer card in a gated acquisition mode. The gated acquisition mode ensures that only data points within the gate signal are sampled, along with a predefined number of pre- and post-trigger samples. Such an ADC event is defined by a pair of timestamps which is written to an exclusive timestamp buffer

on the digitizer card. Within the operation loop, the RX device continuously monitors this buffer and waits for the completion of a new gate event. Based on the timestamps of a new gate event, the ADC duration and the number of sample points can be calculated by the time delta of the timestamp pair. The timestamp buffer is immediately reset, and the acquired sample points are transferred from card memory to the main memory. Once the transfer is complete, the card memory is released and can be filled with new data points. Data is received as 16-bit integer values which correspond to a voltage amplitude dependent on the maximum input voltage of channels. All the channel data of the readout is stored in a single *Numpy* array with *Fortran* order. The operation loop of the RX device is running until all the expected gate events were received, or until the loop is interrupted by a timeout or error.

Signal post-processing

Once the acquisition is finished, the acquisition control triggers the post-processing workflow. In a first step, the *Numpy* array with the received sample points is rearranged into separate dimensions for averaging, receive coils, gate events and readout samples. The gate events may already be correctly sorted, however, in general the sorting depends on the trajectory which is defined by the sequence. After sorting the data, the phase reference signal is separated from the primary input at receive channel zero. The pure receive signal is subsequently converted to a voltage amplitude. Both, the receive signal and the phase reference signal are then demodulated and decimated using a FIR filter from *SciPy* [100], where the decimation factor depends on the receive bandwidth. This workflow remains fully digital, what enables access to the unprocessed raw data.

To encapsulate the acquisition raw data with the *Pulseq* sequence, the specific acquisition parameters, device configuration and additional metadata, an acquisition data object is generated at the end of the processing workflow. The acquisition data object provides a method for non-standardized saving of the acquisition data and for the export to ISMRMRD format [93]. Non-standardized saving stores the restructured *Numpy* array which is advantageous for development work, while the standardized ISMRMRD format ensures seamless integration into existing reconstruction frameworks [94]. If required, the unprocessed raw data, meaning the acquired data points before demodulation, can be stored as well. However, this needs to be enabled explicitly due to the high memory demand. The meta information, which can be extended by custom user inputs, is stored in a *JavaScript Object Notation* (JSON) file which is human-readable and can be easily imported to a python dictionary if needed.

3.2 Experimental setup

Based on the implemented hardware and software, different experiments are conducted to explore the Nexus console performance. Demonstrations include three-dimensional imaging in vivo, the detection and mitigation of EMI using additional sensing coils, monitoring of temperature and B_0 drifts, as well as a self-monitoring experiment which utilized a feedback loop that includes the hardware components to obtain calibration data of the system.

3.2.1 System configuration

Experiments with the Nexus console are conducted at two facilities: LUMC and PTB in Berlin. On both sites, identical console configurations are used, however, some of the hardware components differ. The PTB setup utilizes a 1 kW open-source RFPA [101] and GPA [102], while LUMC employs a 250 W RFPA (Barthel HF-Technik, Aachen, Germany) and a B-AFPA 40 GPA (Bruker Corporation, Billerica, USA). On the receive side, both facilities utilize the ABL0050-00-4510 pre-amplifier (Wenteq Microwave Corporation, Monrovia, USA) to amplify the signal obtained by the MR coil. Within the device configuration, the gradient channels of the console are configured to provide the maximum output at both systems, utilizing as much of the GPA input range as possible. The channel which outputs the RF waveform is configured to a maximum of 200 mV, maintaining a higher resolution. Additionally, the sampling rate for the AWG and digitizer card must be configured. As all channels operate with the same sampling rate, it is determined by the channel with the highest sampling requirements. For a Larmor frequency that is at approximately 2 MHz a sampling rate of 20 MS/s is chosen which yields an oversampling factor of 10 to avoid aliasing. The same sampling rate is set for the receive device. Under consideration of a 20 MS/s sampling rate, the current implementation limits the sampling time to 25.6 s per ADC gate with a single receive channel, and 3.2 s with eight channels. This is because the receive device expects a complete ADC event. However, usually the ADC durations does not exceed 10 ms in practice.

3.2.2 Sequence development

Due to the relatively high field inhomogeneity of portable low-field MRI systems build with permanent magnets, the use of sequences based on spin-echos are favorable. The relaxation of transversal magnetization within a sample or subject is dominated by T_2^* effects, once the magnetization is tilted into the transversal plane. Using a 180° refocussing pulse, the magnetization can be pancake-flipped in the transversal plane, so that the vanishing

magnetization refocusses after Echo Time (TE) [54]. By refocussing the magnetization, the dephasing due to T_2^* can mostly be recovered. An interpreter for *Pulseq* sequences was implemented which translates them into waveforms that can be replayed by the console. Figure 3.3 illustrates the transition from sequence definition to calculated waveforms and output measurements of a digital oscilloscope. Here, the first phase encoding step of a two-dimensional spin-echo sequence with Cartesian sampling is presented as an example, comprising excitation and refocusing pulses followed by a ADC event. As demonstrated by the Figure, the measurement waveforms can be traced back directly to the calculated waveforms and the description in *PyPulseq*. Calculated and measured waveforms are well-matched.


When considering three-dimensional imaging experiments, sequences with a single readout per excitation result in long acquisition times. Thus, a Turbo Spin-Echo (TSE) sequence or more precisely a multi-shot Rapid acquisition with relaxation enhancement (RARE) sequence [103] is implemented for the imaging experiments. The RARE sequence allows acquiring multiple Phase Encoding (PE) steps with a single excitation using multiple refocussing pulses, arranged in an echo train [54]. To achieve three-dimensional k -space encoding, two phase encoding gradients are used, where each echo corresponds to a different phase encoding step. Techniques have been developed to obtain an image by using only a single echo train, referred to as single-shot RARE [104]. It allows fast image acquisition, but noticeably decreases SNR.

During the echo train, the MR signal amplitude decays with $\frac{1}{T_2}$. The peak amplitude of an echo is defined as follows and depends on the number n of the echo in the train and the Echo Spacing (ESP) time.


$$S(n) = S_0 e^{-\frac{nt_{\text{ESP}}}{T_2}} \quad (3.1)$$

The TE of the echo train is defined by the total number N_{etl} of the echos and the ESP by $N_{\text{etl}} \cdot t_{\text{ETL}}$. Because of the T_2 -induced signal decay, the k -space data has non-uniform T_2 -weighting what allows to systematically modulate k -space. An effective TE is defined by the time at which the k -space center is sampled. As the contrast is predominantly determined by the contribution of lower spatial frequencies, dedicated image contrasts are obtained by k -space sorting. Figure 3.4a and 3.4b illustrate different ordering of the k -space PE steps to obtain Proton Density (PD) and T_2 contrast.

Based on these ordering techniques, different k -space trajectories with Cartesian sampling are implemented with *PyPulseq*. The sequences with 3D k -space encoding utilizes two PE and one Readout (RO) gradient. PD-weighted contrast is obtained by using an inside-out trajectory with an Echo-train length (ETL) of 6 and TE/Repetition Time (TR) of 20/500 ms. An interleaved linear trajectory with the ETL increased to 22 and a ratio of TE/TR of 20/2000 ms, yields strongly T_2 -weighted image contrast with an effective TE of



(a) Comparison of the sequence defined with *PyPulseq* and the calculated sequence waveforms, including the digital RF enable and ADC gate signals.



(b) Waveform output at the Nexus console measured with a digital oscilloscope.

Figure 3.3: Comparison of the calculated and replayed sequence waveforms at the Nexus console (Reproduced from [96] CC BY 4.0, 2025).

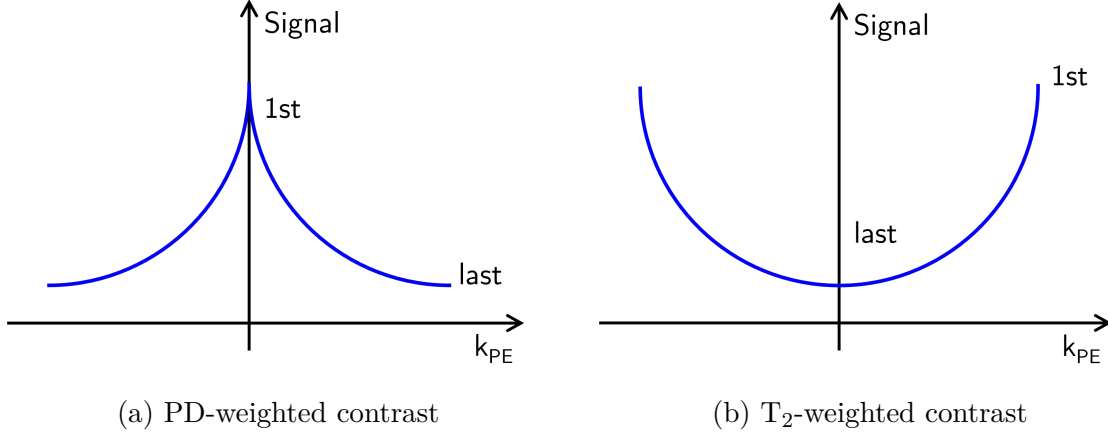


Figure 3.4: Different phase encoding techniques within an echo train of the RARE sequence to obtain PD and T₂ weighted image contrasts (adopted from [54]). The sub-figures plot the signal intensity over the PE location in k -space.

220 ms. Both contrasts are acquired with a FOV of $240 \times 200 \times 210$ mm (RO, PE1, PE2) and a spatial resolution of $2 \times 2 \times 5$ mm (RO, PE1, PE2). The readout bandwidth was set to 20 kHz for both contrasts, resulting in an ADC duration of 6 ms with 120 readout samples and a FIR decimation factor of 1000.

The sequence design was kept as simple as possible to allow easy debugging and to guarantee that the sequence can be executed. For excitation and refocussing, RF block pulses with pulse durations ≤ 200 μ s are used. To compensate for gradient imperfections, an adjustable time offset is added to the flat time of the frequency encoding gradient before the ADC gate. The offset is empirically tuned and depends on the GPA. To minimize saturation effects, three dummy shots were executed before image acquisition.

3.2.3 In vivo imaging

For a comparative reference, almost identical in vivo imaging experiments are performed with a reference console, the Kea2 spectrometer (Magritek Inc., Aachen, Germany) which has previously already been employed for imaging on the low-field MR system [7]. On the hardware side, the key architectural difference between the Kea2 and Nexus consoles lies in the signal processing chain. The Kea2 console integrates a pre-amplifier with 36 dB gain and a noise figure of 1.2 dB prior to a second stage pre-amplifier with 22 dB gain. It uses a passive T/R switch and a 40 MHz antialiasing filter. The down-conversion process utilizes a combination of Cascaded-Integrator-Comb (CIC) and FIR filters. Unfortunately, not all specifications of the signal processing chain were available. Given that the Kea2 console lacks *Pulseq* compatibility, the three-dimensional TSE sequences were replicated with identical acquisition parameters and correction terms using the native sequence description language of the Kea2 console software.

To assess the imaging performance of the Nexus console, experiments on a healthy volunteer were conducted who provided informed consent, as approved by the LUMC ethics committee. The PD- and T₂-weighted 3D image contrasts are acquired with each console for the same volunteer on two consecutive days. During acquisition, a Faraday cage is installed around the volunteer to prevent any distortions from EMI. Image reconstruction is performed using a simple Fast Fourier Transform (FFT) without further post-processing after decimation. Therefore, the acquired k -space data is exported from the reference console and reconstructed using the identical steps as with the Nexus console to guarantee comparability with regard to the image reconstruction. For further comparison between the Nexus console and the reference console, the SNR is calculated based on distinct regions of signal and noise. This approach is inspired by the SNR assessment methodology outlined in the National Electrical Manufacturers Association (NEMA) standard [105], which utilizes a dedicated phantom for measurement rather than in vivo images. Additionally, the imaging results from both consoles are evaluated using a diagonal profile plot, which covers different anatomical structures for a comparative analysis.

Before image acquisition, calibration protocols were executed via the console to adjust essential system parameters to the sample or the subject. The Larmor frequency was calibrated by centering the frequency spectrum of an Free Induction Decay (FID) signal, while a power calibration is conducted by maximizing the FID signal integral for the 90° flip angle. A shimming algorithm is employed to determine offsets for each of the gradient axis, which are superimposed to the gradient waveforms, to achieve first-order active B₀ shimming, compensating for static linear field inhomogeneities. The algorithm acquires repeated FID signals to iteratively find the best combination of gradient offset values that maximize the FID signal amplitude. Given that eddy currents may disturb the MR signal, the refocussed signal may not be perfectly centered in k -space during readout. To mitigate this effect, a correction time of 160 μ s is added to the flat time of the readout gradient before the ADC gate is enabled. This correction time, empirically determined from prior experiments using a spherical phantom, is assumed to be invariant across different subjects. The calibration steps which are individually performed for the Nexus console and the reference system can be summarized in the following calibration protocol:

- Adjustment of the Larmor frequency
- Verification of the receive coil's resonance frequency
- 90° flip-angle calibration
- Calibration of gradient offsets for active shimming
- Readjustment of the Larmor frequency

3.2.4 Integration of auxiliary sensors

In addition to image acquisition, the utilization of auxiliary sensors play a crucial role as they can help to identify sources of distortion that impact the image quality. What is particular advantageous about the Nexus console is the versatile use of multiple different sensors and their easy setup. To demonstrate the versatility of the console, different setups involving auxiliary sensor and feedback from important system components setups are demonstrated. Namely, system self-monitoring, EMI mitigation, temperature and B_0 field drifts are investigated.

Self-monitoring of system components

The Nexus console's multiple receive channels are employed to record various system parameters during data acquisition, demonstrating the benefits of additional channels and their versatile utilization for system calibration. To monitor the transmitted RF waveform, the RF output signal is fed back to the console input. The B_1 field of the MR coil is monitored using a pickup coil, while simultaneously, the forward and reflected power of the RFPA are measured. To verify the GPA output, its current monitor is used to measure the current equivalent of the gradient output voltage for each gradient channel.

A specialized sequence is designed within the *PyPulseq* framework to conduct the monitoring experiment. The sequence consists of a block RF pulse with a duration of 400 μ s, followed by consecutive trapezoidal gradient pulses on each of the three gradient axis. The gradient amplitude is set to 80 kHz/m and the gradient duration to 2 ms. Delays at the sequence start, sequence end and in between consecutive sequence events are all set to 1 ms. To acquire the entire sequence of test signals for monitoring, the ADC event covers the whole sequence, beginning after the initial delay and concluding just before the dead time. This monitoring sequence serves to visualize phenomena such as ringing at the RF coil or performance of the T/R switch. It also enables to estimate a correction factor which compensates for imperfect impedance mismatches between the outputs of the console and the inputs of the GPA. It can result in a deviation of the calculated gradient amplitudes what effects the FOV in imaging. Using the monitored signals provide information about the time response of the GPA output, which can potentially be used to correct the timing of the readout gradient.

Electromagnetic interference detection and mitigation

The versatility of the Nexus console is further demonstrated by integrating sensing coils which acquire EMI during the image acquisition. For an experimental setup at PTB, a

simple sphere phantom is imaged in the presence of different EMIs. Three EMI sensing coils are constructed with 10 turns each and a diameter of 80 mm. Each sensing coil connects to a ZFL-500LN preamplifier (Mini-Circuits, Brooklyn, USA) which provides 24 dB gain and a noise figure of 2.9 dB. The idea is that the EMI measured at the sensing coils become more distinguishable.

For the experimental setup, the three EMI coils are distributed around the MRI system. Two coils are positioned parallel to the bore, while one is placed at a 90-degree rotation, as depicted in Figure 3.1. Four receive channels are enabled on the console to acquire the MR signal simultaneously to the EMI. The original implementation of the External Dynamic InTerference Estimation and Removal (EDITER) [88] algorithm is translated to *Python* to mitigate the EMI in the acquired MR data. The scanner is operated in a shielded environment without any meaningful EMI. Thus, artificial interferences are generated by an external antenna to investigate distinct EMI sources. Here, the two cases of a single frequency and broad band frequency EMI are distinguished to evaluate the detection and mitigation capabilities with the Nexus console. For imaging, a two-dimensional TSE sequence with PD-weighting is used, where only two gradients are employed. The TR/TE ratio is 600/14 ms, the ETL is 18 and an isotropic FOV with 150 mm is used. Readout and phase encoding dimensions of 120×120 , and a readout bandwidth of 20 kHz, resulting in a total acquisition duration of 4.2 s. The MR signal is captured using a solenoid coil with 140 mm in diameter. A spherical phantom with 70 mm in diameter is positioned in the isocenter of the receive coil.

Temperature and B_0 field sensors

Another example of changing environmental conditions is temperature. Due to the temperature dependency of the permanent magnets which are used in the OSI²One scanner, the B_0 field changes depending on temperature. Using seven glass fiber sensors (imc Test & Measurement GmbH, Berlin, Germany) which are attached at several points on the Halbach magnet, as shown in Figure 3.10, the option of carrying out temperature measurements is presented. Six out of the seven sensors are attached to the spaces between the rings of the magnets and one probe is positioned inside the bore. The temperature sensors are connected to a measurement module that transmits data to the console at predefined time intervals via a network connection. An additional processing step ensures synchronization of the temperature data with the MRI measurement. While this approach introduces some complexity, it could be streamlined by employing probes which can be directly attached to the digitizer card of the console enabling synchronized measurements. Nevertheless, this setup highlights the high degree of flexibility offered by the console. Concurrently, the static magnetic B_0 field is measured next to the temperature probe inside

the magnet bore using a PT2026 NMR probe (Metrolab, Plan-les-Ouates, Switzerland). Temperature and B_0 field are tracked over various stages of a fictitious examination. In the first stage, the system is monitored for approximately 20 min after power on. At the beginning of the second stage, a subject is positioned in the scanner. It is to be noted, that during the experiment the RF coil was disconnected from the system. After another 10 min interval at rest but with a subject in the scanner, the three-dimensional TSE sequence with a duration of 6 min 28 s is executed. Finally, the system is monitored again at rest for 20 min. The experiment is conducted in a closed room without air conditioning, artificial heating or significant airflow.


3.3 Results

To evaluate the versatility and the imaging performance of the Nexus console, different experiments have been introduced in the previous section. This section presents the results obtained from the experiments. First, the in vivo images of the healthy volunteer are shown and compared to a reference console. Second, the results from multiple different auxiliary sensors are provided.


3.3.1 In vivo imaging

Figures 3.5a and 3.5b show the PD- and T_2 -weighted image slices of the three-dimensional data acquisition, which was acquired with the Nexus console at LUMC. The images depict multiple transversal slices of the brain from a healthy volunteer. Both contrasts exhibit an in-plane pixel resolution of 2×2 mm and a through-plane resolution of 5 mm. The imaging FOV in the plane is 240×200 mm for RO and the first PE dimensions, while the through-plane FOV, i.e. the FOV in the second PE dimension, is 210 mm. Due to a static position of the FOV most of the slices does not contain any meaningful information. Thus, only 15 out of the in total 42 slices are presented in axial orientation for each of the image contrasts. The acquisition times were 5 min 52 s for the PD-weighted image and 6 min 28 s for the T_2 -weighted image.

Figure 3.6 compares similar image slice from the PD- and T_2 -weighted contrast from each consoles. The scans were conducted on consecutive days, what introduces minor differences in subject positioning, leading to minor differences in the selected slices. The reconstructed three-dimensional magnitude images are normalized to their maximum value to enhance comparability. Inspired by the NEMA standard [105], SNR is calculated for each image slice using the same signal and noise regions. The selected signal and noise regions are overlaid on the images in Figure 3.6 to verify correct positioning. For the PD-weighted



(a) PD-weighted image contrast with TE/TR of 20/500 ms, ETL of 5 and an acquisition duration of 5 min 52 s.



(b) T₂-weighted image contrast with TE/TR of 20/2000 ms, ETL of 22 and an acquisition duration of 6 min 28 s.

Figure 3.5: In vivo brain images of a healthy volunteer acquired with the Nexus console on a 50 mT Halbach-based system. For each contrast, 15 image slices of a 3D TSE sequence with $240 \times 200 \times 210$ mm (RO \times PE1 \times PE2) FOV and $2 \times 2 \times 5$ mm resolution are displayed. (Reproduced from [96] CC BY 4.0, 2025).

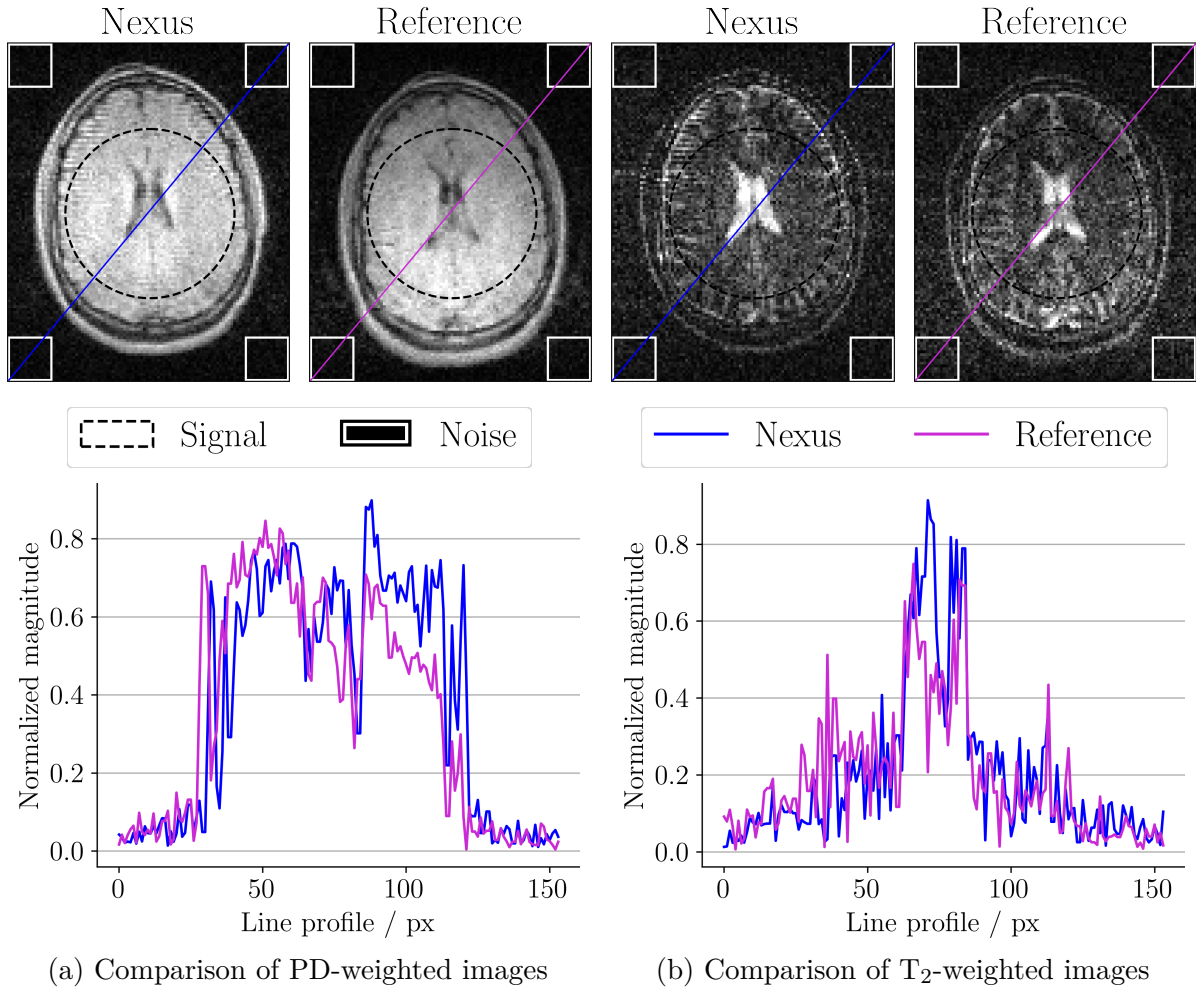


Figure 3.6: Comparison of the image quality obtainable by a three-dimensional TSE sequence across consoles. The proposed Nexus console is compared to a reference console which is commercially available. SNR is calculated from signal (S) and noise (N) regions as denoted in the images, and reported at the bottom of each image. The profiles below compare the intensity values along the diagonals in the images. (Adaptated from [96] CC BY 4.0, 2025).

image, the SNR achieved with the Nexus console is approximately 8.4% higher than that of the reference console. In the case of T₂-weighted image contrast, the SNR improvement with the Nexus console reaches 23.2%. While both sets of images demonstrate high quality, a slight zebra artifacts along the readout dimension in the upper part of the brain is visible in the Nexus images. This artifact is most probably attributed to miscalibration of the B₁ scaling factor, i.e. the flip angle, leading to signal leakage into the outer regions of k -space. However, using simple post-processing methods, the causative k -space regions can be filtered or cropped as demonstrated by Figure 3.7. Images from the reference console in contrast exhibit marginally lower intensity, as evidenced by the comparison of the magnitude profile. For each image contrast, the magnitude profile is compared along the diagonal which is plotted on top of the images in Figure 3.6. Compared to the Nexus console, the reference console exhibits an intensity drop in the right half of the profile in




Figure 3.7: Comparison of the image reconstruction (top row) with full (left) and a truncated (right) k -space (bottom row). The truncation and reconstruction has been done in 3D along the readout dimension. Each readout is truncated by 12 samples in the beginning and 12 samples at the end. For illustration, the magnitude data of a single slice is shown.

the PD-weighted image.

In contrast to the reference system, the proposed console facilitates data export in ISMRMRD format, allowing integration with reconstruction frameworks like Gadgetron [94]. For image reconstruction, a Gadgetron client running in a Docker container was employed, performing FFT-based reconstructions and saving the resulting images in *DICOM* format.

3.3.2 Auxiliary sensors

Beyond imaging capabilities, the Nexus console allows to directly integrate different auxiliary sensors. In this subsection, system monitoring, EMI mitigation, temperature and B_0 tracking results are shown.




Figure 3.8: Monitoring results of RF and gradient waveforms, demonstrating system characterization with the console. Left: Calculated RF waveform (black) compared to measurements using a pickup coil (purple), with corresponding forward (blue) and reflected (yellow) power at the RFPA. Right: Comparison of calculated (black) and GPA-monitored (purple) gradient waveforms. (Adapted from [96] CC BY 4.0, 2025).

System monitoring

In a feedback-based measurement setup, the Nexus console's additional input channels were utilized to monitor the RF performance and compare gradient GPA monitor outputs against calculated values, as demonstrated in Figure 3.8. This setup enables to identify important system characteristics by a direct comparison of monitoring signals to the calculated waveforms. All waveforms were replayed and sampled at 20 MS/s without applying any decimation. The left plots in Figure 3.8 show the calculated and monitored RF signal. Monitoring the transmitted RF signal was conducted using a pickup coil positioned in front of the MR coil. Due to this measurement setup, identical amplitudes are not expected in the comparison plot. Below, the corresponding forward and reflected power, measured at the RFPA, is shown. Deviations from the ideal RF waveform were detected in the pickup coil signal, corresponding to spikes in the monitoring signals from the RFPA. Likewise, the measured gradient signals displayed on the right side exhibit a slight delay compared to the calculated waveforms. The discrepancies between the monitored and calculated amplitudes of the gradient waveforms are likely attributable to the lower input impedance of the GPA. To fully utilize the analog output range of the Nexus console, a higher impedance at the GPA input might be necessary. However, the deviation which has an impact on the FOV can be compensated by a scaling factor which can be derived from the monitored signals. The demonstrated monitoring setup allows for the characterization of certain system components without the need for additional measurement equipment.




Figure 3.9: Experimental results for EMI mitigation using three additional noise sensing coils with the Nexus console. Compared to the image obtainable without any interferences (a), single frequency (b) and broad band (c) EMI could be mitigated successfully using the EDITER algorithm [88]. (Adapted from [96] CC BY 4.0, 2025).

Mitigation of electromagnetic interferences

Figure 3.9 showcases the effect of EMI contamination and mitigation using the EDITER algorithm. For the detection of EMI, three coils were positioned around the magnet of the MRI system. The experiments for EMI mitigation were conducted at the PTB system which suffered from much higher B_0 field inhomogeneities at the time of the measurements. This leads to noticeable geometric distortion of the phantom, which is a spherical ball, and less SNR compared to the in vivo experiments as can be seen in the image with no EMI. Two types of EMI were evaluated throughout the experiments: a constant-frequency signal and a broad-frequency band signal with a bandwidth of 100 kHz. The interferences are clearly visible in both cases and further degrade the obtainable image quality. In contrast to the main MR coil, which acquires the MRI signal, the three sensing coils only detect the EMI signals but not the MRI signal. Using the EDITER algorithm this information is used to significantly mitigate the EMI artifacts in the images. This works for the single frequency distortion, as well as for the broad band EMI as shown in Figure 3.9. It should be noted that in comparison to the setup used for the in vivo measurements this experimental setup exhibited greater B_0 field inhomogeneity (> 4000 ppm) at the time of the experiments, resulting in some geometric distortions in the sphere phantom.

Temperature and B_0 field monitoring

To further demonstrate the versatility of the Nexus console, temperature and B_0 field measurements were performed over different stages. Figure 3.10 illustrates the positioning of the seven temperature probes (T1-T7) and the NMR probe which is located inside the bore. The plot to the right side show the change of the B_0 field, and below the measured temperature change of each individual probe. In the course of the experiment, the temperature constantly increases what leads to a decrease in the B_0 field strength of $20 \mu\text{T}$. A change in temperature caused by the subject is not observed by the temperature

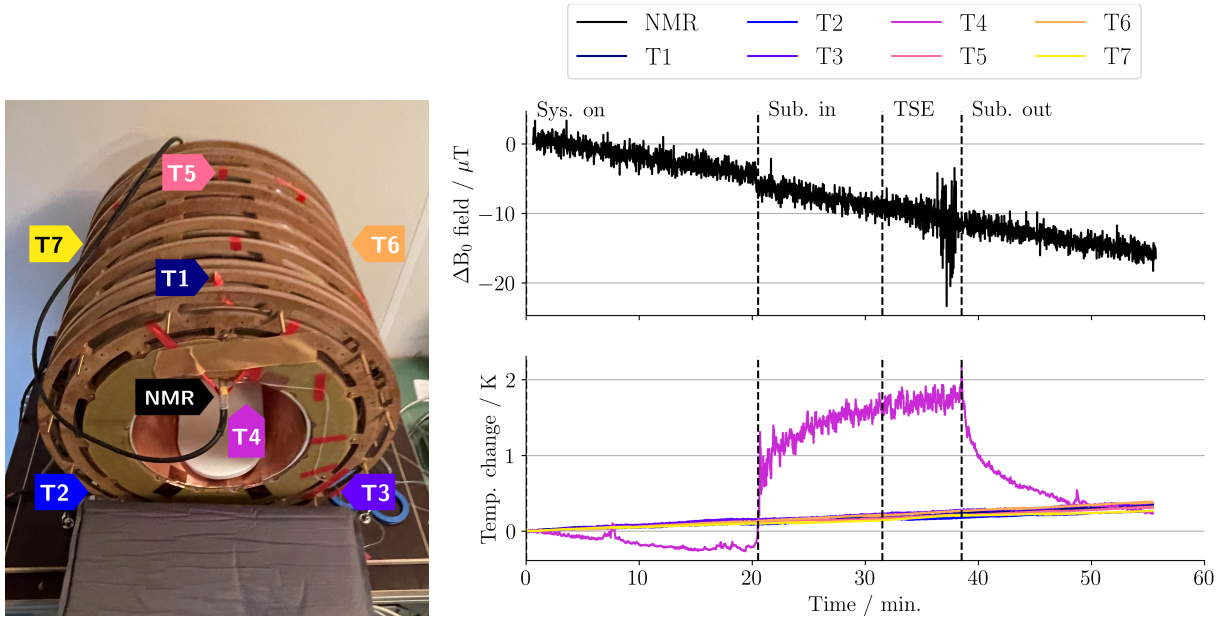


Figure 3.10: Positioning of the fiber Bragg grating temperature sensors (T1-T7) and the NMR probe (left). Measurement of the room temperature increase and the correlated decrease of the B_0 field strength at different stages: System on (Sys. on), subject in the scanner (Sub. in), execution of a TSE sequence (TSE), and subject removed (Sub. out). (Adapted from [96] CC BY 4.0, 2025).

probes positioned at the outside of the magnet. Only the temperature measured by probe T4 which is located inside the bore rises by approximately 2 K. However, this increase can be attributed to the subject's breathing, though it does not significantly affect the mean B_0 field. The slight drop observed when the subject enters the magnets (Sub. in) is likely caused by changes in the magnetic load. Execution of the three-dimensional TSE sequence, also doesn't really affect the magnet temperature. During sequence execution, the measured temperature increases by less than 0.5 K, what corresponding to around 112 Hz. The change in temperature throughout the overall experiment is most probably caused by an increase of the room temperature, indicating that the Larmor frequency should be updated in between multiple measurements.

3.4 Discussion

This chapter presented the Nexus console, which addresses the limitations of existing low-field MRI consoles by a transparent, high-performance, and more flexible design which is easy to adapt. The system, as demonstrated, comprises four transmit and eight receive channels, with the potential to be straightforwardly extended by incorporating additional measurement cards. Architecturally, the motherboard features seven PCI Express slots, of which five remain available for further expansion. Each card supports between four and eight channels, allowing for a total system capacity ranging from 32 to

52 channels, depending on the configuration of the measurement cards. Since the cards support synchronization across up to 16 measurement units, an alternative hardware setup could enable an even higher number of measurement cards and channels.

Without altering the imaging protocol, the console can integrate up to seven additional EMI sensing coils. Effective EMI mitigation in different scenarios was demonstrated using the EDITER algorithm [88], leveraging the additional sensor data to enhance the MR signal quality. To further highlight the adaptability of the system, temperature and B_0 field monitoring were conducted during an MRI examination, utilizing auxiliary measurement hardware. This data can be leveraged to automatically track variations in the Larmor frequency, which shifts due to the temperature dependency of the magnet, as could be observed throughout the experiment. The receive and transmit channels of the console are not restricted to the applications demonstrated in this work but can be easily configured to a wide range of other auxiliary sensors [106]:

- B_0 field probes and current drivers for B_0 shim coils [107] to enhance the field homogeneity under changing environmental conditions.
- Motion-tracking sensors for artifact correction [108].
- Physiological monitoring devices such as electroencephalography (EEG) [109] and electrocardiography (ECG) [110].
- RF pulse and gradient field monitoring [111].
- Parallel imaging with additional receive coils [112].

The implementation in Python, combined with native support for *PyPulseq* sequences, facilitated the seamless development and execution of TSE-based sequences for three-dimensional in vivo imaging. Furthermore, this allows to compare imaging results at different field strengths using the same sequence without much additional effort. The acquired PD- and T_2 -weighted images exhibit image quality comparable to that of a commercially available state-of-the-art console. An increase in SNR was observed with the Nexus console, which may be attributed to differences in the receive chain, variations in post-processing, or environmental factors between experimental sessions. Unlike the proprietary reference system, the Nexus console is fully digital and operates with an open-source software framework, providing unrestricted access to unprocessed raw data before down-conversion, as well as complete control over the data processing pipeline. With the current implementation, waveform calculation from the *Pulseq* sequence is performed prior to acquisition, introducing a delay before sequence execution and leading to high memory consumption in case of longer sequences. However, this latency could be minimized using parallel computing, allowing waveform computation to run in parallel with the sequence

execution. An alternative approach involves the configuration of different transmit cards with distinct sampling rates for RF and gradient waveforms or employing external DAC hardware, such as the *OCRA1* board [113], to output gradient waveforms through digital channels at a lower sampling rate. Within the operational constraints of the Nexus console's measurement cards, which support a maximum sampling rate of 40 MS/s, the sampling rate remains freely configurable, allowing adaptation to different experimental frequency ranges. For example, given Nyquist sampling requirements, MR experiments at frequencies below 20 MHz, such as 15 MHz ($B_0 \sim 0.35$ T), are feasible. Direct access to unprocessed raw data at up to 40 MS/s enables the exploration of oversampling techniques that could further enhance SNR [114]. If higher frequencies are required, the system can accommodate AWG and digitizer cards with greater sampling rates, requiring only minor modifications to the existing open-source software framework.

The integration of an ISMRMRD interface ensures compatibility with established reconstruction frameworks, such as Gadgetron [94]. Additionally, the ability to directly streamline data from the measurement cards to a GPU facilitates seamless integration of deep learning-based image reconstruction techniques which can be directly employed at the scanner.

With an overall material cost of approximately 12 000 € and no requirement for specialized tools or expertise, the Nexus console offers a cost-effective and easily replicable alternative given its performance. It bridges the gap between existing low-cost open-source designs, which are underpowered for advanced applications, and expensive proprietary systems. Assembly is as straightforward as building a PC, using standardized computer hardware housed in a 19" enclosure. This enhances reproducibility for low-field MRI research while also simplifying adherence to standard compliance procedures for electrical safety, as specified in IEC 60601-157. Such conformity streamlines the approval process for in vivo applications, fostering broader adoption of the Nexus console and potentially enabling access to more transparent and comprehensible data.

3.5 Summary

With the Nexus console, a novel versatile and high-performance console designed for advanced POC low-field MRI has been presented. Due to the implementation in *Python* it is easily adaptable to integrate information from different auxiliary sensor, which can help to provide robust imaging under changing environmental conditions. By providing interfaces to established frameworks such as *PyPulseq* and ISMRMRD, the console enables straightforward integration into existing workflows while supporting flexible customization for rapid prototyping. The system has demonstrated reliable image acquisition using TSE

sequences on two portable low-field MRI scanners operating at 50 mT, producing PD- and T₂-weighted image contrasts in vivo comparable to those obtained with a proprietary reference console.

Throughout this chapter, image reconstructions were performed using FFT, demonstrating promising results while also revealing limitations, particularly the impact of noise in T₂-weighted in vivo images. The challenges of imaging in the presence of strong B₀ field inhomogeneities further underscore the need for dedicated reconstruction techniques. With its transparent acquisition pipeline and the capability to seamlessly integrate GPU-powered techniques, the Nexus console establishes a solid foundation for deep learning-driven image reconstruction and optimization in POC low-field MRI.

4 Image and B_0 field reconstruction

Based on the acquisition results from the previous chapter, it can be concluded that the limitations in low-field MRI are dominated by low SNR and high B_0 field inhomogeneities. Two similar low-field MRI scanners at different sites with different shimming were utilized for the imaging experiments. In any of the measurements, the acquired data was affected by noise which is noticeable in the image domain. Principally, the SNR can be improved by averaging, but this leads to very long scan times what prevents the clinical use. Since the system used at PTB had significant higher B_0 field inhomogeneities at the time of the experiments, the images suffer from geometric distortions. The system at LUMC, which underwent an extensive shimming procedure to homogenize the B_0 field, is significantly less affected by such artifacts. Nevertheless, B_0 field distortions need to be avoided as much as possible in order to provide a correct representation of the anatomy and to allow for accurate post-processing steps such as automatic segmentations. Changes in environmental temperatures will affect the B_0 field and portable POC systems are in particular susceptible to these. Shimming a magnet involves multiple iterations of precise field mapping, field optimization with thousands of small neodymium magnets, whose field is subject to tolerances and precise manufacturing of magnet holders. Consequently, effective shimming requires substantial workload and specialized equipment [115], making ad-hoc application unfeasible. If the B_0 field is known, the information can be used to compensate inhomogeneity artifacts in the images [116], what motivates the investigation of a B_0 -informed reconstruction approach.

Chapter 1 provided an overview of existing deep learning-based techniques which improve the image quality by denoising and B_0 field correction. None of the methods which involve machine learning to estimate the B_0 field map have been applied to a low-field system yet and were developed for high field systems, where the relative field deviations are much smaller. Consequently, these techniques are not directly applicable to low-field MRI. However, a method targeting the estimation of image and B_0 field map at low field strength of about 50 mT using conventional regularization was proposed in [116]. The authors propose different models for B_0 field reconstruction based on the relative phase difference of two acquisitions and B_0 -informed image reconstruction. Image and field map estimation is carried out using an alternating optimization scheme. Despite promising results, the reconstruction is, however, computationally expensive.

To address the lack of a model-based reconstruction algorithm, targeting the image and B_0 field related problems, this chapter introduces a novel reconstruction network which is fully end-to-end trainable and explicitly incorporates the physics of the data acquisition process. Inspired by the algorithm unrolling techniques which have already been successfully demonstrated for cine MRI [50], model-driven and data-driven reconstruction are combined to yield noise-free images with corrected geometric distortions.

4.1 Formulation of the joint reconstruction problem

Based on the formulation of the forward problem (2.7), this section first introduces the operator of the method, which takes as input the image as well as the B_0 field map. Because the operator is used in an end-to-end training later on, it must be defined in such a way that it is differentiable. In a second step, the reconstruction problem is defined as a joint minimization problem with a variable splitting approach, which divides the problem into two sub-problems for the image and the B_0 field map, respectively.

4.1.1 Approximation of the forward operator

Due to the dependency on space and time, the operator \mathbf{E} in (2.7) becomes computationally expensive to evaluate. To overcome this problem, an approximation of the term $e^{-i\omega_n t_m}$ based on time segmentation is used, as proposed in [117] which is computationally more favorable. The approximation is given by

$$e^{-i\omega_n t_m} \approx \sum_{l=1}^L c_{l,m} e^{-i\omega_n \mathbf{t}'_l}, \quad (4.1)$$

where the vector $\mathbf{t}' \in \mathbb{R}^L$ is defined by $\mathbf{t}'_l = l \frac{T_{\text{acq}}}{L}$, $l = 1, \dots, L$ and $n, m = 1, \dots, N$. L defines the number of time segments and N accounts for the total number of image and k-space samples. The approximation is obtained by a linear combination with the coefficients $c_{l,m}$, which are introduced later in this subsection. For a fixed image \mathbf{x} , the function $g_{\mathbf{x}} : \mathbb{C}^N \rightarrow \mathbb{C}^{LN}$ is defined by

$$g_{\mathbf{x}}(\boldsymbol{\omega}) = (\mathbf{1}_L \otimes \mathbf{x}) \circ e^{-i(\mathbf{t}' \otimes \boldsymbol{\omega})}, \quad (4.2)$$

where $\mathbf{1}_L$ denotes an L -dimensional vector of ones and the exponential function is applied component-wise. The Kronecker product is denoted by " \otimes " and the Hadamard product by " \circ ". The operator $\text{diag} : \mathbb{C}^N \rightarrow \mathbb{C}^{N \times N}$, $\mathbf{x} \mapsto \text{diag}(\mathbf{x}) = [x_1 \mathbf{e}_1, \dots, x_N \mathbf{e}_N]$, where $\{\mathbf{e}_j\}_{j=1}^N$ denotes the standard basis of the space \mathbb{R}^N , which transforms a vector into a diagonal matrix.

For fixed $\boldsymbol{\omega}$, the linear operator \mathbf{B}_ω is defined by the matrix

$$\mathbf{B}_\omega = [\text{diag}(e^{-i\omega t_1}), \dots, \text{diag}(e^{-i\omega t_L})]^\mathbf{T} \in \mathbb{C}^{LN \times N}. \quad (4.3)$$

It should be noted that the linear operator \mathbf{B}_ω applied to \mathbf{x} equals the evaluation of $g_{\mathbf{x}}$ in $\boldsymbol{\omega}$, i.e., $g_{\mathbf{x}}(\boldsymbol{\omega}) = \mathbf{B}_\omega \mathbf{x}$. This notational aspect serves to highlight the fact that the function $g_{\mathbf{x}}$ is nonlinear with respect to the field map $\boldsymbol{\omega}$, but \mathbf{B}_ω defines a linear operator with respect to the image \mathbf{x} . Transformation from image to Fourier space is achieved by the construction of a block diagonal Fourier operator $\mathbf{F}_L := \mathbf{I}_L \otimes \mathbf{F}$, where \mathbf{I}_L denotes the $L \times L$ -identity matrix.

In (2.8), the field inhomogeneities $\boldsymbol{\omega}$ cause phase accumulation. In the impulse response function, this phase term can be neglected if the field map varies slowly enough. If the number of L is chosen sufficiently small, the phase term in the impulse response function can be neglected [118], [119]. To compute the coefficients $c_{l,m}$ in (4.1), the functions

$$\gamma : \mathbb{R}^N \rightarrow \mathbb{C}^{N \times N}, \boldsymbol{\omega} \mapsto \gamma(\boldsymbol{\omega}) = e^{-i\boldsymbol{\omega} \cdot \mathbf{t}^\mathbf{T}}, \quad \text{and} \quad (4.4a)$$

$$\gamma^L : \mathbb{R}^N \rightarrow \mathbb{C}^{N \times L}, \boldsymbol{\omega} \mapsto \gamma^L(\boldsymbol{\omega}) = e^{-i\boldsymbol{\omega} \cdot (\mathbf{t}')^\mathbf{T}} \quad (4.4b)$$

are defined. For fixed $\boldsymbol{\omega}$, $\boldsymbol{\Gamma}_\omega \in \mathbb{C}^{N \times N}$ is defined by $\boldsymbol{\Gamma}_\omega = \gamma(\boldsymbol{\omega})$ and $\boldsymbol{\Gamma}_\omega^L \in \mathbb{C}^{N \times L}$ by $\boldsymbol{\Gamma}_\omega^L = \gamma^L(\boldsymbol{\omega})$, respectively. Both matrices contain the exponential expression from (4.1) whereby $\boldsymbol{\Gamma}_\omega$ contains all the time points, while $\boldsymbol{\Gamma}_\omega^L$ only contains segmented time points. The approximation in (4.1) for all $n = 1, \dots, N$ and $m = 1, \dots, L$ can be compactly written in the form of a system $\boldsymbol{\Gamma}_\omega^L \mathbf{C}_\omega = \boldsymbol{\Gamma}_\omega$, whose solution is calculated as

$$\mathbf{C}_\omega = (\boldsymbol{\Gamma}_\omega^L)^\dagger \boldsymbol{\Gamma}_\omega, \quad (4.5)$$

where $(\boldsymbol{\Gamma}_\omega^L)^\dagger$ is the Moore-Penrose inverse given by

$$(\boldsymbol{\Gamma}_\omega^L)^\dagger := ((\boldsymbol{\Gamma}_\omega^L)^\mathbf{H} \boldsymbol{\Gamma}_\omega^L)^{-1} (\boldsymbol{\Gamma}_\omega^L)^\mathbf{H}. \quad (4.6)$$

The matrix \mathbf{C}_ω from equation (4.5) contains L row vectors with N coefficient each, i.e. $\mathbf{C}_\omega = [\mathbf{c}_1, \dots, \mathbf{c}_L]^\mathbf{T}$ with $\mathbf{c}_l \in \mathbb{C}^N, l = 1, \dots, L$. To apply the coefficients to the Fourier transformed vector of L images, $\tilde{\mathbf{C}}_\omega = [\text{diag}(\mathbf{c}_1), \dots, \text{diag}(\mathbf{c}_L)] \in \mathbb{C}^{N \times LN}$ is defined.

The time-segmented operator $\mathbf{E}^L(\mathbf{x}, \boldsymbol{\omega})$ depends on two variables. For fixed $\boldsymbol{\omega}$, the operator is denoted by $\mathbf{E}_\omega^L : \mathbb{C}^N \rightarrow \mathbb{C}^N$ and for fixed \mathbf{x} by $\mathbf{E}_\mathbf{x}^L : \mathbb{R}^N \rightarrow \mathbb{C}^N$. They are defined by

$$\mathbf{E}_\omega^L(\mathbf{x}) := (\tilde{\mathbf{C}}_\omega \mathbf{F}_L \mathbf{B}_\omega)(\mathbf{x}) \quad (4.7a)$$

$$\mathbf{E}_\mathbf{x}^L(\boldsymbol{\omega}) := v(\boldsymbol{\omega}) \mathbf{F}_L g_{\mathbf{x}}(\boldsymbol{\omega}) \quad (4.7b)$$

respectively, where $v(\boldsymbol{\omega}) = \tilde{\mathbf{C}}_{\boldsymbol{\omega}}$. $\mathbf{E}_{\boldsymbol{\omega}}^L$ is linear with respect to \mathbf{x} , while $\mathbf{E}_{\mathbf{x}}^L$ is nonlinear with respect to $\boldsymbol{\omega}$ due to the functions $g_{\mathbf{x}}(\boldsymbol{\omega})$ and $v(\boldsymbol{\omega})$. Both of the operators from equations (4.7a) and (4.7b) will be used in the formulation of two sub-problems. The proposed reconstruction operator will be utilized in end-to-end training, requiring backpropagation through the entire network, which underscores the need for a detailed notation.

4.1.2 Reconstruction problem

The reconstruction problem can be defined as a joint minimization task on the field map $\boldsymbol{\omega}$ and the image \mathbf{x} . This is expressed mathematically as:

$$\min_{\mathbf{x}, \boldsymbol{\omega}} \frac{1}{2} \left\| \mathbf{E}^L(\mathbf{x}, \boldsymbol{\omega}) - \mathbf{y} \right\|_2^2 + \mathcal{R}_{\mathbf{x}}(\mathbf{x}) + \mathcal{R}_{\boldsymbol{\omega}}(\boldsymbol{\omega}), \quad (4.8)$$

where the objective function consists of a data fidelity term and two regularization terms. The regularization terms $\mathcal{R}_{\mathbf{x}}(\mathbf{x})$ and $\mathcal{R}_{\boldsymbol{\omega}}(\boldsymbol{\omega})$, enforce constraints on the image \mathbf{x} and the field map $\boldsymbol{\omega}$, respectively. The regularization for the image \mathbf{x} is defined as:

$$\mathcal{R}_{\mathbf{x}}(\mathbf{x}) = \frac{\lambda_{\mathbf{x}}}{2} \left\| \mathbf{x} - \mathbf{x}_{\text{NN}} \right\|_2^2, \quad (4.9)$$

where $\lambda_{\mathbf{x}} > 0$ is a regularization parameter. This term encourages the reconstructed image \mathbf{x} to remain close to a previous estimate \mathbf{x}_{NN} , which is provided by a NN. Similarly, the regularization for the field map $\boldsymbol{\omega}$ is given by:

$$\mathcal{R}_{\boldsymbol{\omega}}(\boldsymbol{\omega}) = \frac{\lambda_{\boldsymbol{\omega}}}{2} \left\| \boldsymbol{\omega} - \boldsymbol{\omega}_{\text{NN}} \right\|_2^2, \quad (4.10)$$

where $\lambda_{\boldsymbol{\omega}} > 0$ is a regularization parameter and $\boldsymbol{\omega}_{\text{NN}}$ is the field map estimate provided by a NN. Substituting the regularization terms into the joint minimization problem, the objective can be rewritten as:

$$\min_{\mathbf{x}, \boldsymbol{\omega}} \frac{1}{2} \left\| \mathbf{E}^L(\mathbf{x}, \boldsymbol{\omega}) - \mathbf{y} \right\|_2^2 + \frac{\lambda_{\mathbf{x}}}{2} \left\| \mathbf{x} - \mathbf{x}_{\text{NN}} \right\|_2^2 + \frac{\lambda_{\boldsymbol{\omega}}}{2} \left\| \boldsymbol{\omega} - \boldsymbol{\omega}_{\text{NN}} \right\|_2^2. \quad (4.11)$$

Due to the dependency on two variables, directly solving the above problem is challenging. To address this complexity, an approximation of the forward operator is utilized alongside a variable-splitting approach. The approximation of the forward operator reduces the number of Fourier transforms from M to L by segmenting the acquisition time into uniform intervals.

4.1.3 Sub-problem for the image

When the field map $\boldsymbol{\omega}$ is fixed, the reconstruction problem reduces to the following sub-problem:

$$\min_{\mathbf{x}} \frac{1}{2} \left\| \mathbf{E}_{\boldsymbol{\omega}}^L(\mathbf{x}) - \mathbf{y} \right\|_2^2 + \frac{\lambda_{\mathbf{x}}}{2} \left\| \mathbf{x} - \mathbf{x}_{\text{NN}} \right\|_2^2, \quad (4.12)$$

where $\mathbf{E}_{\boldsymbol{\omega}}^L$ is defined as in (4.7a). Since $\mathbf{E}_{\boldsymbol{\omega}}^L$ is linear with respect to \mathbf{x} , the objective function in (4.12) is convex. Consequently, the unique solution to this sub-problem can be obtained by solving the linear system:

$$\mathbf{H}_{\boldsymbol{\omega}} \mathbf{x} = \mathbf{b}_{\boldsymbol{\omega}}, \quad (4.13)$$

where the system matrices are defined as:

$$\mathbf{H}_{\boldsymbol{\omega}} = (\mathbf{E}_{\boldsymbol{\omega}}^L)^H \mathbf{E}_{\boldsymbol{\omega}}^L + \lambda_{\mathbf{x}} \mathbf{I}_N, \quad (4.14)$$

$$\mathbf{b}_{\boldsymbol{\omega}} = (\mathbf{E}_{\boldsymbol{\omega}}^L)^H \mathbf{y} + \lambda_{\mathbf{x}} \mathbf{x}_{\text{NN}}. \quad (4.15)$$

The solution to this system can be efficiently computed using the CG method.

4.1.4 Sub-problem for the field map

For a fixed image \mathbf{x} , the corresponding sub-problem becomes:

$$\min_{\boldsymbol{\omega}} \frac{1}{2} \left\| \mathbf{E}_{\mathbf{x}}^L(\boldsymbol{\omega}) - \mathbf{y} \right\|_2^2 + \frac{\lambda_{\boldsymbol{\omega}}}{2} \left\| \boldsymbol{\omega} - \boldsymbol{\omega}_{\text{NN}} \right\|_2^2. \quad (4.16)$$

Due to the exponential correlation between $\boldsymbol{\omega}$ and \mathbf{y} this nonlinear sub-problem is much more complex compared to sub-problem (4.12), which is why a first-linearize-then-solve approach is used here. To simplify the problem, a first-order Taylor approximation is applied to the forward operator $\mathbf{E}_{\mathbf{x}}^L$, following a similar approach to that described in [120]. The Jacobian of $\mathbf{E}_{\mathbf{x}}^L$, evaluated at a linearization point $\boldsymbol{\omega}_0$, is denoted by $\mathbf{J}_{\mathbf{E}_{\mathbf{x}}^L}|_{\boldsymbol{\omega}=\boldsymbol{\omega}_0}$. Using this approximation, the forward operator can be expressed as:

$$\tilde{\mathbf{E}}_{\mathbf{x}}^L(\boldsymbol{\omega}) = \mathbf{E}_{\mathbf{x}}^L(\boldsymbol{\omega}_0) + \mathbf{J}_{\mathbf{E}_{\mathbf{x}}^L}|_{\boldsymbol{\omega}=\boldsymbol{\omega}_0}(\boldsymbol{\omega} - \boldsymbol{\omega}_0). \quad (4.17)$$

By defining the modified data term as:

$$\tilde{\mathbf{y}} = \mathbf{E}_{\mathbf{x}}^L(\boldsymbol{\omega}_0) + \mathbf{J}_{\mathbf{E}_{\mathbf{x}}^L}|_{\boldsymbol{\omega}=\boldsymbol{\omega}_0}(\boldsymbol{\omega}_0), \quad (4.18)$$

and substituting \mathbf{E}_x^L with $\tilde{\mathbf{E}}_x^L$ in (4.16), the problem simplifies to the following linearized sub-problem:

$$\min_{\omega} \frac{1}{2} \left\| \mathbf{J}_{\mathbf{E}_x^L} |_{\omega=\omega_0}(\omega) - \tilde{\mathbf{y}} \right\|_2^2 + \frac{\lambda_{\omega}}{2} \left\| \omega - \omega_{\text{NN}} \right\|_2^2. \quad (4.19)$$

The solution to this sub-problem can be found by solving the linear system: The solution to this sub-problem can be found by solving the linear system:

$$\mathbf{G}_x \omega = \mathbf{d}_x, \quad (4.20)$$

where the matrices are defined as:

$$\mathbf{G}_x = (\mathbf{J}_{\mathbf{E}_x^L} |_{\omega=\omega_0})^H \mathbf{J}_{\mathbf{E}_x^L} |_{\omega=\omega_0} + \lambda_{\omega} \mathbf{I}_N \quad (4.21)$$

$$\mathbf{d}_x = (\mathbf{J}_{\mathbf{E}_x^L} |_{\omega=\omega_0})^H \tilde{\mathbf{y}} + \lambda_{\omega} \omega_{\text{NN}}. \quad (4.22)$$

This linear system can also be solved efficiently using the CG method.

4.2 Neural network architecture

The previous Section introduced the reconstruction problem which is divided to solve the image and field map related sub-problems alternatively. Each of the sub-problems utilize a different NN for regularization. This Section introduces a Convolutional Neural Network (CNN) specifically designed to estimate B_0 field maps and a U-Net architecture which is capable of denoising complex-valued low-field MRI data. With the integration of the CNNs for regularization, the end-to-end network architecture is presented.

4.2.1 Spherical harmonic network for B_0 regularization

For the field map estimation, a novel NN architecture, referred to as Spherical Harmonics (SH)-Net, is proposed. This approach follows a methodology similar to the Off-Resonance Network (ORN) architecture, introduced in [46] and originally proposed in [45], estimates a B_0 field map from a single complex-valued distorted image. In the ORN, the estimated field map is interpreted as a non-stationary convolution kernel, which can be learned by a NN. Rather than learning the entire mapping from the complex-valued image to the field map distribution, the proposed approach employs a physically motivated NN structure. This design leverages the inherent spatial smoothness of B_0 field maps, achieved by parameterizing the field map as a linear combination of SH coefficients and basis functions. Using SH harmonics has the advantage that an entire three-dimensional B_0 field can be described by only a few coefficients [54].

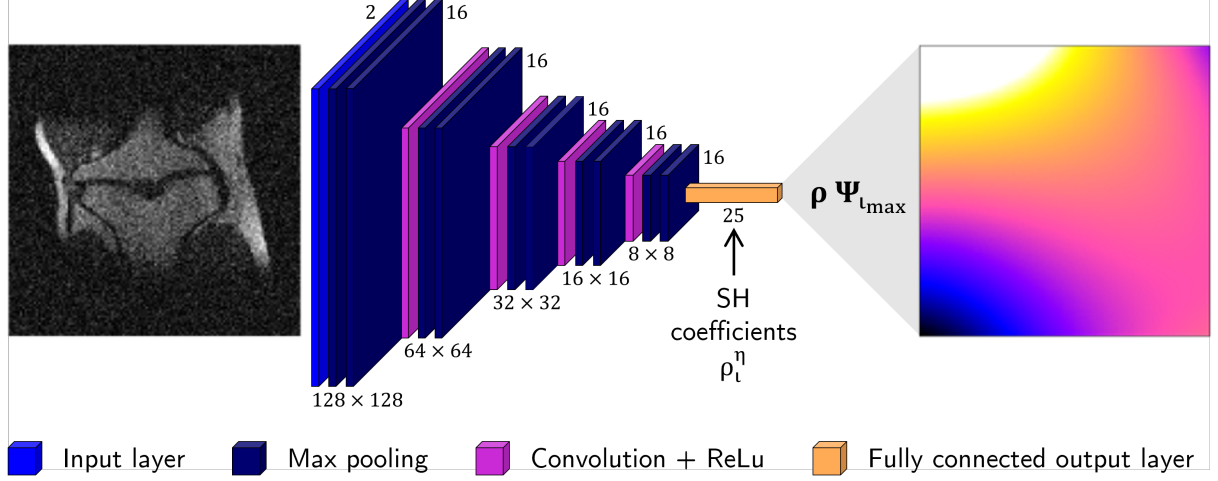


Figure 4.1: SH-Net to estimate a field map prior ω_{NN} from a complex-valued B_0 -uninformed FFT reconstruction. The network is built of an encoding path with two convolutional layers per stage and a constant number of 16 filters. The last layer is a fully connected layer that maps the values to 25 SH coefficients, which are used to calculate a spatially smooth field map by (4.24). (Reproduced from [121] CC BY 4.0, 2024).

A B_0 field map can be expressed as a weighted sum of SH functions, where the basis functions $Y_l^\eta(\theta, \phi)$ are defined by degree ι and order η . These functions depend on the polar angle θ and the azimuthal angle ϕ . In the two-dimensional case considered here, the azimuthal angle is fixed at 90° , simplifying the basis functions to $Y_l^\eta(\theta)$. The field distribution, parameterized by radius r , is described as a weighted sum of real SH functions and coefficients in the polar coordinate system. The field map ω , expressed in polar coordinates and parameterized by the SH coefficients \mathbf{h}_l^η , is defined as:

$$\omega_{r,\theta} = \sum_{\iota=0}^{\iota_{\max}} \sum_{\eta=-\iota}^{\iota} r^\iota \mathbf{h}_l^\eta Y_l^\eta(\theta). \quad (4.23)$$

Expressed in vectorized form, this equation can be written in a more compact way as

$$\omega = \Psi_{\iota_{\max}} \rho, \quad (4.24)$$

where $\Psi_{\iota_{\max}}$ denotes the set of SH basis functions for orders up to ι_{\max} and degrees η , and ρ represents the corresponding coefficients.

As illustrated in Figure 4.1, the network consists of an encoder path that processes the input image and outputs the SH coefficients. Instead of also learning the decoder path of the network, SH basis functions are used to decode the output state of the encoder path as denoted in (4.24). Thus, the proposed SH-Net architecture is designed to estimate the SH coefficients of the underlying field distribution from the input image with complex values. Empirical evaluations determined that the use of a fixed number of 16 feature channels for each convolutional layer and five encoding stages provides optimal performance. In the

final stage of the network, the features are flattened into a vector and a fully connected layer is applied to estimate the SH coefficients. Here, SH coefficients up to degree five are generated, which has been shown to sufficiently describe the field map, as supported by [115]. The input to the network is an initial B_0 -uninformed image reconstruction, calculated as $\mathbf{x}_0 = \mathbf{F}^H \mathbf{y}$. The network, which estimates the SH coefficients to calculate the B_0 field map according to (4.24), is denoted as $\boldsymbol{\omega}_{\text{NN}} = \text{SH}_\Omega(\mathbf{x}_0)$.

4.2.2 U-Net for image regularization

Due to the low magnetic field used for data acquisition, noise in the acquired k -space data \mathbf{y} is dominant. As a result, in (4.11), the NN-image prior to \mathbf{x}_{NN} should represent an estimate of the noise-free image corresponding to the data \mathbf{y} . For the denoising task, a NN is employed, operating in the image domain to produce a denoised image \mathbf{x}_{NN} .

Assuming the availability of a reliable field map estimate $\boldsymbol{\omega}$, such as $\boldsymbol{\omega}_{\text{NN}}$, the operator $\mathbf{E}_{\boldsymbol{\omega}_{\text{NN}}}^H$ can be defined to reconstruct an image and correct for geometric distortion. In the image domain, a U-Net architecture [71], [122], [123] is applied for denoising. This architecture consists of an encoder and a decoder path, where the real and imaginary components of the complex-valued image are processed as separate channels. The denoised and geometrically corrected image is then estimated as $\mathbf{x}_{\text{NN}} = u_\Theta((\mathbf{E}_{\boldsymbol{\omega}_{\text{NN}}}^L)^H \mathbf{y})$. In this context, the NN architecture comprises four encoding stages, each featuring two convolutional layers, with an initial filter size of 32 and a kernel size of three.

4.2.3 Algorithm unrolling with alternating reconstruction scheme

As introduced in Section 4.1, a variable splitting approach is applied which solves the problem utilizing two sub-problems. The solution to problem (4.11) is approached by solving (4.12) and (4.19) in an alternating pattern for a fixed number of iterations T , yielding and end-to-end trainable NN denoted as $f_{\Theta, \Omega}^T$. Formally, $f_{\Theta, \Omega}^T$ is defined as:

$$\begin{aligned} f_{\Theta, \Omega}^T : \mathbb{C}^N &\rightarrow \mathbb{R}^N \times \mathbb{C}^N \\ \mathbf{y} &\mapsto \mathbf{z}^T, \end{aligned} \quad (4.25)$$

where \mathbf{z}^T is the T -th element of the sequence $(\mathbf{z}^t)_{t=1}^T$ defined as:

$$\begin{cases} \mathbf{z}^{t+1} := (\boldsymbol{\omega}^{t+1}, \mathbf{x}^{t+1}), \\ \boldsymbol{\omega}^{t+1} := \arg \min_{\boldsymbol{\omega}} \frac{1}{2} \left\| \mathbf{J}_{\mathbf{E}_{\mathbf{x}^t}^L}(\boldsymbol{\omega}) - \tilde{\mathbf{y}} \right\|_2^2 + \frac{\lambda_{\boldsymbol{\omega}}}{2} \left\| \boldsymbol{\omega} - \boldsymbol{\omega}_{\text{NN}} \right\|_2^2 \\ \mathbf{x}^{t+1} := \arg \min_{\mathbf{x}} \frac{1}{2} \left\| \mathbf{E}_{\boldsymbol{\omega}^t}^L(\mathbf{x}) - \mathbf{y} \right\|_2^2 + \frac{\lambda_{\mathbf{x}}}{2} \left\| \mathbf{x} - \mathbf{x}_{\text{NN}} \right\|_2^2, \end{cases} \quad (4.26)$$

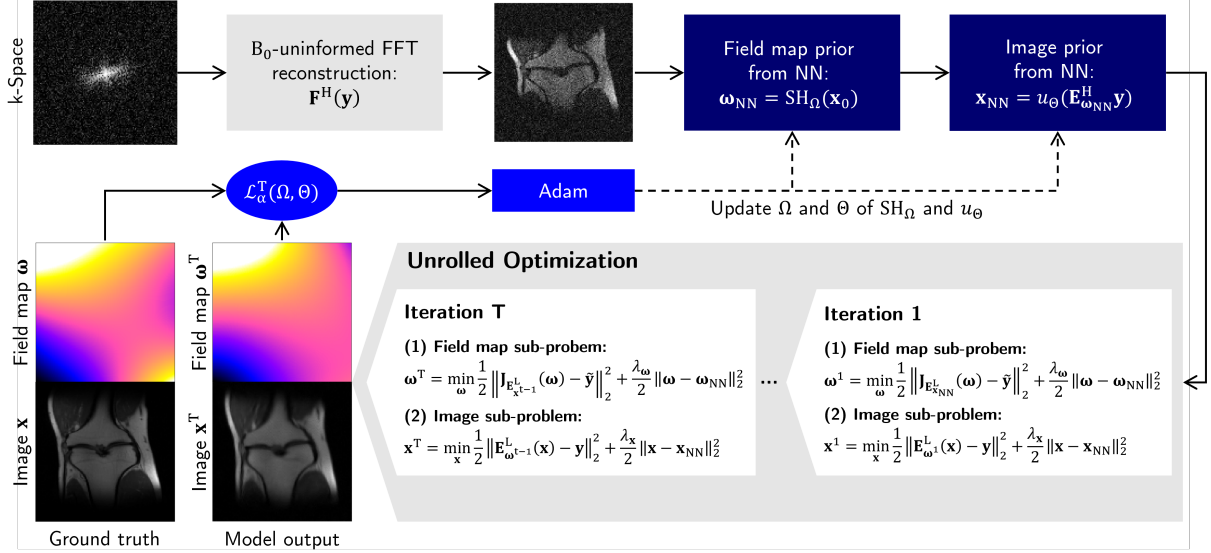


Figure 4.2: Overview of the end-to-end reconstruction network. The initial NNs estimates of the B_0 field map and the denoised image are obtained from B_0 -uninformed FFT reconstruction. Using $(E_{\omega_{NN}}^L)^H$ the noisy data is reconstructed under consideration of the estimate ω prior to denoising. Unrolling the end-to-end reconstruction network yields the final results x^T and ω^T . Although the CNNs SH_{Ω} and u_{Θ} are applied only once, their parameters as well as the regularization parameters λ_{ω} and λ_x are optimized during training. (Reproduced from [121] CC BY 4.0, 2024).

with initial values $x^1 := x_{NN}$ and $\omega^1 := \omega_{NN}$. The network $f_{\Theta, \Omega}^T$ maps the measured k -space data y to an estimate of the ground truth image and the ground truth field map. The NNs are applied only once before unrolling the algorithm to obtain the initial priors.

The optimization process, including both sub-problems, is described in Algorithm 1 and illustrated in Figure 4.2. The initial noisy image reconstruction, performed using FFT, is used to estimate a spatially smooth field map with the proposed SH-Net. Using the field map estimate, a B_0 -informed image reconstruction is performed, which yields a geometrically corrected but noisy image. By applying the denoising U-Net, a noise-free estimate of the reconstructed and distortion-corrected image is obtained. Starting from the NN priors, the method iteratively refines the estimates of the field map and the image by alternating between solving the two sub-problems. The strategy is to always define the operators E_{ω}^L and E_x^L with the best available solutions for ω and x , respectively. This means that at the very first iteration, the NN priors ω_{NN} and x_{NN} are used to define the operators $E_{\omega_{NN}}^L$ and $E_{x_{NN}}^L$. After T iterations, the process results in an optimized field map and a denoised, geometrically corrected image. During the end-to-end training process, the final estimates for the image and the field map are used to update the parameters of the underlying NNs involved in the regularization. The backpropagation from the final estimates to the trainable parameters of the NNs is performed through the entire unrolled optimization process.

Algorithm 1 Alternating optimization scheme to jointly update field map and image.

Require: Measured k-space data \mathbf{y}

- 1: Compute initial priors:
 - 2: $\boldsymbol{\omega}_{\text{NN}} \leftarrow \text{SH}_{\Omega}(\mathbf{F}^{\mathbf{H}}(\mathbf{y}))$ ▷ Estimate B_0 field map using SH-Net
 - 3: $\mathbf{x}_{\text{NN}} \leftarrow u_{\Theta}((\mathbf{E}_{\boldsymbol{\omega}_{\text{NN}}}^L)^{\mathbf{H}}(\mathbf{y}))$ ▷ Denoise geometrically corrected image by U-Net
 - 4: Initialize variables:
 - 5: $\boldsymbol{\omega}^1 \leftarrow \boldsymbol{\omega}_{\text{NN}}, \quad \mathbf{x}^1 \leftarrow \mathbf{x}_{\text{NN}}$
 - 6: **for** $t = 1, \dots, T - 1$ **do**
 - 7: **Field map update (sub-problem B):**
 - 8: $\tilde{\mathbf{y}}^t \leftarrow \mathbf{y} - \mathbf{E}_{\mathbf{x}^t}^L(\boldsymbol{\omega}^t) + \mathbf{J}_{\mathbf{E}_{\mathbf{x}^t}^L}(\boldsymbol{\omega}^t)$ ▷ Update residual
 - 9: $\mathbf{b}_{\boldsymbol{\omega}^t} \leftarrow \mathbf{J}_{\mathbf{E}_{\mathbf{x}^t}^L}^{\mathbf{H}}(\tilde{\mathbf{y}}^t) + \lambda_{\omega} \boldsymbol{\omega}_{\text{NN}}$ ▷ Compute gradient term
 - 10: $\mathbf{H}_{\boldsymbol{\omega}^t} \leftarrow \mathbf{J}_{\mathbf{E}_{\mathbf{x}^t}^L}^{\mathbf{H}} \mathbf{J}_{\mathbf{E}_{\mathbf{x}^t}^L} + \lambda_{\omega} \mathbf{I}$ ▷ Construct system matrix
 - 11: $\boldsymbol{\omega}^{t+1} \leftarrow \text{CG}(\mathbf{H}_{\boldsymbol{\omega}^t}, \mathbf{b}_{\boldsymbol{\omega}^t})$ ▷ Solve with conjugate gradient
 - 12: **Image update (sub-problem A):**
 - 13: $\mathbf{b}_{\mathbf{x}^t} \leftarrow (\mathbf{E}_{\boldsymbol{\omega}^{t+1}}^L)^{\mathbf{H}}(\mathbf{y}) + \lambda_{\mathbf{x}} \mathbf{x}_{\text{NN}}$ ▷ Compute gradient term
 - 14: $\mathbf{H}_{\mathbf{x}^t} \leftarrow (\mathbf{E}_{\boldsymbol{\omega}^{t+1}}^L)^{\mathbf{H}} \mathbf{E}_{\boldsymbol{\omega}^{t+1}}^L + \lambda_{\mathbf{x}} \mathbf{I}$ ▷ Construct system matrix
 - 15: $\mathbf{x}^{t+1} \leftarrow \text{CG}(\mathbf{H}_{\mathbf{x}^t}, \mathbf{b}_{\mathbf{x}^t})$ ▷ Solve with conjugate gradient
 - 16: **end for**
 - 17: **return** Final estimates: $\mathbf{z}^T := (\boldsymbol{\omega}^T, \mathbf{x}^T)$
-

4.3 Computational experiments

To evaluate the proposed method, different computational experiments are conducted. For the training of the proposed methods, a supervised approach is applied which requires ground-truth data for each sample. In the first part of this section, the simulation process is explained to obtain training data and, subsequently, the training of the NNs. Two comparison methods are introduced for the B_0 field estimation which are used to evaluate the novel SH-Net. In the next step, different algorithm unrolling schemes are evaluated for the end-to-end training. Finally, the proposed joint method is compared with purely model-based approaches.

4.3.1 Simulated low-field MRI data

For supervised learning of NNs, a relatively large amount of data is usually required for training. In addition, for supervised learning, a clean and undistorted reference is also required for the training process. Such ground truth data does not exist, which is why often simulated low-field data are used for supervised training. As an alternative to this, also self-supervised techniques could be employed, however, they still require huge amounts of data. Since only a fraction of low-field in vivo data is available at this point, synthetic low-field MRI data was simulated based on single-coil spin-echo acquisitions of the knee at 3 T from the *FastMRI* dataset [34]. By selecting a single

two-dimensional slice from each three-dimensional sample, 973 different complex-valued samples are obtained. Retrospective simulation of low-field MRI data from ground truth samples was implemented according to (2.7).

To obtain a variety of smooth random field maps, random SH coefficients are generated. The B_0 field of a low-field Halbach magnet can be relatively inhomogeneous, unless the field has been optimizing over several iterations of shimming. However, the inhomogeneities are rather slowly varying globally such that the low frequencies dominate. In order to take this into account, the entries of the generated coefficient vector in (4.24) are weighted exponentially decreasing. According to equation (4.24), ground truth field maps can then be generated from random SH coefficients. The weighting function is chosen so that field inhomogeneities in the range of ± 6000 ppm are obtained. This corresponds to the inhomogeneity of the PTB Halbach systems with only a few shim iteration used in Chapter 3 and to the inhomogeneity reported for other Halbach-based systems in the literature [116]. The mean and standard deviation of the correlation coefficient between the field maps showed no correlation or weak correlation, from which it can be concluded that the field maps are sufficiently diverse. For the generation of training samples, SH coefficients up to the order of 10 are used. The order of the simulation step is explicitly higher compared to the order used in the network from Section 4.2.1 to mimic the continuous character of the real field distributions.

Given ground truth field maps $\boldsymbol{\omega}$ and images \mathbf{x} , the k-space data \mathbf{y} is retrospectively computed according to (2.7). The retrospective simulation does not utilize the approximation of the forward operator, which is only used to accelerate the reconstruction for training and inference of the model. This also counteracts the inverse crime of using the identical operator in simulation and reconstruction. During training, the standard deviation of the Gaussian noise σ_y is randomly chosen from the interval $[0.2, 0.6]$. The test set was created using three distinct values $\sigma_y \in \{0.2, 0.4, 0.6\}$.

The simulation of SNR and B_0 field inhomogeneities is highly inspired by the characteristics of real low-field MRI data, so that the methods trained on these data should in principle also be applicable to data from actual low-field MR systems. The 973 unique image slices are randomly combined with a total number of 1800 randomly generated and weighted field maps. For training, validation and testing the data is split into fractions of 1260 (70%), 360 (20%) and 180 (10%), respectively. The resulting dataset is denoted by:

$$\mathcal{D} = \{(\mathbf{y}, \boldsymbol{\omega}, \mathbf{x}) | \mathbf{y} = \mathbf{E}(\mathbf{x}, \boldsymbol{\omega}) + \boldsymbol{\varepsilon}\}. \quad (4.27)$$

4.3.2 Pretraining of neural networks

The network pretraining was performed by minimizing the MSE between the predicted and ground truth field maps. With the previously defined dataset from (4.27), the corresponding loss function is expressed as:

$$\mathcal{L}^{\text{SH}}(\Omega) = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{y}, \boldsymbol{\omega}, \mathbf{x}) \in \mathcal{D}} \|\boldsymbol{\omega} - \text{SH}_{\Omega}(\mathbf{F}^{\text{H}}\mathbf{y})\|_2^2. \quad (4.28)$$

The network architecture, illustrated in Figure 4.1, consists of five encoder stages, each comprising two convolutional layers with a kernel size of 3 and a fixed filter size of 16. It is designed to predict 25 SH coefficients, corresponding to a maximum spherical harmonic order of 4. With 858 trainable parameters, the model strikes a balance between complexity and generalizability while mitigating overfitting. Pretraining was conducted over 400 epochs with a batch size of 1, using the ADAM optimizer with a fixed learning rate of 10^{-4} . The total pretraining time was approximately 44 min.

For image denoising, a U-Net architecture with residual connections was employed, following [123]. The U-Net featured four encoder steps, each containing two convolutional layers with a kernel size of 3, mirroring the SH-Net encoder structure. The initial filter size for the U-Net was set to 32. To train the denoising U-Net, a slightly different dataset is used, which contains only noise but no geometric distortions:

$$\mathcal{D}_u^{\text{Pre}} := \{(\mathbf{x}_0, \mathbf{x}) \mid \mathbf{x}_0 = \mathbf{F}^{\text{H}}\mathbf{y}, \mathbf{y} = \mathbf{F}\mathbf{x} + \boldsymbol{\varepsilon}\}. \quad (4.29)$$

By not using the previously defined dataset with geometric distortions caused by B_0 field inhomogeneities, it is ensured that the U-Net solely covers the denoising task. The correction of geometric distortions will be handled entirely by the SH-Net. The loss function for the denoising U-Net is defined by

$$\mathcal{L}^u(\Theta) = \frac{1}{|\mathcal{D}_u^{\text{Pre}}|} \sum_{(\mathbf{x}_0, \mathbf{x}) \in \mathcal{D}_u^{\text{Pre}}} \|\mathbf{x} - u_{\Theta}(\mathbf{x}_0)\|_2^2. \quad (4.30)$$

Pretraining is carried out over 400 epochs, with a batch size of 1, using the ADAM optimizer with a learning rate of 10^{-4} . The duration of training for U-Net was measured at approximately 21 min.

4.3.3 Comparison of B_0 field map estimation methods

Spatially B_0 field inhomogeneities cause relative phase differences dependent on the echo time. A common way to reconstruct the B_0 field map is to acquire multiple shifted

echos and calculate the B_0 map from the relative phase difference map. Therefore, the discrete signal equation (2.8) is extended by a non-zero positive echo shift to simulate the corresponding k -space data. The calculated B_0 field map from the phase difference of these two simulations serves as the first comparison method.

The ORN, published by [45] is a CNN with equally shaped convolutional layers and residual connections that estimates an undistorted image from a single complex-valued distorted image. In [46], the approach was adapted to estimate a field map from the initial distorted reconstruction by the same architecture. In both cases, the CNN architecture consists of 8 convolutional layers with 3 residual connections, a kernel size of 5×5 and 128 filters. The CNN approach is implemented as proposed by the authors and compared to the proposed SH-Net. To provide a fair comparison, the training was carried out as described in Section 4.3.2 by replacing SH_Ω with ORN_Ξ in loss-function (4.28).

4.3.4 Algorithm unrolling for end-to-end training

The end-to-end training unrolls the optimization algorithm as described in Section 4.2.3. This process is evaluated by three different cases for the iteration number T :

- SP-A Only the image sub-problem is minimized once, using the field map provided by the pretrained NN.
- $T = 1$ Each sub-problem is solved once within a single iteration. This case investigates the impact of the B_0 field map sub-problem.
- $T = 3$ Sub-problems A and B are alternated and solved over three iterations. Assuming the image or field map result after one of the minimization steps can be improved, this case investigates the interdependence of both sub-problems.

For all CG blocks, five iterations are performed to improve the data consistency of the outputs for each sub-problem (4.12) and (4.19). The hyperparameters of the pretrained NNs are further refined during the alternating optimization process using the ADAM optimizer. The regularization parameters $\lambda_{\mathbf{x}}$ and λ_{ω} , initially selected based on empirical evaluations, are included as training parameters and adjusted during training.

The learning rates are set to 10^{-6} for the NN hyperparameters and 10^{-5} for the regularization parameters. With a batch size of one, the training is pursued over 400 epochs. The duration of training varied according to the number of iterations T , with longer durations for higher values of T due to the increased complexity of the unrolling.

The end-to-end training is performed by minimizing the following loss function:

$$\mathcal{L}_\alpha^T(\Omega, \Theta) = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{y}, \omega, \mathbf{x}) \in \mathcal{D}} \left[\alpha \|\omega - \omega^T(\mathbf{y})\|_2^2 + (1 - \alpha) \|\mathbf{x} - \mathbf{x}^T(\mathbf{y})\|_2^2 \right], \quad (4.31)$$

where $\alpha > 0$ is a weighting factor chosen to balance the contributions of the errors associated with the field map and the reconstructed image. The network is implemented in PyTorch, leveraging its automatic differentiation engine to compute gradients of the trainable parameters.

4.3.5 Comparison to conventional approaches for joint image and B_0 reconstruction

Based on the formulation presented in equation (4.8), a similar joint reconstruction method was introduced in [120]. The approach uses a quadratic regularization term $\mathcal{R} = \|\nabla \mathbf{x}\|_2^2$ instead of the quadratic penalty terms derived from NNs. Although this method was not originally developed for low-field MRI, it is included for comparison due to its structural similarity to the proposed Algorithm 1. The quadratic penalty term is replaced by the ℓ_1 norm in this work, resulting in a TV minimization approach, which has been widely adopted for MRI reconstruction [124]. To solve the sub-problem with respect to the image, the PDHG method [67] is applied. The initialization for the field map is obtained from the relative phase difference, while the starting value for the image is given by $\mathbf{x}_0 := \mathbf{F}^H \mathbf{y}$. $T = 10$ iterations are used in total to alternate between the optimization of the sub-problems. Herein, the B_0 field map sub-problem is solved using the CG method for up to 100 iterations or until a tolerance of 10^{-5} is reached. The image sub-problem is solved with 128 iterations of the PDHG method. The regularization parameters λ_w and λ_x are selected by a grid search, ensuring the lowest MSE with respect to the ground truth. Although this grid search is not feasible in practical scenarios, it provides a fair basis for comparison by identifying the optimal parameters for this method. In essence, this comparison reveals the impact of using NNs for regularization instead of conventional approaches like TV.

A second comparison method, called the bi-model approach, is also based on conventional techniques to jointly estimate the image and B_0 field map [116]. This method uses separate models to optimize the field map and the image and was specifically designed for heavily distorted low-field MRI images, which is why it is chosen for comparison. The initial field map is derived from the relative phase difference of two acquisitions and smoothed using Tikhonov regularization. The optimization uses the original forward operator \mathbf{E} , i.e. without the approximation operator, with the extension by the echo time shift. Reconstruction of both measured data, \mathbf{y}_1 and \mathbf{y}_2 , is performed with \mathbf{E}_{TE_1} and \mathbf{E}_{TE_2} , respectively, and the resulting phases of the reconstructed images \mathbf{x}_1 and \mathbf{x}_2 are used to update the field map. The Split-Bregman algorithm is used for image reconstruction with two outer iterations, each involving up to 100 CG steps with a tolerance of 10^{-5} . Empirical evaluations indicated that three alternations between the sub-problems yielded

the best results for the simulated dataset. A grid search over the test dataset was used to determine the regularization parameters that minimized the image MSE.

4.3.6 Evaluation and metrics

To evaluate network performance, test data was generated with distinct noise levels by simulating the forward problem (2.7) using Gaussian noise with standard deviations $\sigma_y \in \{0.2, 0.4, 0.6\}$. The field maps estimated by the proposed SH-Net and the ORN [46] were compared to the field map obtained from the relative phase difference of two measurements with an echo time difference of 20 ms and the simulated ground truth.

For evaluation of the field map, the Mean Absolute Error (MAE) is computed. The evaluation was performed on the entire field map and within a circular Region of Interest (ROI) at the center of the map to exclude phase noise from areas outside the object of interest. The ROI was defined with a radius of 32 pixels, encompassing approximately 20 % of the total pixels. Data augmentation was applied to ensure that all test slices contained meaningful information within this ROI. The reconstructed images were assessed using Root Mean Squared Error (RMSE), Peak Signal-to-Noise Ratio (PSNR), and Structural Similarity (SSIM) [125]. While RMSE and PSNR are global metrics, SSIM evaluates image quality based on contrast, luminance, and structural comparison, making it more sensitive to distortions such as blurring [125].


The regularization parameters for the comparative methods were optimized separately for each noise level using a sample-wise grid search with the ground truth. Usually a ground truth is not available, which is why this technique cannot be employed in practice.

4.4 Results


The preceding section delineated an overview of the computational experiments that were conducted to systematically assess the efficacy of the proposed reconstruction methodology. Initially, the section presents the outcomes of the novel B_0 field estimation method, followed by the results of the different algorithm unrolling schemes and subsequently, the comparison of the proposed joint reconstruction approach against conventional methods.

4.4.1 Performance comparison of field map estimation methods

The results of the field map of SH-Net and ORN, which are obtained from a single echo, are compared to the field map obtained from the relative phase difference of two measurements.



(a) Boxplot comparing the MAE of the field maps obtained from relative phase difference of two measurements with different echo times, our proposed SH-Net and the ORN, whereby the latter are based on a single echo. The comparison is carried out at distinct noise levels $\sigma \in \{0.2, 0.4, 0.6\}$ for the entire field maps (left) and within the ROI (right).



(b) Field maps obtained from our proposed SH-Net and the ORN compared to the phase difference map and the ground truth at $\sigma = 0.4$.

Figure 4.3: Comparison of SH-Net and ORN on the test dataset after pretraining. (Reproduced from [121] CC BY 4.0, 2024).

Figure 4.3 compares the MAE of the results at three distinct noise levels given by the standard deviation of the simulated k -space data, that is, $\sigma_y \in \{0.2, 0.4, 0.6\}$. For both the entire field map and for the ROI, the SH-Net outperforms the other methods by showing lower MAE values. For the entire field map the MAE of the SH-Net is 324.7 ± 145.3 ($\sigma_y = 0.2$) Hz, 324.6 ± 144.1 ($\sigma_y = 0.4$) Hz and 327.2 ± 142.4 ($\sigma_y = 0.6$) Hz, while the ORN leads to an MAE of 711.3 ± 231.2 ($\sigma_y = 0.2$) Hz, 726.5 ± 240.7 ($\sigma_y = 0.4$) Hz and 742.8 ± 247.5 ($\sigma_y = 0.6$) Hz. The highest error was found for the phase difference maps with 852.3 ± 326.2 ($\sigma_y = 0.2$) Hz, 945.7 ± 325.0 ($\sigma_y = 0.4$) Hz and 1107.8 ± 350.4 ($\sigma_y = 0.6$) Hz.

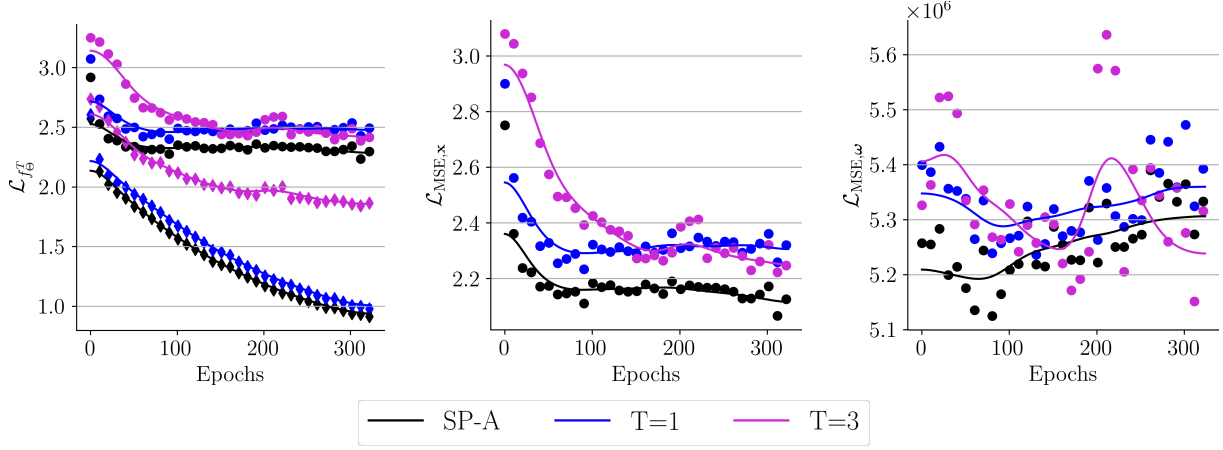


Figure 4.4: Comparison of different metrics during further fine-tuning of the pretrained u_Θ and SH_Ω by unrolling the proposed scheme in Algorithm 1 with SP-A, $T = 1$ and $T = 3$ alternations. For SP-A, the field map is entirely defined by the output of the SH-Net and only the image is ensured to be data-consistent by solving problem (4.12). (Reproduced from [121] CC BY 4.0, 2024).

4.4.2 Comparison of unrolling strategies for the proposed algorithm

The proposed end-to-end network addresses two sub-problems in an alternating fashion, with each sub-problem optimized by the CG method. In resolving these sub-problems, the operator is formulated with regard to one specific variable, with the other variable kept constant. Between consecutive steps, i.e. sub-problem A or B, the operator is updated based on the most recent solution.

Before evaluating the different algorithm unrolling strategies, the time-segmented approximation of the reconstruction operator is validated. A comparison of the point spread function between the full encoding operator \mathbf{E} and its approximation \mathbf{E}^L is performed. The field map used for the comparison exhibits a linear increase, with a maximum of 7517 ppm. With the full encoding operator \mathbf{E} , the maximum value of the reconstructed point reaches 99.99% of the original value. Using the approximation \mathbf{E}^L with $L = 24$ results in 98.91% of the original magnitude. Based on this evaluation, the approximation is deemed sufficiently accurate.

Using the approximated operator, this section examines the various algorithm unrolling strategies introduced in Section 4.3.4. Figure 4.4 illustrates the performance of the training in the different strategies. The methodologies $T = 1$, $T = 3$, as well as a particular case SP-A, where only sub-problem A is considered and the field map is defined entirely by the SH-Net, are compared. Figure 4.4 shows the behavior of the components of the loss function defined in (4.31). The left box plot shows the weighted combination for $\beta = 10^{-6}$, the center box plot shows the validation error of the image, and the right box plot shows the validation error of the field map, respectively. The validation error is represented by

dots, while the training error is represented by crosses for each epoch. Upon comparing the three different cases, the NNs converge at similar levels. With a larger T , the error increases at the beginning of training. Based on the error courses, the gap between the training and validation errors is the lowest for case $T = 3$. The case SP-A tends to have the lowest error over the combined and image-specific MSE. In contrast to the image MSE, which decreases by -25.12% (SP-A), -22.07% ($T = 1$) and -28.49% ($T = 3$), only small improvements of MSE can be achieved for the field map. For the field map-specific MSE, SP-A slightly increases by 0.32% over the epochs, while $T = 1$ and $T = 3$ slightly decrease by -1.26% and -3.98% . Within the ROI, $T = 3$ has the highest error, while $T = 1$ has the lowest. Here, the validation error decreases relatively by -5.83% (SP-A), -6.71% ($T = 1$) and -5.15% ($T = 3$).

In terms of performance, all the approaches work well and the validation error with respect to the field map in Figure 4.4 shows that the nonlinear SP-B is challenging. In contrast to performance, the behavior with regard to overfitting can be distinguished more clearly and $T = 3$ shows a considerably lower tendency to overfitting, as can be seen in the left box plot of Figure 4.4. Consequently, $T = 3$ was used in the further analysis due to the small gap between training and validation error, performance comparable to $T = 1$ and the least tendency to overfitting.

4.4.3 Comparison to conventional approaches for joint image and B_0 reconstruction

The proposed method employs a variable splitting strategy to address the B_0 -informed image reconstruction and denoising problem through alternating optimization between two sub-problems. To regularize these sub-problems, two NNs are utilized to estimate the noise-free image and the B_0 field map. During the end-to-end training process, the parameters of these networks are refined by backpropagating through the unrolled algorithm. The proposed combination of model-based reconstruction and deep learning-based regularization is compared to the conventional methods introduced in Section 4.3.5.

All methods are compared on the simulated test data introduced in Section 4.3.1 with Gaussian noise at three different standard deviation levels $\sigma_y \in \{0.2, 0.4, 0.6\}$. Figure 4.5 illustrates the comparison of RMSE, PSNR, and SSIM between the proposed method with $T = 3$, the bi-model, the joint-model, and a B_0 -uninformed FFT reconstruction. Since the final image quality is of particular interest, only the image metrics of the different methods are compared. Complementary to these results, median and Interquartile Range (IQR) of RMSE, PSNR, SSIM and MAE are reported in Table 4.1 for both the image and the field map at the distinct noise levels $\sigma_y \in \{0.2, 0.4, 0.6\}$. The table of results (Table 4.1) provides median and IQR values for RMSE, PSNR, SSIM, and MAE for both the image and the field map across all noise levels. Median and IQR are used instead of mean




Figure 4.5: Boxplot comparison of image RMSE, PSNR and SSIM for B_0 -uninformed FFT reconstruction, joint-model, bi-model, and end-to-end trained unrolled optimization with $T = 3$. (Reproduced from [121] CC BY 4.0, 2024).

σ_y		B_0 -uninformed FFT/ Phase difference map		Joint-Model		Bi-Model		Proposed, T = 3	
		Median	IQR	Median	IQR	Median	IQR	Median	IQR
0.2	Image RMSE (a.u.)	1.7369	2.1287	0.9204	1.1941	0.8465	1.1148	0.1386	0.1851
	Image PSNR (a.u.)	22.4169	3.6544	27.2592	3.9025	27.7094	3.0306	28.058	3.8212
	Image SSIM (a.u.)	0.1955	0.0987	0.4204	0.094	0.4332	0.1082	0.6906	0.0832
	Field Map MAE (Hz)	813.1	440.2	630.6	364.2	201.5	115.8	301.0	157.6
0.4	Image RMSE (a.u.)	2.1545	2.7612	1.0836	1.8792	0.9616	1.2318	0.1393	0.1857
	Image PSNR (a.u.)	20.4635	2.8819	25.6011	3.4752	26.4848	2.7191	28.0603	3.5063
	Image SSIM (a.u.)	0.1254	0.0643	0.3163	0.089	0.3746	0.0679	0.6661	0.0734
	Field Map MAE (Hz)	932.3	420.7	657.7	370.6	205.6	121.7	306.4	154.1
0.6	Image RMSE (a.u.)	2.6819	3.6108	1.5388	2.8575	1.1324	1.5082	0.1449	0.1888
	Image PSNR (a.u.)	18.4514	2.3879	22.9965	3.8283	24.7952	2.4025	27.7169	3.3339
	Image SSIM (a.u.)	0.0869	0.0456	0.2228	0.0999	0.3139	0.0601	0.6293	0.0878
	Field Map MAE (Hz)	1089.0	436.4	672.6	362.8	210.9	127.0	307.1	149.2

Table 4.1: Median and IQR of MSE, PSNR and SSIM for the reconstructed image \mathbf{x} with joint-model, bi-model, and unrolled optimization $T = 3$. The second table shows the MAE for the field map of the corresponding methods. The data is shown for the noise levels given by $\sigma_y \in \{0.2, 0.4, 0.6\}$. (Reproduced from [121] CC BY 4.0, 2024).

and standard deviation, because the observable error distributions are rather non-normal according to the boxplots. Both values were calculated per method and metric for the overall test dataset.


The results show that, compared to the B_0 -uninformed FFT reconstruction, all other methods perform better on the reported metrics. In terms of noise robustness, errors increase progressively with increasing noise levels σ_y of the input data for the B_0 -uninformed FFT, joint-model, and bi-model approaches. In contrast, the proposed method maintains consistent performance in different values of σ_y , with only a minor reduction in SSIM observed. Moreover, an examination of the whisker graphs for RMSE in Figure 4.5 and the corresponding IQR reveals that the proposed method produces fewer and less pronounced outliers, further underscoring its robustness.

The MAE values for the field map confirm these findings. Field maps derived from phase difference maps without additional regularization yield the highest MAE, followed by the joint-model approach. Although the bi-model approach achieves the lowest MAE for field maps, it marginally outperforms the proposed method. It is worth noting that both the bi-model and joint-model approaches rely on relative phase difference maps derived from two acquisitions, whereas the SH-Net employs a single acquisition, offering a more efficient solution. Within the ROI, all methods show comparable performance, as summarized in Table 4.1.

Figure 4.6 illustrates a comparison of image reconstructions for two representative samples from the test dataset, obtained using B_0 -uninformed FFT, joint-model, bi-model and the proposed unrolled algorithm with $T = 3$. Reconstructions are evaluated at two distinct noise levels, $\sigma_y = 0.3$ and $\sigma_y = 0.6$, with all intensities scaled relative to the ground-truth reference image. This comparison highlights the consistent performance that can be achieved at varying noise levels by using NN-based regularization.

Figure 4.6 shows exemplary image and B_0 field reconstruction results for noise levels $\sigma_y = 0.3$ and $\sigma_y = 0.6$. The joint-model and bi-model approaches deliver comparable performances in terms of SNR, while the proposed method demonstrates a slightly lower SNR. At both noise levels and for both samples, the unrolled reconstruction algorithm produces smooth images with minimal blurring artifacts. In contrast, the joint-model and bi-model approaches exhibit a substantial increase in blurring. The bi-model approach reconstructs geometric distortions with greater accuracy, although it is significantly affected by blurring. The joint-model approach exhibits the highest residual error after correction, which is particularly evident near the object boundaries. Within the object center, the joint-model approach achieves higher accuracy compared to the bi-model approach. The proposed method achieves relatively low overall error and high correction accuracy, particularly in the central regions of the image. However, all methods retain a residual error that remains relatively consistent between the noise levels evaluated, $\sigma_y \in \{0.3, 0.6\}$. Figure 4.6b depicts the corresponding field maps for the images shown in Figure 4.6a. Compared to the field maps derived from the relative phase difference map, all evaluated methods produce smooth representations of the spatial field distribution. The field map from the joint-model reconstruction exhibits substantial errors near the edges but lower errors at the center. The bi-model approach and the proposed method achieve similarly low overall errors, which align with the metrics reported in Table 4.1. As noise levels increase, the error remains constant. This robustness in field map estimation under varying noise conditions, as well as the spatial accuracy of the field maps, is consistent with the observations from Figure 4.6.

The bi-model approach requires approximately 15s for reconstruction per sample. The



(a) Reconstructed magnitude images, error in arbitrary units.


(b) Reconstructed B_0 -field maps, error in Hz.

Figure 4.6: Comparison of reconstruction results from B_0 -uninformed FFT or phase difference map, joint-model, bi-model and the proposed end-to-end trained unrolled optimization. The first and the third row show the magnitude images or the field maps respectively for noise levels $\sigma_y = 0.3$ and $\sigma_y = 0.6$. Rows two and four depict the absolute error map with respect to the ground truth image or the field map respectively. (Reproduced from [121] CC BY 4.0, 2024).

joint-model approach achieves reconstruction within approximately 5 s. The proposed model, once trained, requires less than 1 s for a sample size of 128×128 . It should be noted that none of the implementations were optimized for fast GPU computation, and therefore these results should be interpreted as relative performance estimates.

4.5 Discussion

This chapter introduced a novel end-to-end training approach for the joint reconstruction of image and B_0 field maps. Utilizing a variable splitting technique the proposed method integrates model-based reconstruction and deep learning in an end-to-end training, yielding the first algorithm of this kind for low-field MRI.

For estimating the B_0 field map, a novel SH-Net architecture was proposed, which estimates SH coefficients. The spatial distribution of the field is derived via a linear combination of the SH basis functions and the estimated coefficients, which are learned from B_0 -uninformed FFT reconstructions. The design of the SH-Net is based on the physical property that the B_0 field maps exhibit inherent smoothness. This smoothness is ensured by estimating the SH coefficients instead of the field map directly, thereby reducing the risk of introducing local artifacts. Although this approach is effective for representing global field variations, it presents a limitation when dealing with highly localized field inhomogeneities. In contrast, alternative architectures, such as the ORN, are capable of resolving very localized changes. However, given that the network is specifically designed to model variations in the background B_0 field, this restriction provides an advantage. The enforced smoothness is expected to yield more reliable and robust field map estimations, particularly under the condition of poor SNR, which are used for distortion correction by time-segmented approximation of conjugate phase reconstruction [118], [126], [127]. Previous methods for joint reconstruction of images and field maps have been limited to 3 T MRI systems and relatively homogeneous B_0 field distributions of < 3 ppm [120], [127], [128]. In this work, low-cost MR systems are considered, which exhibit significantly higher relative field inhomogeneities [52]. At the time of the measurements, the PTB system had B_0 field inhomogeneities above 4000 ppm. In addition, such POC systems may experience further fluctuations due to environmental factors such as temperature variations that further degrade the B_0 field homogeneity [115]. The proposed network architecture demonstrated the ability to estimate B_0 field maps with strong field inhomogeneities from the data which can be used to mitigate geometric distortions caused by it. The performance of SH-Net was evaluated by comparing it to field maps derived from relative phase difference and the ORN method, as proposed by [45], [46]. The results indicate that incorporating SH coefficients yields the most accurate and smooth field maps. In contrast, ORN lacks prior knowledge of spatial smoothness and requires 829 185 hyperparameters. The SH-Net

model, with only 858 hyperparameters, demonstrates superior performance, generating more realistic field maps.

In combination with a U-Net architecture, trained to denoise the low-field MRI data, the NNs are used as regularizers when solving the joint reconstruction problem. The alternating optimization is unrolled and used to fine-tune the NNs by end-to-end training, with different unrolling schemes, i.e. number of alternations, being evaluated. Considerations included the image-related sub-problem alone, as well as one and three alternations between the two sub-problems. Although the overall performance in all of these configurations was similar, the initial validation error was influenced by the number of iterations, T , used in the approach. Since sub-problem B is nonlinear, it is unclear whether performance can be further improved. Also, it should be noticed that the reconstruction operator itself is an approximation and that the nonlinear optimization required Taylor approximation to obtain a solution. However, the inclusion of data consistency contributes to the stability of the solution and enhanced robustness. Despite a comparable final performance, differences in validation error at the beginning of training suggest that a higher number of iterations may be beneficial. For improved robustness, $T = 3$ was selected, accepting minor performance trade-offs. Future work could reduce training durations by manually implementing the backpropagation step, bypassing the automatic differentiation engine [129].

Training with simulated knee images was evaluated at epoch 320, as spontaneous error deflections were observed near epoch 400 and the validation error had already converged. The loss function used during training combined the image and the field map MSE, with a constant weighting factor. However, different convergence rates were observed in the validation errors, suggesting that static weighting might not be the optimal choice for joint end-to-end training, potentially leading to training instabilities after epoch 320. Perhaps, gradient normalization and rotation techniques could improve the simultaneous optimization of both networks [130].

To classify the performance of the approach, the method with $T = 3$ alternations between the sub-problems was compared to both the joint-model and the bi-model approach, using the test data of the simulated low-field MRI knee images. The joint-model approach does not incorporate SH coefficients, resulting in insufficient information in areas with no acquired signal. This leads to lower field map accuracy at the object center, similar to the performance of the bi-model approach and the proposed method. The bi-model approach also uses SH coefficients, providing comparable field map accuracy, but requires reconstruction times on the order of minutes and at least two acquisitions with different echo times. In contrast, the proposed method achieves reconstruction times on the order of milliseconds and requires only a single acquisition. The more extensive spatial description of the field distribution allows for more accurate image reconstructions, especially at the

object boundaries. However, residual errors from imperfect geometric distortion correction remain visible in all approaches. For model-based methods, a reasonable search space was determined for regularization parameters and the best coefficient combination was selected to reconstruct the test data. The parameter configuration that minimized the image MSE resulted in excessively blurry images at a noise level of $\sigma_y = 0.6$. While this empirical parameter search is not applicable in practice, it demonstrates that the proposed method outperforms the comparison methods, even when their hyperparameters are optimized using ground-truth data. Regardless of noise level, the NNs-based approach consistently outperforms classical regularization methods, providing sharper images and significantly better SNR. This is reflected in the improved SSIM scores, as shown in Figure 4.4, compared to PSNR. The k -space samples used in this study were retrospectively simulated from a ground truth to obtain data for training, validation, and testing. This type of data cannot be generated in practice, but it facilitated comparison between different SNR levels and field map configurations. Using a wide range of system characteristics, it covers various experimental conditions.

4.6 Summary

The novel approach for joint image and B_0 field reconstruction in low-field MRI, addresses the main challenges of low SNR and high B_0 field inhomogeneities. The proposed method integrates physics-informed modeling with deep learning in an end-to-end trainable network. For each of the problems, i.e. image and field map, separate networks are trained and used as regularizers in an unrolled algorithm. Image denoising is achieved by a U-Net architecture, while the novel SH-Net architecture was proposed to estimate spatially smooth B_0 field maps from the complex-valued image data. The end-to-end reconstruction alternates between updates of the image and the field map throughout the unrolled algorithm. In accordance with the results from the Nexus console, introduced in Chapter 3, realistic training data could be simulated, which was used throughout this chapter to evaluate the performance of the method and compare it to existing methods. Based on quantitative evaluations on the simulated data, superior performance of the proposed algorithm could be demonstrated by higher SNR and more accurate geometric distortion correction. The combination of model-based reconstruction with data-driven regularization leads to more robust performance, especially for varying noise levels. While this chapter focussed on the detailed methodology and the quantitative evaluation of the method, the application on measured data, acquired with the Nexus console, is evaluated in the subsequent chapter.

5 Application of image and B_0 field reconstruction in vivo

The previous chapter introduced a method that integrates model-based reconstruction with deep learning-based regularization for joint image and B_0 field reconstruction in low-field MRI. A time-segmented reconstruction operator was incorporated in an iterative reconstruction algorithm which was unrolled for end-to-end training, to correcting geometric distortions and denoise the image. Using simulated data, the optimal network configuration was determined through quantitative evaluation against the ground truth. While the method outperformed conventional approaches, its training and validation were limited to simulated knee data.

Applying it directly to measured data, such as in vivo brain scans acquired with the Nexus console, is challenging because the supervised model was trained on a data distribution that differs significantly from real measurements [42]. To address this limitation, method and training are adjusted to improve the applicability to measured data. This chapter presents the modifications, followed by an overview of the computational experiments and the results obtained from measured in vivo brain scans acquired with the Nexus console.

5.1 Method adaption

In the proposed joint reconstruction approach, two distinct NNs were introduced to estimate the regularizers for both the image and the B_0 field sub-problems. The SH-net architecture, described in Section 4.2.1, estimates the field map using SH coefficients, ensuring spatial smoothness of the network output. As described by (2.8), the B_0 field influences the phase of the MR signal. By explicitly incorporating phase information, the network’s sensitivity to phase variations is expected to improve, leading to more accurate B_0 field map estimation. While the overall architecture remains largely unchanged, it is adapted to process magnitude and phase data instead of real and imaginary components. As a preprocessing step, the phase of the B_0 -uninformed FFT reconstruction is unwrapped before being processed by the SH-net.

For image denoising, the U-Net architecture, introduced in Section 4.2.2, is employed. Since both networks contribute to optimizing the image and B_0 field map alternately, preserving accurate phase information motivates the use of magnitude and phase representations in

the denoising network as well. Instead of processing real and imaginary components, the complex-valued image is decomposed into magnitude and phase, which are handled by separate network branches. Each branch integrates a residual connection of its respective input, i.e. magnitude or phase, which is added to the network output through a convex combination. The outputs of the network branches are then recombined to form a complex-valued image. Utilizing magnitude and phase information necessitates different activation functions within the network. The magnitude, being strictly positive, is processed using a ReLU activation function [131], while the phase, which can take negative values, employs a leaky ReLU activation [132]. All other architectural parameters, including the number of convolutional layers, initial filter sizes, and the depth of encoding stages, remain unchanged.

5.2 Computational experiments

The goal of this chapter is to test the model performance on the measured in vivo brain data, acquired by the Nexus console and presented in Chapter 3. Therefore, some adaptations were presented in the previous section regarding the NNs which estimate the regularization terms that necessitate a new training of the models. Additionally, some changes are applied to the training process, which are described by the following. Subsequently, the testing on the measured low-field MRI data is described.

5.3 Model training

In addition to the knee data, the *FastMRI* data set contains 2678 samples of complex-valued T_2 -weighted brain data samples, acquired with a TSE sequence [34]. Due to the higher similarity to the measured in vivo data at the low-field MRI scanner, the simulated training data is adapted to the available brain samples. Since here only multicoil data is provided, the coil data is combined in a first step using an implementation of the Walsh algorithm [133], to obtain equivalent complex-valued single coil data. From the three-dimensional coil-combined data, two-dimensional samples are extracted from the center. The FOV of the *FastMRI* data remains unchanged, because it compares to the FOV of the acquisition with low-field MRI scanner. To obtain comparable resolution, the coil-combined data is further augmented by downsampling to the dimensions 100×120 (PE \times RO). The downsampled data are normalized in terms of magnitude and phase, such that magnitude is within the interval $[0, 1]$ and phase within $[-\pi, \pi]$. Retrospective simulation as explained in Section 4.3.1, is used to add geometric distortions by randomly




Figure 5.1: Photograph of the ACR phantom with a diameter of 203mm and a height of 173mm, mounted in a solenoid RF coil in side the OSI²One low-field MRI scanner. The phantom contains a ring and a grid structure as can be seen on the photograph.

generated B_0 field maps and Gaussian noise. The simulated low-field MRI brain data is split into 1874/535/269 samples for training, validation and testing.

With the simulated low-field MRI brain dataset, pretraining is performed for each of the networks, while the SH-net estimates a B_0 field map from the phase information of the Fourier transformed input, and the U-Net denoises complex-valued images. Both networks are pretrained for 300 epochs using a ℓ_1 loss function and the AdamW optimizer with a weight decay of 10^{-2} . To ensure better generalizability, the SH-net is pretrained with a batch size of 20 and a learning rate of 10^{-3} , while the U-Net is pretrained with a batch size of 8 and an initial learning rate of 10^{-4} . Using the pretrained NNs, the joint reconstruction network is constructed with $T = 3$ alternations.

5.4 Model testing

Before applying the model which was trained with the simulated brain data on the measured data from the low-field MRI scanner, the model is initially evaluated on the simulated test data. Because a ground truth is available for the simulated data, the results can be quantified in terms of Normalized Root Mean Squared Error (NRMSE) and PSNR. Following the evaluation on the simulated test dataset, the method is applied to data which was measured with the Nexus console, introduced in Chapter 3. Two different TSE-based sequences were employed, as introduced in Section 3.2.2, yielding PD- and T_2 -weighted three-dimensional images.

Measurements have been carried out at two different systems at LUMC and at PTB, where the PTB system suffered from higher magnetic field inhomogeneities ~ 4000 ppm in the measurement FOV at the time of the measurement. To assess the method’s capabilities to correct geometric distortions under conditions similar to the simulations, images are acquired using a phantom from the American College of Radiology (ACR). Figure 5.1 shows a photograph of the ACR phantom in the scanner, containing a ring structure and a grid structure. The structures in the phantom are designed for quality assessment of MRI systems. Unfortunately, a B_0 -mapping sequence was not available at the point of the measurement, which is why the geometric distortion correction is evaluated based on the images obtained from the structures in the phantom. After applying the method to measured data from an ACR phantom, it is further tested on in vivo images of a healthy volunteer, as presented in Section 3.3.1. This data was acquired at the LUMC system with a B_0 field inhomogeneity below 1000 ppm. Clearly identifiable geometric distortions in the shown B_0 -uninformed reconstructions are not visible. Nevertheless, it is questionable, if the method can detect the differences in the B_0 field map and adapt to it.

For each of the T_2 - and PD-weighted data, including in vivo and phantom measurements, a single slice of the measured three-dimensional data is used to test the joint reconstruction approach. It undergoes magnitude and phase normalization as described for the simulated low-field MRI data. Due to a lack of sufficient PD-weighted data for the simulation of PD-weighted low-field MRI data, the training is solely based on T_2 -weighted data. Therefore, the direct application to the PD-weighted data from the brain and the phantom remains challenging. However, compared to the simulated knee data, the adapted training data facilitates a higher structural similarity to the measured data. Beside the quantitative comparison for the simulated test data, the results for the measured data are only evaluated qualitatively due to a lack of references.

5.5 Results

With the introduced adaptation of the NNs and the supervised training, the following section first validates the modifications by evaluating the simulated test data. Furthermore, the results from the phantom study and the application in vivo are evaluated.


5.5.1 Simulated low-field MRI brain data

Prior to the application of the trained model to real low-field MRI data, its performance is assessed using the simulated test dataset. To demonstrate the performance, the end-to-end trained joint model $f_{\Theta, \Omega}^T(\mathbf{y})$ with $T = 3$ alternations is applied to a representative sample


from the test dataset. The output of the network $\mathbf{z}^T := (\boldsymbol{\omega}^T, \mathbf{x}^T)$ encompasses both the image and the field map results, which are referred to as $\boldsymbol{\omega}^{T=3}$ and $\mathbf{x}^{T=3}$, respectively. Figure 5.2a compares the magnitude images from B_0 -uninformed FFT of the input data denoted as $\text{FT}(\mathbf{y})$, the application of the NNs denoted by \mathbf{x}_{NN} , the end-to-end network denoted by $\mathbf{x}^{T=3}$ and the ground truth. Their corresponding field map results are displayed in Figure 5.2b. All the compared magnitude images and B_0 field maps in Figure 5.2 are presented on the same scale. To obtain \mathbf{x}_{NN} , a B_0 -informed image reconstruction is performed with the field map $\boldsymbol{\omega}_{\text{NN}}$ estimated by SH-Net, before the denoising U-Net is applied. In comparison to the end-to-end network output $\mathbf{x}^{T=3}$, this step only utilizes the network priors and does not consider any optimization. For each image and field map, a magnified segment is displayed underneath to provide a more detailed view on the reconstruction results.

Without correction, the input image exhibits noticeable geometric distortions, particularly in the upper right and lower left regions of the brain. Using the field map estimated by SH-Net within a B_0 -informed image reconstruction which is denoised by the U-Net already yields a significant improvement over the input image. Employing the joint model further refines the reconstruction, as can be seen in the zoomed-in plot at the upper boundary of the left ventricle. Notably, a small dark spot present in the \mathbf{x}_{NN} image above the left ventricle is corrected in the optimized image, $\mathbf{x}^{T=3}$. While the model effectively suppresses noise, it also introduces some blurring compared to the ground truth, particularly where stronger geometric distortions have been corrected. The overall shape of the estimated field map closely matches the ground truth, with no substantial differences between $\boldsymbol{\omega}_{\text{NN}}$ and $\boldsymbol{\omega}^{T=3}$. Both B_0 field map results show discrepancies in the upper right corner and at the bottom in comparison to the ground truth. These regions, however, contain mostly noise and lack meaningful information, making accurate B_0 field map determination challenging.

To assess the generalizability of the method across different samples, NRMSE and PSNR of the complex-valued output image are evaluated on the test dataset. The box plots in Figure 5.3 compare the B_0 -uninformed noisy reconstruction obtained by FFT against the denoised image \mathbf{x}_{NN} after reconstruction with the estimated B_0 field map and the final output after alternating optimization, $\mathbf{x}_{T=3}$. The slight improvement observed in the example from Figure 5.2 is also reflected in the box plot evaluation. The most substantial enhancement of the B_0 -uninformed FFT reconstruction is achieved by the cascaded application of the denoising U-Net and the SH-Net, yielding a reduction of the mean NRMSE by 32.1 % and an increase of the mean PSNR by 29.4 % respectively. Further refinement is obtained with the end-to-end network, which integrates the alternating model-based reconstruction, leading to an additional 2.1 % decrease in mean NRMSE and a 1.1 % increase in mean PSNR.



(a) Magnitude images of the input, i.e. $\text{FT}(\mathbf{y})$, the U-Net output and the optimized output in comparison to the ground truth.



(b) Field map estimations from the SH-Net and joint optimization in comparison to the ground truth field map.

Figure 5.2: Image and field map results of the joint optimization with intermediate NN results for $T = 3$ alternations of the two sub-problems.

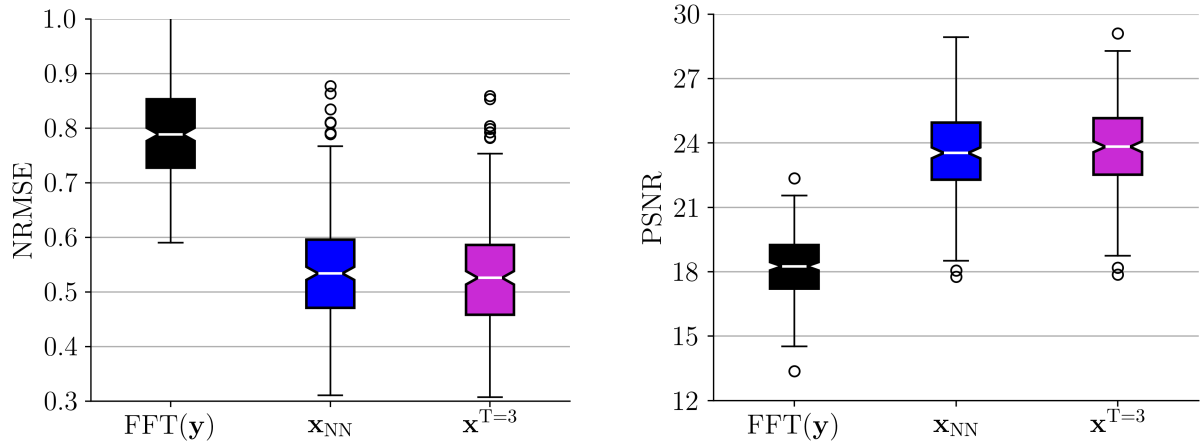


Figure 5.3: Evaluation of input, NN and optimized complex-valued images from the test dataset with respect to the ground truth.




Figure 5.4: Image magnitude and B_0 field map reconstruction of the joint-model reconstruction with $T = 3$ compared to B_0 -uninformed FFT reconstruction, when applied to the PD-weighted data from an ACR phantom measurement at PTB. The closeup of the circular structure illustrates the correction capabilities of geometric distortion.

5.5.2 ACR phantom measurements


To evaluate the capabilities of correcting the geometric distortions in measured data with the proposed end-to-end reconstruction method, PD-weighted contrast of the ACR phantom was acquired using the low-field scanner MRI at PTB in Berlin. Figure 5.4 compares the B_0 -uninformed FFT reconstruction of the measured data to the model output with $T = 3$ alternations between the sub-problems. The resulting magnitude image $x^{T=3}$

and the B_0 field map $\omega^{T=3}$ are shown. The ring structure was chosen for the comparison in Figure 5.4 due to its improved visibility in contrast to the grid structure. Due to the high B_0 field inhomogeneities, the B_0 -uninformed reconstruction leads to geometric distortions which are particularly visible on the left boundary of the phantom. Also, the inner ring structure is distorted which can be seen in the zoomed-in section of Figure 5.4. Unfortunately, there is an additional artifact superimposed to the geometric distortion from B_0 field inhomogeneities. It appears as a halo-like artifact in the vertical PE dimension of the image at the upper and lower boundaries of the phantom. Recent experiments indicate that it is attributable to the gradient system, and not dependent of the B_0 field. However, at the time of this work, the artifact is still under investigation. As it was only observed at the PTB system, the origin in the console or the sequence can most likely be excluded. Regarding the denoising performance of the method, the background noise could be removed successfully, however, in comparison to the B_0 -uninformed FFT reconstruction, the resulting image is relatively blurry. Using the proposed unrolled network for the reconstruction, yields a partially corrected image of the ACR phantom. Particularly the left boundary of the phantom is corrected in comparison to the B_0 -uninformed FFT reconstruction. Comparing the closeup of the ring structure inside the ACR phantom, the joint model reconstruction successfully corrects for the inner distortion.


5.5.3 Application in vivo

Figure 5.5 illustrates the application of the model to two different slices of the T_2 -weighted in vivo data acquired at LUMC, as presented in Section 3.3.1. The scanner used for the data acquisition exhibits a homogeneous B_0 field, with inhomogeneities below 1000 ppm within the FOV, thus noticeable geometric distortions are not to be expected. This means that no significant geometric corrections are expected within the joint reconstruction, which can be confirmed by the compared images in Figure 5.5. Only a slight compression in the vertical direction can be recognized. The estimated field maps for the two slices of T_2 -weighted in vivo data are consistent, displaying similar field distributions with local differences, particularly in the lower and upper right regions. Compared to the ACR phantom data, the shape of the B_0 inhomogeneity map differs considerably, as expected due to the use of different scanners. However, the observed inhomogeneity scale is similar to that of the PTB system and does not reflect the inhomogeneity differences between the two systems. Due to the lack of measured B_0 field maps, a qualified statement about the correctness cannot be made. Regarding denoising, the model effectively reduces noise compared to the B_0 -uninformed FFT reconstruction, but also introduces some blurring in the reconstructed images.

In addition to the T_2 -weighted data, Figure 5.6 compares the PD-weighted image magni-



(a) Comparison for slice number 8.



(b) Comparison for slice number 16.

Figure 5.5: Image magnitude and B₀ field map from the joint-model reconstruction with $T = 3$ compared to FFT of the unprocessed k -space data for two different image slices of the T₂-weighted in vivo brain data, acquired at LUMC. The cutout below each image sample shows a zoomed in section for a more detailed comparison.




Figure 5.6: Image magnitude and B_0 field map from the joint-model reconstruction with $T = 3$ compared to FFT of the unprocessed k -space data for image slice 16 of the PD-weighted in vivo brain data, acquired at LUMC. The cutout below each image sample shows a zoomed in section for a more detailed comparison.

tude and the corresponding B_0 field map reconstructed with the proposed model to the magnitude image obtained from B_0 -uninformed FFT reconstruction. Since the training data consisted exclusively of T_2 -weighted images, processing the PD-weighted data is more challenging, as this type of contrast is unknown for the model. While noise is effectively suppressed by the proposed reconstruction network, the resulting image appears blurred. Details which are clearly visible in the B_0 -uninformed reconstruction, are lost throughout the proposed reconstruction. The accuracy of the estimated field map remains uncertain. However, it is consistent with the B_0 field map obtained from the T_2 -weighted data, despite the different image contrast.

5.6 Discussion

To apply the proposed network for image and B_0 field reconstruction to measured data from the Nexus console, the model was adapted to enhance phase sensitivity. For supervised training, low-field MRI brain data with T_2 -weighted contrast was simulated to better match the measured in vivo data. Before applying the model in vivo, these adaptations were validated using a T_2 -weighted test dataset simulated under the same conditions as

the training and validation data. In this evaluation, the joint reconstruction approach improved NRMSE and PSNR by 2.1 % and 1.1 %, respectively. For qualitative assessment, the proposed method was tested on data acquired from two different low-field MRI systems. First, joint reconstruction was applied to PD-weighted data acquired from an ACR phantom, allowing evaluation of geometric distortion correction. Subsequently, the method was tested on PD- and T_2 -weighted in vivo brain data. It should be noted, that this evaluation does not replace a quantitative assessment, but rather demonstrates the general applicability of the method to measured low-field MRI data. A systematic investigation using B_0 field mapping sequences for instance is required to quantify the accuracy of the reconstructed B_0 field maps.

The geometric distortion correction was qualitatively assessed using the ring structure of the ACR phantom, which provides high image contrast in B_0 -uninformed FFT reconstruction. With the joint reconstruction network, both the inner ring structure and parts of the phantom boundary were successfully corrected, despite the presence of additional artifacts in the measured data. As concomitant gradient fields [55] or gradient non-linearities [116] may also contribute to geometric distortions in the reconstructed images, a measurement of the B_0 field map is essential to be able to conclusively evaluate the model performance with respect to the geometric distortion correction. When applied to in vivo brain data acquired on a system with a magnet approximately four times more homogeneous, less significant geometric corrections were observed as expected, and the estimated B_0 field profiles remained consistent across slices and image contrasts. This suggests that the model applies reasonable corrections to some extent. However, the fact that the B_0 field map is always scaled identically raises concerns about the method’s reliability. All training samples contain B_0 field maps with inhomogeneities similar to those observed in the PTB scanner, while none of the simulated data includes fully homogeneous B_0 field maps. As a result, the SH-Net was trained to estimate B_0 field maps within a specific inhomogeneity range, which may limit its ability to generalize to more homogeneous cases. This issue could be addressed by expanding the training dataset to include samples with more homogeneous B_0 fields. Alternatively, incorporating additional sensor inputs or integrating data from a second scan could improve the model’s ability to estimate B_0 field maps more reliably.

The proposed joint reconstruction network effectively reduces noise across different examples. However, for data with PD-weighted image contrast, the output appears over-smoothed, reducing the contrast compared to B_0 -uninformed FFT reconstruction. This effect can be attributed to the higher signal intensity in PD-weighted data, which the model was not trained on, as only T_2 -weighted samples were used during training. In contrast, the results for the measured T_2 -weighted data exhibit sharper features, underscoring the importance of data similarity between training and model inference.

The proposed method employs NNs to regularize a joint reconstruction problem. For both the image and field map sub-problems, the alternating optimization introduces only minor improvements, with the output largely driven by the NNs estimates, as could be observed by the evaluation on simulated test data in Figure 5.2. This suggests that the balance between model-based and data-driven contributions may require further refinement. Additionally, the data acquired with the Nexus console in Chapter 3 consists of three-dimensional volumes, whereas the evaluation was performed on individual two-dimensional slices due to GPU memory limitations. By manually implementing the backward pass of an end-to-end network for training, as proposed in [129], the method could be optimized for greater efficiency, enabling application to full three-dimensional data. This could, in turn, enhance the B_0 field estimation.

5.7 Summary

Building on the evaluation of simulated low-field MRI knee images, the joint method for image and B_0 field map reconstruction, which combines model-based and data-driven approaches, was further extended for measured low-field MRI data. Supervised training was conducted using simulated low-field MRI brain images with T_2 -weighted contrast, and performance was validated on a simulated test dataset first. To assess the applicability of the proposed method, a qualitative evaluation was performed using measured low-field images with PD- and T_2 -weighted contrast. Partial correction of reasonable geometric distortions caused by B_0 inhomogeneities was demonstrated, and effective denoising was achieved when the test sample closely matched the training data. However, the NNs used to regularize the joint reconstruction network still have a significant impact on image quality. Therefore, the following chapter explores an alternative regularization approach and further investigates the alignment between simulations and measurements.

6 Learning regularization parameter maps

As elaborated in the previous chapters, model-based reconstruction with regularization by deep NNs enables joint image and B_0 field reconstruction, outperforming conventional model-based approaches. However, after training with simulated data, performance losses can be observed when the model is applied directly to measured low-field MRI data. Compared to the T_2 -weighted images, the results for the PD-weighted data were over-smoothed and blurry. If the distribution of data used for training significantly differs the distribution used at model inference, the model fails to maintain its performance. This could be observed in the previous chapter, but has also been reported for other problems in the literature [42].

The proposed method which was investigated throughout the previous chapters involved a physical model for the reconstruction of images and field maps and two NNs for the regularization. Instead of entirely leaving the regularization to the network, what still requires very accurate model results, another viable approach is to rely on conventional non-data-dependent regularization methods. In [134], an approach that learns a regularization parameter map for image denoising was investigated, whose application to low-field MRI data will be explored in the following. The approach relies on conventional regularization and addresses the shortcoming of selecting an optimal regularization parameter. Thereby an improved transferability of the method is expected. The reconstruction problem previously considered was complex and required a time-segmented approximation of the reconstruction operator in combination with first-order Taylor approximation. These approximations were mainly introduced due to the consideration of the B_0 field. For well-shimmed systems, as in the case of the LUMC system which was used for in vivo measurements, the impact of the B_0 field inhomogeneities on image quality is negligible. Instead, SNR is the dominant component which degrades the quality of the image. To reduce complexity and focus on SNR, the following approach only considers the image-related denoising sub-problem and the possibility of employing image denoising as a post-processing step. In turn, the reduced complexity of the problem enables the application to three-dimensional image data.

6.1 Denoising network

This section presents the denoising problem, which is addressed by an iterative algorithm using conventional regularization, weighted by a regularization parameter map estimated through a NN. In principle, this step can be applied to the entire image reconstruction problem addressed in Chapter 4. However, here the problem is simplified and restricted to denoising as a post-processing step which does not involve image reconstruction. A more general description of the method has been published in [134].

6.1.1 Denoising problem

To reduce complexity, only the denoising problem is considered for learning a regularization parameter map. Considering a fully sampled k -space \mathbf{y} , this simplifies the forward problem 2.7 to

$$\mathbf{y} = \mathbf{F}\mathbf{x}_{\text{True}} + \boldsymbol{\varepsilon} \quad (6.1)$$

where \mathbf{F} is the Fourier operator and $\boldsymbol{\varepsilon}$ the noise, which corrupts the measured data. By applying the inverse Fourier operator on both sides, the problem can be written as

$$\mathbf{x}_{\text{Noisy}} = \mathbf{x}_{\text{True}} + \mathbf{F}^H \boldsymbol{\varepsilon}, \quad \text{with} \quad \mathbf{x}_{\text{Noisy}} = \mathbf{F}^H \mathbf{y}, \quad (6.2)$$

where \mathbf{x} is the desired clean image and $\mathbf{x}_{\text{Noisy}}$ is the image corrupted by noise. Thus, the simplification allows solving the problem solely in the image space, which is why it can be considered as a post-processing step. To obtain a denoised version of the image, a possible approach for regularization is TV denoising [64], [65], [135]. Using the ℓ_1 norm, the minimization problem can be defined by

$$\arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{x}_{\text{Noisy}}\|_2^2 + \lambda \|\nabla \mathbf{x}\|_1. \quad (6.3)$$

Regularization is achieved by promoting sparsity in the image's gradient domain, as ∇ denotes the finite difference operator for three spatial dimensions. The sparsity is quantified using the ℓ_1 norm, yielding the TV of \mathbf{x} :

$$\text{TV}(\mathbf{x}) = \|\nabla \mathbf{x}\|_1. \quad (6.4)$$

To obtain a good result for the denoised image, the choice of the global scalar regularization parameter λ is crucial. If λ is unfavorable, the regularization might be too poor or the result is too smooth. The idea is to replace the scalar value with a parameter map $\boldsymbol{\Lambda} \in \mathbb{R}_+^{qn}$, where q represents the directions of the partial derivatives and n the number of pixels or voxels. This enables pixel- or voxel-dependent regularization and allows adaption of the

regularization to individual samples. Thus, using a parameter map for the regularization yields the following.

$$\arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{x}_{\text{Noisy}}\|_2^2 + \|\mathbf{\Lambda} \nabla \mathbf{x}\|_1. \quad (6.5)$$

This flexibility ensures that areas requiring more smoothing receive it, while areas with important structural details are preserved. The end-to-end network is denoted by $f_{\Theta}^T(\mathbf{x}_{\text{Noisy}})$. It estimates the regularization parameter map $\mathbf{\Lambda}$ from the noisy image $\mathbf{x}_{\text{Noisy}}$ once in the beginning, using some CNN with trainable parameters Θ . The regularization parameter map is applied within the regularization of Problem (6.5), which is solved by the PDHG algorithm [67]. The number of iterations is defined by T .

6.1.2 Learning the regularization parameter map

The overall denoising network consists of two parts, the CNN to estimate the regularization parameter map and the iterative PDHG algorithm. In this context, the network u_{Θ} predicts the parameter map $\mathbf{\Lambda}$ from the input $\mathbf{x}_{\text{Noisy}}$, that is, $\mathbf{\Lambda}_{\Theta} = u_{\Theta}(\mathbf{x}_{\text{Noisy}})$, with Θ denoting the trainable parameter of the network. In this setting, $\mathbf{x}_{\text{Noisy}}$ is a real-valued image with three spatial dimensions, contaminated by noise. A U-Net architecture is used to estimate the regularization parameter map from the input image. To constrain the regularization to a positive value, a soft-plus activation function is applied to the network output. In addition, the output, i.e. the regularization parameter map, is scaled by a multiplicative factor, given that the arbitrary initialization of the network parameters Θ is deemed suboptimal. All input images have an isotropic in-plane resolution of 2×2 mm and a resolution of 5 mm in the third dimension. This motivates the choice of two different regularization parameter maps. Thus, the network is constructed to provide a two-channel output: The first channel is replicated to serve as $\mathbf{\Lambda}_{\Theta}^{xy}$, while the second channel is utilized for $\mathbf{\Lambda}_{\Theta}^z$, resulting in

$$u_{\Theta}(\mathbf{x}_{\text{Noisy}}) = \mathbf{\Lambda}_{\Theta} = (\mathbf{\Lambda}_{\Theta}^{xy}, \mathbf{\Lambda}_{\Theta}^{xy}, \mathbf{\Lambda}_{\Theta}^z). \quad (6.6)$$

6.1.3 Primal dual hybrid gradient with total variation regularization

To solve the denoising problem defined in Equation (6.5), the PDHG algorithm [67] is applied. The primal-dual construct involves two related optimization problems: the primal problem, which directly solves for the unknowns in an optimization task, and the dual problem, which focuses on optimizing the trade-offs between the objective and constraints of the primal. Throughout the iterations, the algorithm alternates between updating the primal and dual components, with each iteration refining the balance between noise reduction and detail preservation. The variable T represents for the total number of

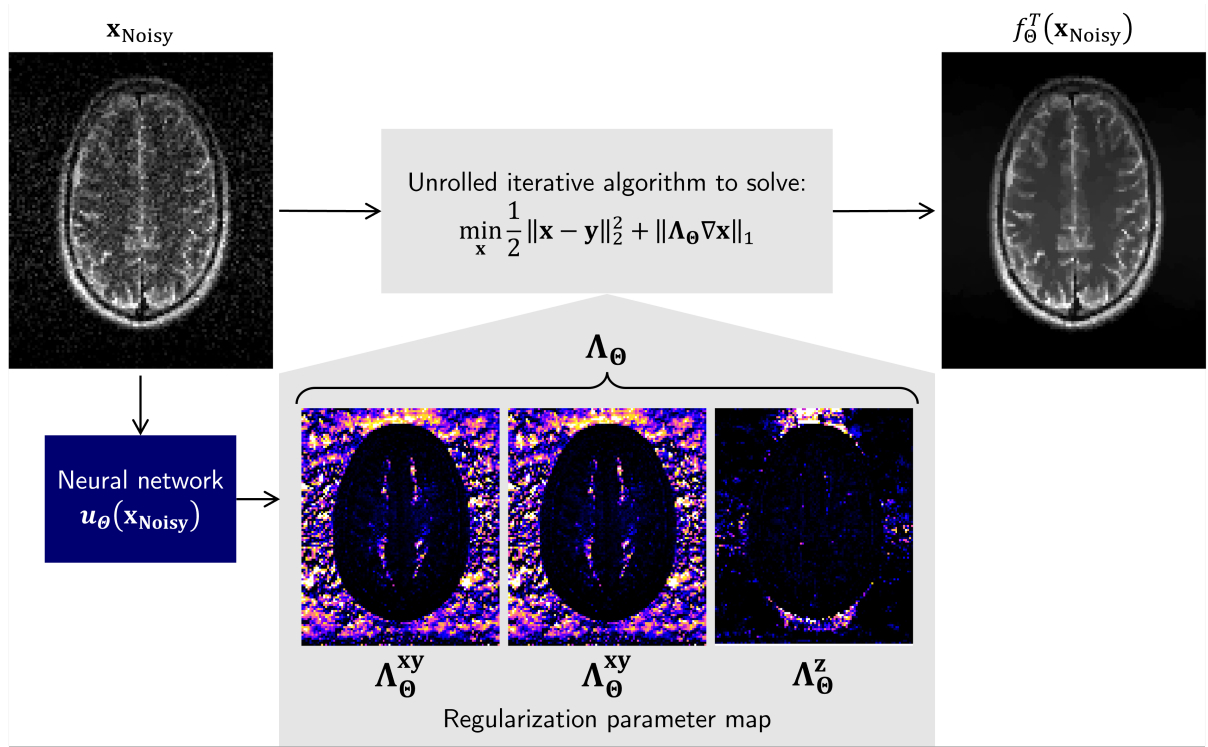


Figure 6.1: Overview of the network architecture facilitating regularization parameter map estimation by CNN and unrolling of iterative denoising scheme with TV regularization.

iterations or alternations between primal and dual update. Figure 6.1 illustrates the role of the NN in solving the denoising problem with the PDHG algorithm. It shows the noisy input image $\mathbf{x}_{\text{Noisy}}$ and the regularization parameter map Λ_{Θ} estimated by $\text{NN}\Theta$. During the training of $\text{NN}\Theta$, the PDHG algorithm is unrolled for T iterations, leading to the final denoised output $f_{\Theta}^T(\mathbf{x}_{\text{Noisy}})$.

6.2 Computational experiments

To investigate how image quality can be enhanced by estimated regularization parameter maps, supervised learning with simulated low-field MRI data is used to train a NN which estimates the parameter maps. The following section first addresses the data simulation and supervised training, before introducing the different comparisons.

6.2.1 Simulation of low-Field MRI samples

For the supervised training, low-field MRI data is simulated based on high-field MRI samples from the *FastMRI* dataset [34] which are adapted to match the image quality that is typically observed on low-field MRI scanners. The simulation process is very similar to the one described in Section 5.1. However, here B_0 field inhomogeneities are neglected as

the method targets the denoising of the in vivo images which do not suffer from severe geometric distortions. Throughout this Chapter, real-valued three-dimensional data is considered, as the method is restricted to denoising. Starting from the real-valued ground truth, the 3 T images are down-sampled to dimensions $16 \times 100 \times 120$ for (PE $2 \times$ PE $1 \times$ RO) which corresponds to the dimension of the measured in vivo data as shown in Chapter 3. Only samples with 16 or more slices are considered. Any sample exceeding this limit is reduced to 16 slices located in the center of the slice encoding dimension. To simulate the low SNR observed at low-field systems, the real-valued samples are augmented by Rician noise. Using random numbers from a two-channel normal distribution, Rician noise is defined by the root of the squared sum of both channels. The noise is scaled by a factor that is manually tuned to match the SNR observed in vivo. During training, the scaling factor is randomly chosen from a predefined interval to simulate varying SNR conditions and to achieve better generalizability of the model at inference. In the in vivo images, a characteristic receive profile is observable which can be explained by the narrow bandwidth of the receive coil. Because the methodology focuses on denoising, this specific profile is considered for the data simulation here. Using pure noise measurement data, the readout profiles are obtained by calculating the standard deviation along the PE dimension of multiple noise scans. A polynomial of degree 4 is fitted to the averaged readout profile and normalized to weight the readout dimension of the simulated low-field MRI image. The overall simulation process is illustrated in Figure 6.2. As most of the samples are T_2 -weighted and because T_2 is less dependent on the field strength which provides a more realistic contrast in the simulated low-field regime, the supervised training is restricted to the T_2 -weighted images. The total number of 2678 is divided into fractions of 70 %, 20 %, and 10 % for training, validation, and testing.

6.2.2 Supervised training

The regularization parameter maps are estimated from single channel real-valued three-dimensional low-field MRI images using a U-Net architecture [71]. It incorporates residual connections between the encoder and decoder and consists of three encoding stages with two convolutional layers per stage and the number of initial filters set to 32. The output of the network architecture comprises two channels for the in-plane and through-plane regularization parameter maps, i.e. Λ_{Θ}^{xy} and Λ_{Θ}^z , which are motivated by the anisotropic voxel resolution. The NN is applied once in the beginning to obtain the regularization parameter maps. Within the unrolled network $f_{\Theta}^T(\mathbf{x}_{\text{Noisy}})$, the NN estimates the regularization parameter map once at the beginning. The PDHG algorithm then iteratively solves the denoising problem for T iterations. During training, a moderate $T = 100$ iterations are used, while testing employs a maximum of $T = 1000$ iterations.

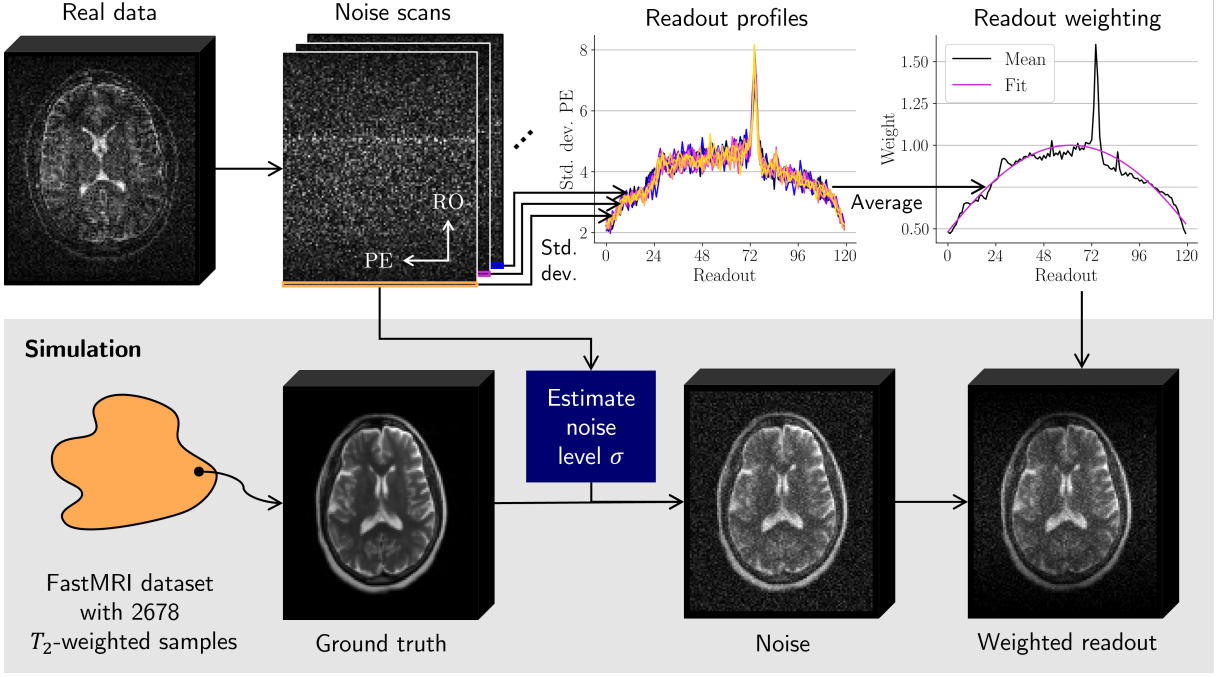


Figure 6.2: Simulation process to obtain pairs of noise-free ground truth and noisy low-field MRI samples: Noise standard deviation and coil profile are estimated from real data. A smooth coil profile is fitted to the readout profiles derived from pure noise measurements along the PE dimension, using a 4th-degree polynomial before applying it to the simulated low-field MRI image.

The unrolled network f_{Θ}^T is trained in a supervised manner by optimizing its weights Θ , using the objective function

$$\mathcal{L}(\Theta) = \frac{1}{N_{\text{Sample}}} \sum_{i=1}^{N_{\text{Sample}}} \ell(f_{\Theta}^T(\mathbf{x}_{\text{Noisy}}^i), \mathbf{x}_{\text{True}}^i), \quad (6.7)$$

where ℓ denotes some loss-function. N_{Sample} accounts for the total number of training or validation samples, respectively, depending on the current stage in training. With the data introduced in Section 6.2.1, $N_{\text{Sample}} = 1874$ is used for the training phase and $N_{\text{Sample}} = 535$ for the validation phase. For any of the experiments the \mathcal{L}_1 -norm is used as the loss function ℓ . To find the best Θ which minimizes (6.7), the AdamW optimizer [136] is used. It is configured with an initial learning rate of 10^{-4} and a weight decay of 10^{-2} . The learning rate determines the magnitude of the updates of the model parameters during optimization, with a value of 10^{-4} ensuring stable convergence while avoiding abrupt changes that could hinder the training process. The weight decay parameter serves as a regularization term, preventing overfitting by penalizing large parameter values. In particular, AdamW decouples weight decay from gradient updates, enabling more effective control over model regularization without interfering with the adaptive learning rate mechanism. This configuration facilitates both effective learning and improved generalization.

6.2.3 Comparisons

With the network trained on simulated low-field MRI data, different experiments are conducted on the simulated test data and the actual low-field MRI data. Initially, the proposed regularization methodology is assessed using the simulated test dataset. The utilization of simulated data facilitates an objective quantification of errors relative to the ground truth images, enabling a comparison with conventional approaches. Conventionally, a scalar value is chosen at the global level to weight the regularization term. With a scalar value, the problem to be solved corresponds to (6.3). To provide a fair comparison, the best scalar is identified by setting λ^{xyz} a learnable parameter, which is optimized by supervised learning as described in the previous section. Furthermore, the comparison is supplemented by the case of two different scalars for each of the different spatial resolutions. This results in the scalar λ^{xy} designated for the 2 mm in-plane resolution in RO and PE1 dimension, and the scalar λ^z designated for the 5 mm through-plane resolution in PE2 dimension. Again, supervised learning is used to find the best combination of scalars. The first experiment observes the training loss and convergence of the different scalar values. By evaluating the NRMSE and PSNR distributions over the test dataset, the scalar-based weighting of the regularization term is contrasted with the proposed regularization parameter map approach.

Subsequently, the different regularization approaches are then applied to the measured in vivo low-field MRI data, which were acquired using the Nexus console introduced in Chapter 3. The different regularization approaches are compared across both PD- and T₂-weighted three-dimensional data. For the network that estimates the regularization parameter maps, the model weights trained on simulated T₂-weighted low-field MRI data are used. Since no ground truth is available for the in vivo data, reference-free metrics are employed for quantitative comparison. Specifically, a no-reference blur metric [137] is calculated, along with the SNR introduced in Section 3.3.1, which uses signal and noise regions as depicted in Figure 3.6 to calculate the SNR.

6.3 Results

The evaluation of the learned regularization parameter maps is carried out in two distinct stages. First, the model is tested on simulated low-field MRI data, allowing a direct comparison with ground truth data for quantitative evaluation. Second, the trained model is transferred and applied to in vivo data, demonstrating its generalizability and practical utility. For both cases, identical models are considered, trained on simulated low-field MRI data.




Figure 6.3: Convergence of training and validation ℓ_1 -loss over epochs (left) and convergence of λ^{xyz} , λ^{xy} and λ^z over epochs during training (right).

6.3.1 Simulated low-field data

Figure 6.3 shows the training and validation losses over 300 epochs on the left, evaluated every second epoch. The solid lines represent the training loss, while the dashed lines correspond to the validation loss. The learned regularization parameter map $\mathbf{\Lambda}_\Theta$ is compared against a global scalar λ , with distinctions made between λ^{xyz} and the combination of λ^{xy} and λ^z . Due to the similarity between the training and validation losses, the curves are nearly indistinguishable. All training processes converge rapidly, reaching a minimum around epoch 30. The training was continued until epoch 300 to exclude any overfitting, which would be noticeable by diverging training and validation loss. The course of the scalar lambda values λ^{xyz} , λ^{xy} , and λ^z is plotted on the right side of Figure 6.3. It can be noticed that λ^{xy} and λ^z clearly converge to different values, which emphasizes the spatial dependence of the regularization parameter. As expected, the joint regularization scalar λ^{xyz} is in between λ^{xy} and λ^z , converging toward a compromise.

10 % of the simulated low-field MRI data was reserved as an independent test dataset, unseen by the model during training. After training, the model parameters were fixed, and performance was evaluated on the test set using the NRMSE and PSNR metrics. The trained model was tested on noisy samples from the test dataset, with performance measured against their corresponding ground truth. Figure 6.4 compares the performance for learned global scalar values and learned regularization parameter maps.

The performance, obtainable with one or two global scalars, i.e. λ^{xyz} or λ^z and λ^{xy} , for regularization yields similar performance. However, splitting the scalar into two distinct values, representing through-plane and in-plane resolutions, respectively, yielded a minimal improved NRMSE and PSNR values, over the test dataset. The mean NRMSE decreases by 0.6 % and the mean PSNR increases by 0.2 %. The varying performance with two




Figure 6.4: Evaluation of NRMSE and PSNR for the denoising performance on the test dataset comparing λ^{xyz} , λ^{xy} and λ^z , and the proposed Λ_{Θ} regularization.

different scalar values aligns with the findings from Figure 6.3, suggesting that lambdas with different spatial dependence converge to different values. Introducing learned spatially varying regularization parameter maps, which further employ this dependency, lead to an enhancement of the results obtainable by the two different scalar values. Here, the mean NRMSE noticeably decrease by additional 4.7% and the mean PSNR increases by 1.6%. This underscores the advantage of accommodating spatial dependencies in the regularization process.

A comparison by means of a single two-dimensional slice of a three-dimensional test samples is shown in Figure 6.5. For both samples, the input is compared with the different regularization scenarios and the ground truth. The in-plane and through-plane regularization parameter maps are shown next to the image results obtainable with the maps. Due to the use of TV regularization, all the results exhibit a more subdued grayscale appearance compared to the ground truth. The example clearly emphasizes the importance of the spatial dependency of the regularization parameter map. Not only do the in-plane and through-plane regularization parameter maps differ significantly, but there is also a clear distinction between the imaged object and the background in the Λ_{cnn}^{xy} maps of Figure 6.5. Compared to the brain structures, the regularization of the background is much stronger and more coarse. Finer structures inside the brain are weighted more individually, resulting in a much sharper result. In comparison to regularization with global scalars, the background noise can be removed much better with the estimated regularization parameter map. The image slice shown in Figure 6.5 clearly shows a pathological bright spot in the upper right part of the cortex, which is highlighted in the zoomed in section. None of the compared methods fails to depict the structure correctly. All methods manage to delineate the structure more clearly than in the original image.




Figure 6.5: Comparison of reconstruction results for a representative sample from the test data set of simulated low-field images with T_2 -weighted contrast. The first row compares the magnitude images of the input, i.e. $\mathbf{x}_{\text{Noisy}}$, regularization with λ^{xyz} , the combination of λ^z and λ^{xy} , and the map of regularization parameters against the ground-truth. The second row shows the regularization parameter maps Λ_{Θ}^{xy} and Λ_{Θ}^z , which allow interpreting the impact of the model.

6.3.2 In vivo data

Figures 6.6 and 6.7 each show a slice from the center of the T_2 - and the PD-weighted images, comparing the image quality obtained by the different regularization methods. A qualitative comparison across the T_2 -weighted results in Figures 6.6 shows that the model can be successfully transferred to actual in vivo images and maintains its denoising performance. The T_2 -weighted data closely matches the training data, yielding an overall decent performance. Noise is effectively removed in all the compared cases of the T_2 -weighted images, while important information in the denoised output is retained. The

reference-free quantitative metrics, i.e. blur and SNR, were added to the denoising results in Figures 6.6. TV regularization weighted by global scalar values, increases the SNR compared to the input, which is a simple FFT reconstruction of the measured data. By using two separate scalars, SNR increases by 6.7 %. The local regularization obtained by estimating a regularization parameter map further increases the SNR by 3.8 %. The regularization by TV introduces blurring to the images what increases the no-reference blur metric calculated for the scalar-based approaches by 42 % with respect to the input. By applying the estimated regularization parameter map the blurring only increases by 20.7 %, yielding sharper results compared to the global regularization parameter, confirming the previous observations. The spatially-dependent regularization which is achieved by the estimated regularization parameter maps can be observed by the shown parameter maps in Figures 6.6. In the center of the in-plane regularization map $\Lambda_{\text{cnn}}^{xy}$, the regularization is low and a structure similar to the ventricles is visible. Here, the input image already facilitates a relatively high contrast and less noise, thus less regularization is required.

In addition to the T_2 -weighted images, the results for the PD-weighted measurements are shown in Figure 6.7. It should be noted that this example exhibits a different contrast, which is unknown to the model, since it was trained exclusively on simulated T_2 -weighted data. Given that the network only estimates the strength and location of the regularization by the parameter map, the methods are still transferable and capable of effectively eliminating the noise from the images. Compared to the T_2 -weighted results, the PD-weighted reconstructions tend to be a bit over-smoothed, as can be seen by a qualitative comparison of the results in Figure 6.7. Some finer details in the denoised image appear less contrastive or even vanish, as can be seen in the zoomed-in section of Figure 6.7. This effect is observed in both scalar-based regularization and the proposed method using a regularization parameter map. Nevertheless, all methods demonstrate a clear improvement in contrast compared to the background. According to the no-reference blur metric, the regularization with λ^{xyz} and the combination of λ^z and λ^{xy} equally increase the blurring by 30.5 % compared to the input. The image denoised using the estimated regularization parameter map experiences an 11.1 % increase in blurring with respect to the input. With the scalar-based regularization approaches, SNR of the noisy input can be increased by 61.2 % with λ^{xyz} and by 71.8 % with the combination of λ^z and λ^{xy} . Employing the regularization parameter map, SNR increases by 82.0 %. The parameter maps shown in Figure 6.7 show a strong local change in regularization in the transition between head and background, emphasizing the high contrast to the background in the denoised image. The differences in regularization between the left and right hemispheres can not be directly related to the shown image slice. However, it should be noted that the approach estimates only one regularization parameter map Λ_{Θ}^{xy} for in-plane regularization, which is applied across all the different slices.




Figure 6.6: Reconstruction of the T₂ weighted in vivo results obtained in Chapter 3 using the unrolled network $f_{\Theta}^T(\mathbf{x}_{\text{Noisy}})$. For each slice, the original noisy input image, and the denoised model output is shown along with the regularization parameter maps Λ_{Θ}^{xy} and Λ_{Θ}^z .




Figure 6.7: Reconstruction of the PD-weighted in vivo results obtained in Chapter 3 using the unrolled network $f_{\Theta}^T(\mathbf{x}_{\text{Noisy}})$. For each slice, the original noisy input image, and the denoised model output is shown along with the regularization parameter maps Λ_{Θ}^{xy} and Λ_{Θ}^z . The training was performed on simulated low-field data with T_2 weighted contrast.

6.4 Discussion

This chapter introduced an approach to low-field MRI denoising that combines conventional regularization by TV with a spatially dependent regularization parameter map estimated by a NN. The algorithm was trained on simulated three-dimensional data and evaluated on a simulated test dataset, as well as on measured low-field MRI data with PD- and T_2 -weighted contrast. Quantitative evaluations on the simulated test data, using PSNR and NRMSE, demonstrated that the regularization parameter map outperforms conventional scalar-based approaches for weighting the regularization term. This emphasizes that a spatially dependent regularization yields much better denoising performance in comparison to global scalar values. In addition, the estimated regularization parameter map allows for an interpretation of the network output. The application to the actual low-field MRI data yields a significant improvement of SNR what could be shown by a reference-independent measure. Employing TV regularization requires a trade-off between blurring and denoising performance, which causes the reference-free blur metric to increase. Using the regularization parameter map, the blurring introduced is only half as strong compared to the global scalar values. However, since the model weights are transferred from simulated data to actual low-field MRI data, the estimation of regularization is not optimal. This applies for both the regularization parameter map, and for the scalar-based weighting. By manually adjusting the strength of regularization, the regularization could be further fine-tuned to the different data. Alternatively, the training data could be extended by more diverse data, which also cover different image contrasts. This could in particular benefit the results obtained for the PD-weighted images. For the PD-weighted data, the lower increase in blurring compared to the T_2 -weighted data is a bit counterintuitive to the qualitative assessment of the image quality after denoising. The improvement in brain-to-background contrast appears to have a greater impact on the blur metric than the enhancement of finer details, which inherently exhibit lower contrast. Because the network only estimates two two-dimensional maps for in-plane and through-plane regularization, spatial adaptability is limited. By estimating three-dimensional regularization parameter maps, this problem could be addressed. Alternatively, the strength of the estimated regularization parameter map could be fine-tuned to yield improved denoising results, but would require a manual adjustment of the method which is usually unfavorable.

6.5 Summary

The joint image and B_0 field reconstruction approach which was utilized in the previous chapter combined model-based reconstruction and deep learning-based regularization. It showed promising results but also revealed challenges when applying the method to

measured data. Consequently, this chapter investigated an alternative approach which employs conventional TV regularization in combination with spatially varying regularization parameter maps estimated by a NN. These constraints limit the NN to estimating the strength and location of the regularization, making its output more interpretable and improving the denoising results for the acquired in vivo data.

By using spatially adaptive in-plane and through-plane maps to weight the regularization instead of conventional global scalar values, improved denoising performance is obtained while less blurring is introduced. The method could be successfully applied to measured low-field MRI in vivo data. However, an optimal choice for the values in the regularization parameter maps could not be obtained, as an overestimating of the regularization strength introduced significant blurring, which was not observed in the test data. Simulations always contain approximations which limit their ability to replicate actual low-field MRI data for supervised training. Thus, the next chapter investigates the optimization of models trained with supervision on individual samples to overcome the need for highly realistic simulations, which are challenging to obtain given the diversity of low-field systems available and varying operational conditions.

7 Model adaptation by test time training

In the previous chapter, the combination of CNN and conventional regularization demonstrated promising performance for low-field MRI denoising. The key strategy involved estimating regularization parameter maps from noisy inputs using a NN, which are then combined with conventional regularization for iterative denoising. Training was conducted on simulated low-field data, and although the model performed well on the measured MRI data, its effectiveness depended on the similarity between the inference data and the original training data. A significant shift between the training and inference data distributions may lead to a performance gap at model inference [42], [43], [48]. Previous methods, which rely entirely on simulated data during supervised training, introduced blurring in the denoised images when applied to measured low-field MRI data due to their deviation from the training data.

Alternative self-supervised approaches have shown promise in denoising low-field MRI data [44], [138]–[140]. However, their performance often falls short of supervised techniques, which typically employ retrospective simulation of low-field MRI data from higher-field scans [36], [37], [40]. This approach is common due to the wide availability of high-field datasets [34], [82], [83] and the unfeasibility of acquiring noise-free reference data which is essential for supervised training. Due to the dependency of relaxation coefficients on the magnetic field strength, the simulated data often differs inherently from real low-field acquisitions. The differences between training and test data complicate the direct translation of the trained models from simulations to real data. Using more sophisticated simulation techniques, the training data could be further improved [38]. But this requires a large set of tissue models to generate the training data. Rather than bridging the gap between simulated and measured low-field MRI data, this chapter investigates a novel approach to adapt pretrained models through a two-step training process. First, supervised learning is employed to pretrain a model on a large set of retrospectively simulated low-field MRI images. Then, the pretrained model is adapted to individual in vivo images using self-supervised learning, which does not require any reference data. The goal is to relax the demands on the simulation, while still profit from the outstanding performance of supervised learning.




Figure 7.1: The two-stage training approach integrates an initial phase of supervised learning (top) with a subsequent phase of unsupervised refinement (bottom). In this framework, the pretrained model weights undergo optimization for individual samples utilizing the unsupervised *Noise2Self* training strategy. The objective is to exploit supervised learning during the testing phase through adaptations of the denoising model tailored to individual samples.

7.1 Two-stage training approach

The objective of the proposed method is to divide the overall training process into two stages. In the initial stage, supervised learning is utilized to train a denoising model using simulated pairs of noisy and ground-truth images. The methodologies presented in the preceding chapters have demonstrated that, while supervised learning is to some extent transferable to real in vivo data, the models exhibit optimal performance only when the training data closely resembles the data encountered during inference. Consequently, the aim of this approach is to extend training into a second stage taking place after the supervised training. Upon convergence of supervised training, the model parameters undergo further refinement through unsupervised learning. The underlying presumption of this approach is that both training procedures optimize towards similar minima, and these minima are sufficiently proximate to each other, thereby enabling the model to leverage the experience acquired during supervised training in the course of unsupervised refinement. Section 2.4.4 introduced various methodologies for unsupervised learning. In this context, the objective is to optimize the model specifically for individual samples. Hence, the *Noise2Self* [87] training strategy is adopted for unsupervised refinement. Unlike the *Noise2Noise* and *Noise2Noise* strategies, *Noise2Self* operates on a single image and

does not rely on assumptions about the noise distribution, making it a suitable choice for this approach. Instead, it assumes that the noise in an image is statistically independent of the underlying signal. This independence means that the noise is uncorrelated and random, whereas the true image content exhibits structured spatial relationships. The strategy employs a masking process during training, in which a subset of pixels in the image is deliberately masked. Their true values are hidden from the network, which is trained to predict the masked values based on the structural information of the surrounding unmasked pixels. The masking process enforces that the model learns to infer the image values by focussing on the spatial context provided by neighboring pixels what effectively filters out noise.

Figure 7.1 illustrates the two-stage methodology. The gray box represents the pretraining phase, where a denoising model is optimized through supervised learning using ground truth data. Subsequently, the model weights obtained during supervised training undergo a second training step, which processes individual images as input, specifically measured images from the low-field MRI scanner. While this step may be repeated for each individual sample, it always starts at the pretrained weights. The second training stage aims to further refine the model optimizing its performance for the measured in vivo data. If a substantial amount of measured low-field data is available, this step can potentially be replaced by methods such as *Noise2Noise*. This two-stage training approach enables pretraining on an extensive simulated dataset while allowing model optimization that is tailored to individual samples at inference.

7.1.1 Network architecture

For image denoising, a three-dimensional U-Net architecture [71] was chosen due to its proven effectiveness in image denoising tasks, which could be observed in the previous chapters. In this context, NNs were used to regularize image reconstruction and optimization problems. It was observed that the denoised image tends to become blurry when the NN overestimates the regularization term. The underlying issue arises when the model struggles to make an accurate estimate from the input, particularly when the data exhibit contrasts not encountered during training. This motivates to investigate the proposed two-stage training process, which is considered here for a simple denoising task to isolate the problem. However, the procedure can also be adapted for the method presented in Chapter 6. For this study, a denoising U-Net is considered which is trained to denoise a noisy low-field MRI input image. The model, denoted by u_{Θ} , is given as a weighted combination of the U-Net output and the noisy input image. This corresponds to the minimizer of a denoising problem in which the output of the U-Net is employed in the form of Tikhonov regularization [141]. Similarly to the adapted model used for the image

related sub-problem in Chapter 5, an architecture with multi-head self-attention in the bottleneck is employed. It allows the network to capture long-range dependencies by simultaneously attending to different parts of the input feature map, enhancing its ability to integrate global contextual information [75]. This improves the network’s ability to handle complex spatial dependencies and enhances its robustness. The U-Net architecture consists of three encoding stages, each comprising two convolutional layers, with an initial filter count of 32 and a kernel size of three.

7.2 Computational experiments

To examine the transferability of a denoising operator after a two-stage training process, first the training data and the supervised learning are described. With the subsequently introduced self-supervised refinement step on individual samples, various comparisons are defined to evaluate its impact on denoising performance.

7.2.1 Simulated training data

For the supervised pretraining, three-dimensional low-field MRI data is simulated from T_2 -weighted 3 T brain images [34]. To mimic low-field characteristics, the data are downsampled to a resolution of $16 \times 100 \times 120$ ($PE_2 \times PE_1 \times RO$) and augmented by Rician noise, with additional weighting of the readout dimension based on the receiver coil profile of the actual low-field MRI system. Section 6.2.1 describes the simulation process in more detail. The resulting data set is split into fractions of 70 %, 20 %, and 10 % for training, validation, and testing. This yields 1874 samples for training, 535 for validation, and 269 for testing. The test data of the simulated T_2 -weighted brain samples is used to review the performance of the model for samples originating from the same data distribution as training and test data. Using the identical simulation pipeline, T_1 -weighted brain samples are generated, which exhibit a data distribution shift due to the different image contrast. Because, the dataset contains much more T_2 -weighted samples, supervised training is performed with the simulated T_2 -weighted data only. The T_1 data is used to investigate the consequences of a distribution shift by quantitative means with respect to the ground truth. Last but not least, the two-stage training process is tested on real, three-dimensional T_2 - and PD-weighted in vivo images of the brain acquired in Chapter 3.

7.2.2 Supervised training

The training process is subdivided into two stages, where the first stage of the training process involves supervised training. Using Θ^{pre} to denote the optimized model parameters after supervised pretraining, the objective function for the supervised training is defined by

$$\mathcal{L}(\Theta^{\text{pre}}) = \frac{1}{N_{\text{Sample}}} \sum_{i=1}^{N_{\text{Sample}}} \ell(u_{\Theta^{\text{pre}}}(\mathbf{x}_{\text{Noisy}}^i), \mathbf{x}_{\text{True}}^i). \quad (7.1)$$

The number of training samples which are pairs of simulated noisy data $\mathbf{x}_{\text{Noisy}}$ and noise-free ground truth data \mathbf{x}_{True} , is denoted by N_{Sample} . Supervised training was carried out over 300 epochs using an \mathcal{L}_1 -norm as loss function ℓ in combination with the AdamW optimizer [81]. The learning rate is set to 10^{-4} with a weight decay of 10^{-2} and a batch size of 8 is employed.

7.2.3 Unsupervised test time training

As described by the illustration in Figure 7.1, the model weights from the supervised pretraining Θ^{pre} are further refined in a second training stage, in which self-supervised training at test-time is employed. Test-time training [42] refers to an adaptation of the model, i.e. the model parameters Θ^{pre} , at inference. That is, to employ the *Noise2Self* [87] training strategy, which facilitates single-sample training, to further adapt the model $u_{\Theta^{\text{pre}}}$. Self-supervised training stage at test-time is achieved by minimizing the objective function

$$\mathcal{L}(\Theta) = \left\| \mathbf{M}_i^c u_{\Theta}(\mathbf{M}_i \mathbf{x}_{\text{Noisy}}) - \mathbf{M}_i^c \mathbf{x}_{\text{Noisy}} \right\|_2^2. \quad (7.2)$$

Here, \mathbf{M} denotes the masking operate and \mathbf{M}^c its complement. During training, \mathbf{M} is used to mask the input at voxel locations on a regular three-dimensional grid using local mean values, such that continuity of the network input is guaranteed. The true input value at the masked locations is hidden for the network. Using the complement operator \mathbf{M}^c , the loss is calculated only for the unseen voxels. A mask width of four voxels results in 64 different mask configurations indexed by i . The self-supervised refinement is carried out over 1000 iterations, which repeatedly iterate the mask configurations. To refine the model parameters which are initialized by the parameters Θ^{pre} from the supervised pretraining, the Adam optimizer with an initial learning rate of 10^{-5} is employed. Because the model already experienced supervised pretraining, the learning rate was set lower to prevent the network from forgetting what was learned by supervised training, a phenomenon known as *Catastrophic Forgetting* [142]. The self-supervised refinement step is done individually for each sample.

7.2.4 Comparisons and evaluation metrics

To investigate the impact of test-time training as a second self-supervised training stage, supplementing supervised learning, Across the evaluation, different scenarios of mode training are distinguished, which are defined as:

ST	Supervised training with simulated low-field MRI data.
TTT	Self-supervised test-time training for individual samples during inference.
SPT+TTT	Supervised pretraining with simulated data followed by self-supervised test-time training for individual samples at inference.

First, the performance of the different scenarios is evaluated by quantitative measures on the simulated test data. To validate the denoising capability, the model is applied on the test data which originates from the same data distribution as the training and validation data, i.e. the simulated T_2 -weighted data. Subsequently, the model is evaluated using the simulated T_1 -weighted images. Due to the different contrast, the simulated T_1 -weighted images constitute a shift towards the known distribution of inputs, because the model was trained only by T_2 -weighted images. In the case of ST, the model was trained with T_2 -weighted data and its performance is measured on the T_1 -weighted images without an additional refinement step. The TTT investigates the self-supervised *Noise2Self* training approach, without any pretraining, i.e. the model weights are randomly initialized. Self-supervised TTT is carried out until the model fully converged to provide a fair comparison. Last but not least, the SPT+TTT investigates, if the performance obtainable from TTT alone can be further improved, if ST is used to initialize the model weights for TTT. Since a noise-free ground truth is available for the simulated data, NRMSE and PSNR are reported. The same comparison is carried out for the actual low-field MRI measurements of a healthy volunteer’s brain with T_2 -weighted and PD-weighted contrast. In case of the T_2 -weighted data, a distribution shift is given by the deviations from simulation, and in case of the PD-weighted data, both, contrast and simulation deviations contribute to the data distribution shift. Since here a ground truth is not available, no-reference metrics are utilized for a quantitative comparison. A reference-free blur metric [137] and the SNR which was introduced in Section 3.3.1 are considered.

7.3 Results

To evaluate the impact of the proposed two-stage training process and in particular the possibilities of refining a model by test-time training which has been pretrained with supervision, the results for simulated and real low-field MRI images are presented in the




Figure 7.2: Denoising example for simulated low-field MRI data obtained from the U-Net architecture, when trained by supervised learning. The example shows the center slice of a three-dimensional sample from the test dataset which originates from the same distribution as the training and validation data. The model output in the second column demonstrates the performance of the supervised training.

following. First, a data distribution shift in simulated low-field MRI data is evaluated, then actual in vivo data is considered.

7.3.1 Simulated low-field data

In the first scenario, the model is trained on the simulated T_2 -weighted low-field MRI data. To validate the models denoising performance on this distribution of data, the application for a random sample from the test dataset is shown in Figure 7.2. The model is clearly capable of removing the noise from the input while retaining smaller details in the output. The evaluations of NRMSE and PSNR for the test dataset, reveal that the NRMSE of the input can be decreased by 67.6 % when denoised by the model, and the PSNR can be increased by 90.5 %. These values correlate with the example in Figure 7.2, confirming the denoising capabilities of the model trained with supervision, when applied to T_2 -weighted data from the same distribution.

Using the identical model, the adaptability to unseen data from a different distribution is explored using simulated T_1 -weighted data for which a ground truth is available. The box plots in Figure 7.3 show the quantitative assessment by comparing NRMSE and PSNR




Figure 7.3: Comparison of supervised training (ST) with distributional discrepancy, self-supervised test-time training (TTT) and supervised pretraining with self-supervised test-time training (SPT + TTT) across the test data by NRMSE and PSNR. The results correspond to the observations in Figure 7.4.

the T_1 -weighted test dataset which includes 150 samples. If the denoising model is directly applied to the T_1 -weighted samples without further adjust, the mean NRMSE increases by 105.4 % and the mean PSNR decreases by 4.1 %. With TTT the mean NRMSE over the inputs can be decreased by 10.6 % and the mean PSNR can be increased by 4.9 %. The contribution of supervised pretraining to the self-supervised test-time training yields an additional 5.1 % and 1.5 % improvement, towards the TTT without pretraining. How this translates to the obtainable image quality is being demonstrated by the example in Figure 7.4. Here, the direct application of the different training cases, that is ST, TTT, SPT+TTT, is compared to the original input and the ground truth of an exemplary sample from the T_1 -weighted dataset. Each sample is three-dimensional, but only a single image slice from the center is shown in Figure 7.2. However, the model performance over the difference image slices is consistent. A direct application of the model trained by T_2 -weighted data yields effectively removal of the background noise, but the contrast of the brain structure vanishes compared to the input, whereby the details in the image are no longer recognizable. For the shown example, the model solely trained by supervision with T_2 -weighted data clearly fails. Employing the self-supervised TTT approach, noise is removed from the background and the brain structure, while retaining important details. The background has a slight offset, but the image is still effectively denoised. With additional pretraining, i.e. the combination of SPT+TTT, noise is also effectively removed, while the image features a smoother but still sharp contrast.





Figure 7.4: Comparison of supervised training (ST) with distributional discrepancy, self-supervised test-time training (TTT) and supervised pretraining with self-supervised test-time training (SPT + TTT) on a set of simulated T_1 weighted test data.


7.3.2 In vivo data

Last but not least, the two-stage training process is applied to real in vivo low-field MRI data of a healthy volunteer's brain. The different training scenarios are compared for T_2 -weighted data and for PD-weighted data. Here, the same model used previously, which was trained on simulated T_2 -weighted data is employed. Figures 7.5 and 7.6 present the in vivo denoising for both contrasts, each showing two slices of the data. For the T_2 -weighted data, the direct application of the model trained with supervision maintains high contrast but loses some structural contours. A qualitative comparison of the model output to the input reveals that the result are a bit over-smoothed and hyper intense. This is reflected by the 27.7 % increase of blurring and 130.5 % increases of SNR for the T_2 -weighted image data in Figure 7.5a. The self-supervised training without test-time training exhibits checkerboard artifacts. Referring to Figure 7.5a, blurring increases by 40.4 % and SNR by 42.0 %. Compared to ST, the images obtained by TTT are more blurry but the intensities match the original images much better. If the TTT is combined with supervised pretraining in the case of SPT+TTT, the image quality does not suffer by checkerboard artifacts or hyper intensities, yielding the sharpest output with the highest degree of details compared to the other training cases. In Figure 7.5a, the SNR is increased by 7.6 %, what is less compared to ST, therefor the blurring only increases by 4.8 %. The zoomed-in sections which provide a more detailed comparison confirms that the two-stage training method yields the best compromise between noise removal and conservation of details.

Compared to the T_2 -weighted images, the PD-weighted data exhibits a more significant



(a) Image slice number 8 with a magnified detail view from the upper right image section



(b) Image slice number 16 with a magnified detail view from the lower left image section


Figure 7.5: Denoising results for T2-weighted 3D in vivo brain images obtained by TSE sequence with TE/TR of 20/2000 ms. The FOV of the data is $240 \times 200 \times 210$ mm with a resolution of $2 \times 2 \times 5$ mm for RO \times PE1 \times PE2. Comparison of supervised training (ST), self-supervised test-time training (TTT) and self-supervised test-time training with supervised pretraining (SPT + TTT).

deviation from the original training data. It constitutes a fundamentally different image contrast and is affected by approximation due to the simulation process. Figure 7.6 shows the identical two slices of the healthy volunteers brain as for the T_2 -weighted contrast, but with PD-weighted images. When trained solely by supervision as in the case of ST, the application of the model yields a blurry and over-smoothed result with a mismatch in intensity. Compared to the original image, blurring in Figure 7.6a increases by 16.9 %. Similar effects could be observed for the simulated data with T_1 -weighted contrast. Since the background noise could still be removed effectively, SNR increases by 30.6 %. Test time training without pretraining leads to some similar artifacts as observed in Figure 7.5. However, for the PD-weighted images, it is less pronounced. Here, the blurring increases by 15.1 % and SNR increases by 23.7 %. When supervised training is further refined by test-time training, the denoising results can be clearly enhanced. Particularly the details in the image are shown with a much better contrast, compared to the ST and TTT cases, as can be observed for the delimitation of the ventricles in the center of the brain for instance. In quantitative terms, the improvement by the combination SPT+TTT yields an increase of 6.1 % in SNR while being able to decrease the blurring in the image by 2.87 %.


7.4 Discussion

This chapter investigated a two-stage training approach which refines supervised pretraining by test-time training (SPT+TTT), and compared it to supervised-only training (ST) and self-supervised-only (TTT) training. Different data distribution shifts are considered for the comparison, that is supervised training is performed using simulated T_2 -weighted images, while at inference, the model is applied to simulated T_1 -weighted images and measured low-field MRI data. As shown by the results in this chapter, models trained with supervision perform very well when applied to unseen data with identical characteristics as the training data, i.e. for data from the same distribution. However, due to approximations in the simulation process, e.g. in terms of image contrast, the trained model may not necessarily retain its performance when applied to measured low-field MRI data. The self-supervised refinement step on the individual samples aims to overcome these issues.

When directly applying the denoising model trained by supervised learning to image contrasts that deviate from the contrast of the training data, the denoising performance is strongly reduced, what could be shown by an increase of NRMSE and a decrease in PSNR. Employing the proposed refinement by test-time training, the denoising performance can be regained. According to the metrics, the model weights are not just relearned as the combination of supervised training with test-time training improves the test-time training without pretraining. Similar observations can be made when applying the model to actual



(a) Image slice number 8 with a magnified detail view from the upper right image section



(b) Image slice number 16 with a magnified detail view from the lower left image section

Figure 7.6: Denoising results for PD-weighted 3D in vivo brain images obtained by TSE sequence with TE/TR of 20/500 ms. The FOV of the data is $240 \times 200 \times 210$ mm with a resolution of $2 \times 2 \times 5$ mm for RO \times PE1 \times PE2. Comparison of supervised training (ST), self-supervised test-time training (TTT) and self-supervised test-time training with supervised pretraining (SPT + TTT).

low-field MRI in vivo data. Applying the supervised training without adaption leads to mismatched intensities and over-smoothed results. If the discrepancy between the data is bigger, as in the case of the PD-weighted data, this effect is more severe. Since the model has never encountered PD-weighted low-field MRI data during training, its denoising performance on these samples is lower compared to data which is more similar to data from the training distribution. Self-supervised test-time training can help to overcome this issue, because the model weights are optimized on individual samples. However, if the data is very noisy, as in the case of the T_2 -weighted data, the performance of the test-time training is limited and leads to artifacts in the denoised result. Employing the supervised training as a pretraining stage which is refined by the self-supervised test-time training improves the image quality for T_2 -weighted and PD-weighted data. The results are not entirely noise-free, but the two-stage training approach offers an excellent compromise between SNR enhancement and the amount of blurring which is introduced to the output. The combination of supervised training and self-supervised refinement in a two-stage training strategy provides the most robust denoising performance on the compared in vivo examples. Notably, for PD-weighted data, the proposed approach enhances SNR while simultaneously reducing image blur. The adaptability to new data improves the applicability of the model and is particularly valuable for low-field MRI, where training data is predominantly simulated.

7.5 Summary

The denoising methods proposed in the previous chapters worked well on the simulated data, but suffer from performance losses, when the model is applied to measured low-field MRI data which deviates from the simulated data. This chapter investigated a two-stage training approach to adapt a pretrained NN to real low-field MRI data. The method combines supervised pretraining on simulated data with self-supervised test-time training on individual in vivo samples. Supervised pretraining uses a U-Net with multihead self-attention, achieving significant noise reduction on data similar to the training distribution. Test-time training refines the model weights at inference using the *Noise2Self* framework, adapting to distributional shifts without requiring reference data. The combined approach of supervised pretraining and test-time training outperformed supervised-only and self-supervised-only models, showing improved noise suppression and structural preservation in simulated and measured in vivo low-field data. This is reflected by reduced blurring and enhanced SNR, yield an optimized image quality. The proposed method provides an effective solution to bridge the model performance gap when training and inference data are significantly difference. Thus, it is a good candidate to improve the transferability also

for the joint reconstruction approach and the estimation of the regularization parameter maps.


8 Discussion and conclusion

The presented work demonstrates a successful application of deep learning to enhance image quality in low-field MRI at 50 mT. Initially, a novel console architecture for advanced low-field MRI was developed and tested. With the imaging results from the console, a method was presented that combines model-based reconstruction with deep learning to target image and B_0 field reconstruction for low-field MRI. Subsequently, an alternative regularization approach was investigated, which combined conventional TV regularization with the estimation of a regularization parameter map to improve the interpretability and transferability of the network to measured data. Last but not least, a two-staged training concept was proposed, which allows adapting models trained with supervision to different data by a self-supervised refinement step, while maintaining the model's performance. All the proposed methods were evaluated on simulated data but also tested on in vivo data acquired with the low-field MRI scanner.


Nexus, a versatile console for advanced image acquisition in low-field MRI was developed and demonstrated in Chapter 3. It facilitates the acquisition of additional sensor information, a fully transparent data processing pipeline, and the embedding of GPU, what enables deep learning-based image reconstruction. The modular design allows custom configuration of the channel number and enables easy adaption for use at different field strengths. Thereby, the proposed console overcomes limitations with respect to the interpretability and adaptability of existing proprietary systems and delivers more performance in comparison to existing open-source systems. By default, it enables techniques, such as the integration of auxiliary sensors for EMI mitigation, temperature, or B_0 field measurements, which offers great potential to enhance image quality. As demonstrated by the authors of [90], deep learning helps to further improve EMI mitigation by the EDITER algorithm [88], which was used in Chapter 3. With the proposed Nexus console, such methods could be directly integrated into the console, making them a standard for low-field MRI. The option of transferring the acquired data directly to the GPU, enables faster data processing and seamless integration of deep learning-based techniques. Using TSE-based sequences written in *PyPulseq*, imaging capabilities of the system could be demonstrated at two different sites, yielding three-dimensional in vivo PD- and T_2 -weighted image contrasts on portable 50 mT systems in 5 min 52 s and 6 min 28 s, respectively. A comparison to a state-of-the-art reference console demonstrated image acquisition without losses in data quality.

Different deep learning approaches have been investigated to address the identified main limitations of low-field MRI which could also be observed in Chapter 3: Low SNR and high B_0 field inhomogeneities. The joint model-based approach for image and B_0 field reconstruction with deep learning-based regularization presented in Chapter 4 showed superior performance in comparison to conventional methods by improving the PSNR up to 11.7% and the RMSE up to 86.3%. End-to-end training of the NNs with alternations between the image and B_0 field related sub-problems contributed to improved robustness, while the evaluation was limited to simulated low-field MRI data. Chapter 5 evaluated the proposed method on data acquired with the Nexus console at a 50 mT low-field MRI system, including phantom and in vivo brain scans. Effective denoising capabilities could be demonstrated, especially for the T_2 -weighted in vivo brain data. The correction of distortions, caused by the B_0 field inhomogeneities, could be validated for an ACR phantom. However, a quantitative evaluate of the B_0 field, e.g. by a B_0 mapping sequence [116], needs to be done to fully validate the results. If needed, the model input could also be extended by additional sensor information [106] to increase reliability. For well-shimmed systems, as in the case of the in vivo acquisitions corrections of the B_0 field inhomogeneities might not necessarily be required. However, it remains to be demonstrated whether field distortions occur outside well-regulated lab conditions, such as in portable applications where temperature fluctuations may arise, and whether the proposed distortion correction method can effectively compensate for them.

Denoising approaches to improve SNR can be applied independently to substantially improve image quality. This was shown in particular for low-signal sequences as in the case of the T_2 -weighted contrast. Relatively high blurring in the in vivo reconstructions with the joint model-based approach suggests a performance gap when applying the model trained on simulated low-field data to measured data, particularly if the measured sample features a different image contrast. Consequently, Chapter 6 investigated an approach that combines conventional TV regularization with the prediction of a regularization parameter map by a NN instead of employing the network output directly for regularization. This leads to better denoising performance with less blurring, especially for the application on the PD-weighted in vivo images. As an additional refinement step, Chapter 7 introduced a two-staged training process of supervised learning combined with self-supervised training at test time. The evaluations revealed that a model trained with supervision can be further enhanced by refining the model weights on individual samples, especially in the presence of a data distribution gap between training and inference data. This refinement helped to avoid excessive blurring while effectively improving the SNR for both the PD- and T_2 -weighted image contrasts in vivo. Figure 8.1 provides a comprehensive comparison of all the presented methods as they were proposed in this thesis, including the joint model-based approach, regularization parameter map estimation, and the two-stage training strategy.



(a) PD-weighted image contrast.



(b) T₂-weighted image contrast.

Figure 8.1: Comparison of magnitude images from in vivo brain data acquired with the Nexus console: Images are obtained through FFT reconstruction, joint image and B₀ field reconstruction, and FFT reconstruction with denoising using estimated regularization parameter maps and the two-stage training strategy. The reference-free blur metric and SNR is calculated for each image.

Among these, the two-stage training approach from Chapter 7 offers the best compromise between image denoising and sharpness, yielding a 13.7 % and 8.8 % increase in SNR for T_2 -weighted and PD-weighted images, respectively. Unlike other methods, it limits the blurring in the T_2 -weighted image to 6.5 % and even reduces the blur in PD-weighted images by 2.2 %. A qualitative comparison further confirms that the two-stage training strategy delivers the highest perceived image quality.

The findings of this thesis suggest that no single method alone provides an optimal solution to enhance image quality in a low field MRI. Instead, combining multiple approaches is key to obtain optimized image quality, with an accurate acquisition model at its core. As suggested in [143], a “plug-and-play” approach could be considered, where image and B_0 field map estimations are used as priors, while the learned regularization parameter maps provide adaptive regularization. By incorporating multiple training stages, including both supervised and self-supervised learning, the output could be refined for individual samples directly at the scanner, further improving robustness and image quality in real-world applications. This enables reliable image reconstruction in a modular fashion, facilitating adaptability to different hardware setups. However, the training process and particularly the selection of training data is crucial and can be enhanced by more sophisticated simulations [38] or training strategies that fully rely on self-supervision [44]. Using the information of auxiliary sensors, available from the Nexus console, EMI sensing coils can be used to enrich the acquisition model or preprocess the acquisition data, while field probes provide additional information for the B_0 field estimation or gradient field correction. Such additional sensor information helps to further improve the robustness of the scanner under uncontrolled environmental conditions, and contributes to improve the overall image quality [90], [144], [145]. Towards acceleration of the imaging process, different undersampling techniques can be employed and integrated into the acquisition model [35], [39], [40]. As this further reduces the amount of information available, the importance of model-based reconstruction increases. Integrating all these techniques into the Nexus console is technically feasible. However, more advanced methods may require additional computational resources, which warrants an analysis of their cost effectiveness. Therefore, it is advisable to tailor the reconstruction model to the specific requirements of the deployment environment. Although the evaluations in this thesis are limited to head and knee images, the methods can in principle also be extended to other extremities or even whole body MRI [29].

The methodologies presented in this thesis not only demonstrate the effectiveness of deep learning in enhancing image quality for low-field MRI but also establish a framework that seamlessly integrates data acquisition with these techniques for image reconstruction and processing. Emphasizing robustness and reliability, the proposed methods enhance the performance of low-field MRI and offer a promising foundation for clinical integration.

With their compact design and affordability, low-field MRI scanners are well-suited for diverse applications, from resource-limited settings and underserved regions to stroke units and smaller hospitals in the countryside. Deep learning further expands the capabilities of these accessible systems by improving image quality. However, ensuring robustness and reliability remains critical for their adoption in radiological practice. Transitioning these methodologies into clinical practice is essential to fully explore the potential of low-field MRI for value-based diagnostic imaging, but remains an open challenge.

Bibliography

- [1] B. M. Dale, M. A. Brown, R. C. Semelka, and M. A. Brown, *MRI: Basic Principles and Applications*, 5th ed. John Wiley & Sons, Ltd, 2015.
- [2] L. C. Maas and P. Mukherjee, „Diffusion MRI: Overview and clinical applications in neuroradiology“, *Applied Radiology*, pp. 44–60, Nov. 2005. DOI: 10.37549/AR1383. (visited on 01/24/2025).
- [3] T. C. Arnold, C. W. Freeman, B. Litt, and J. M. Stein, „Low-field MRI: Clinical promise and challenges“, *Journal of Magnetic Resonance Imaging*, vol. 57, no. 1, pp. 25–44, Jan. 2023. DOI: 10.1002/jmri.28408. (visited on 11/02/2023).
- [4] J. Hennig, „An evolution of low-field strength MRI“, *Magnetic Resonance Materials in Physics, Biology and Medicine*, vol. 36, no. 3, pp. 335–346, Jun. 2023. DOI: 10.1007/s10334-023-01104-z. (visited on 10/17/2024).
- [5] M. Hori, A. Hagiwara, M. Goto, A. Wada, and S. Aoki, „Low-Field Magnetic Resonance Imaging: Its History and Renaissance“, *Investigative Radiology*, vol. 56, no. 11, pp. 669–679, Nov. 2021. DOI: 10.1097/RLI.0000000000000810. (visited on 02/14/2022).
- [6] M. Sarraçanie and N. Salameh, „Low-Field MRI: How Low Can We Go? A Fresh View on an Old Debate“, *Frontiers in Physics*, vol. 8, p. 172, Jun. 2020. DOI: 10.3389/fphy.2020.00172. (visited on 06/04/2021).
- [7] T. O'Reilly, W. Teeuwisse, and A. Webb, „Three-dimensional MRI in a homogenous 27 cm diameter bore Halbach array magnet“, *Journal of Magnetic Resonance*, vol. 307, p. 106578, Oct. 2019. DOI: 10.1016/j.jmr.2019.106578. (visited on 11/09/2021).
- [8] T. Pires and J. M.M, „The Rise and Efficiency of Low Field Portable MRI Scanners“, *Journal of Health and Allied Sciences NU*, vol. 14, no. 02, pp. 163–168, Apr. 2024. DOI: 10.1055/s-0043-1769759. (visited on 10/17/2024).
- [9] L. L. Wald, P. C. McDaniel, T. Witzel, J. P. Stockmann, and C. Z. Cooley, „Low-cost and portable MRI“, *Journal of Magnetic Resonance Imaging*, vol. 52, no. 3, pp. 686–696, Sep. 2020. DOI: 10.1002/jmri.26942. (visited on 11/09/2021).

- [10] A. E. Campbell-Washburn, R. Ramasawmy, M. C. Restivo, *et al.*, „Opportunities in Interventional and Diagnostic Imaging by Using High-Performance Low-Field-Strength MRI“, *Radiology*, vol. 293, no. 2, pp. 384–393, Nov. 2019. DOI: 10.1148/radiol.2019190452. (visited on 07/14/2023).
- [11] E. Kuoy, J. Glavis-Bloom, G. Hovis, *et al.*, „Point-of-Care Brain MRI: Preliminary Results from a Single-Center Retrospective Study“, *Radiology*, vol. 305, no. 3, pp. 666–671, Dec. 2022. DOI: 10.1148/radiol.211721. (visited on 06/03/2024).
- [12] K. T. Islam, S. Zhong, P. Zakavi, *et al.*, „Improving portable low-field MRI image quality through image-to-image translation using paired low- and high-field images“, *Scientific Reports*, vol. 13, no. 1, p. 21 183, Dec. 2023. DOI: 10.1038/s41598-023-48438-1. (visited on 10/17/2024).
- [13] M. H. Mazurek, B. A. Cahn, M. M. Yuen, *et al.*, „Portable, bedside, low-field magnetic resonance imaging for evaluation of intracerebral hemorrhage“, *Nature Communications*, vol. 12, no. 1, p. 5119, Aug. 2021. DOI: 10.1038/s41467-021-25441-6. (visited on 06/27/2024).
- [14] K. N. Sheth, M. M. Yuen, M. H. Mazurek, *et al.*, „Bedside detection of intracranial midline shift using portable magnetic resonance imaging“, *Scientific Reports*, vol. 12, no. 1, p. 67, Jan. 2022. DOI: 10.1038/s41598-021-03892-7. (visited on 06/03/2024).
- [15] G. Hovis, M. Langdorf, E. Dang, and D. Chow, „MRI at the Bedside: A Case Report Comparing Fixed and Portable Magnetic Resonance Imaging for Suspected Stroke“, *Cureus*, Aug. 2021. DOI: 10.7759/cureus.16904. (visited on 10/17/2024).
- [16] A. M. Prabhat, A. L. Crawford, M. H. Mazurek, *et al.*, „Methodology for Low-Field, Portable Magnetic Resonance Neuroimaging at the Bedside“, *Frontiers in Neurology*, vol. 12, p. 760 321, Dec. 2021. DOI: 10.3389/fneur.2021.760321. (visited on 01/24/2025).
- [17] S. C. L. Deoni, P. Medeiros, A. T. Deoni, *et al.*, „Development of a mobile low-field MRI scanner“, *Scientific Reports*, vol. 12, no. 1, p. 5690, Apr. 2022. DOI: 10.1038/s41598-022-09760-2. (visited on 04/18/2023).
- [18] J. M. Algarín, T. Guallart-Naval, E. Gastaldi-Orquín, *et al.*, *Portable MRI for major sporting events – a case study on the MotoGP World Championship*, 2023. DOI: 10.48550/arxiv.2303.09264. arXiv: 2312.10549. (visited on 05/03/2023).
- [19] C. M. Rogers, P. S. Jones, and J. S. Weinberg, „Intraoperative MRI for Brain Tumors“, *Journal of Neuro-Oncology*, vol. 151, no. 3, pp. 479–490, Feb. 2021. DOI: 10.1007/s11060-020-03667-6. (visited on 02/12/2025).

-
- [20] J. Otero, J. M. Pearce, D. Gozal, and R. Farré, „Open-source design of medical devices“, *Nature Reviews Bioengineering*, vol. 2, no. 4, pp. 280–281, Feb. 2024. DOI: 10.1038/s44222-024-00162-9. (visited on 11/08/2024).
- [21] C. De Maria, L. Di Pietro, A. Ravizza, A. D. Lantada, and A. D. Ahluwalia, „Open-source medical devices: Healthcare solutions for low-, middle-, and high-resource settings“, in *Clinical Engineering Handbook*, Elsevier, 2020, pp. 7–14. DOI: 10.1016/B978-0-12-813467-2.00002-X. (visited on 11/08/2024).
- [22] L. Winter, J. Periquito, C. Kolbitsch, *et al.*, „Open-source magnetic resonance imaging: Improving access, science, and education through global collaboration“, *NMR in Biomedicine*, vol. 37, no. n/a, e5052, Nov. 2023. DOI: 10.1002/nbm.5052. (visited on 06/03/2024).
- [23] L. Winter, R. Pellicer-Guridi, L. Broche, *et al.*, „Open Source Medical Devices for Innovation, Education and Global Health: Case Study of Open Source Magnetic Resonance Imaging“, in *Co-Creation*, T. Redlich, M. Moritz, and J. P. Wulfsberg, Eds., Springer International Publishing, 2019, pp. 147–163. DOI: 10.1007/978-3-319-97788-1_12. (visited on 10/19/2021).
- [24] T. Guallart-Naval, T. O'Reilly, J. M. Algarín, *et al.*, „Benchmarking the performance of a low-cost magnetic resonance control system at multiple sites in the open MaRCoS community“, *NMR in Biomedicine*, vol. 36, no. 1, e4825, Jan. 2023. DOI: 10.1002/nbm.4825. (visited on 11/06/2023).
- [25] S. S. Bhat, T. T. Fernandes, P. Poojar, *et al.*, „Low-Field MRI of Stroke: Challenges and Opportunities“, *Journal of Magnetic Resonance Imaging*, vol. 54, no. 2, pp. 372–390, Aug. 2021. DOI: 10.1002/jmri.27324. (visited on 02/14/2022).
- [26] O. Islam, A. W. Lin, and A. Bharatha, „Potential application of ultra-low field portable MRI in the ICU to improve CT and MRI access in Canadian hospitals: A multi-center retrospective analysis“, *Frontiers in Neurology*, vol. 14, p. 1220091, Sep. 2023. DOI: 10.3389/fneur.2023.1220091. (visited on 10/17/2024).
- [27] T. Pogarell, R. Heiss, R. Janka, *et al.*, „Modern low-field MRI“, *Skeletal Radiology*, vol. 53, pp. 1751–1760, Feb. 2024. DOI: 10.1007/s00256-024-04597-4. (visited on 07/22/2024).
- [28] F. F. Schröder, C. E. Post, S. M. Van Raak, *et al.*, „The diagnostic potential of low-field MRI in problematic total knee arthroplasties - a feasibility study“, *Journal of Experimental Orthopaedics*, vol. 7, no. 1, p. 59, Dec. 2020. DOI: 10.1186/s40634-020-00274-2. (visited on 10/17/2024).
- [29] Y. Zhao, Y. Ding, V. Lau, *et al.*, „Whole-body magnetic resonance imaging at 0.05 Tesla“, *Science*, vol. 384, no. 6696, May 2024. DOI: 10.1126/science.adm7168. (visited on 07/22/2024).

- [30] A. Webb, „Artifacts at Low Field“, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 30, 2022, p. 1818. (visited on 07/22/2024).
- [31] J. P. Marques, F. F. Simonis, and A. G. Webb, „Low-field MRI: An MR physics perspective“, *Journal of Magnetic Resonance Imaging*, vol. 49, no. 6, pp. 1528–1542, Jun. 2019. DOI: 10.1002/jmri.26637. (visited on 12/17/2021).
- [32] N. Nazir, A. Sarwar, and B. S. Saini, „Recent developments in denoising medical images using deep learning: An overview of models, techniques, and challenges“, *Micron*, vol. 180, p. 103615, May 2024. DOI: 10.1016/j.micron.2024.103615. (visited on 03/24/2025).
- [33] Z. Chen, K. Pawar, M. Ekanayake, *et al.*, „Deep Learning for Image Enhancement and Correction in Magnetic Resonance Imaging—State-of-the-Art and Challenges“, *Journal of Digital Imaging*, vol. 36, no. 1, pp. 204–230, Nov. 2022. DOI: 10.1007/s10278-022-00721-9. (visited on 01/24/2025).
- [34] J. Zbontar, F. Knoll, A. Sriram, *et al.*, „fastMRI: An Open Dataset and Benchmarks for Accelerated MRI“, Nov. 2019. DOI: 10.48550/arXiv.1811.08839. (visited on 11/17/2021).
- [35] M. L. De Leeuw Den Bouter, G. Ippolito, T. P. A. O’Reilly, *et al.*, „Deep learning-based single image super-resolution for low-field MR brain images“, *Scientific Reports*, vol. 12, no. 1, p. 6362, Apr. 2022. DOI: 10.1038/s41598-022-10298-6. (visited on 10/17/2024).
- [36] N. Koonjoo, B. Zhu, G. C. Bagnall, D. Bhutto, and M. S. Rosen, „Boosting the signal-to-noise of low-field MRI with deep learning image reconstruction“, *Scientific Reports*, vol. 11, no. 1, p. 8248, Dec. 2021. DOI: 10.1038/s41598-021-87482-7. (visited on 10/20/2021).
- [37] A. G. Hernandez, P. Fau, S. Rapacchi, *et al.*, „Improving Image Quality In Low-Field MRI With Deep Learning“, in *2021 IEEE International Conference on Image Processing (ICIP)*, IEEE, Sep. 2021, pp. 260–263. DOI: 10.1109/ICIP42928.2021.9506659. (visited on 01/24/2025).
- [38] A. Salehi, M. Mach, C. Najac, *et al.*, „Denoising low-field MR images with a deep learning algorithm based on simulated data from easily accessible open-source software“, *Journal of Magnetic Resonance*, vol. 370, p. 107812, Jan. 2025. DOI: 10.1016/j.jmr.2024.107812. (visited on 01/24/2025).
- [39] R. Ayde, T. Senft, N. Salameh, and M. Sarracanie, „Deep learning for fast low-field MRI acquisitions“, *Scientific Reports*, vol. 12, no. 1, p. 11394, Jul. 2022. DOI: 10.1038/s41598-022-14039-7. (visited on 07/13/2023).

-
- [40] C. Man, V. Lau, S. Su, *et al.*, „Deep learning enabled fast 3D brain MRI at 0.055 tesla“, *Science Advances*, vol. 9, no. 38, Sep. 2022. DOI: 10.1126/sciadv.adi9327. (visited on 01/24/2025).
- [41] V. Lau, L. Xiao, Y. Zhao, *et al.*, „Pushing the limits of low-cost ultralow-field MRI by dual-acquisition deep learning 3D superresolution“, *Magnetic Resonance in Medicine*, mrm.29642, Apr. 2023. DOI: 10.1002/mrm.29642. (visited on 04/18/2023).
- [42] M. Z. Darestani, J. Liu, and R. Heckel, *Test-Time Training Can Close the Natural Distribution Shift Performance Gap in Deep Learning Based Compressed Sensing*, 2022. DOI: 10.48550/arXiv.2204.07204. arXiv: 2204.07204.
- [43] Y. Sun, X. Wang, Z. Liu, *et al.*, *Test-Time Training with Self-Supervision for Generalization under Distribution Shifts*, 2020. DOI: 10.48550/arXiv.1909.13231. arXiv: 1909.13231.
- [44] B. Zhou, J. Schlemper, N. Dey, *et al.*, „Dual-domain self-supervised learning for accelerated non-Cartesian MRI reconstruction“, *Medical Image Analysis*, vol. 81, p. 102538, Oct. 2022. DOI: 10.1016/j.media.2022.102538. (visited on 02/10/2025).
- [45] D. Y. Zeng, J. Shaikh, S. Holmes, *et al.*, „Deep Residual Network for Off-Resonance Artifact Correction with Application to Pediatric Body Magnetic Resonance Angiography with 3D Cones“, *Magnetic Resonance in Medicine*, p. 24, 2020. DOI: 10.1002/mrm.27825.
- [46] M. W. Haskell, A. A. Cao, D. C. Noll, and J. A. Fessler, „Deep learning field map estimation with model-based image reconstruction for off resonance correction of brain images using a spiral acquisition“, in *ISMRM Workshop on Data Sampling & Image Reconstruction*, Jan. 2020.
- [47] D. B. T. Le, M. Sadinski, A. Nacev, R. Narayanan, and D. Kumar, „Deep Learning-based Method for Denoising and Image Enhancement in Low-Field MRI“, in *2021 IEEE International Conference on Imaging Systems and Techniques (IST)*, IEEE, Aug. 2021, pp. 1–6. DOI: 10.1109/IST50367.2021.9651441. (visited on 01/24/2025).
- [48] V. Antun, F. Renna, C. Poon, B. Adcock, and A. C. Hansen, „On instabilities of deep learning in image reconstruction - Does AI come at a cost?“, *Proceedings of the National Academy of Sciences*, vol. 117, no. 48, pp. 30088–30095, Dec. 2020. DOI: 10.1073/pnas.1907377117. arXiv: 1902.05300. (visited on 11/17/2021).
- [49] V. Monga, Y. Li, and Y. C. Eldar, „Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing“, *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, Mar. 2021. DOI: 10.1109/MSP.2020.3016905. (visited on 04/20/2023).
-

- [50] A. Kofler, M. Haltmeier, T. Schaeffter, and C. Kolbitsch, „An end-to-end-trainable iterative network architecture for accelerated radial multi-coil 2D cine MR image reconstruction“, *Medical Physics*, vol. 48, no. 5, pp. 2412–2425, May 2021. DOI: 10.1002/mp.14809. (visited on 11/08/2021).
- [51] T. O’Reilly and A. G. Webb, „In vivo T1 and T2 relaxation time maps of brain tissue, skeletal muscle, and lipid measured in healthy volunteers at 50 mT“, *Magnetic Resonance in Medicine*, mrm.29009, Sep. 2021. DOI: 10.1002/mrm.29009. (visited on 09/27/2021).
- [52] T. O’Reilly, W. M. Teeuwisse, D. Gans, K. Koolstra, and A. G. Webb, „In vivo 3D brain and extremity MRI at 50 mT using a permanent magnet Halbach array“, *Magnetic Resonance in Medicine*, vol. 85, no. 1, pp. 495–505, Jan. 2021. DOI: 10.1002/mrm.28396. (visited on 06/18/2021).
- [53] R. W. Brown, Y.-C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan, *Magnetic Resonance Imaging: Physical Principles and Sequence Design*, 2nd ed. John Wiley & Sons, Inc, 2014.
- [54] M. A. Bernstein, K. F. King, and Z. J. Zhou, *Handbook of MRI Pulse Sequences*. Academic Press, 2004.
- [55] B. De Vos, R. F. Remis, and A. G. Webb, „Characterization of concomitant gradient fields and their effects on image distortions using a low-field point-of-care Halbach-based MRI system“, *Magnetic Resonance in Medicine*, vol. 91, no. 2, pp. 828–841, Feb. 2024. DOI: 10.1002/mrm.29879. (visited on 03/21/2025).
- [56] C. A. Michal, K. Broughton, and E. Hansen, „A high performance digital receiver for home-built nuclear magnetic resonance spectrometers“, *Review of Scientific Instruments*, vol. 73, no. 2, pp. 453–458, Feb. 2002. DOI: 10.1063/1.1433950. (visited on 01/23/2025).
- [57] A. Valori, J. Mitchell, and E. J. Fordham, „Digital filters for low-field NMR“, *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, vol. 46B, no. 4, pp. 202–220, Oct. 2016. DOI: 10.1002/cmr.b.21346. (visited on 10/15/2024).
- [58] M. W. Haskell, J.-F. Nielsen, and D. C. Noll, *Off-resonance artifact correction for magnetic resonance imaging: A review*, May 2022. DOI: 10.48550/arXiv.2205.01028. arXiv: 2205.01028. (visited on 08/19/2022).
- [59] J.-P. Korb and R. G. Bryant, „Magnetic field dependence of proton spin-lattice relaxation times“, *Magnetic Resonance in Medicine*, vol. 48, no. 1, pp. 21–26, Jul. 2002. DOI: 10.1002/mrm.10185. (visited on 02/12/2025).
- [60] A. Macovski, „Noise in MRI“, *Magnetic Resonance in Medicine*, vol. 36, no. 3, pp. 494–497, Sep. 1996. DOI: 10.1002/mrm.1910360327. (visited on 01/07/2025).

-
- [61] E. R. McVeigh, R. M. Henkelman, and M. J. Bronskill, „Noise and filtration in magnetic resonance imaging“, *Medical Physics*, vol. 12, no. 5, pp. 586–591, Sep. 1985. DOI: 10.1118/1.595679. (visited on 01/07/2025).
 - [62] J. Fessler, „Model-Based Image Reconstruction for MRI“, *IEEE Signal Processing Magazine*, vol. 27, no. 4, pp. 81–89, Jul. 2010. DOI: 10.1109/MSP.2010.936726. (visited on 01/06/2025).
 - [63] R. A. Willoughby, „Solutions of Ill-Posed Problems (A. N. Tikhonov and V. Y. Arsenin)“, *SIAM Review*, vol. 21, no. 2, pp. 266–267, Apr. 1979. DOI: 10.1137/1021044. (visited on 02/11/2025).
 - [64] L. I. Rudin, S. Osher, and E. Fatemi, „Nonlinear total variation based noise removal algorithms“, *Physica D: Nonlinear Phenomena*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992. DOI: 10.1016/0167-2789(92)90242-F. (visited on 01/22/2025).
 - [65] A. Chambolle, V. Caselles, D. Cremers, M. Novaga, and T. Pock, „An Introduction to Total Variation for Image Analysis“, in *Theoretical Foundations and Numerical Methods for Sparse Recovery*, M. Fornasier, Ed., De Gruyter, Jul. 2010, pp. 263–340. DOI: 10.1515/9783110226157.263. (visited on 03/28/2025).
 - [66] M. Hestenes and E. Stiefel, „Methods of conjugate gradients for solving linear systems“, *Journal of Research of the National Bureau of Standards*, vol. 49, no. 6, p. 409, Dec. 1952. DOI: 10.6028/jres.049.044. (visited on 01/08/2025).
 - [67] A. Chambolle and T. Pock, „A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging“, *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 120–145, May 2011. DOI: 10.1007/s10851-010-0251-1. (visited on 02/14/2023).
 - [68] M. Genzel, J. Macdonald, and M. März, *Solving Inverse Problems With Deep Neural Networks – Robustness Included?*, Nov. 2020. DOI: 10.48550/arXiv.2011.04268. arXiv: 2011.04268. (visited on 11/17/2021).
 - [69] K. M. Koch, M. Sherafati, V. E. Arpinar, *et al.*, „Analysis and Evaluation of a Deep Learning Reconstruction Approach with Denoising for Orthopedic MRI“, *Radiology: Artificial Intelligence*, vol. 3, no. 6, Nov. 2021. DOI: 10.1148/ryai.2021200278. (visited on 02/13/2025).
 - [70] A. Kofler, D. Si, D. Schote, *et al.*, *MR imaging in the low-field: Leveraging the power of machine learning*, Jan. 2025. DOI: 10.48550/arXiv.2501.17211. arXiv: 2501.17211. (visited on 02/03/2025).
 - [71] O. Ronneberger, P. Fischer, and T. Brox, *U-Net: Convolutional Networks for Biomedical Image Segmentation*, May 2015. DOI: 10.48550/arXiv.1505.04597. arXiv: 1505.04597. (visited on 10/18/2022).
-

- [72] C. Tian, L. Fei, W. Zheng, *et al.*, „Deep learning on image denoising: An overview“, *Neural Networks*, vol. 131, pp. 251–275, Nov. 2020. DOI: 10.1016/j.neunet.2020.07.025. (visited on 02/15/2023).
- [73] K. He, X. Zhang, S. Ren, and J. Sun, „Deep Residual Learning for Image Recognition“, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Jun. 2016, pp. 770–778. DOI: 10.1109/CVPR.2016.90. (visited on 01/08/2025).
- [74] J. Hu, L. Shen, and G. Sun, „Squeeze-and-Excitation Networks“, in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Jun. 2018, pp. 7132–7141. DOI: 10.1109/CVPR.2018.00745. (visited on 01/08/2025).
- [75] A. Vaswani, N. Shazeer, N. Parmar, *et al.*, *Attention Is All You Need*, Aug. 2023. DOI: 10.48550/arXiv.1706.03762. arXiv: 1706.03762. (visited on 07/22/2024).
- [76] J. Long, E. Shelhamer, and T. Darrell, „Fully convolutional networks for semantic segmentation“, in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Jun. 2015, pp. 3431–3440. DOI: 10.1109/CVPR.2015.7298965. (visited on 01/08/2025).
- [77] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, *Loss Functions for Neural Networks for Image Processing*, Apr. 2018. DOI: 10.48550/arXiv.1511.08861. arXiv: 1511.08861. (visited on 04/18/2023).
- [78] U. Michelucci, *Advanced Applied Deep Learning: Convolutional Neural Networks and Object Detection*. Apress, 2019. DOI: 10.1007/978-1-4842-4976-5. (visited on 01/21/2025).
- [79] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, „Learning representations by back-propagating errors“, *Nature*, vol. 323, no. 6088, pp. 533–536, Oct. 1986. DOI: 10.1038/323533a0. (visited on 03/28/2025).
- [80] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>.
- [81] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, Jan. 2017. DOI: 10.48550/arXiv.1412.6980. arXiv: 1412.6980. (visited on 01/21/2025).
- [82] D. Van Essen, K. Ugurbil, E. Auerbach, *et al.*, „The Human Connectome Project: A data acquisition perspective“, *NeuroImage*, vol. 62, no. 4, pp. 2222–2231, Oct. 2012. DOI: 10.1016/j.neuroimage.2012.02.018. (visited on 01/22/2025).

-
- [83] D. S. Marcus, T. H. Wang, J. Parker, *et al.*, „Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI Data in Young, Middle Aged, Nondemented, and Demented Older Adults“, *Journal of Cognitive Neuroscience*, vol. 19, no. 9, pp. 1498–1507, Sep. 2007. DOI: 10.1162/jocn.2007.19.9.1498. (visited on 06/05/2024).
 - [84] J. Lehtinen, J. Munkberg, J. Hasselgren, *et al.*, *Noise2Noise: Learning Image Restoration without Clean Data*, Oct. 2018. DOI: 10.48550/arXiv.1803.04189. arXiv: 1803.04189. (visited on 01/22/2025).
 - [85] Y. Lian, X. Ye, H. Luo, Z. Wu, and H. Guo, „A deep learning network for low-field MRI denoising using group sparsity information and a Noise2Noise method“, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, 2022, p. 4345. DOI: 10.58530/2022/4345. (visited on 01/22/2025).
 - [86] N. Moran, D. Schmidt, Y. Zhong, and P. Coady, „Noisier2Noise: Learning to Denoise From Unpaired Noisy Data“, in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Jun. 2020, pp. 12 061–12 069. DOI: 10.1109/CVPR42600.2020.01208. (visited on 01/22/2025).
 - [87] J. Batson and L. Royer, „Noise2Self: Blind Denoising by Self-Supervision“, in *Proceedings of the 36th International Conference on Machine Learning*, PMLR, May 2019, pp. 524–533. [Online]. Available: <https://proceedings.mlr.press/v97/batson19a.html> (visited on 01/17/2024).
 - [88] S. A. Srinivas, S. F. Cauley, J. P. Stockmann, *et al.*, „External Dynamic InTerference Estimation and Removal (EDITER) for low field MRI“, *Magnetic Resonance in Medicine*, vol. 87, no. 2, pp. 614–628, Feb. 2022. DOI: 10.1002/mrm.28992. (visited on 03/28/2025).
 - [89] L. Yang, W. He, Y. He, *et al.*, „Active EMI Suppression System for a 50 mT Unshielded Portable MRI Scanner“, *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 11, pp. 3415–3426, Nov. 2022. DOI: 10.1109/TBME.2022.3170450. (visited on 07/22/2024).
 - [90] Y. Zhao, L. Xiao, Y. Liu, A. T. Leong, and E. X. Wu, „Electromagnetic interference elimination via active sensing and deep learning prediction for radiofrequency shielding-free MRI“, *NMR in Biomedicine*, vol. 37, no. 7, Jul. 2024. DOI: 10.1002/nbm.4956. (visited on 07/11/2024).
 - [91] K. J. Layton, S. Kroboth, F. Jia, *et al.*, „Pulseseq: A rapid and hardware-independent pulse sequence prototyping framework: Rapid Hardware-Independent Pulse Sequence Prototyping“, *Magnetic Resonance in Medicine*, vol. 77, no. 4, pp. 1544–1552, Apr. 2017. DOI: 10.1002/mrm.26235. (visited on 12/21/2020).
-

- [92] K. Ravi, S. Geethanath, and J. Vaughan, „PyPulseq: A Python Package for MRI Pulse Sequence Design“, *Journal of Open Source Software*, vol. 4, no. 42, p. 1725, Oct. 2019. DOI: 10.21105/joss.01725. (visited on 12/21/2020).
- [93] S. J. Inati, J. D. Naegele, N. R. Zwart, *et al.*, „ISMRM Raw data format: A proposed standard for MRI raw datasets“, *Magnetic Resonance in Medicine*, vol. 77, no. 1, pp. 411–421, 2017. DOI: 10.1002/mrm.26089. (visited on 06/03/2024).
- [94] M. S. Hansen and T. S. Sørensen, „Gadgetron: An open source framework for medical image reconstruction“, *Magnetic Resonance in Medicine*, vol. 69, no. 6, pp. 1768–1776, Jun. 2013. DOI: 10.1002/mrm.24389. (visited on 06/27/2024).
- [95] V. Negnevitsky, Y. Vives-Gilabert, J. M. Algarín, *et al.*, „MaRCoS, an open-source electronic control system for low-field MRI“, *Journal of Magnetic Resonance*, vol. 350, p. 107424, May 2023. DOI: 10.1016/j.jmr.2023.107424. arXiv: 2208.01616. (visited on 11/06/2023).
- [96] D. Schote, B. Silemek, T. O’Reilly, *et al.*, „Nexus: A versatile console for advanced low-field MRI“, *Magnetic Resonance in Medicine*, vol. 93, no. 5, pp. 2224–2238, May 2025. DOI: 10.1002/mrm.30406. (visited on 03/28/2025).
- [97] D. Schote, B. Silemek, F. Seifert, *et al.*, „Beyond Boundaries – A versatile Console for Advanced Low-Field MRI“, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 33, 2024, p. 2838.
- [98] L. Winter, B. Silemek, J. Petzold, *et al.*, „Parallel transmission medical implant safety testbed: Real-time mitigation of RF induced tip heating using time-domain E-field sensors“, *Magnetic Resonance in Medicine*, vol. 84, no. 6, pp. 3468–3484, Dec. 2020. DOI: 10.1002/mrm.28379. (visited on 07/11/2024).
- [99] D. Schote, J. Behrens, L. Winter, C. Kolbitsch, and C. Dinh, „ScanHub: Open-Source Platform for MR Scanner Control, Acquisitions and Postprocessing“, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 31, 2023, p. 2391.
- [100] P. Virtanen, R. Gommers, T. E. Oliphant, *et al.*, „SciPy 1.0: Fundamental algorithms for scientific computing in Python“, *Nature Methods*, vol. 17, no. 3, pp. 261–272, Mar. 2020. DOI: 10.1038/s41592-019-0686-2. (visited on 06/03/2024).
- [101] D. de Gans, *Radio-Frequency Power Amplifier with 1kW peak power*, GitLab: Open Source Imaging Initiative, Mar. 2022. [Online]. Available: <https://gitlab.com/osii/rf-system/rf-power-amplifier/1kw-peak-rfpa> (visited on 07/11/2024).
- [102] D. de Gans, *TU Delft Gradient Amplifier*, GitLab: Open Source Imaging Initiative, Mar. 2022. [Online]. Available: <https://gitlab.com/osii/gradient-system/gradient-power-amplifier/tu-delft-ga2> (visited on 07/11/2024).

-
- [103] J. Hennig, A. Nauerth, and H. Friedburg, „RARE imaging: A fast imaging method for clinical MR“, *Magnetic Resonance in Medicine*, vol. 3, no. 6, pp. 823–833, Dec. 1986. DOI: 10.1002/mrm.1910030602. (visited on 03/26/2025).
- [104] D. C. Alsop, „Phase insensitive preparation of single-shot RARE: Application to diffusion imaging in humans“, *Magnetic Resonance in Medicine*, vol. 38, no. 4, pp. 527–533, Oct. 1997. DOI: 10.1002/mrm.1910380404. (visited on 02/13/2025).
- [105] NEMA, *Determination of Signal-to-Noise Ratio and Image Uniformity for Single-Channel Non-Volume Coils in Diagnostic MR Imaging*, Standards Publication, 2008.
- [106] B. Madore, A. T. Hess, A. M. J. Van Nierkerk, *et al.*, „External Hardware and Sensors, for Improved MRI“, *Journal of Magnetic Resonance Imaging*, vol. 57, no. 3, pp. 690–705, Mar. 2023. DOI: 10.1002/jmri.28472. (visited on 03/27/2025).
- [107] Y. Duerst, B. J. Wilm, B. E. Dietrich, *et al.*, „Real-time feedback for spatiotemporal field stabilization in MR systems“, *Magnetic Resonance in Medicine*, vol. 73, no. 2, pp. 884–893, Feb. 2015. DOI: 10.1002/mrm.25167. (visited on 07/11/2024).
- [108] E. Boskamp, M. Twieg, and R. O’Halloran, „Dipole position sensors integrated in receiver array for motion detection and correction“, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 29, 2021, p. 4038. (visited on 07/22/2024).
- [109] T. Warbrick, „Simultaneous EEG-fMRI: What Have We Learned and What Does the Future Hold?“, *Sensors*, vol. 22, no. 6, p. 2262, Mar. 2022. DOI: 10.3390/s22062262. (visited on 07/11/2024).
- [110] D. Stäb, J. Roessler, K. O’Brien, C. Hamilton-Craig, and M. Barth, „ECG Triggering in Ultra-High Field Cardiovascular MRI“, *Tomography*, vol. 2, no. 3, pp. 167–174, Sep. 2016. DOI: 10.18383/j.tom.2016.00193. (visited on 07/22/2024).
- [111] D. O. Brunner, B. E. Dietrich, M. Çavuşoğlu, *et al.*, „Concurrent recording of RF pulses and gradient fields – comprehensive field monitoring for MRI“, *NMR in Biomedicine*, vol. 29, no. 9, pp. 1162–1172, Sep. 2016. DOI: 10.1002/nbm.3359. (visited on 07/11/2024).
- [112] D. Waddington, E. Shimron, N. Hindley, N. Koonjoo, and M. Rosen, „Accelerating Ultra-Low Field MRI with Compressed Sensing“, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 31, 2022, p. 1818. DOI: 10.58530/2022/1818. (visited on 07/22/2024).
- [113] M. Prier, D. Schote, and I. Vogt, *OCRA1 – SPI controlled 4 channel 18bit DAC and RF Attenuator – OCRA*, Blog, Nov. 2020. [Online]. Available: <https://zeugmatographix.org/ocra/2020/11/27/ocra1-spi-controlled-4-channel-18bit-dac-and-rf-attenuator/> (visited on 07/22/2024).
-

- [114] F. Galve, J. Alonso, J. M. Algarín, and J. M. Benlloch, *Model-driven reconstruction with phase-constrained highly-oversampled MRI*, Dec. 2020. DOI: 10.48550/arXiv.2007.15674. arXiv: 2007.15674. (visited on 11/26/2024).
- [115] K. Wenzel, H. Alhamwey, T. O'Reilly, *et al.*, „B0-Shimming Methodology for Affordable and Compact Low-Field Magnetic Resonance Imaging Magnets“, *Frontiers in Physics*, vol. 9, p. 704566, Jul. 2021. DOI: 10.3389/fphy.2021.704566. (visited on 11/09/2021).
- [116] K. Koolstra, T. O'Reilly, P. Börnert, and A. Webb, „Image distortion correction for MRI in low field permanent magnet systems with strong B0 inhomogeneity and gradient field nonlinearities“, *Magnetic Resonance Materials in Physics, Biology and Medicine*, Jan. 2021. DOI: 10.1007/s10334-021-00907-2. (visited on 06/18/2021).
- [117] J. Fessler and B. Sutton, „Nonuniform fast fourier transforms using min-max interpolation“, *IEEE Transactions on Signal Processing*, vol. 51, no. 2, pp. 560–574, Feb. 2003. DOI: 10.1109/TSP.2002.807005. (visited on 12/21/2020).
- [118] D. Noll, C. Meyer, J. Pauly, D. Nishimura, and A. Macovski, „A homogeneity correction method for magnetic resonance imaging with time-varying gradients“, *IEEE Transactions on Medical Imaging*, vol. 10, no. 4, pp. 629–637, Dec. 1991. DOI: 10.1109/42.108599. (visited on 11/15/2022).
- [119] L.-C. Man, J. M. Pauly, and A. Macovski, „Multifrequency interpolation for fast off-resonance correction“, *Magnetic Resonance in Medicine*, vol. 37, no. 5, pp. 785–792, May 1997. DOI: 10.1002/mrm.1910370523. (visited on 06/23/2021).
- [120] M. Usman, L. Kakkar, A. Matakos, *et al.*, „Joint B0 and image estimation integrated with model based reconstruction for field map update and distortion correction in prostate diffusion MRI“, *Magnetic Resonance Imaging*, vol. 65, pp. 90–99, Jan. 2020. DOI: 10.1016/j.mri.2019.09.008. (visited on 08/20/2021).
- [121] D. Schote, L. Winter, C. Kolbitsch, *et al.*, „Joint B0 and Image Reconstruction in Low-Field MRI by Physics-Informed Deep-Learning“, *IEEE Transactions on Biomedical Engineering*, pp. 1–12, 2024. DOI: 10.1109/TBME.2024.3396223. (visited on 06/03/2024).
- [122] M. P. Heinrich, M. Stille, and T. M. Buzug, „Residual U-Net Convolutional Neural Network Architecture for Low-Dose CT Denoising“, *Current Directions in Biomedical Engineering*, vol. 4, no. 1, pp. 297–300, Sep. 2018. DOI: 10.1515/cdbme-2018-0072. (visited on 10/18/2022).
- [123] S. Nasrin, M. Z. Alom, R. Burada, T. M. Taha, and V. K. Asari, „Medical Image Denoising with Recurrent Residual U-Net (R2U-Net) base Auto-Encoder“, in *2019 IEEE National Aerospace and Electronics Conference (NAECON)*, IEEE, Jul. 2019, pp. 345–350. DOI: 10.1109/NAECON46414.2019.9057834. (visited on 02/03/2023).

-
- [124] K. T. Block, M. Uecker, and J. Frahm, „Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint“, *Magnetic Resonance in Medicine*, vol. 57, no. 6, pp. 1086–1098, Jun. 2007. DOI: 10.1002/mrm.21236. (visited on 05/23/2023).
 - [125] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, „Image Quality Assessment: From Error Visibility to Structural Similarity“, *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, Apr. 2004. DOI: 10.1109/TIP.2003.819861. (visited on 12/06/2023).
 - [126] B. Sutton, D. Noll, and J. Fessler, „Fast, iterative, field-corrected image reconstruction for MRI“, in *Proceedings IEEE International Symposium on Biomedical Imaging*, IEEE, 2002, pp. 489–492. DOI: 10.1109/ISBI.2002.1029301. (visited on 11/10/2022).
 - [127] J. Fessler, Sangwoo Lee, V. Olafsson, H. Shi, and D. Noll, „Toeplitz-based iterative image reconstruction for MRI with correction for magnetic field inhomogeneity“, *IEEE Transactions on Signal Processing*, vol. 53, no. 9, pp. 3393–3402, Sep. 2005. DOI: 10.1109/TSP.2005.853152. (visited on 08/19/2022).
 - [128] B. P. Sutton, D. C. Noll, and J. A. Fessler, „Simultaneous Estimation of Image and Inhomogeneity Field Map“, in *ISMRM Workshop on Minimum MR Data Acquisition Methods: Making More with Less*, 2001, pp. 15–18.
 - [129] F. F. Zimmermann, C. Kolbitsch, P. Schuenke, and A. Kofler, „PINQI: An End-to-End Physics-Informed Approach to Learned Quantitative MRI Reconstruction“, *IEEE Transactions on Computational Imaging*, vol. 10, pp. 628–639, 2024. DOI: 10.1109/TCI.2024.3388869. arXiv: 2306.11023. (visited on 01/23/2025).
 - [130] A. Javaloy and I. Valera, „RotoGrad: Gradient Homogenization in Multitask Learning“, in *International Conference on Learning Representations*, Oct. 2021. [Online]. Available: <https://openreview.net/forum?id=T8wHz4rnuGL> (visited on 07/12/2023).
 - [131] A. F. Agarap, *Deep Learning using Rectified Linear Units (ReLU)*, Feb. 2019. DOI: 10.48550/arXiv.1803.08375. arXiv: 1803.08375. (visited on 03/20/2025).
 - [132] A. L. Maas, A. Y. Hannun, and A. Y. Ng, „Rectifier Nonlinearities Improve Neural Network Acoustic Models“, in *Proceedings of the 30th International Conference on Machine Learning*, vol. 28, Journal of Machine Learning Research, 2013.
 - [133] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin, „Adaptive reconstruction of phased array MR imagery“, *Magnetic Resonance in Medicine*, vol. 43, no. 5, pp. 682–690, 2000. DOI: 10.1002/(SICI)1522-2594(200005)43:5<682::AID-MRM10>3.0.CO;2-G. (visited on 01/23/2025).
-

- [134] A. Kofler, F. Altekruiger, F. Antaro Ba, *et al.*, „Learning Regularization Parameter-Maps for Variational Image Reconstruction Using Deep Neural Networks and Algorithm Unrolling“, *SIAM Journal on Imaging Sciences*, vol. 16, no. 4, pp. 2202–2246, Dec. 2023. DOI: 10.1137/23M1552486. (visited on 10/31/2024).
- [135] A. Chambolle and P.-L. Lions, „Image recovery via total variation minimization and related problems“, *Numerische Mathematik*, vol. 76, no. 2, pp. 167–188, Apr. 1997. DOI: 10.1007/s002110050258. (visited on 03/27/2025).
- [136] I. Loshchilov and F. Hutter, *Decoupled Weight Decay Regularization*, Apr. 2019. DOI: 10.48550/arXiv.1711.05101. arXiv: 1711.05101. (visited on 02/25/2025).
- [137] F. Crete, T. Dolmiere, P. Ladret, and M. Nicolas, „The blur effect: Perception and estimation with a new no-reference perceptual blur metric“, in *Electronic Imaging 2007*, B. E. Rogowitz, T. N. Pappas, and S. J. Daly, Eds., Feb. 2007, p. 64920I. DOI: 10.1117/12.702790. (visited on 10/22/2024).
- [138] B. Yaman, S. A. H. Hosseini, and M. Akcakaya, *Zero-Shot Self-Supervised Learning for MRI Reconstruction*, 2022. DOI: 10.48550/arXiv.2102.07737. arXiv: 2102.07737.
- [139] C. Millard and M. Chiew, „Clean Self-Supervised MRI Reconstruction from Noisy, Sub-Sampled Training Data with Robust SSDU“, *Bioengineering*, vol. 11, no. 12, p. 1305, Dec. 2024. DOI: 10.3390/bioengineering11121305. (visited on 01/22/2025).
- [140] L. Pfaff, J. Hossbach, E. Preuhs, *et al.*, „Self-supervised MRI denoising: Leveraging Stein’s unbiased risk estimator and spatially resolved noise maps“, *Scientific Reports*, vol. 13, no. 1, p. 22 629, Dec. 2023. DOI: 10.1038/s41598-023-49023-2. (visited on 07/22/2024).
- [141] Z.-Y. Hou and Q.-N. Jin, „Tikhonov regularization for nonlinear ill-posed problems“, *Nonlinear Analysis: Theory, Methods & Applications*, vol. 28, no. 11, pp. 1799–1809, Jun. 1997. DOI: 10.1016/S0362-546X(95)00235-N. (visited on 03/28/2025).
- [142] E. L. Aleixo, J. G. Colonna, M. Cristo, and E. Fernandes, *Catastrophic Forgetting in Deep Learning: A Comprehensive Taxonomy*, Dec. 2023. DOI: 10.48550/arXiv.2312.10549. arXiv: 2312.10549. (visited on 02/11/2025).
- [143] K. Zhang, Y. Li, W. Zuo, *et al.*, *Plug-and-Play Image Restoration with Deep Denoiser Prior*, Jul. 2021. DOI: 10.48550/arXiv.2008.13751. arXiv: 2008.13751. (visited on 01/23/2025).
- [144] Y. Liu, A. T. L. Leong, Y. Zhao, *et al.*, „A low-cost and shielding-free ultra-low-field brain MRI scanner“, *Nature Communications*, vol. 12, no. 1, p. 7238, Dec. 2021. DOI: 10.1038/s41467-021-27317-1. (visited on 12/17/2021).

- [145] Y. Zhao, L. Xiao, J. Hu, and E. X. Wu, „Robust EMI elimination for RF shielding-free MRI through deep learning direct MR signal prediction“, *Magnetic Resonance in Medicine*, vol. 92, no. 1, pp. 112–127, Jul. 2024. DOI: 10.1002/mrm.30046. (visited on 01/24/2025).

Publications

1. **D. Schote**, L. Winter, C. Kolbitsch, G. Rose, O. Speck, and A. Kofler, “Joint B0 and Image Reconstruction in Low-Field MRI by Physics-Informed Deep-Learning”, *IEEE Transactions on Biomedical Engineering*, pp. 1-12, 2024.
DOI: 10.1109/TBME.2024.3396223
2. **D. Schote**, B. Silemek, T. O’Reilly, F. Seifert J.-L. Assmy, C. Kolbitsch, A. Webb, and L. Winter, “Nexus: A versatile console for advanced low-field MRI,” *Magnetic Resonance in Medicine*, vol. 93, no. 5, pp. 2224–2238, May 2025.
DOI: 10.1002/mrm.30406
3. **D. Schote**, L. Winter, C. Kolbitsch, F. Zimmermann, T. O’Reilly, A. Webb, and A. Kofler, “Physics-Informed Deep Learning for Image Distortion Correction from B₀-inhomogeneities in Low-Field MRI”, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 31, abstract no. 1816, London, 2022.
DOI: 10.58530/2022/1816
4. **D. Schote**, L. Winter, C. Kolbitsch, and A. Kofler, “Reliable Low-Field B0-Maps by Deep Learning with Physical Constraints”, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 32, abstract no. 4344, Toronto, 2023.
5. **D. Schote**, J. Behrens, L. Winter, C. Kolbitsch, and C. Dinh, “ScanHub: Open-Source Platform for MR Scanner Control, Acquisitions and Postprocessing”, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 32, abstract no. 2391, Toronto, 2023.
6. **D. Schote**, B. Silemek, F. Seifert, C. Kolbitsch, T. O’Reilly, A. Kofler, A. Webb, and L. Winter, “Beyond Boundaries - A versatile Consolee for Advanced Low-Field MRI”, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 33, abstract no. 2838, Singapore, 2024.
7. **D. Schote**, C. Kolbitsch, L. Winter, and A. Kofler, “Deep Learning Image Denoising for In-Vivo Low-Field MRI Using Test-Time Training”, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 34, abstract no. 3838, Honolulu, 2025 (accepted).

8. A. Kofler, F. Altekürer, F. Antarou Ba, C. Kolbitsch, E. Papoutsellis, **D. Schote**, C. Sirotenko, F. Zimmermann, and K. Papafitsoros, “Learning Regularization Parameter-Maps for Variational Image Reconstruction Using Deep Neural Networks and Algorithm Unrolling”, *SIAM Journal on Imaging Sciences*, vol. 16, no. 4, pp. 2202-2246, 2023.
DOI: 10.1137/23M1552486
9. H. Rabus, N. Agarwal, A. Boubekki, Y. Bozkurt Varolgünes, B. Clark, C. Elster, J. Faller, N. Hegemann, S. Haufe, S. Heidenreich, R.-P. Kapsch, A. Kofler, C. Kolbitsch, A. Märten, J. Martin, G. O’Connor, D. Panknin, O. Pfeffer, S. Pojtinger, A.-B. Schönfeld, **D. Schote**, P. Schünke, M. Schwarze, P. Tappayuthpijarn, L. Thomas, L. Winter, R. Zhumagambetov and F. Zimmermann, “Metrology for artificial intelligence in medicine”, *Measurement: Sensors*, p. 101789, 2025.
DOI: 10.1016/j.measen.2024.101789
10. A. Kofler, F. Altekürer, F. Antarou Ba, C. Kolbitsch, E. Papoutsellis, **D. Schote**, C. Sirotenko, F. Zimmermann, and K. Papafitsoros, “Unrolled three-operator splitting for parameter-map learning in Low Dose X-ray CT reconstruction,” 2023.
DOI: 10.48550/arXiv.2304.08350, arxiv: 2304.08350
11. A. Kofler, C. Sirotenko, F. Zimmermann, **D. Schote**, C. Kolbitsch, F. Antarou Ba, F. Altekürer, E. Papoutsellis, and K. Papafitsoros, “CNN-based Estimation of Spatio-Temporal Regularization Parameter-Maps for TV-Reconstruction in Dynamic Cardiac MRI,” in *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 32, abstract no. 3076, Toronto, 2023.
DOI: 10.58530/2023/3076
12. A. Kofler, D. Si, **D. Schote**, R. M. Botnar, C. Kolbitsch, and C. Prieto, “MR imaging in the low-field: Leveraging the power of machine learning”, 2025.
DOI: 10.48550/arXiv.2501.17211, arxiv: 2501.17211

A Appendix

A.1 Sequence calculation and execution with the Nexus console




Figure A.1: Diagram which illustrates the sequence calculation by the implemented sequence provider within the Nexus console. The gradient and RF waveforms of a *Pulseseq* sequence are calculated dependent on adjustable acquisition parameters and a static device configuration. Each sample point of the gradient waveform encodes a digital signal, replayed synchronously with the analog waveform. The computed sequence is stored as a 16-bit *Numpy* array in *Fortran* order.




Figure A.2: Schematic of the sequence execution in the Nexus console after computation, utilizing the implemented transmit and receive device classes. The sequence is streamed to the AWG card's memory and subsequently replayed. Whenever new timestamp data becomes available, the digitizer card transfers it from the card to the buffer. Digital post-processing includes demodulation, decimation, and phase correction.

Ehrenerklärung

Ich versichere hiermit, dass ich die vorliegende Arbeit vollständig ohne unzulässige Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe.

Alle verwendeten fremden und eigenen Quellen sind als solche kenntlich gemacht und im Falle einer Ko-Autorenschaft, insbesondere im Rahmen einer kumulativen Dissertation, ist der Eigenanteil richtig und vollständig ausgewiesen. Insbesondere habe ich nicht die Hilfe einer kommerziellen Promotionsberaterin/eines kommerziellen Promotionsberaters in Anspruch genommen. Dritte haben von mir weder unmittelbar noch mittelbar geldwerte Leistungen für Arbeiten erhalten, die im Zusammenhang mit dem Inhalt der vorgelegten Dissertation stehen.

Ich habe insbesondere nicht wissentlich:

- Ergebnisse erfunden oder widersprüchliche Ergebnisse verschwiegen,
- statistische Verfahren absichtlich missbraucht, um Daten in ungerechtfertigter Weise zu interpretieren,
- fremde Ergebnisse oder Veröffentlichungen plagiiert,
- fremde Forschungsergebnisse verzerrt wiedergegeben.

Mir ist bekannt, dass Verstöße gegen das Urheberrecht Unterlassungs- und Schadensersatzansprüche der Urheberin/des Urhebers sowie eine strafrechtliche Ahndung durch die Strafverfolgungsbehörden begründen können.

Ich erkläre mich damit einverstanden, dass die Dissertation ggf. mit Mitteln der elektronischen Datenverarbeitung auf Plagiate überprüft werden kann.

Die Arbeit wurde bisher weder im Inland noch im Ausland in gleicher oder ähnlicher Form als Dissertation eingereicht und ist als Ganzes auch noch nicht veröffentlicht.

Ort, Datum

Unterschrift