NNT : 2024AIXM0402

THÈSE DE DOCTORAT

Soutenance prévue à Aix-Marseille Université
le 22 Novembre 2024 par

Swetali NIMJE

An Evaluation of Scan-Specific Deep-Learning Strategies for Brain MRI Reconstruction

Discipline
Informatique

École doctorale
ED 184 Mathématiques et Informatique

Laboratoire/Partenaires de recherche
Laboratoire d'Informatique et Systèmes
Centre de Résonance Magnétique Biologique et Médicale

Composition du jury

Marie POIRIER-QUINOT Professeure, Université Paris-Saclay	Rapporteure
Saïd MOUSSAOUI Professeur, Centrale Nantes	Rapporteur
Vincent GUIGUE Professeur, Agroparistech	Examinateur
Alexandre VIGNAUD Dir. de recherche, HDR, CEA	Président du jury
Thierry ARTIERES Professeur, Centrale Méditerranée	Directeur de thèse
Ludovic DE ROCHEFORT Chargé de recherche, HDR, Aix-Marseille Université	Co-directeur de thèse

Affidavit

I, undersigned, Swetali Nimje, hereby declare that the work presented in this manuscript is my own work, carried out under the scientific supervision of Thierry Artières and Ludovic de Rochefort, in accordance with the principles of honesty, integrity and responsibility inherent to the research mission. The research work and the writing of this manuscript have been carried out in compliance with both the french national charter for Research Integrity and the Aix-Marseille University charter on the fight against plagiarism.

This work has not been submitted previously either in this country or in another country in the same or in a similar version to any other examination body.

Marseille, 2024 July 9



This work is licensed under Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International Public License

Scientific production

List of publications and/or patents produced as part of the project:

1. **Nimje S**, Artières T, Guye M, de Rochefort L, *Improved Brain Image Quality in Accelerated Parallel MR Imaging with A Deep Complex Residual Convolutional Neural Network: 2D/3D and Multi-Contrast Proof-of-Concept*, (Under review).
2. Patent application (21/09/2023): FR2309986: *Procédé de mesure de la qualité d'une image reconstruite* SATT-SE, inventors: **Nimje S**, de Rochefort L., Artières T. Owners: AMU/CNRS/Ecole Centrale Méditerranée

Conferences and workshops during the PhD:

1. **Nimje S**, Artieres T, de Rochefort, *COrelation-Based Artifact Index (COBRAI), an image quality metric to quantify structured artifacts for MR parallel imaging reconstruction evaluation*, Colloque Français d'Intelligence Artificielle en Imagerie Biomédicale - IABM, March 2024, Grenoble, France.
2. **Nimje S**, Artieres T, de Rochefort, *Accelerating Phase and Quantitative susceptibility mapping with Scan-Specific Complex Convolutional Neural Networks*, June 2023, Proc. International Society for Magnetic Resonance in Medicine, Toronto, Canada. https://hal.science/hal-04269563v1
3. **Nimje S**, Beaumont, J, Artieres T, de Rochefort L, *Accelerating Phase Mapping with Scan-Specific Complex Convolutional Neural Networks*, QMR Lucca, Italy 2022 Joint Workshop on MR phase, magnetic susceptibility and electrical properties mapping.
4. **Nimje S**, Artieres T, de Rochefort L, *SANGRIA: one-Shot learning super-ResolutIon with Adversarial training for accelerated Magnetic Resonance Imaging*, CAP - RFIAP : French conference on machine learning July 2022 (Conference sur l'Apprentissage automatique), Vannes, France. https://hal.science/hal-04050162v1
5. Beaumont J, Troalen T, **Nimje S**, Rapacchi S, de Rochefort L, *Energy dependent z-scores improve parallel imaging motion correction using trimmed autocalibrating k-space estimation (TAKE)*, June 2022, Proc. International Society for Magnetic Resonance in Medicine, London, UK. https://hal.science/hal-03859860v1
6. **Nimje S**, Artieres T, de Rochefort L, *Reconstruction, Réseau de neurones convolutifs complexes et scan-spécifiques pour accélérer l'IRM : faisabilité sur le cerveau pour l'imagerie 2D pondérée T2*, RITS May'22 (Recherche en Imagerie et Technologies pour la Sante), Brest, France.
7. Beaumont J, Troalen T, **Nimje S**, Rapacchi S, de Rochefort L, *Correction robuste d'artefacts de mouvement en imagerie par résonance magnétique*, RITS May'22 (Recherche en Imagerie et Technologies pour la Sante), Brest, France.

8. **Nimje S**, Beaumont J, Artieres T, de Rochefort L, *Accelerating Parallel MRI Acquisition and Reconstruction using Machine Learning*, Colloque Numérique et IA pour la Santé à Aix-Marseille Université, 2021, Marseille, France.

Co-supervision of internships on related work

1. Ayayi Ayilo JE, *Accelerating MRI acquisition with deep learning and reinforcement learning*, M2 IAAA 2021-2022, Aix-Marseille Université, France.
2. Faraji M, *Fitting Multi-Parametric MRI Images Using Deep Learning*, M2 IAAA 2021-2022, Aix-Marseille Université, France.

Open source software

1. **Nimje S**, Artieres T, de Rochefort L, Scan-specific MRI reconstruction / COBRAI repository (Under finalization).

Résumé

L'IRM est une technologie polyvalente fournissant une multitude de contrastes, mais elle est réputée être lente. Réduire la durée d'un examen IRM tout en conservant une qualité diagnostic est un enjeu important. Récemment, les approches basées sur de l'apprentissage automatique ont été proposées. Parmi ces approches, certaines focalisent sur la reconstruction de données sous-échantillonnées dans la ligne de l'imagerie parallèle avec des apprentissages spécifiques à la séquence d'imagerie. La contribution de ce travail est d'évaluer différentes implémentations, tant sur les options d'acquisition expérimentales que sur l'apprentissage des modèles de reconstruction, tout en proposant une métrique originale pour quantifier le niveau d'artefact résiduel.
Après une introduction de la physique de l'IRM et de l'apprentissage automatique associé, le premier point abordé concerne les métriques de qualité d'image pour lequel la métrique COrreletion-Based Residual Artifact Index (COBRAI) est présentée. C'est une nouvelle mesure basée sur la corrélation qui est sensible aux résidus structurés dans les images reconstruites. Le principe de conception de cette mesure est présenté, ainsi que plusieurs exemples pour illustrer son intérêt.
Par la suite, les approches d'apprentissage spécifique à la séquence d'imagerie employée sont évaluées. La méthode prend pour famille d'architecture de modèles l'approche RAKI (robust artificial-neural-networks for k-space interpolation). C'est au départ un réseau neuronal convolutif qui est entraîné pour la tâche de reconstruction en imagerie parallèle qui effectue l'interpolation de données manquantes dans l'espace-k pour des sous-échantillonnages structurés. Initialement proposé avec des réseau contenant des couches de convolution à valeurs réelles pour l'imagerie 2D et avec des données d'apprentissage en mode dit 'intégré', i.e. qui ont le même contraste que les données à reconstruire, il est considéré dans ce travail l'extension des modèles à valeurs complexes et avec une connexion résiduelle, ainsi que l'évaluation et l'optimisation des hyperparamètres. L'influence des non-linéarités est étudiée, et ce, en particulier en lien avec le mode d'acquisition des données d'apprentissage, mode intégré versus mode séparé, qui est propice à produire des artefacts résiduels du type hallucination. La faisabilité de ces approches est démontrée pour plusieurs contrastes standards en 2D, pour finir avec une extension en 3D. Enfin, une extension de la méthode scan-spécifique, SANGRIA, utilisant une stratégie d'apprentissage adversarial en tant que régularisation basée sur les données a été explorée. Un discriminateur par patch est mis en oeuvre dans l'espace image, ce qui signifie que des patches sont extraits d'une seule image entièrement échantillonnée utilisé pour discriminé des patchs réalistes. Cette approche améliore la qualité visuelle et les performances des images reconstruites avec des taux de sous-échantillonnage plus élevés.
Ces travaux interpellent notamment sur l'utilisation de modèles neuronaux utilisés pour la reconstruction d'image IRM entraînés avec peu de données d'apprentissage du type spécifique à la séquence. La caractérisation complémentaire permise par COBRAI laisse envisager des stratégies d'acquisition plus optimales et des modèles de reconstruction optimisés.

Mots clés: Reconstruction IRM, imagerie parallèle, réseau neuronal convolutif, scan-spécifique, métriques de qualité d'image

Abstract

MRI is a versatile medical imaging technology that provides a multitude of contrasts; however, it is notoriously slow. Reducing the duration of an MRI examination while maintaining the diagnostic quality is a major challenge. Recently, approaches based on machine learning have been proposed. Among these approaches, some focus on the reconstruction of undersampled data in parallel imaging with learning specific to the imaging sequence employed (said scan-specific). The contribution of this study is to evaluate different strategies for the practical implementation of these approaches, both in terms of experimental acquisition options and the training of reconstruction models, while proposing an original metric for quantifying the level of artifacts in reconstructions. After a state-of-the-art introduction to the physics of Magnetic Resonance Imaging (MRI) and associated machine learning, the first point addressed in this work concerns image quality metrics, for which the COrreletion-Based Residual Artifact Index (COBRAI) metric is proposed. This is a new correlation-based metric that is sensitive to structured residuals in the reconstructed images. The design principle of this metric is presented, along with several examples to illustrate its interest.
Subsequently, deep-learning scan-specific approaches were evaluated. This method takes the RAKI (robust artificial neural networks for the k-space interpolation (RAKI) approach as its model architecture family. Initially, a convolutional neural network was trained for the parallel imaging reconstruction task by performing k-space interpolation of missing data for structured undersampling. Initially proposed with networks containing real-valued convolution layers for two-dimensional imaging and with training data in so-called integrated mode, i.e. having the same contrast as the data to be reconstructed, this work considers the extension of models to complex values and with a residual connection, as well as the evaluation and optimization of hyperparameters. The influence of nonlinearities is investigated, particularly in relation to the mode of acquisition of training data, integrated versus separated mode, which is prone to producing residual artifacts assimilated to hallucinations. The feasibility of these approaches is demonstrated for several standard 2D contrasts, ending with an extension to 3D contrasts. Finally, an extension of the scan-specific reconstruction method, SANGRIA, using an adversarial training strategy as data-driven regularization, was explored. A patch-based discriminator is implemented in the image space, meaning that patches are extracted from a single fully sampled image and used to discriminate realistic patches. This approach improves the quality of the reconstructed images both visually and in terms of performance in achieving higher undersampling rates.
In particular, this work calls into question the use of neural networks for MRI image reconstruction trained with a low amount of scan-specific training data. The characterization made possible by COBRAI suggests a complementary characterization to state-of-the-art metrics, enabling more optimal MRI data acquisition strategies and optimized reconstruction models to be envisaged.

Keywords: MRI reconstruction, Parallel imaging, convolutional neural network, scan-specific, image quality metrics

Acknowledgements

I would like to express my gratitude to my doctoral supervisors, Prof. Artieres Thierry and Dr. de Rochefort Ludovic of Aix Marseille University, for their support, invaluable guidance, and profound insights throughout my doctoral journey. Their unwavering commitment to nurturing critical thinking skills has been instrumental in shaping my path as an independent researcher. I am deeply appreciative of their encouragement, patience, and dedication to promoting my intellectual growth.

I am grateful for the contributions of the CRMBM and QARMA LIS teams, whose expertise and collaborative spirit have strengthened the foundations of my research. Their efforts in experimentation, data analysis, and interpretation are essential. I extend my heartfelt gratitude to each member of these esteemed labs for their unwavering support and collaboration. Furthermore, I thank the CRMBM and QARMA LIS lab assistants for their tireless efforts and expertise, which have been essential to the success of my research.

I express my gratitude to the Mathematics Doctoral School at Aix Marseille University for offering a supportive academic environment and encouraging interdisciplinary collaboration. The extensive range of resources and possibilities provided by doctoral schools have significantly influenced my academic and professional development. We are grateful for the financial support provided by ANR-AAP doctoral contract in Artficial intelligence and Institut Marseille Imaging. Their invaluable contributions have significantly aided our research. Finally, I would like to extend my heartfelt regards to the project leader, Lavoute Cécile, for managing the administrative tasks, coordinating the conferences, and extending an invitation for me to present my work at several scientific events. Also, Special thanks to A. Vignaud, P.Ciuciu and GR Chaithya, Neurospin, CEA/CNRS/Université Paris-Saclay for thier valuable comments.

In addition to the academic support I have received, I would like to extend my heartfelt gratitude to my fiancé Leo for his unwavering love, patience, and understanding throughout this demanding journey. His encouragement and belief in my abilities have provided me with the strength and motivation to persevere, even during the most challenging times. His unwavering support has been a constant source of inspiration, and I am deeply grateful for his presence in my life. I would also like to acknowledge the friendships I have formed during my doctoral journey. The camaraderie and support of my fellow colleagues and friends have been invaluable. Their companionship, shared experiences, and words of encouragement have made the journey more enriching and enjoyable. In particular, I would like to mention Hamed, Charly, Hanwei, Malek, Mimoun, Khushboo, Rohit, Manvi, Chandrashekar, Julia, Felipe, Mounir, Maya, Tangi, Andrea, Pierre, Jeremy, whose friendship and support have been a source of strength and inspiration. To my fiancé and friends, thank you for being the pillars of support in my life. Your love, encouragement, and friendship have made all the difference, and I am immensely grateful for your presence in my life.

I am deeply indebted to all those who supported and inspired me along this transformative path. Their belief in my abilities and unwavering encouragement have been a constant source of strength and motivation, for which I am truly grateful.

“Tell me and I forget,
Teach me and I may remember,
Involve me and I learn.”

- Benjamin Franklin

“Paṛh paṛh ‘ilm hazaar kitaabaan
Kadi apne aap noon paṛhiya naeen
Ja ja waṛda een mandar maseeti
Kadi man apne wich waṛiya naeen
Aeween laṛna en shaitaan de naal bandiya
Kadi nafs apne naal laṛya naeen
Aakhe peer bullhe shah aasmaani phaṛda en
Jihṛa man wich wasda ohnoon phaṛiya naeen,”

- Bulleh Shah

Contents

Affidavit 2

Scientific production 3

Résumé 5

Abstract 6

Acknowledgements 7

Contents 9

Introduction 10

1 Magnetic Resonance Imaging basis 13

1.1 A brief historical perspective . . . 13
1.2 From magnetization to the Free Induction Decay . . . 15
1.2.1 Nuclear magnetization . . . 15
1.2.2 Resonance and flip angle . . . 16
1.2.3 Precession and relaxation . . . 17
1.3 Localization . . . 19
1.3.1 Gradients and k-space . . . 19
1.3.2 Gradient-echo pulse sequence . . . 19
1.3.2.1 Slice selection . . . 20
1.3.2.2 Frequency and phase encoding . . . 20
1.3.3 Spin-echo pulse sequence . . . 21
1.3.4 Coil array and parallel imaging . . . 21
1.3.4.1 Image combination in coil arrays . . . 22
1.3.4.2 Reducing scan time with parallel imaging . . . 23
1.4 Overview of MRI contrasts and quantitative imaging . . . 24
1.4.1 Basic MRI contrasts . . . 24
1.4.1.1 Proton density contrast . . . 24
1.4.1.2 Susceptibility contrast . . . 25
1.4.1.3 T1-weighted . . . 25
1.4.1.4 T2-weighted . . . 25
1.4.2 Quantitative imaging . . . 26
1.4.2.1 Relaxometry . . . 26
1.4.2.2 Magnetic susceptibility quantification . . . 27
1.5 Challenges . . . 27

2 Machine learning in MRI **29**
2.1 Introduction . . . 29
2.1.1 Primer to Deep Learning . . . 29
2.1.1.1 Basics: one neuron at a time . . . 30
2.1.1.2 Activation functions . . . 33
2.1.1.3 Convolutional Neural Networks . . . 33
2.1.1.4 Generative Models . . . 36
2.2 MRI reconstruction and Machine Learning . . . 40
2.2.1 Image-based super-resolution . . . 42
2.2.2 K-space-based super-resolution . . . 43
2.2.2.1 MR Reconstruction using fixed sampling . . . 45
2.2.2.1.1 Scan-Specific MR reconstruction . . . 45
2.2.2.1.2 Database-driven MR reconstruction . . . 52
2.2.2.1.3 Hybrid MR reconstruction . . . 54
2.2.2.2 Optimizing sampling patterns . . . 54
2.2.2.3 ML for qMRI . . . 57
2.2.2.3.1 Traditional methods . . . 57
2.2.2.3.2 DL-based MR parametric mapping . . . 58
2.3 Conclusion . . . 59

3 COBRAI: A Contribution to Image-Quality Assessment **61**
3.1 Image Quality Assessment in MRI . . . 61
3.2 Standard metrics used in this work . . . 62
3.2.1 MSE and its variant . . . 63
3.2.2 SSIM . . . 64
3.2.3 Blur Metric . . . 65
3.2.4 limitations . . . 65
3.3 COBRAI . . . 66
3.3.1 Metrics evaluation for various degradations . . . 67
3.3.1.1 Shepp-Logan phantom with various degradations . . . 67
3.3.1.2 Subtle natural image degradations . . . 68
3.3.1.3 COBRAI and SSIM in the context of MRI Reconstruction . . . 68
3.4 Discussion and Conclusion . . . 69

4 Focus on scan-specific convolutional models for MRI reconstruction **75**
4.1 Introduction . . . 75
4.2 Method and Material . . . 76
4.2.1 Train-Validation Partition (TVP) . . . 76
4.2.2 Models . . . 77
4.2.2.1 Real and complex layers . . . 77
4.2.2.2 Residual RAKI and complex residual RAKI . . . 77
4.2.2.3 Grid-search: K-fold cross validation . . . 78
4.2.3 Implementation details . . . 79
4.2.3.1 Code and computing resources . . . 79
4.2.3.2 Model training . . . 80
4.2.3.3 Network architecture hyperparameters . . . 80
4.2.3.4 Statistics . . . 81

4.2.3.5 fastMRI and separated-ACS mode data . . . 81
4.3 Results . . . 82
4.3.1 Grid-search in 2D and separated ACS mode . . . 83
4.3.2 2D results . . . 83
4.3.2.1 fastMRI retrospective undersampling (integrated ACS mode) . . . 83
4.3.2.2 T2 Weighted in separated ACS mode . . . 89
4.3.2.3 2D multi-echo GRE in separated mode . . . 95
4.3.2.4 Effect of the number of ACS lines . . . 95
4.3.3 3D RESULTS in separated mode . . . 100
4.3.3.1 3D MPRAGE . . . 100
4.3.3.2 3D Multi-Echo GRE . . . 100
4.4 Discussion . . . 104
4.4.1 Main findings . . . 104
4.4.2 Non-linearity . . . 104
4.4.3 Real vs complex . . . 105
4.4.4 TVP and K-fold cross-validation . . . 106
4.4.5 Kernel sizes and depths . . . 106
4.4.6 ACS and Timing . . . 106
4.4.7 COBRAI . . . 107
4.4.8 Improvements and practical implications . . . 107
4.5 Conclusions . . . 108

5 SANGRIA: one-Shot leArNinG super-ResolutIon with Adversarial training for accelerated Magnetic Resonance Imaging 109
5.1 Introduction . . . 109
5.2 Background and Related works . . . 109
5.2.1 One-shot learning . . . 109
5.2.2 Single-image patch-based generative models . . . 110
5.2.3 Data-driven regularization . . . 111
5.3 SANGRIA . . . 112
5.3.1 Model . . . 113
5.3.2 Model description . . . 113
5.4 Experiments . . . 114
5.4.1 Data and experimental setting . . . 114
5.4.2 Quantitative evaluation of the reconstruction quality . . . 115
5.4.2.1 Synthetic data . . . 115
5.4.2.2 In-vivo data . . . 115
5.4.3 Qualitative evaluation of reconstruction quality . . . 118
5.5 Discussion and Conclusion . . . 118

Conclusion 122

Bibliography 126

List of Figures 163

List of Tables 164

Introduction

MRI is a non-invasive medical imaging technique that has emerged as a powerful tool for generating detailed anatomical and functional information without relying on ionizing radiation. Unlike other imaging modalities, such as radiography, Computational Tomography (CT), and Positron Emission Tomography (PET), MRI utilizes a strong magnetic field and radiofrequency pulses to produce images. The inherent advantages of MRI make it particularly suitable for various clinical applications. MRI is a valuable diagnostic tool in modern healthcare, offering high-resolution visualization of anatomical structures and the ability to generate three-dimensional reconstructions. With the ability to differentiate between various types of tissues and employ specialized pulse sequences, MRI can provide detailed information on organs, blood vessels, and the central nervous system. However, the lengthy acquisition time required for scanning poses a significant challenge. Long scan times can result in patient discomfort, reduced throughput, and increased cost. MRI scanners sequentially collect raw data in k-space measurements which represent spatial frequency in Magnetic Resonance (MR) image. MR image is reconstructed by applying an inverse Fourier transform. MRI data is inherently multidimensional, 3D in space, but also has a depth corresponding to different contrasts and coils. The most common way to reduce the acquisition time is to acquire fewer measurements i.e. subsampling k-space and contrasts. Although the reconstructed image from the partially observed k-space measurements results in image blurring or aliasing artifacts. To address this issue, it is important to develop strategies and techniques that can enhance the quality of MRI reconstructions while maintaining optimal diagnostic quality. Accelerating MRI without compromising image quality has been a growing interest in recent years. Several deep-learning approaches have been proposed in the line of parallel imaging, i.e. focussing on the reconstruction task of undersampled data. Some of these approaches are learning from data that are specific to the imaging sequence employed (said scan-specific). The contribution of this work is to evaluate different strategies for their practical implementation, both in terms of experimental acquisition options and the training of reconstruction models, while proposing an original metric for quantifying the level of artifact in reconstructions.

Contributions

The problem statement of this thesis concerns the evaluation of deep learning models, particularly in terms of state-of-the-art (SOTA) performance, metrics, and generalizability, with a focus on adversarial loss and the lack of efficient metrics for assessing model robustness and training effectiveness. There are three main contributions of this thesis.

First, this work which begins with an overview of MRI physics and Machine learning (ML), applied ML to MRI, starts with a discussion on image quality metrics, including the Correlation-Based Residual Artifact Index (COBRAI) metric. The evaluation of the image quality in MRI reconstruction includes both subjective and objective methods. The subjective approach depends on expert reader who visually inspect and qualitatively assess the images based on their experience and perceptual judgment. However, subjective assessments are susceptible to inter- and intra-observer variabilities, which

can limit their reliability and reproducibility. Despite the availability of numerous objective metrics, challenges persist in developing universally accepted and reliable evaluation frameworks. Objective evaluations using commonly used metrics such as Normalized Root Mean Square Error (NRMSE), Normalized Mean Absolute Error (NMAE), Peak Signal-to-Noise Ratio (PSNR), and Structural Similarity Index (SSIM) often fail to account for local variations and fine details, resulting in suboptimal assessments. To address these limitations, we propose COrreletion-Based Residual Artifact Index (COBRAI), a novel correlation-based metric that is sensitive to structured residuals in reconstructed images, and we compare it with different scan-specific Deep learning (DL) reconstruction methods. Qualitative and quantitative comparisons were performed for multiple acceleration rates, ACS size and type (integrated versus separated), as well as sequence types (2D T2W, MPRAGE, FLAIR, T1W, 2D/3D gradient-echo magnitude and phase images).

Secondly, the development of the scan- and patient-specific MRI reconstruction methods where the goal is to generalize in the context of multi-parametric (or multi-contrast) setting. In recent years, DL methods have been successful in performing MRI image reconstructions. Although most studies have focused on magnitude images, the MRI phase is crucial for numerous applications such as field mapping and Quantitative Susceptibility Mapping (QSM). Deep neural networks with complex convolutions have shown high representational power and accuracy, making them suitable for MRI reconstruction, including phase-based applications, with highly undersampled data. Thus, we propose complex residual RAKI (crRAKI), a novel complex-valued residual Convolutional Neural Network (CNN) designed for scan-specific 2D MRI data, as well as its extension to 3D data. The proposed method was compared with Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) and residual RAKI (rRAKI) for accelerated MRI reconstruction. This work focussed on the time-efficient separated-mode AutoCalibration Signals (ACS) acquisition and investigated the effects of architecture hyperparameters, non-linearity or real versus complex implementations for generalizability. We proposed a Training–Validation Partition (TVP) scheme used for early stopping to avoid overfitting. To ensure optimal performance, we recommend implementing a K-fold cross-validation strategy that enables the selection of the most suitable hyperparameters and adapts the network size according to the configuration of the receiver coil. To validate these methods, a publicly available dataset, fastMRI, was tested by simulating the integrated Autocalibration Signal (ACS) mode, and in-house data was used to evaluate the separated ACS mode.

Lastly, the extension of scan-specific methods, SANGRIA, which is trained using an adversarial learning strategy as data-driven regularization, has been explored. We focus on uniform undersampled acquisitions for which SANGRIA performs one-shot superresolution in k-space. A CNN operating in the k-space domain is combined with patch-based adversarial learning in the image domain and trained for every new single image. Image enhancement is presented as well as potential drawbacks.

Thesis Outline

Chapter 1: Magnetic Resonance Imaging basis provides a brief historical timeline of the various technological advancements in MRI along with the introduction of the fundamentals of MRI physics. This chapter introduces the various MRI contrasts and quantitative imaging used in this thesis, including proton density and susceptibility contrasts as well as T1-Weighted (T1W)/T2-Weighted (T2W). Despite MRI's numerous benefits, it faces technical and practical challenges. This chapter also examines the limitations and areas for improvement in current MRI techniques.

Chapter 2: Machine learning in MRI focuses on the integration of ML in MRI, which involves around introducing relevant machine learning and deep learning methods. The chapter also provides an overview of neural network architectures such as CNNs. The application of ML in MRI reconstruction is reviewed, including reconstruction methods that utilize fixed cartesian sampling patterns and the optimization of sampling patterns i.e. adaptive sampling. Furthermore, ML techniques are discussed for Quantitative Magnetic Resonance Imaging (qMRI). The chapter concludes by summarizing the significance, potential implications, and future directions of integrating machine learning into MRI acquisition and reconstruction processes.

Chapter 3: COBRAI: A Contribution to Image-Quality Assessment offers a review of the traditional image quality metrics that are widely used in the field of MRI and their inherent limitations. Additionally, this chapter introduces COBRAI, a novel metric and its advantages over conventional metrics. The discussion in this chapter focuses on the applicability and limitations of COBRAI in the context of MRI image quality assessment on phantom, natural and MR images.

Chapter 4: Focus on scan-specific convolutional models for MRI reconstruction is dedicated to the optimization of scan-specific convolutional models for MRI reconstruction. An overview of scan-specific RAKI-based models is provided. Various variants are evaluated. Models are trained on the ACS using a Train-Validation Partition (TVP) scheme to prevent overfitting. The hyperparameters were optimized using a grid-search strategy and K-fold cross-validation. The performance of real/complex/residual variants was evaluated using both qualitative and quantitative metrics, including the Correlation-Based Residual Artifact Index (COBRAI), and was compared for different acceleration rates, ACS sizes and types (integrated versus separated), and sequence types (2D T2W, MPRAGE, FLAIR, T1W, 2D/3D gradient-echo magnitude and phase images).

Chapter 5: SANGRIA: one-Shot leArNinG super-ResolutIon with Adversarial training for accelerated Magnetic Resonance Imaging presents a comprehensive overview of one-shot learning in the context of MRI and provides an evaluation of the benefits of using adversarial methods in this context. Additionally, it offers a thorough introduction to the background of single-image-patch-based generative models and provides an overview of the architecture and training methodology of the SANGRIA model. SANGRIA is an extension of a scan-specific model that has been trained using an adversarial approach and a patch discriminator for solving Super-Resolution (SR) in k-space for accelerated MRI.

1 Magnetic Resonance Imaging basis

Summary

1.1 A brief historical perspective . . . 13
1.2 From magnetization to the Free Induction Decay . . . 15
1.2.1 Nuclear magnetization . . . 15
1.2.2 Resonance and flip angle . . . 16
1.2.3 Precession and relaxation . . . 17
1.3 Localization . . . 19
1.3.1 Gradients and k-space . . . 19
1.3.2 Gradient-echo pulse sequence . . . 19
1.3.2.1 Slice selection . . . 20
1.3.2.2 Frequency and phase encoding . . . 20
1.3.3 Spin-echo pulse sequence . . . 21
1.3.4 Coil array and parallel imaging . . . 21
1.3.4.1 Image combination in coil arrays . . . 22
1.3.4.2 Reducing scan time with parallel imaging . . . 23
1.4 Overview of MRI contrasts and quantitative imaging . . . 24
1.4.1 Basic MRI contrasts . . . 24
1.4.1.1 Proton density contrast . . . 24
1.4.1.2 Susceptibility contrast . . . 25
1.4.1.3 T1-weighted . . . 25
1.4.1.4 T2-weighted . . . 25
1.4.2 Quantitative imaging . . . 26
1.4.2.1 Relaxometry . . . 26
1.4.2.2 Magnetic susceptibility quantification . . . 27
1.5 Challenges . . . 27

1.1 A brief historical perspective

The foundations of Nuclear Magnetic Resonance (NMR) were established in the early 1930s by Otto Stern and his colleagues, who discovered the magnetic moment of protons, a fundamental constituent of atomic nuclei. This remarkable finding paved the way for further exploration of the magnetic properties of nuclei and set the groundwork for subsequent advancements in NMR. In 1938, Isidor Rabi[214] conducted groundbreaking experiments demonstrating the interactions between atomic nuclei and electromagnetic waves. He showed that when subjected to specific frequencies of electromagnetic radiation, atomic nuclei absorb and subsequently emit energy. Rabi coined the term "NMR" to describe the resonant behaviour of atomic nuclei in the presence of a magnetic field. The true breakthrough in NMR research occurred in 1946 when two independent groups of scientists, Felix Bloch[32] and Edward Purcell[241], explained the phenomenon of NMR on water and paraffin samples . They developed techniques to measure nuclear magnetic moments and gained invaluable insights into atomic structures and molecular dynamics. Their exceptional contributions to this field were recognized with the Nobel Prize in Physics in 1952. The 1950s saw further advancements in NMR technology, as Proctor and Yu[41, 237] discovered the frequency shift of NMR signals based on the molecular environment.

In the 1970s, Paul Lauterbur[167] and Sir Peter Mansfield[203] independently introduced a groundbreaking concept that transformed the field of NMR and paved the way for medical imaging. They proposed using magnetic-field gradients to encode spatial information within the NMR signal. This innovative idea led to the development of multi-planar image formation and slice selection, enabling the visualization of anatomical structures. Their remarkable achievements were recognized with the Nobel Prize in Physiology or Medicine in 2003, which honored their pioneering contributions to the growth of modern MRI.

Since its inception, MRI technology has made remarkable progress, leading to improved imaging capabilities and a wider range of clinical applications[144]. Key advancements in magnet design, gradient systems, and pulse sequences have significantly enhanced image quality and reduced acquisition time. The development and optimization of pulse sequences, including T1-weighted, T2-weighted, and Gradient echo (GRE), have further improved signal-to-noise ratio and reduced total acquisition time[28]. One of MRI's main strengths is its ability to provide multiple contrasts, enabling exceptional soft tissue differentiation. This versatility allows clinicians to visualize intricate anatomical structures with greater sensitivity and detect minute pathological changes. However, it is important to acknowledge the challenges posed by the prolonged scan duration.

From patient-centric perspective, patients undergoing MRI examinations often find themselves in a difficult situation as they must remain still for extended periods, as even minor involuntary movements can introduce motion artifacts and compromise the quality and reliability of the acquired images[113, 304, 331]. This issue is particularly significant when dealing with pediatric or claustrophobic patients, who may struggle to endure lengthy scans. Furthermore, the extended image acquisition time limits MRI's capacity to capture dynamic physiological processes in real time. Such as cardiac imaging, where the heart's rapid motion requires both high spatial and temporal resolutions, the restriction of time can prove to be a significant challenge. This limitation hinders the ability to accurately visualize and diagnose cardiac conditions in a timely manner. From a healthcare management perspective, extended scan durations reduce the overall patient throughput of MRI facilities, resulting in longer wait times for patients and lowering the economic efficiency of healthcare institutions. Additionally, specific diagnostic protocols often depend on promptly capturing

images after the administration of contrast agents. Delays in acquisition can lead to missing the optimal imaging "window," thereby undermining the diagnostic precision required in time-sensitive scenarios.

The ongoing pursuit of accelerating MRI acquisition times without compromising the image quality has fueled an active field of research. In recent years, significant advancements have been made in MRI hardware and sequence designs. Hardware improvements, including higher-field-strength magnets[84] and more efficient gradient systems[80, 302], have enhanced the sensitivity and speed of image acquisition. Similarly, As illustrated in Figure1.1, the current trend in MRI reconstruction techniques exemplifies the continuous efforts to improve image quality and reduce scan times. Innovations, such as parallel imaging[98] and compressed sensing[76], have emerged to address challenges such as motion artifacts and faster imaging. In the early 2000s, a significant milestone in the advancement of MRI was achieved with the introduction of parallel imaging techniques, as notably pioneered by Griswold and Pruessmann [98, 239]. These strategies harnessed the spatial information garnered from multi-coil MRI systems, effectively reducing acquisition times. In parallel, the emergence of compressive sensing algorithms capitalized on the sparsity of MRI data, enabling the reconstruction of images from undersampled data. Nonetheless, these techniques have limitations, including the presence of residual aliasing artifacts and a reliance on specific sampling patterns[269]. This has catalyzed the shift towards assessing the capabilities of neural networks in MRI. In MRI reconstruction, the integration of neural networks with parallel imaging and compressive sensing has led to the development of hybrid methodologies. These combined techniques optimize data acquisition rates and enhance image quality while preserving image fidelity. Additionally, the incorporation of neural networks into fingerprinting methods[107] in MRI presents advanced imaging strategies. Specifically, employing neural networks in Magnetic Resonance Fingerprinting (MRF) facilitates the correlation between the observed signal dynamics and a library of the simulated signals. This advancement allows the simultaneous characterization of different tissue in a singular acquisition, reflecting significant progress in quantitative MRI techniques[196].

In this chapter, we delve into the fundamental principles of Free Induction Decay (FID), including nuclear magnetization, resonance, flip angles, precession, and relaxation. We explore spatial localization in MRI techniques by examining k-space, gradient- and spin-echo pulse sequences. Additionally, we provide an overview of various MRI contrasts and quantitative imaging methods, including basic contrasts such as proton density, susceptibility, T1-weighted, and T2-weighted imaging, as well as advanced techniques like relaxometry and magnetic susceptibility quantification.

1.2 From magnetization to the Free Induction Decay

MRI is a non-invasive medical imaging technique rooted in NMR, facilitating the capture of detailed soft tissue contrasts, providing anatomical, structural or functional information in vivo. It capitalizes on the inherent magnetic properties of atomic nuclei, particularly the abundance of hydrogen protons present in biological tissues. Hydrogen nuclei, or protons, are the principal subjects of clinical imaging MRI, constituting an estimated 63% of atoms within the human body[9]. MRI foundation lies in the principles of spin physics, which involve the manipulation of nuclear spin and its response to magnetic fields[228]. By applying magnetic fields, MRI can control the alignment and precession of these nuclear spins, thereby generating signals that can be used to reconstruct detailed images. In this section, the basic principles of MRI physics are discussed.

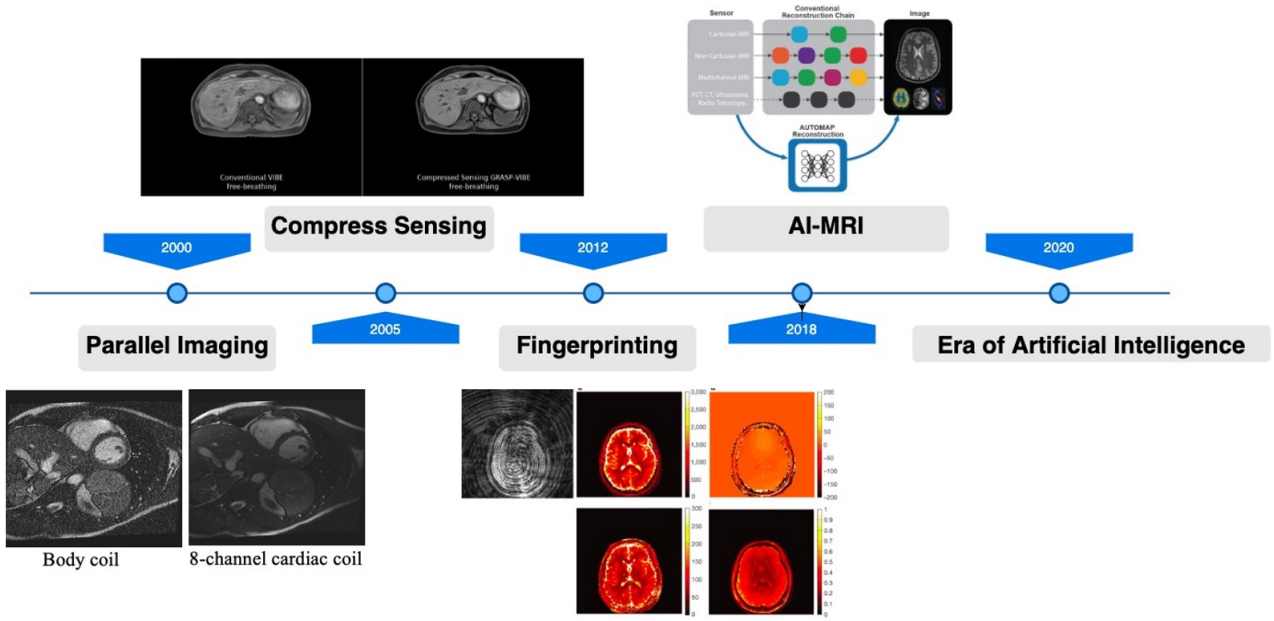

Figure 1.1: Emerging Trends in MR Image Reconstruction: highlights current advancements in MRI reconstruction methods, such as parallel imaging[98, 239], compressive sensing[76], MR fingerprinting[196], and AI-based techniques[375]. Further details on these methods will be provided in the subsequent sections.

1.2.1 Nuclear magnetization

To understand the principles of MRI and its image-acquisition process, it is vital to first present what MRI measures: the magnetization of nuclear spin ensembles. Nuclear spin refers to the inherent quantum property of the nucleus, which arises from its nonzero net angular momentum. Consequently, these nuclei possess a net magnetic moment, endowing them with characteristics similar to those of small magnets. Only the case of 1/2-spins is presented herein, which is the case for the hydrogen nucleus, for which only two energy levels are possible. In the context of MRI, direct measurement of individual spins is not feasible; instead, an ensemble of spins is observed, leading to the notion of magnetization.

Magnetization is represented by vector $\vec{M}(\vec{r},t)$, where $\vec{r}$ is the position and t is the time. The net magnetization denotes the average magnetic moment per unit volume. In the absence of an external magnetic field $\vec{B}_0(\vec{r}) = \vec{0}$, the nuclear spin magnetic moments have random orientations, resulting in a net magnetization of zero. However, when subjected to an external magnetic field, $\vec{B}_0(\vec{r}) = B_0(\vec{r})\vec{z}$, the spins are polarized by a small fraction, that is, they have a higher probability of being in the low-energy state (parallel to the field). In the above expression, the magnetic field B_0 is assumed to be invariant in time but inhomogeneous in space. Note that this expression is an approximation, as it does not satisfy Maxwell's equations. However, this is a realistic approximation, as in practice, the relative spatial variation of B_0 is of the order of 10-to-100 part-per-million (ppm), such that components along other axes are often neglected for high-field systems. The typical magnetic field used for clinical MRI range from 1.5 to 7 T (Telsa). On average, the net magnetization in the direction of the field is proportional to B_0 in magnitude. This phenomenon is known as thermal polarization (as it depends on the temperature), leading to thermal magnetization $\vec{M}_0(\vec{r}) = M_0(\vec{r})\vec{z}$ proportional to the average number of nuclei at a given location $\vec{r}$, called the proton density $\rho(\vec{r})$. Estimating the

spatial distribution $M_0(\vec{r})$ is thus refered to as "proton density imaging."

Measuring the magnetization requires bringing the magnetization $\vec{M}(\vec{r},t)$ out of its thermal equilibrium $\vec{M}_0(\vec{r})$, which is perfomed using the NMR phenomenon, that will 'rotate' the magnetization and enables detection through induction. Given the preferential axis of $\vec{B}_0$, the magnetization is often separated into its component along $\vec{z}$, which is denoted as $M_Z(\vec{r},t)$, and in the transverse XY plane that will be noted $M_{XY}(\vec{r},t)$, the latter being expressed as a complex field, where the real and imaginary parts denote the x and y components, respectively.

1.2.2 Resonance and flip angle

The energy difference between the two levels of a 1/2 nuclear spin is equivalent to a frequency specific to the nucleus. The Larmor angular frequency for a given nucleus is $\omega_0 = \gamma B_0$ where γ is the gyromagnetic ratio specific to the nucleus. Because $\gamma_{^1\mathrm{H}}/2\pi = 42.57\mathrm{MHz/T}$, this typically lies in the Radiofrequence (RF) range. The application of precisely tuned RF can thus interact with the nuclei and induce "resonance," i.e a coherent transition, which on average flips the magnetization from its thermal equilibirum. Classically, the average magnetization undergoes precession, that is, rotation, around the magnetic field direction at the Larmor angular frequency and then becomes detectable as a FID.

The Bloch equation describes the behavior of the net magnetization, $\vec{M}(\vec{r},t)$, in the presence of a time-varying magnetic field $\vec{B}(\vec{r},t)$:

$$\frac{d\vec{M}(\vec{r},t)}{dt} = \gamma\vec{M}(\vec{r},t) \times \vec{B}(\vec{r},t) \tag{1.1}$$

The RF excitation is represented as an oscillating magnetic field applied orthogonally to the main magnetic field, that is, in the transverse plane XY (the third component has no effect and is thus omitted). This RF pulse is produced by RF coils. It is often expressed as a magnetic field $\vec{B}_1(\vec{r},t)$ or equivalently, as an angular velocity $\vec{\omega}_1(\vec{r},t) = \gamma\vec{B}_1(\vec{r},t)$. In practice, a radiofrequency pulse induces an inhomogeneous spatial profile owing to the transmit radiofrequency coil design, which can vary from several % to tens of % in some situations. To simplify the notations, the effect of a transmit pulse is often expressed in a complex notation $\omega_1(\vec{r},t) = |\omega_1(\vec{r},t)|\exp\big(i\phi_1(\vec{r},t)\big)$, with i, the imaginary unit with $i^2 = 1$.

If we apply an oscillating radiofrequency pulse at the angular frequency ω, or equivalently with a phase $\phi_1(\vec{r},t) = \phi_1(\vec{r},0) + \omega t$, it is often simpler to describe the motion in a frame of reference that rotates at the same angular frequency, as if demodulation is performed. In this frame of reference, referred to as the "rotating frame," we can rewrite the Bloch equations as follows:

$$\frac{d\vec{M}_r(\vec{r},t)}{dt} = \vec{M}_r(\vec{r},t) \times (\vec{\omega}_1(\vec{r},t) + \delta\vec{\omega}_0(\vec{r})) \tag{1.2}$$

where $\vec{M}_r$ denotes the magnetization in the rotating frame (demodulated), and with the following angular velocity vectors:

$$\vec{\omega}_1(\vec{r},t) = \begin{bmatrix} |\omega_1(\vec{r},t)|\cos\big(\phi_1(\vec{r},0)\big) \\ |\omega_1(\vec{r},t)|\sin\big(\phi_1(\vec{r},0)\big) \\ 0 \end{bmatrix} \tag{1.3}$$

$$\delta\vec{\omega}_0(\vec{r}) = \begin{bmatrix} 0 \\ 0 \\ \delta\omega_0(\vec{r}) \end{bmatrix} \tag{1.4}$$

where $\delta\omega_0(\vec{r}) = \omega_0(\vec{r}) - \omega$ is the spatial inhomogeneity of the Larmor frequency in reference to the RF pulse one.

As can be seen from the Bloch equations in the rotating frame above, close to the resonance condition ($\omega \approx \omega_0$), the magnetization rotates around the axis defined by vector $\vec{\omega}_1$ at angular velocity ω_1.

This perturbation, known as "excitation," causes the net magnetization to deviate from its thermal equilibrium alignment with the static magnetic field. Assuming that the magnetization starts from its thermal equilibrium $\vec{M}_0(\vec{r})$ and that a radiofrequency pulse with constant angular frequency $\vec{\omega}_1(\vec{r})$ is applied for a duration τ, the magnetization is flipped toward the transverse plane by the flip angle $\alpha(\vec{r}) = \omega_1(\vec{r})\tau$ and with a phase $\phi_1(\vec{r},0)$ which depends on the RF pulse phase and coil design.

1.2.3 Precession and relaxation

Upon turning off the RF pulse, the perturbed magnetization gradually returns to its thermal equilibrium state in a process called "relaxation." During relaxation, the net magnetization precesses, which can be detected by induction. Relaxation is often described by two time constants: the longitudinal relaxation time constant (T_1) or spin-lattice relaxation time and the transverse relaxation time constant (T_2) or spin-spin relaxation time. These relaxation time constants vary depending on the composition and structure of the tissue of interest and both contribute to the generation of MRI contrasts. The inverses of these time constants are known as the longitudinal and transverse relaxation rates, R_1 and R_2, respectively.

The behavior of magnetization in these precession and relaxation phases is described using the Bloch equations[32][1] with relaxation terms. Without RF pulses, the magnetization evolution in the rotating frame can be described as:

$$\frac{d\vec{M}_r(\vec{r},t)}{dt} = \vec{M}_r(\vec{r},t) \times \delta\vec{\omega}_0(\vec{r}) - \begin{bmatrix} R_2(\vec{r}) & 0 & 0 \\ 0 & R_2(\vec{r}) & 0 \\ 0 & 0 & R_1(\vec{r}) \end{bmatrix} \vec{M}_r(\vec{r},t) + \begin{bmatrix} 0 \\ 0 \\ M_0(\vec{r})R_1(\vec{r}) \end{bmatrix} \tag{1.5}$$

where $R_1(\vec{r})$ and $R_2(\vec{r})$ vary through space depending on the tissue type.

This equation describes the dynamics of the net magnetization $\vec{M}_r(\vec{r},t)$ over time in the rotating frame taking into account the relaxation processes and the thermal equilibrium magnetization. Starting from an initial value (after, e.g., an exitation) $\vec{M}_r(\vec{r},0) = \vec{M}_{XY}(\vec{r},0) + \vec{M}_Z(\vec{r},0)$ in which transverse and longitudinal parts have been separated, the evolution of the net magnetization $\vec{M}_r(\vec{r},t)$ over time is then:

$$\vec{M}_r(\vec{r},t) = e^{-R_2(\vec{r})t} e^{i\delta\omega_0(\vec{r})t} \vec{M}_{XY}(\vec{r},0) + e^{-R_1(\vec{r})t} \vec{M}_Z(\vec{r},0) + (1 - e^{-R_1(\vec{r})t}) \vec{M}_0(\vec{r},0) \tag{1.6}$$

These equations describe the time evolution of the transverse and longitudinal components of the net magnetization, considering the relaxation processes with respect to the relaxation rates, as well as the initial magnetization state at time $t = 0$:

[1] https://www.drcmr.dk/bloch

$$M_Z(\vec{r},t) = M_Z(\vec{r},0)e^{-R_1(\vec{r}),t} + M_0(\vec{r})(1-e^{-R_1(\vec{r}),t}) \tag{1.7}$$

$$M_{XY}(\vec{r},t) = M_{XY}(\vec{r},0)e^{-R_2(\vec{r})t}e^{i\delta\omega_0(\vec{r})t} \tag{1.8}$$

The longitudinal component progressively returns to its thermal equilibrium during exponential decay. Additionally, it can be seen that the transverse magnetization precesses at the frequency $\delta\omega_0(\vec{r})$ (in the rotating frame) while its magnitude is decaying at the rate $R_2(\vec{r})$. This time-varying magnetization can then be detected via induction in the tuned receiver coils. Let $C(\vec{r})$ be the sensitivity of a tuned receiver coil (C is a complex function). In practice, the induced signal originates from a volume element (referred to as a voxel in the context of imaging) in which $\delta\omega_0(\vec{r})$ is not homogeneous. By introducing the mean angular frequency $\Delta_{\omega_0}(\vec{r})$ inside this volume element, the induced signal can be approximated as

$$s(\vec{r},t) = C(\vec{r})M_{XY}(\vec{r},0)e^{-R_2^*(\vec{r})t}e^{i\Delta_{\omega_0}(\vec{r})t} + \mathcal{N}(0,\sigma^2) + i\mathcal{N}(0,\sigma^2) \tag{1.9}$$

where the apparent transverse relaxation rate R_2^* is introduced to account for the faster decay resulting from the intra-volume frequency inhomogeneity. The NMR signal induced inside the coil is referred to as FID. As can be seen, it bears information about the proton density and transmits B_1, as well as receiving the coil sensitivity profile, magnetic field B_0 (through its frequency), and the local magnetic environment (through the apparent relaxation rate). $\mathcal{N}(0,\sigma^2) + i\mathcal{N}(0,\sigma^2)$ is used to account for Gaussian noise, which is generally assumed to have 0-mean and variance σ^2 identical for the real and imaginary parts.

1.3 Localization

1.3.1 Gradients and k-space

The MR signal is generated after excitation and precession. However, it is not possible to determine which part of the body the signal originates from without additional information about the location of the spins. Therefore, it is necessary to spatially encode the origin of the FID signal. This is achieved through the use of "gradient" coils. Gradient coils are used to modify spatialy the value of the static magnetic field component along $\vec{z}$ in the following way:

$$B_z(\vec{r},t) = B_0(\vec{r}) + \vec{G}(\vec{r},t)\cdot\vec{r} \tag{1.10}$$

Consequently, the angular frequency of magnetization can be rendered dependant on position in a controlable manner. Typical rise times are in the order of hundreds of microseconds, while typical maximum gradients are few tens of mT/m.

K-space is caracterized at time t by the vector $\vec{k}$, which can be expressed as:

$$\vec{k}(t) = \gamma/2\pi \int_0^t \vec{G}(\tau)d\tau \tag{1.11}$$

where $t = 0$ related to the RF pulse application used as a reference. As can be seen, the dimension

of each component of the vector $\vec{k}$ is the inverse of a distance, thus they correspond to spatial frequencies.

The effect of applying the gradient modulates the signal $s(\vec{r},t)$ with the term $\exp\left(2i\pi\vec{k}(t)\cdot\vec{r}\right)$. Assuming that $s(\vec{r},t)$ does not vary much during the application of the gradients, the signal (in the rotation frame) arising from the entire imaging volume V is then the sum of all signals arising from individual volumes, which is expressed as:

$$S(\vec{k}) = \int_V s(\vec{r},t)e^{i2\pi\vec{k}\cdot\vec{r}}\,dv \tag{1.12}$$

This relationship sums up the whole principle of MRI localization. This sum over the entire magnetization volume V is the Fourier Transform (FT) of the elementary signal spatial distribution, which is directly linked to the magnetization spatial distribution. If enough information on the value of S at various positions in k-space (for different spatial frequencies), the spatial magnetization distribution $s(\vec{r},t)$ can be estimated using a FT. In practice, the function $S(\vec{k})$ is sampled by varying the gradients through time, i.e., by changing $\vec{k}(t)$ in a controlled manner.

1.3.2 Gradient-echo pulse sequence

The main localization concepts are slice-selection, frequency encoding, and phase encoding realized by combining RF excitations, gradients, acquisition and relaxation periods. The chronogram of these events is called a sequence. Three main principles are used that will be presented using the example of a gradient-echo sequence, as depicted in Figure 1.2.

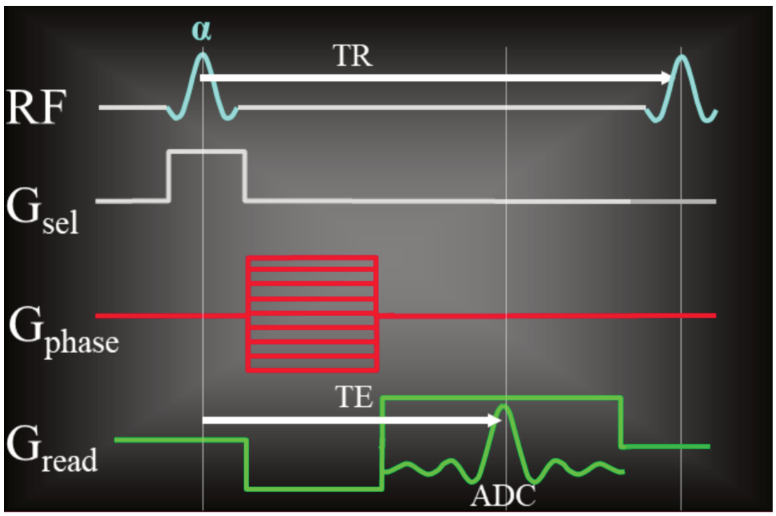

Figure 1.2: Schematic illustration of a gradient-echo sequence presenting the basic ingredients for localization in MRI. A frequency selective RF pulse is applied together with a slice selection gradient G_{sel} producing a flip angle of α. Frequency encoding is obtained using a readout gradient G_{read} producing an echo at the echo time TE. The sequence is repeated (every repetition time TR) with different phase encoding gradients G_{phase}. Note that the gradients presented here are applied in any direction enabling a full latitude on the slice location and orientation.

1.3.2.1 Slice selection

Slice selection is usually the initial step in MR image formation in 2D imaging, where a specific two-dimensional plane, or slice, within the three-dimensional object of interest is targeted for examination. The slice selection gradient (denoted as G_{sel} in Figure 1.2) is applied in conjunction with a radiofrequency excitation pulse with a selective frequency profile. This enables magnetization to be selectively flipped within this frequency range by an angle α. By playing in the gradient direction and RF pulse excitation frequency, the location can be selected in any orientation and position. Slice thickness can also be modulated by varying the gradient strength and RF pulse bandwidth. At this step, the signal can be written in 2D for the selected slice as

$$S_{slice}(k_x, k_y) = \int_{XY} s_{slice}(x, y, t) e^{-i2\pi(xk_x + yk_y)} dxdy \tag{1.13}$$

1.3.2.2 Frequency and phase encoding

Frequency encoding is performed by applying a gradient (denoted G_{read} in Figure 1.2) during the recording of the signal (Analog-to-Digital Converter (ADC)). In the first step, the gradient is negative, dephasing the magnetization along the direction of the gradient. The gradient is then reversed and acquisition (or readout) is performed. Assuming that the gradient is applied along the axis x, the spatial frequency k_x varies during the readout as

$$k_x(t) = \gamma/2\pi G_{read} \times (t - TE). \tag{1.14}$$

A line in k-space is sampled during acquisition with sampling time t_s, starting from negative to positive values. At $t = TE$, the phase accumulated by magnetization along x cancels, leading to an echo for which the signal is the following sum over the slice: $S_{slice}(0,0) = \int_{XY} s_{slice}(x, y, t) dxdy$.

During the preparation phase (before the acquisition), a gradient is also applied along y, leading to a dephasing in this direction. The sequence is repeated with a repetition time TR. The area of the phase encoding gradient is stepped by different values Δk_y at each repetition j, e.g., $k_{y,j} = (j - j_0)\Delta_{k_y}$. By doing so, the signal $S(k_x, k_y)$ is sampled in 2D, and a 2D Fourier transform can be done to recover a 2D image for the selected slice. The image is then approximated as:

$$s_{gre,slice}(x, y, t) = C_{slice}(x, y) M_{XY,slice}(x, y, 0) e^{-R^*_{2,slice}(x,y)TE} e^{i\Delta_{\omega_0,slice}(x,y)TE}, \tag{1.15}$$

where the coil sensitivity profile weights the flipped magnetization in the transverse plane that has apparently relaxed and is dephased. Note that noise has been omitted here to simplify the expression.

For 3D imaging, while some imaging scenarios prefer multi-slice 2D imaging, extension of the phase-encoding principles to the slice direction is also possible, consequently requiring 3D Fourier transforms for reconstruction.

As can be seen, the time required to acquire an image depends on the number of phase-encoding steps, N_y for the first phase-encoding direction, and N_z for the second phase-encoding direction. Thus, the total acquisition time is

$$T_{acq} = N_y \times N_z \times TR \tag{1.16}$$

TR is on the order of 10-1000 ms, and the number of phases encoding is close to 256, an image acquisition can take seconds to minutes.

1.3.3 Spin-echo pulse sequence

The second class of sequences is based on the spin echo principle. In a spin-echo sequence (Figure 1.3), the echo is not produced by gradient reversal, but by flipping the magnetization by 180°, effectively refocusing the transverse magnetization. Using the complex notation for the transverse magnetization and assuming that this 180°pulse is perfect, it transforms the in-plane signal $S(k_x, k_y)$ before the pulse into its complex conjugate $S(-k_x, -k_y)$. Consequently, every B_0 induced dephasing that was present in FID is cancelled at the spin-echo time TE, leading to an image in which the inhomogeneities are removed.

$$s_{se,slice}(x, y, t) = C_{slice}(x, y) M_{XY,slice}(x, y, 0) e^{-R_{2,slice}(x,y)TE}, \tag{1.17}$$

where the phase has been cancelled, as well as intra-volume apparent transverse relaxation, leading to image contrast driven by T_2 effects. Noise is omitted here to simplify the expression.

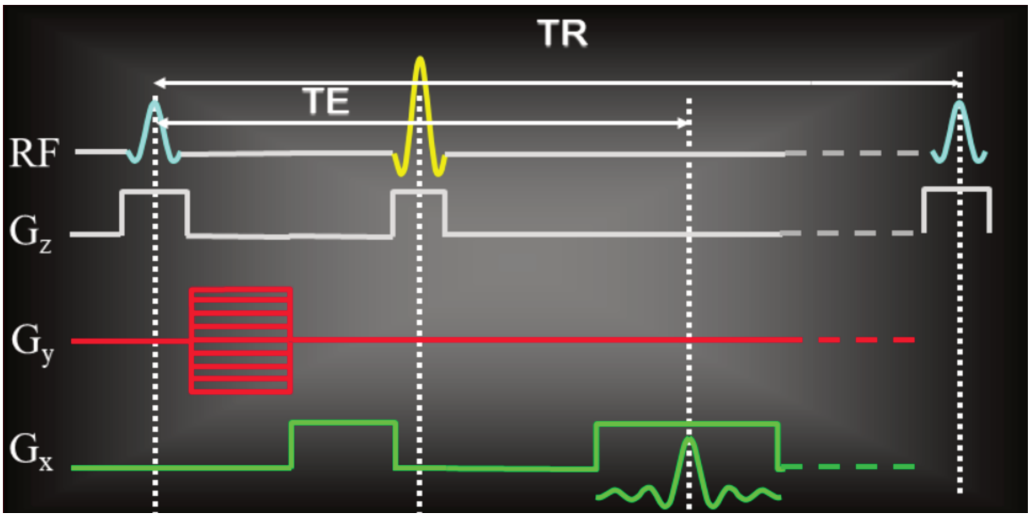

Figure 1.3: Schematic illustration of a spin-echo sequence. A frequency selective RF pulse is applied together with a slice selection gradient producing a flip angle of 90°. Here, a second RF pulse with and angle 180° is also applied, refocussing the magnetization and producing a spin-echo at TE. Frequency and phase encoding are obtained similarly as the gradient-echo sequence. The sequence is repeated (every repetition time TR) with different phase encoding gradients.

1.3.4 Coil array and parallel imaging

The use of multiple smaller receiver coils, known as coil arrays, is a strategy commonly employed in routine MRI to improve the signal-to-noise ratio (Signal-to-Noise Ratio (SNR)) and accelerate image acquisition. Smaller coils produce larger induced signals and can sense less noise. The approach of combining multiple coil images to increase SNR was initially proposed in [262]. Subsequently,

it has been employed to accelerate acquisition using a range of methods collectively referred to as Parallel Imaging (PI), which leverage the inherent localization capabilities of smaller coils. The fundamental principles of this technique will be discussed in greater detail in subsequent sections of the manuscript when relevant to the current discussion.

1.3.4.1 Image combination in coil arrays

When multiple coils are used simultaneously, individual signal from each coil c can be written:

$$s_c(\vec{r},t) = C_c(\vec{r}) M_{XY}(\vec{r},0) e^{-R_2^*(\vec{r})t} e^{i\Delta_{\omega_0}(\vec{r})t} + \mathcal{N}(0,\sigma^2) + i\mathcal{N}(0,\sigma^2), \tag{1.18}$$

assuming that noise has been normalized and is uncorrelated between coils. If k-space is sampled sufficiently, images from each coil can be estimated. They are then combined as a single image. A basic combination is the sum-of-squares reconstruction in which the phase information is lost:

$$s_{sos}(\vec{r},t) = \sqrt{\sum |S_c(\vec{r},t)|^2}, \tag{1.19}$$

More advanced combinations are possible to preserve the phase information, such as if the sensitivity is known or can be estimated[262], such as a noise-uniform combination:

$$s_{phase}(\vec{r},t) = \sum \bar{C}_c(\vec{r}) S_c(\vec{r},t) / \sqrt{\sum \bar{C}_c(\vec{r}) C_c(\vec{r})}, \tag{1.20}$$

where $\bar{C}_c$ is the complex conjugate of C_c.

1.3.4.2 Reducing scan time with parallel imaging

Additionally, by taking advantage of the geometry of the coil array, parallel imaging can accelerate MRI. This technique enables us to work from an under-sampled k-space, which is faster to acquire. The acceleration factor R is introduced. This corresponds to the ratio of the number of k-space lines of a full k-space to the effectively acquired undersampled one.

To introduce the principles of parallel imaging, as illustrated in Figure 1.4, reconstruction methods can be divided into two general categories: those that operate in the image domain and those that operate in the k-space domain. The Sensitivity Encoding (SENSE) method[239] is an example of an image-domain technique that utilizes multiple receiver coil sensitivities to disentangle overlapping signals from different coil elements, resulting in faster image acquisition. SENSE reconstructs images directly from acquired data by exploiting the spatial sensitivity profiles of the coil array.

In contrast, the GRAPPA[98] functions in the k-space domain. GRAPPA is a parallel imaging technique that accelerates image acquisition by undersampling the k-space and then reconstructing the missing data using the information from the fully sampled regions. By exploiting the redundancy inherent in multicoil acquisitions, GRAPPA can reconstruct images from the undersampled k-space data. As we will see, this thesis focuses specifically on GRAPPA-inspired approaches. Furthermore, iterative approaches, such as the Iterative Self-consistent Parallel Imaging Reconstruction (SPIRIT)[193], have been adapted to accommodate non-uniform undersampling patterns. Additional information about these methods can be found in subsequent chapters.

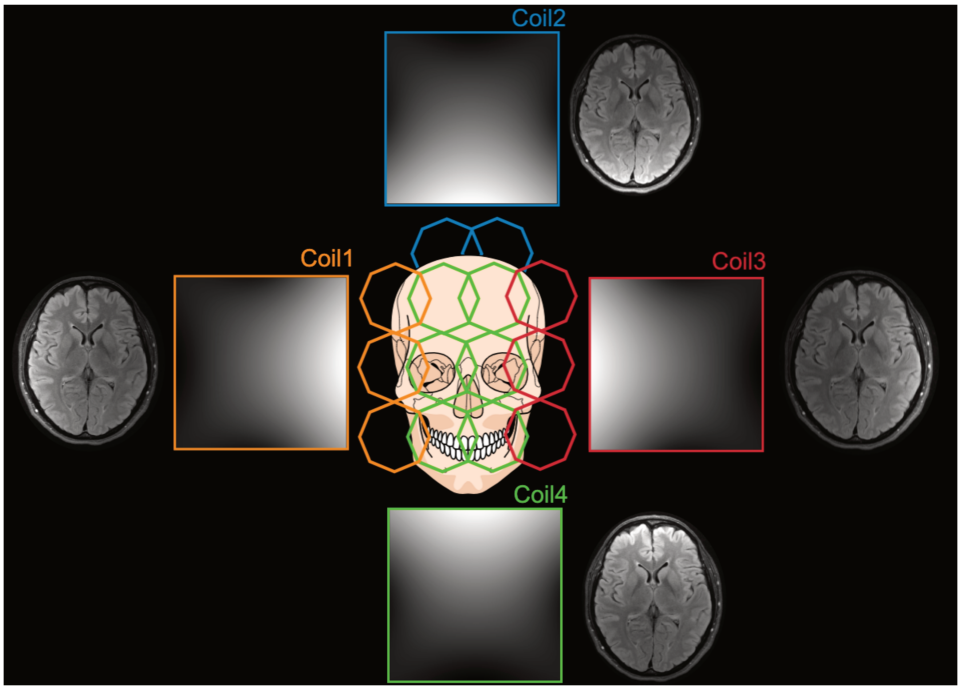

Figure 1.4: Parallel imaging with 4-coil setting. Image to be acquired is shown in the centre along with coil sensitivities and the corresponding reconstructed image.

1.4 Overview of MRI contrasts and quantitative imaging

1.4.1 Basic MRI contrasts

In this work, we explore reconstruction techniques that utilize data acquired through various pulse sequences, resulting in a range of contrasts. MRI offers a multitude of sequences, combining both localization principles and the potential to generate contrast. The timing of the excitation and relaxation periods is crucial for generating the basic contrasts. The sequences relevant to this work are briefly introduced in this section with the intention of providing an overview rather than an exhaustive list. For illustrative purposes, Figure1.5 is shown.

1.4.1.1 Proton density contrast

Proton density-weighted imaging[200, 201] refers to images in which the contrast is dominated by the proton density. Thus, it must be rendered independent of relaxation times. This is generally done using either gradient-echo sequences with small flip angles and/or long TR, as well as short TE, or spin-echo sequences with long TR and short TE. This type of basic contrast offers limited information on soft tissue. However, for the brain, there is a small proton-density contrast that can be observed between grey matter and white matter, resulting from the fact that white matter is quite structured, leading to a fraction of water that has very short transverse relaxation times and cannot be seen, even with the short TE in practice.

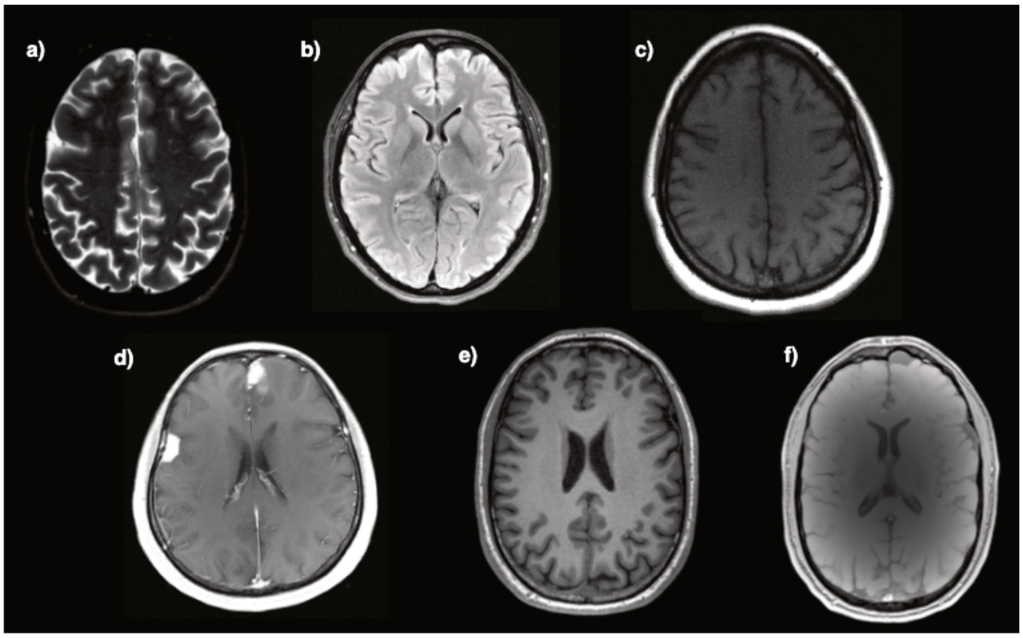

Figure 1.5: Illustration on brain of different contrasts that can be obtained with MRI a) T2-weighted Spin-Echo displaying hypersignals for fluids, b) Fluid-Attenuated Inversion Recovery (FLAIR), c) T1-weighted (using a Gradient echo), d) T1-weighted image after contrast agent injection that changes T1 (said T1-POST), e) T1 contrast using a Magnetization-Prepared Rapid Acquisition using Gradient-Echo (MPRAGE) f) Simple Gradient-Echo (GRE) scan with a limited T1-contrast, thus close to a proton density image.

1.4.1.2 Susceptibility contrast

As previously mentioned, the difference in images generated by GRE sequences is largely influenced by the apparent transverse relaxation time or T_2^*. This parameter is closely related to the intra-voxel field inhomogeneities[199], which can be attributed to the magnetic susceptibility variations within the voxel. These variations may be caused by factors such as hemorrhage, calcifications, or air-tissue interfaces, all of which contribute to the generation of strong deformations in the images[180]. To optimize the image contrast, the echo time, or TE, is typically adjusted to be close to the target tissue's T_2^* value.

1.4.1.3 T1-weighted

T1W imaging techniques capitalize on the properties of the longitudinal relaxation time. Because of the diverse range of T1 relaxation time (T1) values encountered in tissues, this method is commonly employed for anatomical imaging. Several sequences are available to generate the T1 contrast, such as the short echo (SE) and GRE sequences, which feature short Repetition Time (TR) and short Echo Time (TE) values. In practical applications, a Magnetization Prepared Rapid Gradient Echo (MPRAGE)[35] sequence is commonly used. This sequence was designed to generate the desired contrast and was followed by a rapid series of gradient-echo readouts. The preparation involved an inversion pulse of 180°, followed by a relaxation period lasting TI (inversion time). Starting from the inverse thermal magnetization at $t = 0$, the magnetization at TI before the imaging phase is:

$$M_Z(\vec{r}, TI) = M_0(\vec{r}) \times (1 - 2\exp(-R_1(\vec{r})\,TI)) \tag{1.21}$$

When TI is chosen adequately, it increases the contrast between tissues, for example, between white and grey matter. For brain imaging, TI is usually chosen to cancel the signal from the cerebral-spinal fluid with $1 - 2\exp\left(-R_{1,CSF}\,TI\right) = 0$, where $R_{1,CSF}$ denotes the CSF relaxation rate.

1.4.1.4 T2-weighted

T2W imaging primarily utilizes spin-echo sequences, as these are necessary to cancel out the susceptibility effects that occur in gradient-echo sequences. The contrast in T_2 can be manipulated by using TE, which is typically set to match the relaxation time of a given tissue. Transverse relaxation times are highly sensitive to the presence of "free water," which refers to water that can move freely. As a result, tissues with longer T_2 values, such as those found in the cerebrospinal fluid or inflamed areas where cells are less compact, will have longer T_2 values (on the order of seconds). In contrast, healthy, structured tissues have shorter T_2 values (on the order of 100 ms). Owing to these properties, T2W imaging is often used to detect edema, inflammation, and fluid accumulation.

Separating tissues based solely on T2W can be a challenging task. To overcome this, combined contrast can be employed. One such example is the Fluid Attenuated Inversion Recovery (FLAIR)[222] sequence, which utilizes a specific approach to enhance the visualization of inflammation regions. Initially, an inversion is performed to cancel the signal from the CSF, which is based on its T_1 value. Subsequently, a spin-echo imaging sequence is applied, resulting in the more distinct display of inflammation regions due to the attenuation of fluid[17]. This serves as an illustration of how MRI leverages relaxation phenomena to selectively filter tissues for a specific imaging purpose.

1.4.2 Quantitative imaging

The fundamentals of basic MRI, which are sufficient for routine clinical scans, have been discussed. However, reality is more sophisticated, as advanced systems and reconstruction techniques have enabled the exploration of a wider range of phenomena and the extraction of more detailed information. One significant advancement in MRI is the move towards quantification, which allows for the measurement of the physical properties of tissues, rather than just their visualization. qMRI[77] is a rapidly growing field that allows the mapping and estimation of tissue biophysical properties. As established earlier, a variety of physical properties and phenomena, such as relaxation[1, 350], flow[205], perfusion[134], diffusion[148], and magnetic susceptibility[272] can influence MRI signals. Furthermore, the signal intensity of an MR image is influenced not only by the tissue properties but also by the acquisition parameters (e.g., repetition time, echo time, flip angle) and the condition of the MR system being used (e.g., homogeneities of the radiofrequency field B1 and main magnetic field B0, uniformity of receiver coil sensitivity)[100]. In general, the MRI parameter space can be divided into three subspaces corresponding to the tissue relaxation properties, acquisition parameters, and system conditions. Quantitative parametric mapping can be accomplished by sampling the acquisition parameter space and solving the inverse problem by utilizing multiple measurements from the acquisition parameters to estimate the tissue property maps. Typically, an MR image, commonly known as a weighted image, is captured using a predefined pulse sequence with specific imaging parameter settings[261, 347]. The resulting image is qualitative in nature and its signal intensity

is influenced by both tissue characteristics and other factors in the MRI parameter space. It is not possible to directly assess the inherent physical properties of tissues, such as T1 or T2 relaxation times, which present a significant obstacle in image analysis, precision medicine, and tissue microstructure assessment.

qMRI offers several advantages over conventional MRI by enabling direct comparison between imaging sites, time points, and various acquisition methods. This is made possible by the use of calibrated physical units, which allow for standardized diagnoses, multicenter studies, and long-term investigations of brain development and changes. With MRI, we can transition from identifying lesions to assessing their progression and determining the stage in which they are present. Additionally, MRI provides objective interpretations of MRI images and is a valuable tool for assessing levels of iron and myelin in the human brain[166]. Therefore, it is of great interest to study neurological disorders and age-related changes. This technique provides quantitative measurements of specific physical parameters related to the nuclear spin of protons in water that act as intrinsic probes for the microstructure of the surrounding tissue. These parameters include longitudinal and transverse relaxation rates, magnetic susceptibility, proton density, and magnetization transfer, which offer crucial information regarding myelination, iron, and cell membranes in the living brain. Such studies are critical for clinical trials and brain research. We will briefly study about the relaxometry[132, 240] and magnetic susceptibility along with MRF in the next sections.

1.4.2.1 Relaxometry

The basis of relaxometry is the measurement of relaxation times or rates, which translate in imaging to relaxometric maps. These maps can be generated by using spin- or gradient-echo sequences. For transverse relaxation, several images with different echo times are acquired, using gradient-echo sequences for $T2^*$ mapping and spin-echo for T2 mapping. Because T1 contrast depends on TR and flip angle, or on prepartaion time (in magnetization-prepared sequences) both spin-echo or gradient-echo sequences can be used. The accuracy of these methods is influenced by a variety of factors, such as the chosen sequence, TR, TE, the number of images acquired at different TE values, and the fitting model applied to the acquired data. While the equations presented above involve simple exponential decay, tissues are in reality more complex than mono-component water, and phenomena such as magnetization transfer between different chemical species may occur. This deviation from the expected 'simplistic' mono-exponential behaviour complicates the accurate extraction of relaxation times using standard fitting models, while it already provides a quantitative average localized evaluation.

1.4.2.2 Magnetic susceptibility quantification

Magnetic susceptibility (χ) is a fundamental property of materials that measures their magnetization response to an external magnetic field. Susceptibility Weighted Imaging (SWI)[105] is known for its ability to differentiate various tissue types based on their susceptibility properties. However, it can be difficult to differentiate between haemorrhage and calcifications that appear hypointense. QSM[69] overcomes this challenge by providing quantitative information about underlying tissue susceptibility. It differentiates between blood products (paramagnetic) and calcifications (more diamagnetic than tissue). It is sensitive to bio-metals such as iron which is involved in neurodegeneration and neuroinflammation such as in Multiple Sclerosis (MS)[288].

The presence of susceptibility variations within a tissue affects the local magnetic field, which in turn influences the precession frequency of the protons and the received signal. Using several echoes in a gradient echo scan enables to map this frequency very precisely. QSM then exploit the frequency maps to reconstruct the underlying χ distribution solving a field-to-source inverse problem. Several methods exist for generating QSM, each with its own advantages and limitations. A popular iterative methods called MEDI[184] will be used in this work.

1.5 Challenges

Reducing the duration of MRI is a considerable challenge because of various factors that contribute to the length of the acquisition process. Moreover, accurately capturing tissue characteristics can be difficult, particularly in situations that involve tissue heterogeneity, complex structures, and pathological conditions. Each voxel in an MRI image represents a mixture of tissues with distinct relaxation properties. For instance, relaxometry can lead to inaccurate measurements, particularly for smaller structures. To optimize the trade-off between scan time and accuracy, additional data points can be acquired; however, this increases the scan time, which can be a significant drawback in clinical settings, where patient comfort and efficiency are crucial. Patient movement during the scan can introduce artifacts that distort the signal intensity and compromise the accuracy of relaxation time measurements. Ensuring optimal image quality within clinically feasible scan times is a challenge that needs to be addressed. Standardizing and reproducing MRI across different scanners and institutions are crucial for wider clinical adoption. Overall, the challenges associated with data acquisition and quantification highlight the need for innovative approaches to accelerate data acquisition, improve image quality, and enhance clinical utility. To address this issue, research has focused on finding ways to accelerate the acquisition process, known as accelerated MRI. Among the various approaches to accelerated MRI are parallel imaging, the design of fast pulse sequences, and undersampling of the k-space are common methods. In the next chapter section, we investigate potential applications of machine learning in this line.

2 Machine learning in MRI

Summary

2.1 Introduction . . . 29
2.1.1 Primer to Deep Learning . . . 29
2.1.1.1 Basics: one neuron at a time . . . 30
2.1.1.2 Activation functions . . . 33
2.1.1.3 Convolutional Neural Networks . . . 33
2.1.1.4 Generative Models . . . 36
2.2 MRI reconstruction and Machine Learning . . . 40
2.2.1 Image-based super-resolution . . . 42
2.2.2 K-space-based super-resolution . . . 43
2.2.2.1 MR Reconstruction using fixed sampling . . . 45
2.2.2.2 Optimizing sampling patterns . . . 54
2.2.2.3 ML for qMRI . . . 57
2.3 Conclusion . . . 59

2.1 Introduction

ML has become an essential tool across multiple industries, significantly impacting how data is viewed and interpreted. It has gained widespread interest in various sectors such as computer vision, natural language processing (NLP), speech recognition, genomics, automation, finance, and healthcare.[24, 169, 178, 226, 298, 309]. Emerging in the mid-20th century, DL experienced a comeback with the availability of computational resources and extensive datasets. The rise of parallel computing with Graphics Processing Unit (GPU)s, affordable storage, and cloud services has enabled the processing of large amounts of data. Advancements in MRI have significantly increased the availability of data. It is not surprising that ML-based solutions have been investigated.

This chapter discusses the application of ML techniques in MRI. The discussion begins with an overview of deep learning and its components, including activation functions, and various network architectures such as CNN, and Generative Models. Then we shift the focus to accelerated MRI reconstructions, covering its mathematical foundations and methods for MR reconstruction using fixed sampling, where the undersampling pattern is predefined meaning not learned. These include

scan-specific, database-driven, and hybrid approaches. Furthermore, the optimization of sampling patterns and the use of machine learning in qMRI have been covered through brief literature reviews on traditional methods and current state-of-the-art advancements. Finally, the challenges and future trends are discussed in the applied ML into MRI.

2.1.1 Primer to Deep Learning

Deep Neural Network (DNN)s are inspired to mimic the structure and function of biological neural networks. They are composed of interconnected artificial neurons which are stacked into layers, including the input layer, one or more hidden layers, and the output layer. Each neuron in a neural network accepts inputs, applies weights to these inputs, and passes the results through an activation function which introduces non-linearity[215]. The connections between neurons constitute the model parameters. Neural networks are trained using an optimization process that adjusts the weights of the connections between neurons to minimize a loss function, which measures the difference between the predicted output and the actual output.

In this section, we will provide a foundation for the basic building block of a neural network, starting with a single neuron and progressing to a Multi Layer Perceptron (MLP)[59]. We will also explore various concepts, including forward and backward propagations[264], loss functions, and optimization.

2.1.1.1 Basics: one neuron at a time

To understand the various notions of deep learning, we focus on supervised regression tasks. Our primary model of choice is the MLP[59], which is a fundamental deep-learning model. Feedforward neural networks, dense networks, or MLPs are essential elements of deep learning, and the objective of a feedforward network is to approximate the function f that learns a mapping between input and output. The term "feedforward" is used to describe these networks because information flows through the network from input to output, without any feedback connections in which the output of the network is fed back into itself.

The architecture of a neural network is typically represented by a directed acyclic graph known as a computational graph, which represents the dependencies between the neurons and the operations within the network. Computational graphs are directed graphs that represent the dependencies between variables and operations within a model or mathematical expression. To build a computational graph for a neural network, each neuron and operation is represented as a node, and an arrow is drawn from one node to another if the first is the input to the latter. Most deep-learning frameworks such as Pytorch provide an automatic differentiation procedure to compute gradients based on the backpropagation algorithm. The computational graph for the gradients, with respect to the neural network's output, is used to compute the gradients of the weights and biases of the network. During the forward pass, the function may be called multiple times, resulting in the construction of a new graph if the program contains control flow statements such as conditionals and loops. The computational graph is then traversed in the reverse direction during the backward pass, and the gradients are calculated using the chain rule. The gradients are then used to update the weights and biases of the network.

Objective

In a supervised setting, given a distribution p of input data x and target outputs t, the objective is to learn the parameters $\theta = (W, b)$ of a model $f(x, \theta)$ by minimizing the risk (objective, cost, or error) function E.

$$E^*(\theta) := \mathbb{E}_{(x,y)\sim p}\, L(f(x;\theta), t) \tag{2.1}$$

where L is a per-sample loss function that compares predictions $f(x;\theta)$ to targets t.

Since the true distribution p is unknown, the empirical distribution $\hat{p}$ of a training set $x_1, \ldots, x_m$ with associated target outputs $y_1, \ldots, y_n$ is used to minimize instead the empirical risk. Thus converting the learning problem to optimization.

$$E(\theta) := \mathbb{E}_{(x,y)\sim \hat{p}}\, L(f(x;\theta), y) = \frac{1}{n}\sum_{i=1}^{n} L(f(x_i;\theta), y_i) \tag{2.2}$$

Forward propagation

In the simplest case, consider a dense network (Fig2.1) with one hidden layer where all input units are connected to all output units. We call this a fully connected layer. Consider the following notation for a dense network with L layers: a_{l-1} is the input vector for layer l (where $a_0 := x$ is the input vector for the neural network), W_l is the weight matrix, and b_l is the bias vector for layer l (for $l \geq 1$). Let a_l denote the pre-activation vector of layer l and ζ_l represent the post-activation vector output from the activation function g_l. The pre-activations were determined using the equation :

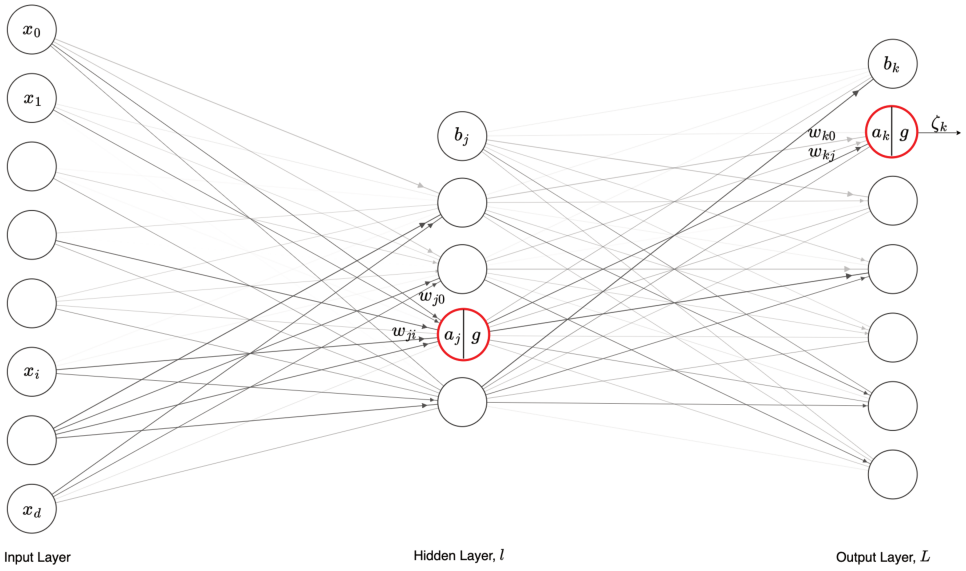

Figure 2.1: MLP: consists of multiple layers of interconnected nodes, where each node in one layer connects to every node in the following layer, resulting in a fully connected network. MLP typically includes an input layer to receive input features, one or more hidden layers to perform nonlinear transformations of the data, and an output layer to produce the final output, which is typically the predicted class or regression value.

$$a_l = \zeta_{l-1} W_l^T + b_l \tag{2.3}$$

whereas the post-activations were calculated with

$$\zeta_l = g_l(a_l) \tag{2.4}$$

The number of neurons in layer l is denoted as n_l. Thus, both vectors a_l and ζ_l are vectors in $\mathbf{R}^{n_l}$; for simplicity, they are treated as row vectors. The weight matrix of layer l, denoted by W_l, has the dimensions $W_l \in \mathbf{R}^{n_l \times n_{l-1}}$.

Loss function

Let E be the cost function that computes the square of the norm of the difference between the network's output ζ_L and target vector t (MSE). We denote $\delta E = \frac{\partial E}{\partial \zeta_L}$ the gradient of the cost function with respect to the output of the network ζ_L. We typically refer to $\delta_L = \frac{\partial E}{\partial a_L}$ as the gradient of the cost function with respect to the activations (before passing through the activation function) of the output layer of the network, and $\delta_l = \frac{\partial E}{\partial a_l}$ as the gradient of the cost function with respect to the activations (before passing through the activation function) of layer number l which is intermediate to the network.

Optimization: Learning with Backpropagation

Forward propagation and backpropagation are interdependent processes in neural network training. Forward propagation involves traversing the computational graph in the direction of dependencies and computing the variables along the path, whereas backpropagation reverses the computation order on the graph to update the model parameters using an optimization algorithm and gradients obtained from backpropagation. Backpropagation requires storing intermediate values from forward propagation to avoid duplicate calculations, but this also increases memory usage during training compared with plain prediction. The sizes of these intermediate values are proportional to the number of network layers and the batch size. Consequently, training deeper networks with larger batch sizes can lead to out-of-memory errors.

The development of efficient optimization strategies is crucial for effectively training neural networks. Gradient descent, a key optimization algorithm, minimizes the loss function and enhances model performance by iteratively adjusting model parameters in the direction of the steepest descent. Although the most basic configuration involves updating the parameters based on the gradient of the entire training dataset, this method is computationally expensive for large datasets. To address this issue, stochastic gradient descent (Stochastic Gradient Descent (SGD)[260]) has been proposed, which calculates gradients and adjusts parameters based on a single training sample at a time. This significantly reduces the computational burden while maintaining the progress towards optimization.

Overfitting and Regularization

A key measure of a model's performance is the generalization gap, which is the difference between its fit on the training and test data. A large gap generally indicates that the model is overfitting the

training data, and in extreme cases, the model may fit the training data perfectly even though the test error remains high. The strategies to avoid overfitting include using early stopping[254, 346] and regularization[92]. Regularization adds a penalty term to the objective function, thus ensuring generalization in the case of the complex model. Traditional methods like Lasso and L2 regularization adjust the model's coefficients to reduce overfitting. Additionally, utilizing dropout can be effective, wherein neurons are randomly dropped during training to prevent an excessive dependence on any one feature.

Early stopping[254, 346] is commonly used in deep neural networks that involve restricting the number of training epochs. This approach is often used to prevent overfitting and improve generalization. The idea is to monitor the validation error throughout the training and stop the training when the error does not decrease by a certain amount for a specified number of epochs. Implementing early stopping can potentially reduce training time and significant costs, particularly for large models that require extensive computing resources. In general, when the data are well-behaved and the classes are easily distinguishable (e.g., classifying cats and dogs), early stopping does not significantly improve the generalization. However, in situations in which the data contain label noise or intrinsic variability, early stopping is essential[91].

Another popular regularization technique is the **dropout**[287]. Dropout randomly drops out the neurons within each layer during forward propagation, thereby enhancing the network's generalisation ability. The idea of training models with input noise to encourage smoothness was first introduced by Bishop in 1995[31], who linked this to Tikhonov regularization. This concept establishes a mathematical relationship between the model's tendency to be simple and its robustness to input perturbations. Subsequently, Srivastava et al.[287] developed a dropout technique, which intentionally adds noise to the internal layers of neural networks during training. The challenge lies in the effective injection of noise into the model. In standard dropout regularization, a fraction of the nodes in each layer is randomly zeroed during training, followed by debiasing through normalization. Specifically, with dropout probability $p_{dropout}$, each intermediate activation h is replaced by a random variable h' that is adjusted to maintain the expected value, thus promoting regularization without compromising the model expressiveness.

$$h' = \begin{cases} 0 & \text{with probability } p_{dropout} \\ \frac{h}{1-p_{dropout}} & \text{otherwise} \end{cases} \tag{2.5}$$

2.1.1.2 Activation functions

Activation functions are essential in neural networks because they enable the network to process complex functions that map inputs to outputs. The most commonly used activation functions are listed in Table 2.1.

2.1.1.3 Convolutional Neural Networks

The development of CNNs began in 1998 with the introduction of LeNet-5[168], which contains convolution and pooling layers, thereby laying the groundwork for contemporary CNNs. In 2012, AlexNet[163], a large-scale network, surpassed traditional computer vision methods in a prominent vision challenge. The VGG network[280] and network-in-network (NiN) architecture[177]

Table 2.1: Activation Functions

Function	**Definition**
Threshold	$g(x) = \begin{cases} 1, & \text{if } x \geq 0 \\ -1, & \text{if } x < 0 \end{cases}$
Sigmoid	$g(x) = \frac{1}{1+e^{-x}}$
Gaussian	$g(x) = e^{-\frac{1}{2}x^2}$
Hyperbolic Tangent (Tanh)	$g(x) = \tanh(x)$
Identity	$g(x) = x$

have also made significant contributions to the field. GoogLeNet[292] introduced multi-branch convolutions, whereas the residual network (ResNet)[114] remains a widely used architecture in computer vision. ResNeXt blocks[340] provide sparser connections and DenseNet[127] generalizes the residual architecture. Various optimizations for efficient networks have been developed over time, such as ShiftNet[332], which uses coordinate shifts, and MobileNetv3[124], which conducts an automatic search for efficient architectures. The semi-automatic design exploration by Radosavovic et al.[249] led to the development of the RegNetX/Y. Liu et al.[190] revisited CNN, ConvNeXts, which outperforms current state-of-the-art hierarchical vision Transformers in various computer vision benchmarks while retaining the standard ConvNet's simplicity and efficiency.

CNN input is high-dimensional such as images, and limits the connectivity between neurons from one hidden layer to another. Unlike fully connected networks, CNNs only allow for connections between a small subset of neurons, which is referred to as the receptive field of the neuron, and is determined by the size of the convolutional filter. The connections are local and extended along the entire depth of the input volume. The CNN architecture is represented by stacking multiple layers such as a convolutional layer, nonlinearity through activation functions, pooling layer, dropout layer, and fully connected layer, as depicted in Figure2.2.

Residual Block

Very deep networks may perform worse because of the vanishing gradient issue that arises during training. This has been addressed by He et al.[114] by the introduction of residual connections, also known as skip connections, in neural networks. This connection enables gradients to flow more freely through the network. The residual block is shown in Fig2.3.

Autoencoders

Autoencoders are neural networks that learn to reconstruct the original input while compressing data to create a more efficient and compressed representation. This unsupervised learning method was first introduced in the 1980s[34] and gained popularity through a seminal paper by Hinton and Salakhutdinov[118]. The network consists of an encoder that transforms the input into a low-dimensional latent code and a decoder that reconstructs the data from the code. The encoder has a smaller output size than the input size, while the decoder has larger output layers. The linear autoencoder performs dimensionality reduction similar to Principal Component Analysis (PCA) or Matrix Factorization (MF), and the parameters are learned together to output a reconstructed data

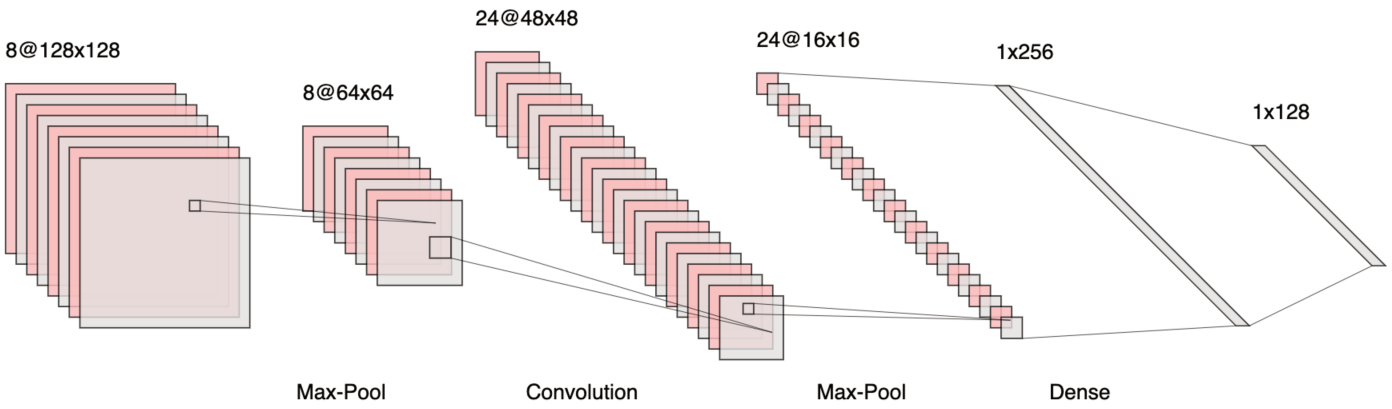

Figure 2.2: CNN: comprises multiple layers, including convolutional, pooling, and fully connected layers. This kind of arrangement of layers is often used in the classification task. The convolutional layers apply filters to the input images and capture spatial patterns and features. The pooling layers then downsampled the feature maps generated by the convolutional layers. Finally, the fully connected layers perform the classification based on the features extracted from the previous layers.

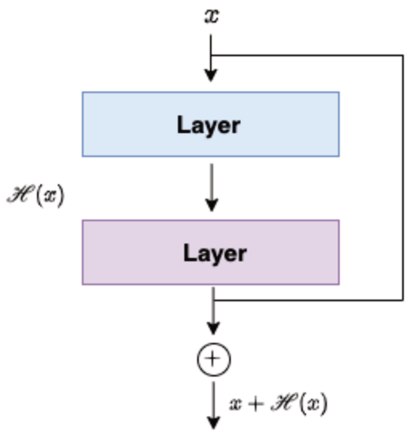

Figure 2.3: A Residual Block in a neural network. Here the residual connection is between two layers.

sample identical to the original input, representing the learning of an identity function. Autoencoders have several types, including Sparse Autoencoder[27], which has sparse activation in the hidden layer, and Variational Autoencoders (VAE)[27], which utilize principles of variational inference from probability theory. Denoising Autoencoders[312] are trained on noisy data to capture the underlying clean data. This can be useful for tasks such as image-denoising. Contractive Autoencoders[259] are trained with an additional penalty term that encourages the encoder to learn a more robust representation of the data. The penalty term is l2 norm of the derivatives of encoder activations in the hidden layer with respect to the input. This helps prevent the autoencoder from simply memorizing the data and enables the learning of more meaningful latent representations.

Complex-valued convolution neural networks

Although MRI data is inherently complex-valued, most current deep learning frameworks are not designed to support complex-valued data. The majority of standard computer vision datasets are indeed real-valued. Typically, reconstruction networks on complex data separate real and imaginary components into real-valued channels, which is suboptimal to represent complex numbers, as well as the derived phase information[108, 217, 285, 286, 375]. Recent research has demonstrated the representational power and accuracy of complex-valued deep neural networks, which has motivated the application of complex-valued CNN to MRI reconstruction[54, 66, 299]. The use of complex numbers allows an accurate representation of both the magnitude and phase of signals. The phase is a crucial aspect of many MRI applications, such as blood flow[206], QSM[323], fat-water separation[93], and chemical shift imaging[234]. Consequently, constructing a network that accurately reconstructs a phase can enhance various MRI applications.

For a given complex input $X = X_{real} + iX_{imag}$ and complex filter $W = W_{real} + iW_{imag}$, the complex convolution Y is calculated as:

$$\begin{aligned} Y_{real} &= X_{real} * W_{real} - X_{imag} * W_{imag}, \\ Y_{imag} &= X_{real} * W_{imag} + X_{imag} * W_{real}. \end{aligned} \tag{2.6}$$

Here, $*$ denotes the convolution operation, and $X_{real}, X_{imag}, W_{real}, W_{imag}$ represent the real and imaginary parts of the complex input and filter, respectively. Similarly, the complex-valued activation functions are defined in table 2.2.

Table 2.2: Complex Activation Functions

Function	**Definition**
modRELU[13]	$\text{modReLU(Z)} = \begin{cases} (\lvert Z\rvert + b)\frac{Z}{\lvert Z\rvert}, & \text{if } \lvert Z\rvert + b \geq 0, \\ 0, & \text{otherwise} \end{cases}$
$\mathbb{C}RELU$[299]	$\text{CReLU(Z)} = \text{ReLU}(\text{Re}(Z)) + i\,ReLU(\text{Im}(Z))$.
zReLU[99]	$\text{zReLU(Z)} = \begin{cases} Z, & \text{if } \theta_z \in [0, \pi/2], \\ 0, & \text{otherwise.} \end{cases}$
Identity	$g(Z) = Z$

$Z \in \mathbb{C}$, θ_z is the phase of Z, and $b \in \mathbb{R}$ is a learnable parameter. θ_z is the phase of Z.

2.1.1.4 Generative Models

Generative modelling is a type of unsupervised machine learning that aims to identify and learn the inherent patterns or regularities present in input data, to generate new examples that are coherent and plausibly derived from the original dataset. The goal of a generative model is to learn to model the distribution of the data. This approach has applications in various fields such as image synthesis[38, 131, 133, 250, 282, 359], music generation[56, 71, 162, 198], and text generation[248]. In this section, we will provide an overview of different generative models, including Generative Adversarial Networks (GAN)s[95], VAEs[151], Flow-based generative models[258], and diffusion models[282].

While GANs models the sampling procedure of the complex distribution learned in an adversarial

manner. VAE optimize the log-likelihood of data by maximizing the Evidence Lower Bound (ELBO). In contrast, flow-based generative models explicitly learn the data distribution through a sequence of invertible transformations. The loss function for flow-based models is the negative log-likelihood. Diffusion models draw inspiration from non-equilibrium thermodynamics and employ a Markov chain of diffusion steps to gradually introduce random noise into the data. The model then learns to reverse the process and generate desired data samples from the noise. Unlike VAE or flow models, diffusion models are learned using a fixed procedure, and the latent variable has the same dimensionality as the original data.

GAN

GANs represent a significant advancement in the field of generative models introduced by Ian Goodfellow et al.[95] which models the sampling procedure of the complex data distribution learned in an adversarial manner. GANs are composed of two networks (Fig2.4): generator G, which produces data, and discriminator D, which distinguishes between real and fake samples. In an original unconditional setting, the generator takes random noise as input and transforms it into fake data, whereas the discriminator attempts to distinguish between the real and generated data. The training objective of an unconditional GAN involves finding a generator that produces realistic images that are difficult for the discriminator to differentiate from real images. The training process is a minimax game in which the generator attempts to minimize the following objective function, whereas the discriminator aims to maximize it.

$$\min_G \max_D V(D,G) = \mathbb{E}_{(\mathbf{x})\sim p_{\text{data}}(\mathbf{x})}\left[\log D(\mathbf{x})\right] + \mathbb{E}_{\mathbf{z}\sim p_z(\mathbf{z})}\left[\log(1-D(G(\mathbf{z})))\right] \tag{2.7}$$

where $\mathbf{x}$ is a real data sample, $\mathbf{z}$ is a latent variable sampled from a prior distribution $p_z(\mathbf{z})$, $G(\mathbf{z})$ is the generated data by the generator, $D(\mathbf{x})$ is the output of the discriminator for real data.

The objective function $V(D,G)$ is a combination of the following two objectives: the first objective, concerning the discriminator's weights, is that the discriminator should accurately discriminate data sampled from the dataset from data samples with the generator, and the second objective, concerning the generator's weights, encourages the generator to produce data that the discriminator identifies as real.

In Conditional Generative Adversarial Networks (cGAN), the generator G and the discriminator D are conditioned on some additional information, denoted as ξ. The objective function for training the cGAN is:

$$\begin{aligned}\min_G \max_D V(D,G) = &\mathbb{E}_{(\mathbf{x},\mathbf{y})\sim p_{\text{data}}(\mathbf{x},\xi)}\left[\log D(\mathbf{x},\xi)\right] \\ &+ \mathbb{E}_{\mathbf{z}\sim p_z}(\mathbf{z})\left[\log(1-D(G(\mathbf{z},\xi),\xi))\right]\end{aligned} \tag{2.8}$$

where x is a real data sample, ξ is the conditional information (e.g., labels), $\mathbf{z}$ is a latent variable sampled from a prior distribution $p_z(\mathbf{z})$, $G(\mathbf{z},\xi)$ is the generated data by the generator conditioned on ξ, $D(\mathbf{x},\mathbf{y})$ is the output of the discriminator for real data and the corresponding conditional information ξ. The objective function $V(D,G)$ is similar to the original GAN formulation but now conditioned on the additional information $\mathbf{y}$.

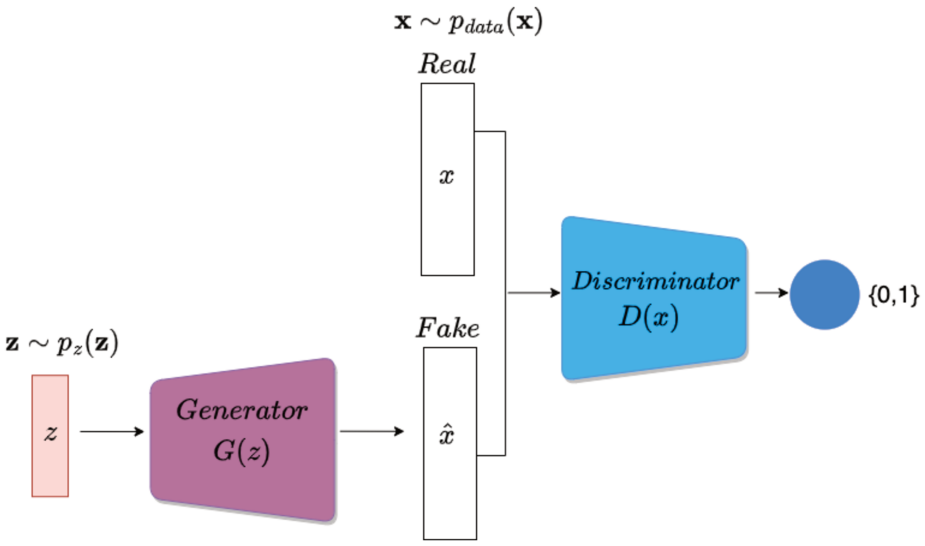

Figure 2.4: GAN: includes two components: generator and discriminator. The generator uses random noise as its input and produces synthetic data samples that are indistinguishable from the real data. The discriminator, on the other hand, receives both real and generated data samples and learns to differentiate between them. During training, the generator and discriminator are trained simultaneously in an adversarial manner, with the generator attempting to fool the discriminator and the discriminator attempting to correctly classify real and generated samples.

Variational autoencoders

VAE employs probabilistic modelling to generate samples. VAEs are composed of an encoder that maps the input data to a latent space and a decoder that reconstructs the input from this latent representation as shown in Fig2.5. The goal is to learn a probabilistic mapping between the input data and the latent space. Let x be the input data and z be the latent variable. The **encoder** maps the input data x to a distribution in the latent space $q_\phi(z|x)$ parameterized by ϕ. The prior distribution $p(z)$ is typically a standard Gaussian distribution.

$$q_\phi(z|x) = \mathcal{N}(\mathbf{z}, \mu_\phi(x), \sigma^2_\phi(x)I) \tag{2.9}$$

$$p_{latent}(z) = \mathcal{N}(\mathbf{z}; 0, I) \tag{2.10}$$

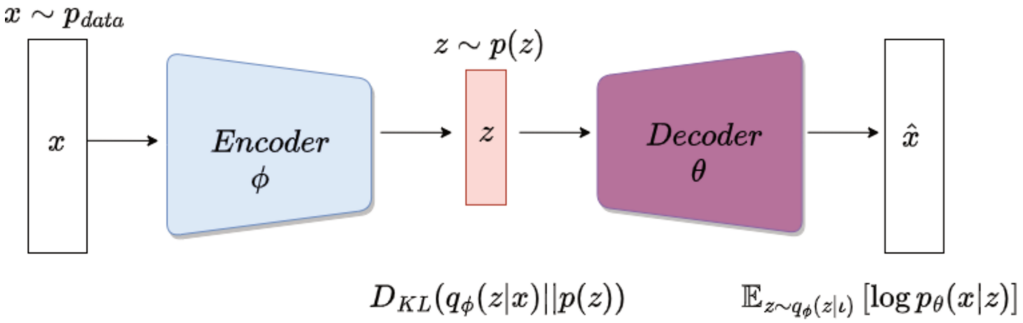

Figure 2.5: VAE: combines neural networks and probabilistic models. it includes an encoder network that maps input data to a latent-space distribution and a decoder network that reconstructs the input data from these latent representations. The primary purpose of training VAEs is to produce new data samples that closely resemble the distribution of training data.

The **decoder** learns a deterministic function $p_\theta(x|z)$ to convert the sampled latent variable z back into an observation x. The overall objective for training a VAE involves maximizing ELBO and more information can be found in [151].

$$\begin{aligned}
\text{ELBO}(\theta,\phi;x) &= \mathbb{E}_{q_\phi(z|x)}\left[\log \frac{p(x,z)}{q_\phi(z|x)}\right] \\
&= \mathbb{E}_{q_\phi(z|x)}\left[\frac{\log p_\theta(x|z)p(z)}{q_\phi(z|x)}\right] \quad \text{(Chain rule of probability)} \\
&= \mathbb{E}_{q_\phi(z|x)}\left[\log p_\theta(x|z)\right] - \mathbb{E}_{q_\phi(z|x)}\left[\log \frac{q_\phi(z|x)}{p(z)}\right] \quad \text{(Split the expectation)} \\
&= \underbrace{\mathbb{E}_{q_\phi(z|x)}\left[\log p_\theta(x|z)\right]}_{\text{Reconstruction Term}} - \underbrace{D_{KL}(q_\phi(z|x)||p(z))}_{\text{Prior Matching}} \quad \text{(Definition of KL divergence)}
\end{aligned} \tag{2.11}$$

In Equation2.11, the two terms have clear and straightforward interpretations. The first term quantifies the ability of the decoder to reconstruct the original data using the variational distribution, while the second term measures the similarity between the learned variational distribution and the prior belief about the latent variables. These terms jointly ensure that the learned distribution effectively captures the underlying representation that generates the original data. The Kullback-Leibler (KL) divergence term can be calculated analytically, while the reconstruction term can be approximated using a Monte Carlo estimate. The objective function is rewritten as:

$$\begin{aligned}
&\arg\max_{\phi,\theta} \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x)\,||\,p(z)) \approx \\
&\arg\max_{\phi,\theta} \sum_{l=1}^{L} \log p_\theta(x|z^{(l)}) - D_{KL}(q_\phi(z|x)\,||\,p(z))
\end{aligned} \tag{2.12}$$

where the latent $\{z^{(l)}\}_{l=1}^{L}$ is sampled from $q_\phi(z|x)$ for every observation x in the dataset. The default

setup, however, presents a challenge. The loss is calculated using $z^{(l)}$, which is generated through a stochastic sampling procedure that is typically non-differentiable. This issue can be resolved by employing the reparameterization trick when $q_\phi(z|x)$ is designed to model specific distributions, such as the multivariate Gaussian distribution. In a VAE, each z is calculated as a deterministic function of input x and auxiliary noise variable ϵ:

$$z = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon \quad \text{with} \quad \epsilon \sim \mathcal{N}(\epsilon; 0, I) \tag{2.13}$$

where $\odot$ represents the element-wise product. With this reparameterized version of z, gradients can be computed with respect to ϕ to optimize μ_ϕ and σ_ϕ.

After training a VAE, new data can be generated by sampling from the latent space $p(z)$ and passing it through the decoder. VAEs are particularly interesting because the dimensionality of z is less than that of the input x. Additionally, when a semantically meaningful latent space is learned, the latent vectors can be edited before being processed by the decoder to generate data more precisely.

To extend the VAE formulation to a Conditional Variational Autoencoders (cVAE), a conditional variable ξ representing additional information, such as class labels or attributes is introduced. The conditional encoder $q_\phi(z|x,\xi)$ now conditions on both input data x and c, and the conditional decoder $p_\theta(x|z,y)$ generates data conditioned on the sampled latent variable z and ξ. The training objective involves maximizing the conditional ELBO that includes both the reconstruction error and the divergence between the approximate and true posteriors conditioned on ξ.

Diffusion model

DIffusion Model (DM) have made significant progress in the field of generative modelling as they produce samples that closely match real-world data. The idea of employing diffusion for generative purposes was first introduced by Sohl-Dickstein et al.[282] and has since been enhanced through the development of noise-conditioned score networks[283], Score-based generative models (SGM)[284] and Diffusion Denoising Probabilistic Models (DDPM)[119]. Compared to other generative models, such as Normalizing Flows[155], GANs, or VAE, diffusion models offer a relatively straightforward approach for converting noise from a simple distribution into a data sample. In this case, the neural network learns to gradually remove noise from the data, starting from a state of pure noise.

The process of image generation involves a predetermined forward diffusion process and a learned reverse denoising diffusion process that collaborates to create images. In the forward process, Gaussian noise is gradually added to a real image sampled from a data distribution over a predetermined number of time steps, T, until the resulting image is composed solely of noise. Using a sufficiently large T and a well-designed schedule for adding noise at each time step, the process results in an isotropic Gaussian distribution at $t = T$ through a gradual transformation. The reverse process, on the other hand, involves gradually denoising an image starting from pure noise sampled from the Gaussian distribution until a real image is produced. This is achieved using a learned reverse denoising diffusion process, $\mathscr{P}_\theta$, in which a neural network is trained to gradually denoise an image starting from pure noise until a real image is obtained.

The field of generative models has witnessed considerable expansion in recent times. A variety of deep generative models have emerged in the past decade, exhibiting performance improvements in terms of sample quality, diversity, or fast sampling[58, 158, 191, 223, 343, 372]. However, achieving all three objectives simultaneously remains a challenge, which is a problem posed by the generative learning

trilemma identified by [339]. Existing models often sacrifice one of these objectives over the others, making it difficult to attain high-quality samples and diversity while also ensuring fast sampling. Denoising diffusion models may not be suitable for practical use due to their high computational cost and slow sampling. The slow sampling is largely a result of the Gaussian assumption in the denoising step, which applies only to small step sizes. To address this issue, Xiao et al.[339] proposed using a complex multimodal distribution to model the denoising distribution, enabling larger step sizes and reducing the total number of denoising steps. They introduced denoising diffusion GANs that use a multimodal cGAN to model each denoising step. Vahdat et al.[303] proposed a Latent Score-based generative models (LSGM) which allows for end-to-end training of SGM in the latent space, with improved expressiveness and reduced sampling time. These models have shown promising results in terms of sample quality, diversity, and fast sampling, thus providing a potential solution to the generative learning trilemma.

2.2 MRI reconstruction and Machine Learning

The medical imaging field has experienced significant changes with the rise of ML[178]. ML has immense potential, particularly for reducing the MRI scan times[129]. Accelerated MRI[157] through undersampling can shorten scan durations by acquiring only a portion of the data typically needed for a full scan. However, it is challenging to reconstruct high-quality images from undersampled data. Undersampled acquisition can lead to artifacts and the loss of diagnostic information. In this context, ML algorithms can be trained to learn the relationships within MRI data and effectively "fill in the gaps," enabling accurate reconstruction from undersampled acquisitions by capitalizing on a large MRI dataset.

Alternatively, accelerated MRI can be considered as a form of super-resolution(SR) that involves reconstructing a High-Resolution (HR) image from a Low-Resolution (LR) image. Either by designing the fast pulse sequences or collecting less raw data, HR images can be reconstructed from LR inputs, such as noisy or blurry images [67, 233, 305, 318, 334]. SR in MRI can be achieved in either an image- or k-space. The data acquisition and reconstruction processes influence the choice between the image- and k-space for SR. Recall that k-space is the Fourier transform of the image-space and contains the spatial frequency information of the MRI signal. SR in k-space manipulates this frequency information before transforming it into the image domain. One advantage of k-space-based SR is that it can directly integrate the MRI data acquisition process, which is inherently based on k-space sampling. A faster scan where a fully sampled lower-resolution noisy image is acquired and super-resolution methods for such images can be classified into image-based SR. The goal of SR image reconstruction is to create a HR image from an LR image. Deep neural networks can be employed to improve image quality by training on HR-LR image pairs, which can be applied to all low-field MR images with the same resolution as network training[67]. On the other hand, the scan time is reduced by undersampling k-space. The approaches where HR images are reconstructed from the undersampled k-space method are known as k-space SR. This thesis focuses on k-space SR reconstruction. In addition, a k-space sampling pattern can be considered adaptive. The research areas related to this topic are generally categorized into three main categories, as depicted in Figure2.6: a) training a "reconstructor" to map LR to HR reconstruction in a fixed sampling scenario, b) employing a "sampler" to learn an optimized sampling pattern, and c) creating a suitable fitting model for qMRI with multi-contrast inputs.

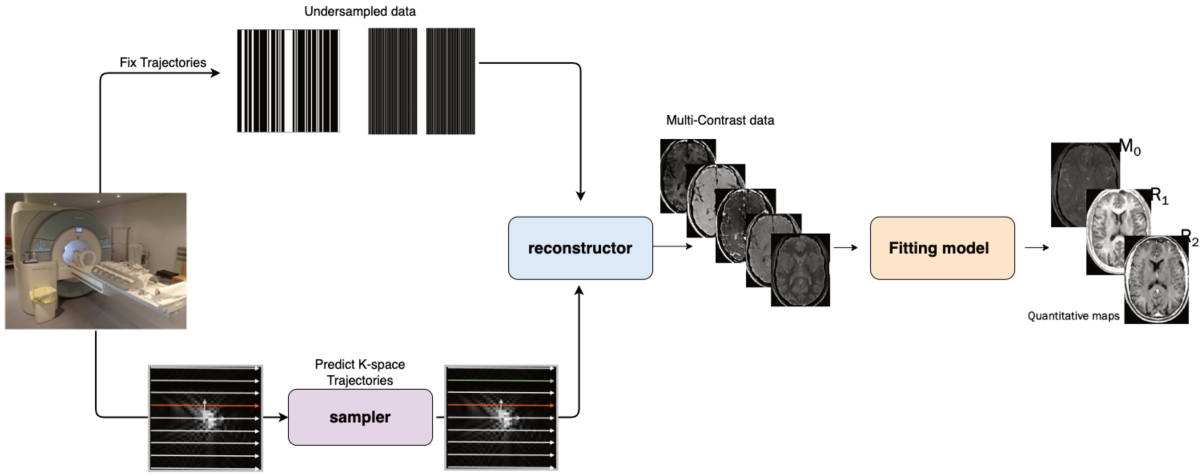

Figure 2.6: Three main research axes: a) Train a reconstructor that can enhance low-resolution images to high-resolution images through fixed sampling; b) sampler: to determine the best sampling pattern; and c) Fitting models: to create a suitable fitting model for MRI

There has been a growing trend towards end-to-end learning in MRI reconstruction, which involves training both the sampler and the reconstructor simultaneously. Pre-training a reconstructor followed by fine-tuning the sampler is another approach that has gained popularity. Another challenge in MRI reconstruction is obtaining quantitative measures after reconstructing weighted images or learning quantitative maps (qmaps) in an end-to-end manner. It remains unclear whether end-to-end approaches are more suitable, or whether learning the sampler, reconstructor, and fitting models individually is more effective.

The next section presents various categories of machine learning-based applications for MRI reconstruction, each with its unique strengths. We start first by briefly discussing the super-resolution literature on computer vision for natural and medical images. Then, we present the current state-of-the-art of machine learning in MRI where k-space is undersampled which covers the three categories of methods discussed above.

2.2.1 Image-based super-resolution

Super-resolution in natural images. The traditional methods of single-image super-resolution (SISR) can be realized through image interpolation such as nearest-neighbor interpolation, bilinear interpolation, bicubic interpolation, and Sinc and Lanczos resampling[39]. Nearest-neighbour interpolation is a simple and fast algorithm that selects the closest pixel value for each position to be interpolated, resulting in low-quality and blocky outputs. Bilinear interpolation, on the other hand, performs linear interpolation, resulting in faster processing time and higher-quality outputs than nearest-neighbour interpolation. Bicubic interpolation produces smoother outputs with fewer artifacts, but at a slower processing speed. However, these methods often have drawbacks such as increased computational complexity, amplified noise, and blurry results. Therefore, the current trend is to replace interpolation-based methods with learnable up-sampling layers. In particular, CNN with transposed convolution layers have been introduced in the field of SR to enable end-to-end learning of upsampling.

The SRCNN[74] and its faster variant, FSRCNN[75], were among the first studies to utilize multi-

layer CNNs to learn mapping from LR images to HR images in an end-to-end manner. Numerous deep networks, such as VDSR[149], DRCN[150], DRRN[293], EDSR[175], Residual Dense Networks (RDN)[363], and Deep Back Projection[111], have been proposed, using techniques such as gradient clipping, skip connections, and recursive supervision to facilitate the training of deep networks. MemNet[294] introduces memory blocks, whereas RDN introduces residual-in-residual connections. Lai et al.[165] introduced a Laplacian pyramid structure. Anwar et al.[10] presented a dense residual Laplacian module. The hybrid attention-based model (HAT)[48] achieved state-of-the-art results in the majority of super-resolution task benchmarks by incorporating an attention mechanism instead of a CNN to effectively capture long-range dependencies. However, transformers face significant computational and memory challenges due to the substantially larger size of image inputs compared to text inputs.

Ledig et al.[170] introduced SRGAN, GAN into SISR, making the way for its widespread adoption. However, training a GAN can be unstable due to the lack of convergence guarantees. Gulrajani et al.[101] addressed this issue by introducing Wasserstein Generative Adversarial Networks with Gradient Penalty (WGAN-GP). Unlike Wasserstein Generative Adversarial Networks (WGAN)[12], which uses weight clipping to enforce 1-Lipschitz continuity on the critic (discriminator), WGAN-GP uses a gradient penalty to enforce the Lipschitz constraint. Isola et al.[131] introduced Patch-GAN, a fully convolutional discriminator network with WGAN-GP. Wang et al.[322] suggested a hybrid approach that combines residual-in-residual and dense network with GAN. Autoregressive Models, normalizing flow, and VAE have been applied to super-resolution tasks, which suffer from the generation of low-quality images compared to GANs. Although powerful, GANs require careful regularization and optimization strategies to address issues, such as instability and mode collapse. Recently, DMs have shown significant advancements in image quality, surpassing GANs by using models such as Dall-E[251] and Stable Diffusion[263]. DMs, such as SR3[265] and SRDiff[171], have achieved high-quality results in the pixel domain; however, their iterative nature and high-dimensional calculations in the RGB space make them computationally intensive. To reduce the computational demands, Latent Diffusion Models (LDM)[263] performs diffusion processes in a low-dimensional latent space, allowing the reuse of latent representations across multiple models. LDM follow a two-stage process, with a pre-trained autoencoder compressing input data to a low-dimensional latent representation and a time-conditional U-Net model trained in the latent space incorporating image-specific biases. For image SR, LR images were concatenated as a condition in the second stage. By contrast, LSGM applies SGM within the latent space and jointly trains SGM with VAE. SRDDGAN[337] achieve faster sampling, sample diversity, and training stability by combining denoising diffusion models with a multimodal cGANs to generate images conditionally.

Superresolution in medical imaging. Unlike 2D natural images, medical imaging modalities, such as MRI, are typically stored in 3D volumes. Traditional MRI image super-resolution methods include bicubic, low-rank, total-variation regularization (LRTV)[277], nonlocal mean upsampling (NMU)[202], similarity of sparse representation[324], and sparsity priors[357]. Several solutions using 2D and 3D CNNs have been proposed for medical applications SISR. Zeng et al.[354] proposed a method for simultaneously estimating single- and multi-contrast images using CNNs. Shi et al.[278] extended SRCNN to solve a 2D MRI SISR problem using local residual blocks and global residual networks. Zhao et al.[371] also extended SRCNN, with a network consisting of three main sub-networks: feature extraction, non-linear mapping, and reconstruction. Liu et al.[179] proposed a multi-scale fusion convolution network (MFCN) that includes several multiscale fusion units (MFU) to help reconstruct SR MR images using different scale features. Sanchez et al.[266] adapted the standard

super-resolution generative adversarial network (SRGAN)[170] framework for 3D brain SISR. Zhao et al.[367] developed a 3D residual network (EDSR)[175] for axial slices in T2 brain images. Chen et al.[49] proposed a multi-layer DenseNet[126] based network for fast and efficient inference in 3D settings and WGAN-GP training for realistic texture recovery on T1W images. Due to the limited memory capacity of modern GPUs, it is often not practical to process an entire 3D volume at full resolution. Instead, images are typically processed patch by patch, [231]. A significant challenge for medical images is the development of a memory-efficient generator that can effectively process patches with a sufficient spatial context while maintaining a reasonable depth. Qu et al. proposed a patch-based nonlocal operator (PANO) to model the sparse representation of similar image patches[246]. Wang et al.[318] proposed a patch-based GAN method for 3D brain MRI super-resolution. Wu et al.[334] apply DDPM on T1- and diffusion-weighted brain MRI dataset. InverseSR[320] extends LDM to training on 3D brain T1W images. Chung et al.[52] applies SGM for denoising and SR using regularized reverse diffusion for MR images.

2.2.2 K-space-based super-resolution

There has been a promising development in image-space-based SR. However, the k-space-based SR has the advantage of working directly with raw data. Akçakaya et al.[6] emphasize that the k-space is directly related to the final image properties, such as resolution and field of view. Thus, k-space-based methods can effectively use spatial frequency information and correlations in k-space to enhance the image quality. In this section, we describe the mathematical formulation of accelerated MRI along with state-of-the-art methods in detail.

Mathematical Formulation of Accelerated MRI

In accelerated MRI, the entire k-space is not acquired, and in practice, multiple phase-encoding lines are skipped, resulting in a proportional reduction in scanning time. This involves undersampling k-space data using an undersampling mask (M_u). The forward model can be expressed as

$$y = M_u \odot Fx + n \tag{2.14}$$

where $\odot$ is the element-wise product, y represents the acquired data, M_u is the undersampling binary mask, F is the Fourier transform, x is the underlying image, and n is the measurement noise.

For multi-coil systems, in which each coil is sensitive to different parts of the imaged volume, the forward model is extended as follows:

$$y_i = M_u \odot FC_i x + n_i, \quad i \in \{1, \ldots, nc\} \tag{2.15}$$

where y_i represents the acquired data for the ith coil, C_i is a diagonal matrix representing the sensitivity of the i_th coil, and n_i is the measurement noise specific to the i_th coil.

The system can be represented compactly by concatenating the measurements across all the coils:

$$y = Ex + n \tag{2.16}$$

where y is the undersampled measurement vector and E is the forward encoding operator that combines the effects of undersampling, Fourier transform, and coil sensitivities across all coils.

The general inverse problem for accelerated MRI is formulated as an optimization problem as follows:

$$\text{argmin}_x \left\| y - Ex \right\|_2^2 + \lambda * R(x) \tag{2.17}$$

Here, $\left\| y - Ex \right\|_2^2$ ensures data consistency and the reconstructed image aligns with the acquired data. The term $R(x)$ serves as a regularizer that includes additional information or constraints into the solution to improve the stability and quality of the reconstruction.

The mask operator, represented by M_u, can be either Cartesian or radial; however, this study solely focuses on the former. Cartesian undersampling typically involves obtaining reduced k-space data points through either randomized or regular patterns. Adaptive sampling strategies enable the dynamic learning of a mask based on real-time feedback during the scan. Multi-echo sampling involves acquiring multiple echoes from a single radiofrequency pulse, with each echo representing a distinct point in time during the decay of transverse magnetization.

2.2.2.1 MR Reconstruction using fixed sampling

In this section, we discuss three primary strategies for MR reconstruction: database-free, database-driven, and hybrid methods. The database-free, also known as patient- or scan-specific learning, focuses on training the reconstruction model on data obtained during the scan without external datasets. This approach is patient-centric and does not rely on external databases. In contrast, the database-driven uses large datasets with diverse anatomical variations, acquisition protocols, and imaging conditions to train neural networks and learn representations and patterns in data. This enables the models to generalize well in various imaging scenarios. The hybrid approach combines the power of representation learning during the training phase with the on-the-fly adaptation of features during the test phase. This approach is particularly useful for scenarios where the training data is limited or the imaging conditions are unknown. These methods can also be broadly categorized into two main groups based on their sampling patterns: regular and pseudorandom methods. PI-based methods involve a uniform sampling pattern in the k-space. The k-space is sampled on a regular grid, forming a Cartesian coordinate system and hence the name. Compressive sensing-like methods use pseudorandom sampling. Pseudo-random sampling involves sampling the k-space in a non-uniform, random manner. This randomness helps spread aliasing artifacts throughout the image, resulting in a noise-like appearance rather than coherent aliasing. This noise-like artifact can be reduced using sparsity constraints in the reconstruction process. In practice, a pseudo-random sampling pattern is often generated using a variable density sampling scheme, where more points are sampled in the centre of the k-space (where the most important image information is located) and fewer points are sampled in the periphery.

2.2.2.1.1 Scan-Specific MR reconstruction

Scan-specific reconstruction models do not rely on databases and are adapted to the unique characteristics of each scan. Model training was launched every time new data were available. These models may incorporate specific scan information, such as coil sensitivity, and use variations in the signals between different receiver coils. Autocalibration signals are often used as training data

that is obtained from fast scans separately or together with undersampled data. Prior knowledge of a patient's anatomy may also be considered[43], although this is not within the scope of this study's framework. Here, we focus on methods that make no assumptions regarding the anatomy of the object scanned in the scan-specific methods. However, in practice, obtaining large amounts of representative training data is often challenging. Furthermore, concerns regarding generalization arise because of discrepancies in image contrast, sampling pattern, SNR, vendor, and anatomy between the training and test datasets can hinder the performance[20, 152, 217]. In this section, we focus on traditional reconstruction techniques, including PI in k-space and image-space, as well as compressed sensing, and discuss other advances in scan-specific reconstruction methods.

PI-based MR reconstruction: PI-based methods, such as GRAPPA and SENSE, rely on the spatial sensitivity profiles of the receiver coils positioned near the area of interest. This PI-based reconstruction can be roughly categorized into image-space and k-space. SENSE is an image space-based PI method and requires explicit knowledge of the sensitivity profiles of the receiver coils used for image acquisition. This information is typically obtained from separate low-resolution calibration scans. SENSE then reconstructs the aliased images by solving an unfolding matrix equation in the spatial domain. Each pixel value in the image is the weighted sum of the contributions from the same pixel in the receiver coil images, with the weights determined by the coil sensitivity profiles. Each coil generates an individual spatially aliased image from undersampled k-space data. The aliased pixels in these images are then unfolded based on the coil sensitivity profiles, resulting in a full field-of-view image.

GRAPPA[98] is a k-space-based PI method and is a linear model that estimates a missing value at a position in a grid by a combination of values in neighbouring locations. It can be used in both 2D and 3D in routine clinical scans[106]. The weights for GRAPPA are determined from a fully sampled, auto-calibration region of k-space located at the centre of k-space, which contains low-frequency image information. These weights are then applied to the undersampled k-space data to interpolate missing lines and reconstruct a full k-space. This k-space can then be transformed into an image using the Fourier transform. In the following, the formalism of GRAPPA in the context of neural networks will be covered. We first discuss the general reconstruction task in k-space. In Cartesian k-space-based PI, the task to be solved by reconstruction models can be summarized as an up-sampling of the data that have been under-sampled during the acquisition. Considering a coil array with N_c coils, full k-spaces are to be reconstructed from k-spaces under-sampled with a structured pattern. Here we focus on uniform undersampling. Moreover, we consider scan-specific reconstruction meaning that the model is learned from ACS and applied to every specific acquisition. In 3D, a *full k-space*, S_f, is a 4D tensor of complex numbers defined as:

$$\begin{aligned} S_f = \big\{ & S(i,j,k,c) \in \mathbb{C}, \\ & i,j,k,c \in [\![1,N_x]\!] \times [\![1,N_y]\!] \times [\![1,N_z]\!] \times [\![1,N_c]\!] \big\} \end{aligned} \tag{2.18}$$

where N_x stands for the number of sampled points along the readout direction, N_y is the number of sampled points in the first phase encoding direction, and N_z is the number of sampled points in the second phase encoding direction (note that $N_z = 1$ in the 2D case). A single complex-valued image I may be computed from the full k-space tensor S_f through Fourier-transforms for each coil and applying a coil combination algorithm such as *adaptive combination* [197, 313].

In the case of uniform undersampling that we consider here, the readout direction (x) is fully sampled while the first and second phase encoding dimensions (y and z) are undersampled. Assuming an

acceleration rate equal to $R_y \in \mathbb{N}$ in the y dimension and $R_z \in \mathbb{N}$ in the z dimension an acquired *undersampled k-space* may then be defined as the following 4D tensor:

$$\begin{aligned} S_u = \{&S(i, 1+(j-1)\times R_y, 1+(k-1)\times R_z, c) \in \mathbb{C}, \\ &i,j,k,coil \in [\![1,N_x]\!] \times [\![1,N_y/R_y]\!] \times [\![1,N_z/R_z]\!] \times [\![1,N_c]\!]\} \end{aligned} \tag{2.19}$$

Here, one line is actually acquired every $R_y \times R_z$ lines so that acquiring such an *undersampled k-space* fastens the *full k-space* acquisition time (which may typically last minutes) by a factor R=$R_y \times R_z$.

Generally speaking, a reconstruction model operating in k-space takes an *undersampled k-space* tensor S_u as input and outputs its *full k-space* approximation $\hat{S}_f$:

$$\hat{S}_f = f_\theta(S_u) \tag{2.20}$$

Approximating S_f, i.e. computing $\hat{S}_f$, often decomposed into (R-1) (independent) subtasks where each subtask corresponds to the approximation of an undersampled view of S_f. Let note $\hat{S}_{f,m,n}$ (with $m,n \in [\![1,N_y/R_y]\!] \times [\![1,N_z/R_z]\!]$) the following approximated *undersampled k-space* tensor:

$$\begin{aligned} \hat{S}_{f,m,n} = \{&\hat{S}(i, m+(j-1)\times R_y, n+(k-1)\times R_z, c) \in \mathbb{C}, \\ &i,j,k,coil \in [\![1,N_x]\!] \times [\![1,R_y]\!] \times [\![1,R_z]\!] \times [\![1,N_c]\!]\} \end{aligned} \tag{2.21}$$

Each of the unacquired *undersampled k-spaces* ($\hat{S}_{f,m,n}$) is estimated with a specific reconstruction model $f_{\theta_{m,n}}$ from the acquired *undersampled k-space* according to:

$$\hat{S}_{f,m,n} = f_{\theta_{m,n}}(S_u) \tag{2.22}$$

As S_u is already acquired, it can readily be used as an estimate of $\hat{S}_{f,1,1}$ here (i.e. setting $f_{\theta_{1,1}}$ as the identity). The *full k-space* reconstruction $\hat{S}_f$ is achieved by combining the $(R_y \times R_z - 1)$ estimated *undersampled k-spaces* $(\hat{S}_{f,m,n})_{m,n}$ by the reconstruction models and the acquired undersample k-space S_u (i.e. $\hat{S}_{f,1,1}$). We make the choice here to use the same architecture for each model $f_{\theta_{m,n}}$ and to train each of them independently.

However, in this DL-MRI formalism, the GRAPPA model network architecture may be viewed as a one-layer neural network that consists of a single complex-valued convolutional layer without bias (see Figure 2.7-(b)). Its architecture is defined by the size of the convolution filter, $(n_{x,g}, n_{y,g})$ in 2D, and $(n_{x,g}, n_{y,g}, n_{z,g})$ in 3D. As stated above $R-1$ GRAPPA models are learnt for each one to infer a particular set of missing lines. The weights for GRAPPA are determined from a fully sampled, auto-calibration region of k-space located at the centre of k-space, which contains essential low-frequency image information. These weights are then applied to the undersampled k-space data to interpolate missing lines and reconstruct a full k-space. This k-space can then be transformed into an image using the Fourier transform. While there exist several ways to solve this linear problem, a regularized strategy is often required when the ACS size is small as compared to the kernel size. This regularization is implemented here in an iterative DL-MRI framework using TVP as well as using Tikhonov regularization[244] that will be respectively referred to as GRAPPA(TVP) and GRAPPA(TR).

The Robust artificial-neural-networks for k-space interpolation (RAKI) model is a slightly deeper CNN whose architecture consists of a few (three) real-valued convolutional layers without bias, each

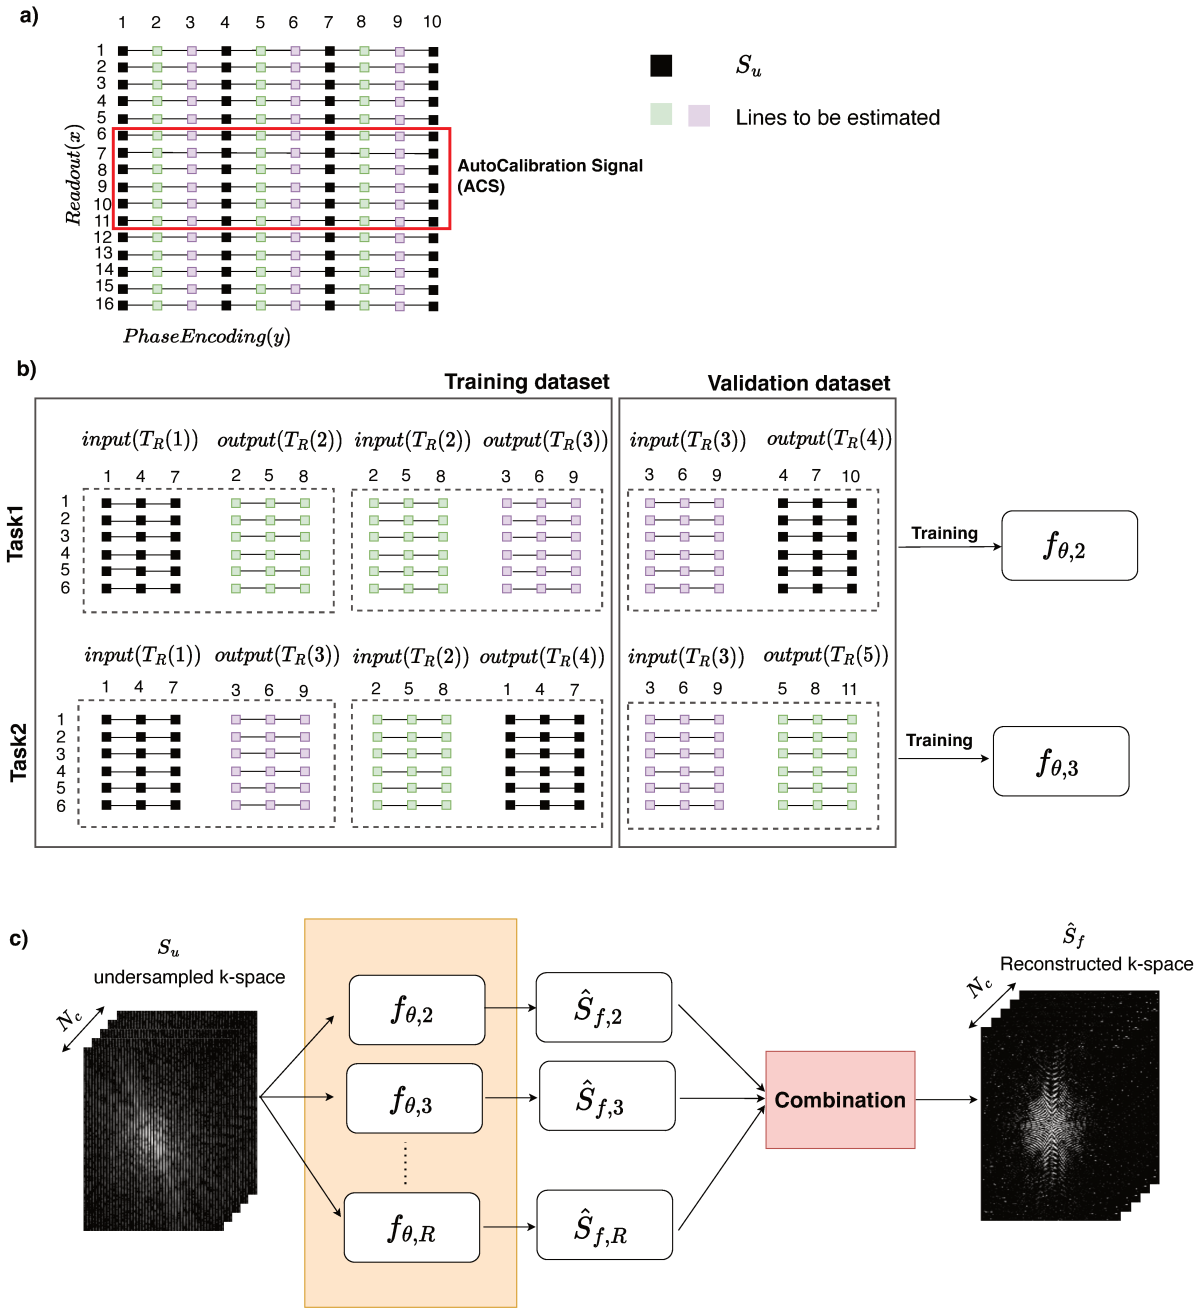

Figure 2.7: Illustration of training strategy in the 2D case for an acceleration factor $R_y = 3$. (a) The acquired *undersampled k-space* includes lines numbered 1, 4, 7... while lines 2, 3, 5, 6, 8, 9... are unobserved. The learning is performed on a small part of the *k-space* called the ACS, which is fully acquired. (b) The fully acquired ACS may be artificially undersampled to build supervised datasets for learning reconstruction models. There are $R-1$ reconstruction tasks for an acceleration factor R, and one reconstruction model is learnt on each one. To select the best model architecture one uses a train/validation split. (c) After the training is performed for each task, the estimations of each model on a new *undersampled k-space* are combined into an estimated *full k-space*.

of which is followed by a non-linear activation such as the Rectified Linear Unit (ReLU)[110]. It is implemented in a way where convolutional layers handle the complex nature of the data with a concatenation of the real and imaginary parts of the input k-space tensor (see Figure2.8). It has been reported for 2D imaging, with either a coil-by-coil implementation[5] or a line-by-line one[355]. We use the latter here, using complex convolution and complex activation functions and replacing the ReLU with LeakyReLU as it has been reported to improve stability[18] and can be used to scale the amount of non-linearity (we refer to this model as cRAKI in the following). This architecture is defined by a number of hyperparameters: in 2D, the size of the convolutional filters of the first layer $(n_{x,1}, n_{y,1})$ and of the second $(n_{x,2}, n_{y,2})$ and third layers $(n_{x,3}, n_{y,3})$, the number of maps in the two first convolutional layers $(n_{d,1}, n_{d,2})$ (extension in 3D also used here is straightforward).

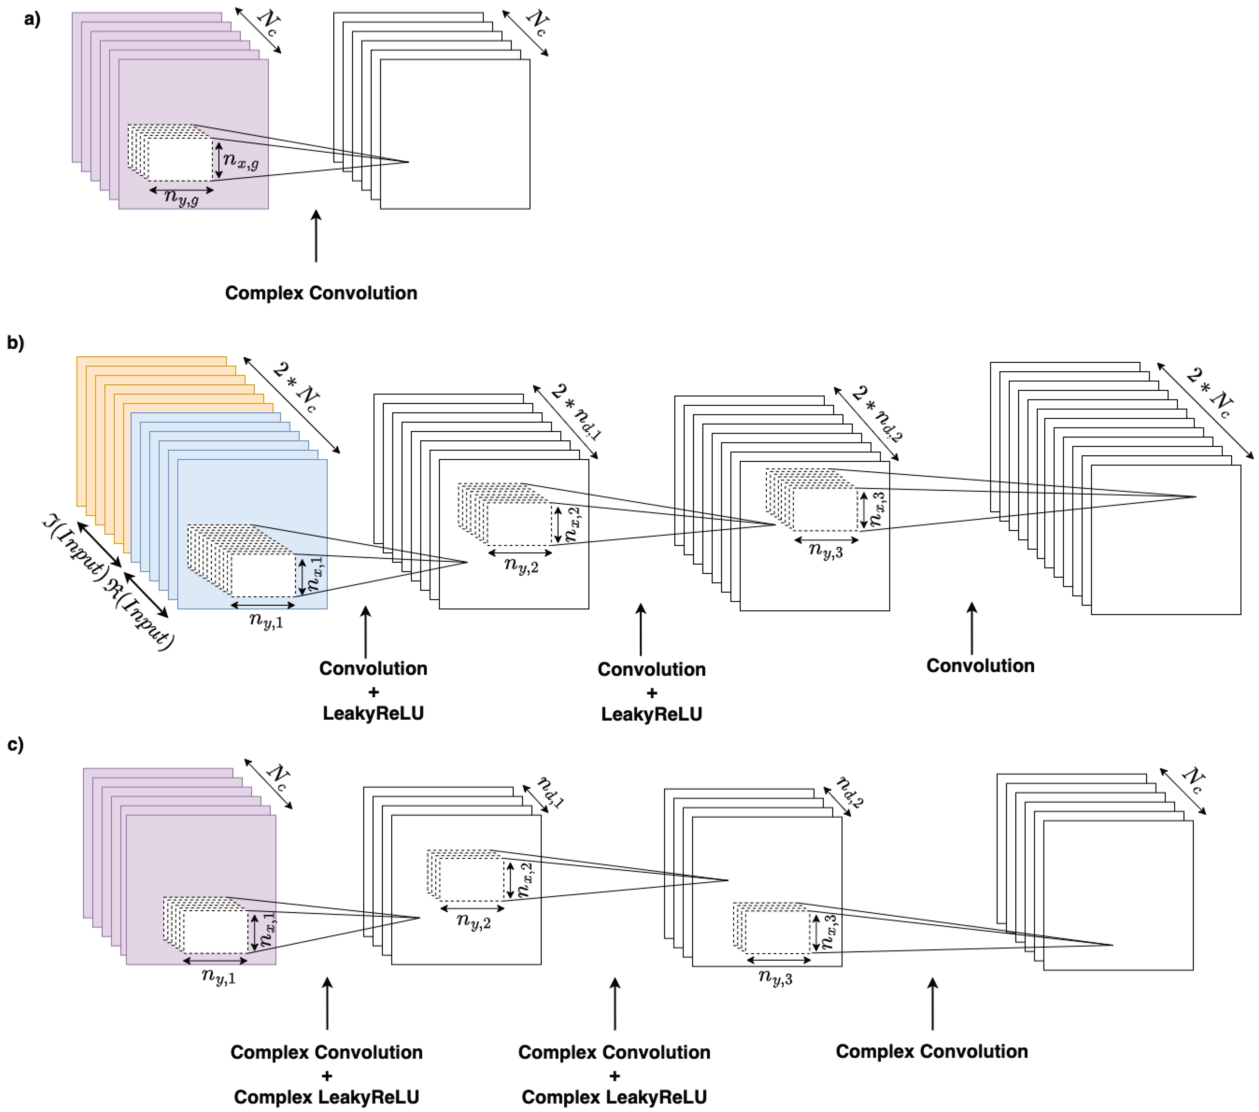

Figure 2.8: Illustration of the a) GRAPPA, b) RAKI, and c) cRAKI models.

Another scan-specific k-space-based approach SPARK[11] involves training a CNN to estimate and correct k-space errors made by an input reconstruction method such as GRAPPA using ACS. First, the difference is calculated between the ACS and reconstructed ACS for each coil. Subsequently,

the reconstructed k-space is inputted across all coils to calculate a correction term, which is then compared to the ACS difference using least-squares loss. To adjust the model parameters, backpropagation can be applied. To apply the corrections, the trained model estimates a correction term using all reconstructed k-space as input and adds the resultant correction to each coil. By repeating this process for all coils, the procedure enables the correction of k-space artifacts in a scan-specific manner for a broad range of reconstruction and acquisition settings.

Compressive Sensing based Reconstruction: In the typical Compressive Sensing (CS) formulation, the reconstruction task is to recover an N-pixel image $x \in \mathbb{C}^N$ from a noisy dataset $y \in \mathbb{C}^M$ obtained via a known linear transform $A \in \mathbb{C}^{M \times N}$. The formation of the dataset is encapsulated by the equation $y = Ax + \eta$, where $\eta \in \mathbb{C}^M$ is additive white Gaussian noise with variance τ^2. The acceleration rate $R = \frac{N}{M} > 1$ poses many solutions for the image x from y. Under this assumption, a standard compressive sensing framework is mathematically expressed as:

$$x := \arg\min_x \left\{ \|Ax - y\|_2^2 + \lambda R(x) \right\} \tag{2.23}$$

In this equation, $R(x)$ is a sparsity-induced regularization term to impose specific constraints in x. Generally, handcrafted regularizers are applied, such as the wavelet transform or total variation. Common regularizers include Total Variation (TV), wavelet, and sparsity constraints. The optimization problem can be solved using iterative algorithms such as the conjugate gradient method, proximal gradient method, and ADMM. The objective is to obtain the optimal solution for v which minimizes the combined effect of the data fidelity term $\|Ax - y\|_2^2$ and the regularization term $R(x)$. The term "data fidelity" refers to the degree of agreement between the reconstructed measurements Ax and the acquired measurements y. Meanwhile, the regularization term encourages sparsity or other preferred features in x. To recover the original image y, three conditions should be specified: incoherent undersampling, sparsifying transformation, and iterative reconstruction process.

To improve image recovery, undersampling should be performed randomly or incoherently, with only noise-like artifacts visible. However, extreme randomness must be avoided because the central region of the k-space contains more important image information than the periphery. Current compressive sensing methods use pseudo-random schemes that preferentially sample near the k-space centre. Numerous algorithms, including gradient-based methods, proximal algorithms, and iterative optimization techniques, can be used to solve the optimization problem in Equation2.23. These algorithms iteratively refine the estimation of x to obtain a solution that satisfies the measurement equation and combines the properties or structures imposed by the regularization term.

Self-consistent Parallel Imaging Reconstruction (SPIRiT): SPIRIT[193] is a generalization of GRAPPA centres around estimating each missing points in k-space by utilizing all neighbour points, irrespective of whether they've been acquired. SPIRIT takes an undersampled k-space $\tilde{y}$, which can follow any arbitrary trajectory (Cartesian or non-Cartesian), and describes the reconstruction as an inverse problem governed by data consistency constraints. The key in the approach is to separate the consistency constraints into (a) consistency with the calibration, and (b) consistency with the data acquisition. The calibration consistency constraint can simply be expressed as:

$$x = Gx \tag{2.24}$$

G denotes a set of convolution kernels. These kernels attempt to learn an identity function to

represent itself and are estimated using ACS.

The second constraint, data consistency, imposes reconstruction has to be consistent with the real data that were acquired by the scanner. This is given by:

$$\tilde{y} = Dy \tag{2.25}$$

In this equation, $\tilde{y}$ refers to the acquired k-space data with non-acquired positions zero-filled across all coils, while D denotes the linear operator conducting undersampling and zero-filling on non-acquired data points. In Cartesian sampling, the operator D selects only acquired k-space locations. The selection can be arbitrary: uniform, variable density, or pseudorandom patterns. In non-Cartesian sampling, the operator D represents an interpolation matrix. It interpolates data from a Cartesian k-space grid onto non-Cartesian k-space locations in which the data were acquired.

SPIRIT approach addresses the following optimization problem, often referred to as ℓ_1-SPIRiT reconstruction:

$$\underset{\hat{y}}{\operatorname{argmin}} \|D\hat{y} - \tilde{y}\|^2 + \lambda_1 \left\|(G-I)\hat{y}\right\|^2 + \lambda_2 \left\|\Psi\left(F^{-1}(x)\right)\right\|_1 \tag{2.26}$$

Here, λ_1 regulates the balance between the two consistency constraints, λ_2 controls the sparsity of the solution, and Ψ signifies the wavelet transform. This equation is solved iteratively using conjugate gradient descent[225] or projection onto convex sets.

Other non-linear variants of spirit have been introduced. SPIRiT-RAKI[122], utilizes scan-specific CNN to enforce nonlinearly through self-consistency. Additionally, regularization terms can also be incorporated. SPIRIT can be applied to both the image- and k-space domain. In the image domain, the reconstruction problem is formulated as a linear inverse problem, where the image is reconstructed from the coil images. The image domain SPIRIT is particularly useful for non-Cartesian acquisitions, where the k-space trajectory is not uniform. Inspired by the success of SPIRIT, several other methods have been proposed, such as iGRAPPA[370] and iRAKI[66] enforces calibration consistency and data consistency iteratively.

Other Advances in Scan-Specific MR Reconstruction: Ulyanov et al. [300] demonstrated that an untrained CNN can provide a strong prior for image reconstruction even without training data. Specifically, given a single measurement instance y to an untrained CNN, the Deep Image Prior (DIP) method involves minimizing a function that combines a fixed, randomly sampled vector z with randomly initialized CNN parameters. The network can generate a high-quality reconstructed image without supervision or training data using an iterative optimization algorithm to fit the weights to the observations (data consistency). This approach was successfully applied to MRI reconstruction by Darestani et al.[62] and Yoo et al.[351]. However, a limitation of DIP is that its flexibility comes at the cost of increased inference time because the network weights must be retrained for each image. Overfitting can be a concern, and heuristics such as early stopping have been used to mitigate this issue[221, 290, 317]. Yaman et al.[342] proposed a zero-shot self-supervised learning (ZS-SSL) approach, which performs scan-specific training. ZS-SSL partitions the acquired measurements into three types of disjoint sets, which are used to define the training loss and to establish an early stopping strategy to avoid overfitting. Thus, the training is self-supervised and self-validated. Feng et al. [86] applies Implicit Neural Representation (INR)[213] to scan-specific MRI reconstructions. INR can learn internal redundancies and correlations within an object due to its continuous representation

function. In addition, it allows for the integration of external prior and constraints by architecture.

Reconstruction using Low-rank Methods. Another category based on the structured low-rank matrix prior has shown promising results. For instance, Shin et al.[279] proposed a calibrationless PI reconstruction method named SAKE, which formulated reconstruction as a structured low-rank matrix completion problem. Jin et al.[137] introduced an annihilating filter-based low-rank Hankel matrix approach (ALOHA)[139] as a general framework for sparsity-driven k-space interpolation, which combined PI with CS. Inspired by a recent mathematical discovery of data-driven framelets linking convolutional neural networks to Hankel matrix decomposition, Han et al.[109] developed a fully data-driven deep learning algorithm for k-space interpolation. Subsequently, Zhang et al.[362] introduced an approach named STDLR-SPIRiT, which combined the simultaneous two-directional low-rankness (STDLR) with the SPIRIT to mine the data correlation from multiple receiver coils. Meanwhile, Haldar et al. developed a flexible framework for constrained image reconstruction that used low-rank matrix modelling of local k-space neighbourhoods (LORAKS)[102]. They subsequently introduced P-LORAKS[104], which extended LORAKS to PI. Chen et al.[47] proposed a locally structured low-rank image reconstruction method by imposing low-rank constraints on submatrices of the Hankel structured k-space data matrix. The Hankel-k-space generative model (HKGM)[229] one-shot learning framework, with a prior learning stage, constructs a large Hankel matrix and extracts structured k-space patches from it to capture the internal distribution among different patches. This enables the generative model to learn from redundant and low-rank data spaces.

2.2.2.1.2 Database-driven MR reconstruction

The scan-specific methods, despite reduced acquisition time, suffer from limitations in speed, data specificity, and reconstruction quality. They require numerous iterations and costly operators, especially for high-resolution 3D images. Sparsity constraints such as total variation or wavelet transform are insufficient to address the complexity of MR images, making achieving high-resolution and clinically acceptable reconstruction quality difficult. In addition, tuning these methods to achieve high reconstruction accuracy on multiple samples with a fixed set of hyperparameters is challenging. To overcome these limitations, DL-based methods have been developed that offer the potential for faster inference and possible acceleration in the acquisition time. These DL-based methods rely on large datasets to extract useful representations. Dishner et al. [72] surveyed the major public MRI datasets in different subfields of radiology, including brain, body, and musculoskeletal, and listed the most important features. In this section, we focus on MR reconstruction based on fixed sampling.

Fully-supervised MRI reconstruction. CNN-based reconstruction methods[242, 271, 321] employ a supervised learning framework to establish a mapping from the undersampled k-space/image to their fully sampled equivalent k-space/images with a fixed undersampling pattern. During testing, this network was fed an under-sampled image/k-space. Despite its potential, this feed-forward network lacks the capability to ensure data consistency during the testing phase. Unrolled neural networks[108, 217, 252, 285] ensure data consistency using unrolling optimization algorithms, such as gradient descent or proximal-dual optimization. Each iteration of the algorithm, represented by a function parameterized by θ_i, is mapped to a single layer of the network. A deep network is formed by stacking a finite number of these layers. When data are fed forward through this network, it is equivalent to executing the iteration L times. The parameters θ_i can be shared across different layers or vary from layer to layer, and they are learned by training the network end-to-end to optimize the performance, which can be viewed as an Recurrent Neural Network (RNN)-like. Varnet[108]

unrolled gradient descent optimization. The filter kernels and activation functions, along with the data term weights, are learned during an offline training process. Aggarwal et al.[3] unroll a conjugate gradients algorithm for single-coil setting. Ramzi et al.[252] provided a benchmark for different unrolled neural networks for two datasets: the fastMRI and OASIS datasets for a single-coil setting. Cross-domain CNNs, such as KIKI-net[82], TEID-Net[352], V-Net, and K-Net[187], comprise three components: a deep CNN operating on the k-space, a deep CNN operating on the image domain, and interleaved data consistency operations. These components were alternated, and each CNN was trained to minimize the loss between the reconstructed and corresponding fully sampled k-spaces. XPDNet[253] unroll the Primal-Dual Hybrid Gradient (PDHG) for a multi-coil setting. It consists of two main components: a denoiser block and a sensitivity map refiner. The denoiser block is constructed using a multiscale wavelet CNN which reduces noise interference in the reconstructed image. The sensitivity map refiner, based on the UNET model, fine-tunes the spatial encoding. E2E-VarNet[285] is an end-to-end framework that simultaneously learns coil sensitivities and maps by processing undersampled k-space data for multi-coil settings. It contains multiple cascades, followed by an inverse Fourier transform and the root sum of a square-coil combined image. The Data Consistency module calculates a correction map to align the intermediate k-space values with the measured k-space values. The Refinement module converts multi-coil k-space data into a single image, applies a U-Net, and then transforms it back into multi-coil k-space data. The Sensitivity Map Estimation module, which is parameterized by a CNN, estimates the sensitivity maps used in the refinement module. Another class of algorithms consists of Plug and Play (PnP)[310] algorithms, which form a non-convex framework that integrates ADMM[45] or other proximal algorithms with advanced denoiser priors. Several studies[4] have emerged from this approach, which can be broadly categorized into four main areas: proximal algorithms[45], imaging applications, denoiser priors[146, 348], and convergence[40, 255]. The PnP framework is rendered tuning-free through reinforcement-learning adaptation[328]. Vision Transformer (ViT)[176] has been utilized for accelerated MRI reconstruction, and although it is shown to be computationally faster than U-Net, its reconstruction accuracy is only slightly better. The main advantage of transformers is their ability to capture long-range dependencies across the entire image, which is useful for situations in which redundant information is distributed throughout the image and dependencies are not limited to local areas. Recently, a new transformer-based method called FRestormer[64] has achieved competitive performance on the fastMRI leaderboard for the brain track by utilizing iterative refinement with neural networks. it addresses the memory inefficiency of transformers by utilizing parameter-free Fourier-based attention modules, resulting in an improved memory efficiency of up to 2x.

Self-supervised/Unsupervised MRI reconstruction. Self-supervised reconstruction approaches that utilize undersampled k-space data for supervision have been introduced for single-contrast MRI reconstructions[70]. HQS-Net[314] presents a neural-network-based approach that decouples the minimization of the data consistency term and regularization term, enabling network training that relies solely on undersampled measurements. Yaman et al.[342] proposed a physics-guided self-supervised learning method that trains the deep cascade CNN by partitioning each undersampled k-space data into a training and validation partition for a large dataset, an approach similar to that discussed in the previous section by ZS-SSL. Korkmaz et al.[160] also proposed a self-supervised transformer GAN for zero-shot instance-specific optimization. Zhou et al.[374] devised a triple-branch-based dual-domain self-supervised reconstruction framework to achieve a promising performance for single-contrast low-field MRI. Cole et al.[55] use cGAN as a reconstruction model, which directly learns a mapping from the input undersampled k-space to the output reconstructed image.

SSDiffRecon[159] is DM that uses cross-attention transformers for reverse diffusion steps and data-consistency blocks. it is trained using self-supervision with only undersampled k-space data. Several DL-based fast multi-contrast MRI reconstruction methods exist. Xiang et al. used fully sampled T1W images as an additional CNN input to T2W reconstructions[335]. Dar et al.[60] added adversarial learning and perceptual loss to further enhance the performance. Recently, DSFormer[373] used a Deep Conditional Cascade Transformer (DCCT) using cascaded Swin Transformer Reconstruction Networks (SwinRN) for multi-contrast settings. The DCCT also utilises a dual-domain self-supervised learning strategy (image and k-space) to reduce the cost of acquiring fully sampled training data.

2.2.2.1.3 Hybrid MR reconstruction

Neural networks are known to experience a drop in performance when applied to images from a distribution that is different from that used during training. For instance, a model trained to reconstruct knee images may not perform well in reconstructing brain images. However, the same network trained on brain images demonstrates perfect performance in reconstructing the brain[61, 140, 152]. Consequently, a distribution-shift performance gap exists for a given neural network. Several studies have proposed approaches for closing this distributional shift using test-time training. Darestani et al.[63] suggested a self-supervised domain-adaptation method for deep-learning-based compressive sensing. This method involves training deep-learning models with self-supervision during the training stage and performing self-supervised test-time training for each new test sample at inference. Yaman at al.[341] extend ZS-SSL to test-time training framework by leveraging transfer learning in cases when a database-pretrained model is available. Recall that CS involves estimating an unknown signal, x, from a fewer set of measurements y. The objective is to determine the "most compressible" signal that fits the available measurements. To formalize the concept of compressibility, CS uses sparsity on some basis, such as the wavelet basis for MRI images. Compressed Sensing using Generative Models (CSGM)[33] proposed the use of generative models, G, instead of sparsity priors. Ideally, when random variable z is drawn from a standard normal distribution, $G(z)$ should follow the desired distribution. In the original CSGM approach, GANs were used; however, because of GANs' limitations of learning the diversity of samples and unstable training, DMs are preferred. Jalal et al.[135] examines the problem of CSGM within the context of a multi-coil MRI setup. The score function is modelled using a deep generative neural network. Tezcan et al.[297] proposed an approach that uses VAE to learn the probability distribution of fully sampled MR images, which is then used as an explicit prior term in reconstruction. This decouples the encoding operation from the prior and results in a powerful image prior to compensating for missing k-space data without requiring paired datasets for training or being prone to the associated sensitivities. The distributional shift can also happen while incorporating parallel imaging in deep learning approaches for reconstruction due to differences in the number of receiver coils at various sites. This could impact the generalization of the model. Capturing the changing relationship between the views captured by each coil for each scan and the object being imaged is particularly difficult for a trained network. Therefore, GrappaNet integrated a scan-specific component into a neural network to improve the generalization of the reconstruction model by adapting it to differences in the sensitivity profile [286]. A GRAPPA layer that learns a scan-specific reconstruction function was jointly trained with the neural network on a fastMRI dataset. The scan-specific layer weights are estimated on ACS at inference while keeping the other parameters fixed.

2.2.2.2 Optimizing sampling patterns

The field of MRI sampling trajectory design has a long history dating back several decades. Notable works in this area include those by Cao et al.[42], Gao et al.[90], Seeger et al.[273], and Haldar et al.[103]. There has been a growing interest in improving MRI reconstruction by utilizing DL-based approaches incorporating formalism to learn the sampling pattern.[2, 3, 14, 15, 232]. These methods produce either non-sequential or sequential sampling patterns. Non-sequential learning generates a dataset-specific acquisition trajectory, whereas sequential policies are conditional on the initial k-space measurements, allowing for the generation of diverse sampling trajectories for each data point[15]. In this section, we focus on investigating DL-based models for optimizing sampling patterns, commonly referred to as active MRI acquisition. These models generally include a sampler and a reconstructor parameterized by a neural network. The learning paradigms for these models are based on supervised, reinforcement, or self-supervised learning, depending on the sampler and reconstructor. The sampler is responsible for generating the sampling pattern, while the reconstructor is responsible for reconstructing the image from the undersampled k-space data, which can be pretrained.

Jointly-trained non-sequential sampler and reconstructor. Several studies have shown that non-sequential policies outperform handcrafted sampling patterns in both single- and multi-coil acquisition settings [14, 128, 316, 329, 358, 376]. In contrast, adaptive policies have shown mixed results and have primarily been developed for single-coil settings[16, 96, 138, 268, 349, 365]. LOUPE[14] co-jointly trains a non-sequential probabilistic sampling pattern and unet as a reconstruction model for a single-coil setting. Extended by Zhang et al.[358] for multi-coil setting. DuDReTLU-net[120] is another end-to-end approach that simultaneously combines learning samplers and reconstructors. It utilizes a transformer-based network that operates in a dual-domain setting, incorporating both image and k-space. This framework extends the LOUPE to undersample dynamic cardiac scans and applies transformer-based networks to a novel dual-domain reconstruction setting. JoJoNet[368] and ASMCR[319] are end-to-end frameworks designed to accelerate multi-contrast setting. JoJoNet[368] consists of a sampling mask generator for each image contrast and a reconstructor that utilizes inter-contrast correlations, allowing information sharing. The sampling mask generator and reconstructor were trained jointly across multiple image contrasts, and the acceleration ratio for each contrast was learnable and could be driven by downstream task performance. Razumov et al.[257] propose pathology-aware sampling patterns based on a pathology's detectability or localizability in reconstructed images. The method uses a CNN to predict the pathology's detectability or localizability in the reconstructed image and uses this information to optimize the sampling pattern.

Jointly-trained sequential sampler and reconstructor. It is beneficial to develop a model capable of adjusting the sampling for each new k-space, because optimizing a single sampling scheme may not always be optimal on a case-by-case basis if the k-space instances show significant differences[96]. Moreover, it is natural to have a mask created sequentially, because MRI acquires k-space sequentially[365]. Jin et al.[138] utilized a Monte-Carlo-Tree-Search (MCTS) formulation, and Yin et al.[349] proposed an end-to-end differential sampler based on the straight-through estimator. Zhang et al.[365] proposed a GAN-based reconstructor, which comprises a cascaded fully connected ResNet and a data consistency layer, takes a zero-filled image as the input and produces a reconstructed image and an uncertainty map that highlights areas with potential errors through an evaluator. They learns the sampling mask through a stochastic sampling layer, which constituted weight parameters that were passed through a sigmoid activation. This produces a probabilistic mask that is then binarized using the Straight-Through Estimator (STE)[25]. During the forward pass, STE compares

the probabilistic mask to a uniform distribution and outputs a value of 1 (or 0) if the probability in the probabilistic mask is greater than the uniform probability. In the backpropagation process, STE returns the identity function. After generating the mask, it was applied to the k-space, and the zero-filled image was fed into the reconstructor. An evaluator, which consists of an embedding layer followed by a CNN, assigns a score to each associated k-space line, denoting whether it is observed. This process encouraged the reconstructor to produce high-quality images. During the inference process, the k-space line with the lowest score is selected for acquisition, and the process continues until a stopping criterion is satisfied. Both the acceleration factor and acquisition trajectory are sequential; however, the method does not choose the next k-space sample based on prior measurements. A different strategy was proposed by Yin et al.[349], which involved selecting the next k-space sample at each step of the process. The reconstructor used in this approach is a U-net model that uses the zero-filled image as the input. Bakker et al.[15] studied adaptive policies for clinically relevant multi-coil MRI setting. They proposed a CNN-based policy network with a reconstruction model and evaluated it on a large-scale multi-coil fastMRI knee dataset. The reconstructor is End-to-End Variational Network (E2E VarNet) which is based on an unrolled network proposed by Sriram et al.[285]. In terms of adaptability, the authors found that sigmoid feature normalization in the sampler resulted in non-adaptive behaviour but better reconstruction performance, whereas softplus normalization was more adaptable.

RL-based. Pineda et al.[232] reformulated the sampling pattern learning as a sequential decision problem as a Partially Observable Markov Decision Process (POMDP) and used a Double Deep Q-Network (DDQN) policy, which was followed by a pre-trained reconstructor. The adoption of Deep Reinforcement Learning (RL) has shown potential due to its capacity to adjust trajectories on-the-fly to patient-specific needs and execute long-horizon planning. In adaptive settings, the DDQN network is the same as the evaluator of [365], whereas in the nonadaptive case, it was a simple MLP. The DDQN learned to iteratively map the current reconstructed image and the current mask to an action, that is, the next column of k-space to acquire, and received an updated reconstruction from the reconstructor. The outputs of the DDQN were the q-values of the different k-space columns, and the next column to be acquired had the highest q-values. Reward was defined as SSIM which is an indicator of improvement in reconstruction quality. Similarly, Liu et al.[189] proposed a lightweight transformer for the DDQN and a U-net for reconstruction, achieving faster and better reconstruction than [365] and [232] on the fastMRI single-coil knee dataset[153].

Compressive Sensing-based. Learning-Based Compressed Sensing (LBCS) and its stochastic variants have shown great potential for addressing clinical challenges, such as multi-coil 3D and dynamic MRI[267]. LBCS uses a greedy, parameter-free approach for image reconstruction, employing an iterative methodology that involves sequentially adding elements to the mask. At each step, LBCS tests all feasible candidates while a budget is available and then fixes the candidate that yields the best performance gain for the sampling mask. This process is repeated until the budget is exhausted, with the goal of incrementally improving image reconstruction performance at each step. Furthermore, compared to RL-based[232], stochastic LBCS performed better and did not adapt or conduct long-term planning. Another approach that has gained significant attention is the use of cGAN for inverse problems[33, 145, 164, 274]. cGAN are used to model the posterior distribution in inverse problems and offer an innate criterion for adaptive sampling[268]. cGAN sequentially identify the location with the highest posterior variance in k-space without the need for explicit training on the task. This adaptive sampling policy is driven by the inherent uncertainty within the data, which optimizes the acquisition process. Another interesting approach is the use of adaptive sequences that enable the

real-time adjustment of sequence parameters. Beracha et al.[26] proposed a framework for adaptive acquisitions includes a signal model, Bayesian estimation, and Cramer-Rao bounds to adjust any sequence parameter dynamically based on incoming measurements. Ravula et al.[256] introduced DM into the CSGM framework to optimize the sampling pattern. This approach utilizes an extension of Tweedie's formula[81] for inverse problems[51], which contains the use of pre-trained generators without the need for further training.

2.2.2.3 ML for qMRI

Extensive research has been conducted to identify efficient methods to explore multiparametric spaces for qMRI. Traditional qMRI methods involve extracting parameters such as tissue relaxation rates from multiple measurements, which can be time-consuming and susceptible to system errors. To address these limitations, the advanced image reconstruction techniques discussed earlier, such as compressed sensing, low-rank processing, or DL-based, for accelerating mri can be combined into qMRI[174, 182, 212, 230, 243, 366]. In addition, advanced multiparametric mapping methods, including MRF[196], syMRI[94, 295], and EPTI[315], have been developed to simultaneously estimate various tissue relaxation properties and field maps from the same input images. However, these methods often require a specific pulse sequence. Recently, deep learning has been applied to accelerate data acquisition and post-processing in qMRI as well as in other medical imaging applications. Although deep learning has shown remarkable speedups, additional scans are necessary, presenting a challenge for the widespread adoption of qMRI in clinical practice given the data-hungry and black-box nature. In this section, we delve into various traditional and deep-learning-based methods for fitting.

2.2.2.3.1 Traditional methods

Classical methods to accelerated parametric mapping techniques is model-based reconstruction, which utilizes physical or physiological signal models as prior knowledge in an iterative reconstruction process. These techniques intrinsically estimate quantitative parameter maps by minimizing a predefined cost function. Classic approaches for retrieving quantitative parameters propose a voxel by voxel curve fitting via an estimator which is commonly Non Linear Least Squares (NLLS) algorithm. Condsider a simple relaxometry case, the signal decay due to relaxation processes is modelled as an exponential decay. For a given voxel, the signal intensity S at time t can be written as

$$S(t) = S_0 \cdot \exp(-t/T2) \tag{2.27}$$

where S_0 is the initial signal intensity, T is the T2 relaxation time. The goal of the exponential decay fitting is to estimate the parameters S_0 and $T2$ from the observed signal intensities at different times. This is typically performed using a least squares (LS) estimator, which minimizes the sum of the squared differences between the observed and predicted signal intensities. Mathematically, the LS estimator solves the following optimization problem:

$$\min_{S_0,T} \sum_{i=1}^{N} (S(t_i) - S_0 \cdot \exp(-t_i/T))^2 \tag{2.28}$$

where N is the number of time points, and t_i and $S(t_i)$ are the time and observed signal intensities at the i-th time point, respectively.

However, because the model is nonlinear in the parameters, the NLLS problem cannot be solved directly and requires iterative methods such as the Levenberg-Marquardt algorithm[185, 204]. These methods can be time-consuming, especially when applied to large 3D+time datasets. The accuracy and precision of the fitting process are influenced by various factors, such as the number of input images, chosen number of parameters for the fitting model, the optimization algorithm used for nonlinear tasks, the selection of a suitable initial guess for the model parameters, and potential presence of artifacts or misregistration between images. Using the mono-exponential decay equation for low-SNR images may cause an unreliable estimation of T2 relaxation times, which could result in their overestimation, as previously reported[207].

CS has been used to reconstruct parametric maps using overcomplete dictionaries[73]. Furthermore, model-based methods can be applied to directly estimate T2 maps and spin density maps from the raw data. MARTINI[289] assumes monoexponential T2 decay and recovers missing samples by directly estimating T2 maps from undersampled k-space data. Additionally, regularization has been suggested to reduce the noise levels in parametric map reconstruction. GRAPPATINI[117] generates T2 maps and synthetic T2W images in less than 2 minutes by extending MARTINI to GRAPPA. SUPER[125] uses blockwise fitting as a model-based acceleration method for parametric mapping and can be combined with SENSE. Poot et al.[236] proposed a framework that simultaneously estimates parametric and noise maps with smooth noise levels using maximum a posteriori estimation. Ben-Eliezer et al.[23] determined T2 maps by employing stepwise Bloch simulations of the pulse sequence, which traced all coherence pathways along the echo train and reproduced all stimulated and indirect echoes. A preprocessing step was performed to generate a database of echo-modulation curves for a range of T2 values. The method can also be extended to multiparametric estimations, such as joint T2 and B1+ fitting. Zhang et al.[361] combined the locally low-rank constraint and parallel imaging to accelerate T2 mapping.

MRF enables multi-parameteric mapping in a single scan[196]. Parameter estimation is achieved by matching the acquired signals to a dictionary of simulated signals that correspond to a grid of parameter combinations. However, increasing the number of estimated parameters in dictionary matching can make the problem more complex due to the exponential increase in the number of simulated signals and the associated computational and memory costs. To address this issue, various strategies have been developed, such as smarter search methods[44] and using singular value decomposition (SVD)[210] to compress simulated signals. While these methods can reduce costs, they have limitations, such as reduced accuracy when only a few singular vectors are used. Yang et al.[344] proposed using a polynomial hyperplane to interpolate the signal, but this method has only been applied to a two-parameter case, limiting its accuracy to an a priori defined refinement factor. Valenberg et al.[306] proposed fitting the acquired signal with a continuous signal model that uses B-spline interpolation of a sparse dictionary.

2.2.2.3.2 DL-based MR parametric mapping

Several studies have investigated the potential of Artificial Neural Network (ANN)s to predict parametric maps of MRI using these advanced models[115, 121]. We refer review article for these emerging deep-learning-based techniques that can provide valuable insights into how to enhance relaxometry measurements[85]. Among the various deep learning-based models, supervised learning is com-

monly used [89, 181, 281, 333]. Supervised frameworks train a neural network to predict various quantitative parameters such as T1 and T2 relaxation times from MRI images. Liu et al.[183] proposed an adversarial learning-based method for image-to-parameter mapping by integrating an end-to-end CNN and MR physical models to enforce data consistency. These models demonstrate comparable accuracy to classical methods with significantly reduced execution times. To address the lack of the ground truth, some studies have proposed to use unsupervised or self-supervised learning methods [141, 142, 182]. These methods are more effective for generalizing unseen data outside the training dataset. For example, an unsupervised framework[307], FLASH, utilizes a series of multi-contrast magnetic resonance images to obtain quantitative parameters. The images were divided into two parts, y_{in} and y_{out}, with the specific scanner parameters Φ_{in} and Φ_{out}. The first part of the images, Φ_{in}, was used to train a CNN. The estimated parameters, including T1, T2*, and Proton Density (PD), along with the scanner parameters Φ_{out}, were then input into a physical model to generate the second batch of multi-contrast images, y_{out}, which were not used during the training process. The difference between the estimated y_{out} and the true y_{out} was used as the loss function to update the CNN parameters. Although the proposed framework appears to be innovative, it is limited to a specific series of magnetic resonance images obtained with a fixed MRI acquisition protocol called FLASH. Consequently, the neural network model cannot be reused for various MRI sequences or physical models. Similarly, RELAX[182], a self-supervised framework, uses two physical models for network training: an inherent MR model and a quantitative model to fit the parameters in qMRI. By enforcing the physical model constraints, RELAX eliminates the need for the full reference datasets required in standard supervised learning. It also allows for direct reconstruction of parameter maps from the undersampled k-space in end-to-end CNN mapping. Recently, Bian et al.[29] proposed a scan-specific self-supervised framework, RELAX-MORE, which uses an optimization algorithm to unroll the learning process. Joint-MAPLE[116] framework comprises a scan-specific self-supervised reconstruction, leveraging multi-contrast data from a multi-echo, multi-flip angle, and gradient echo acquisition. it integrates parallel imaging, signal modelling, and data consistency blocks, which are jointly optimized through a combined loss function. Müller-Franzes et al.[219] proposed a MLP trained on the simulated data to eliminate the need for a dataset to substitute classical methods for Articular Cartilage.

2.3 Conclusion

ML have seen considerable progress; however, challenges remain to be addressed. Among these are the risks of overfitting the training data, balancing the model capacity with the size of the training data, and selecting appropriate hyperparameters for generalization using a validation set. Moreover, optimizing the standard loss function may present difficulties, requires the use of surrogate loss functions. Deep learning models require significant amounts of data for training. However, obtaining such data can be challenging and expensive, particularly in healthcare, where data privacy is a concern. Therefore, ongoing research is being conducted to develop methods for training deep learning models with limited data, such as transfer learning, semi-supervised learning, or scan-specific models. In addition, deep learning models are computationally expensive to train and deploy, particularly for large neural networks with many parameters that require specialized hardware, including GPUs and TPUs. Addressing overfitting and underfitting involves the use of methods such as regularization, dropout, and cross-validation. Deep learning models often lack interpretability, which is crucial, especially in sensitive domains, such as healthcare. Achieving good generalization

involves finding the right balance between model capacity and the size of the training data. The transfer track of the fastMRI 2020 challenge revealed that differences in scanner settings could lead to visible artifacts in the reconstructed images. In a separate study, Darestarni et al.[61] assessed the robustness of three categories of MRI reconstruction methods—trained deep networks, untrained deep networks, and classical sparsity-based approaches—against small adversarial perturbations, distribution shifts, and the recovery of fine details. Moreover, the reconstruction quality is closely tied to small feature recovery, and improving overall reconstruction performance enhances the ability to recover fine details in an image. Another issue is distribution shifts which can impact the performance of DL-based models. As such, scan-specific methods can help bridge this performance gap by accessing a limited amount of data. Training methods can also benefit from larger and more diverse datasets or data augmentation. However, overcoming distribution shifts remains a challenging task.

To address these challenges, we focused on creating scan-specific models. Our primary goal was to leverage the existing clinical MRI acquisition process. We also address the issue of image quality assessment. Throughout this PhD program, two interns conducted preliminary research on optimizing sampling patterns and developing machine-learning-based strategies for qMRI. However, we made the choice not to include their preliminary work in this manuscript by focussing on scan-specific reconstruction implications and image quality metrics.

3 COBRAI: A Contribution to Image-Quality Assessment

Summary

3.1 Image Quality Assessment in MRI . . . 61
3.2 Standard metrics used in this work . . . 62
3.2.1 MSE and its variant . . . 63
3.2.2 SSIM . . . 64
3.2.3 Blur Metric . . . 65
3.2.4 limitations . . . 65
3.3 COBRAI . . . 66
3.3.1 Metrics evaluation for various degradations . . . 67
3.3.1.1 Shepp-Logan phantom with various degradations . . . 67
3.3.1.2 Subtle natural image degradations . . . 68
3.3.1.3 COBRAI and SSIM in the context of MRI Reconstruction . . . 68
3.4 Discussion and Conclusion . . . 69

3.1 Image Quality Assessment in MRI

Ensuring accurate and reliable evaluation of image quality is essential in MRI for purposes such as diagnosis, treatment planning, and research. The assessment of the image quality in MRI relies heavily on both quantitative metrics and expert evaluations. However, several challenges limit the application and efficacy of the existing quality assessment metrics. These challenges include the subjective nature of assessments, lack of a universally accepted ground truth, intricacies of human perception, context sensitivity, multifaceted aspects of image quality, and overall lack of standardization.

Subjective assessments present a significant challenge. Many current metrics are based on the subjective evaluations of radiologists and domain experts, who are invaluable in their expertise but whose evaluations introduce an element of variability. This variability arises because of inherent inconsistencies in individual judgments, making it difficult to establish standardized and objective quality assessment methods. Factors such as varying levels of experience, individual biases, and

the influence of contextual information can lead to divergent opinions, further compromising the reliability of image quality evaluations.

Despite the widespread adoption of various quantitative evaluation metrics, such as NRMSE, PSNR, SSIM, and NMAE, in image-quality assessment, their limitations persist[220, 327]. For instance, PSNR and NRMSE may not align well with human visual perception, as they give equal weight to all pixels and ignore spatial information[270]. SSIM, while better reflecting human visual characteristics, shows weaknesses in assessing the quality of blurred and noisy images[188]. Although subjective measures have limitations, they are widely used due to their non-parametric nature. Variability in acquisition, reconstruction, and post-processing protocols further exacerbates these challenges. These challenges are compounded when dealing with multi-institutional datasets, where differences in equipment and protocols can lead to a wide range of image qualities and artifacts [46, 377]. This lack of standardization hinders the reproducibility of studies and undermines the comparability of results across different investigations. Therefore, there is a need to develop new assessment methods that provide more reliable and standardized results.

Advances in the field of computer vision have led to the utilization of pre-trained deep neural networks for evaluating the perceptual quality of images[360]. This is achieved by determining differences at the representation level or through downstream tasks such as image classification. Nevertheless, pretrained models often struggle with domain-specific tasks owing to generalization issues, and their high computational demands can pose challenges for real-time evaluations[112]. Additionally, their "black-box" nature complicates interpretability, and fine-tuning them can cause overfitting, particularly with limited data. Furthermore, potential biases from the original training dataset and over-reliance on established architectures can limit their efficacy [330].

In this context, the expertise of professionals such as MR physicists and radiologists becomes invaluable. They conduct visual examinations of images, grading them on parameters such as artifact presence, sharpness, and Contrast-to-Noise Ratio (CNR), which have significant implications as they directly affect the end-user's accuracy in interpreting images. However, relying solely on expert evaluations also has drawbacks as it is a resource-intensive process in terms of cost and time.

Recent findings from fastMRI[217] challenge expert evaluations have revealed the nuanced preferences of readers regarding image quality. Interestingly, these experts demonstrated a preference for images with noticeable noise levels. This preference stems from extensive training, which teaches them to identify structures, even in noisy images. By contrast, images that underwent excessive denoising were often deemed less suitable for the interpretation tasks, emphasizing the importance of understanding readers' expertise and preferences in the field of image quality assessment. A notable discrepancy was observed between the readers' subjective visual evaluations and traditional quantitative metrics. This discrepancy highlights the limitations of relying solely on quantitative metrics and underscores the need to incorporate expert evaluations, especially when the ultimate goal is an accurate diagnosis.

In recent years, numerous metrics have been developed for comparing images. These metrics are designed to assess image quality (IQA), which is influenced by several factors, including uniformity of lighting and brightness, resolution and geometry, contrast and sharpness, colour accuracy, and colour discrimination of the observed image[50]. These metrics have been categorized into three groups: Full-reference (FR), Reduced-reference (RR), and No-reference (NR) or blind. The focus of this thesis is on the FR metrics. Taking into account these key insights, several critical focus areas for future research and improvement have been identified: evaluating metrics, characterizing

errors, and addressing the challenges of AI-generated hallucinations. To tackle these challenges, we propose COBRAI, a novel quality assessment metric designed to be more sensitive to visible residual artifacts, offering the advantage of straightforward automatic quantification. In the subsequent sections, we will revise Mean Square Error (MSE) and recall the established standard metrics, such as NRMSE, PSNR, SSIM, and NMAE. Furthermore, we address the inherent limitations of these metrics, emphasizing the necessity for more robust and accurate quality assessment techniques. We then introduce COBRAI and discuss its mathematical foundations and potential applications in MRI reconstruction. Lastly, we present the results of our experiments and explore the implications of these findings.

3.2 Standard metrics used in this work

Traditional metrics for evaluating image quality, such as MSE, have typically been widely used, but are not perfectly aligned with human visual perception. Furthermore, these metrics may not always correspond to subjective evaluations of image quality. To address this issue, SSIM was introduced, which considers aspects such as image structure, luminance, and contrast to simulate human visual processing. In this section, we define MSE, SSIM and Blur metrics in order to understand how COBRAI was constructed.

3.2.1 MSE and its variant

The mean squared error (MSE) was extensively used as a quality criterion and loss function in this work. It is a commonly used metric in scientific disciplines such as statistics and signal processing. Its popularity arises from its mathematical simplicity and versatility in addressing various real-world problems. It is justified to use MSEwhen the underlying noise follows a normal distribution (said Gaussian noise). While numerous studies have explored the complexities and applications of MSE, new perspectives and insights continue to be uncovered. For a comprehensive understanding of MSE, we recommend referring to the review article[327]. This section emphasises the benefits of MSE and highlights its ongoing relevance. Fundamentally, MSE is designed to compare two images, such as a ground truth and a predicted one.

Consider two 2D images of size $M \times N$: the ground truth image x and the predicted image y. Here, N is the number of samples. The MSE between these images is defined as:

$$MSE(x,y) = \frac{1}{M \times N} \sum_{i=1}^{M} \sum_{j=1}^{N} (x_{i,j} - y_{i,j})^2 \tag{3.1}$$

We note Residual Map (RM) and $r_{i,j}$ the difference between x and y and its value at the $i^{th} and j^{th}$ position.

$$r_{i,j} = x_{i,j} - y_{i,j} \tag{3.2}$$

PSNR is derived from the MSE, which inherently ties them together mathematically. The relationship between them is given by:

$$\text{PSNR} = 10\log_{10}\left(\frac{L^2}{\text{MSE}}\right) \tag{3.3}$$

In this equation, L represents the dynamic range of the image intensities. Though PSNR offers notable advantages, especially when contrasting images with varying dynamic ranges, its foundational interpretation predominantly stems from the understanding derived from MSE.

The Normalized Mean Squared Error (Normalized Mean Square Error (NMSE)) is:

$$\text{NMSE} = \frac{\sum_{i=1}^{M}\sum_{j=1}^{N} r_{i,j}^2}{\sum_{i=1}^{M}\sum_{j=1}^{N} x_{i,j}^2} \tag{3.4}$$

The Normalized Root Mean Squared Error (NRMSE) is the square root of (NMSE).

Conversely, the Mean Absolute Error (Mean Absolute Error (MAE)) applies a linear penalty to the error by taking the average of the absolute error:

$$MAE = \frac{1}{M \times N}\sum_{i=1}^{M}\sum_{j=1}^{N} |r_{i,j}| \tag{3.5}$$

Further, the Normalized Mean Absolute Error (NMAE) is a normalized version of the MAE by the ground truth:

$$NMAE = \frac{\sum_{i=1}^{M}\sum_{j=1}^{N} |r_{i,j}|}{\sum_{i=1}^{M}\sum_{j=1}^{N} |x_{i,j}|} \tag{3.6}$$

3.2.2 SSIM

SSIM[326] is a metric designed to evaluate image quality by replicating the human visual system's perception of image degradation. Unlike MSE, which focuses on individual pixels, SSIM emphasizes local patterns, making it more effective for assessing structural degradation. This metric has been widely embraced in the research community for nearly two decades and has established a foundational role in various fields. For a comprehensive understanding of its mathematical foundations and applications in image quality assessment, we recommend referring to the article "Understanding SSIM"[220].

Similar to MSE, SSIM is designed to compare two images, such as a ground truth and a predicted one. It offers computational ease due to its nonparametric and memoryless properties, making it a beneficial metric. Moreover, its differential nature enables it to function effectively as a loss function in optimizing various image-processing tasks. These tasks encompass a wide array of applications, such as image denoising, restoration, contrast enhancement, watermarking, approximation, quantization, compression, inpainting, and super-resolution. In addition, SSIM finds extensive utility in the medical image reconstruction field, particularly in MRI and CT.

Now, let us delve into the core equation of SSIM between two patches of an images x and y of size $M \times N$. The patches are assumed to have a size of $2n+1 \times 2n+1$ with usually $2n+1 = 11$. For a patch

centered at indexes i, j:

$$SSIM_{i,j}(x,y) = l_{i,j}(x,y)^{\alpha} \cdot c_{i,j}(x,y)^{\beta} \cdot s_{i,j}(x,y)^{\gamma} \tag{3.7}$$

The luminance, contrast, and structural comparisons between two images are represented by l, c, and s respectively. α, β and γ are weighting factors usually chosen as equal to 1.

$$l_{i,j}(x,y) = \frac{2\bar{x}_{i,j}\bar{y}_{i,j} + C_1}{\bar{x}_{i,j}^2 + \bar{y}_{i,j}^2 + C_1} \tag{3.8}$$

$$c_{i,j}(x,y) = \frac{2\sigma_{x,i,j}\sigma_{y,i,j} + C_2}{\sigma_{x,i,j}^2 + \sigma_{y,i,j}^2 + C_2} \tag{3.9}$$

$$s_{i,j}(x,y) = \frac{cov_{xy,i,j} + C_3}{\sigma_{x,i,j}\sigma_{y,i,j} + C_3} \tag{3.10}$$

The constants C_1, C_2, and C_3 as described by Wang et al., take into account for stabilization during the division by a very small number. $\bar{x}_{i,j}$ and $\bar{y}_{i,j}$ are estimates of the average values over the patch centered at i, j for images x and y, respectively; $\sigma_{x,i,j}$ and $\sigma_{y,i,j}$ are standard deviation estimates, and $cov_{xy,i,j}$ the covariance between them.

For more efficient computations, a condensed version of SSIM is sometimes used with $C_2 = 2 \times C_3$[326]:

$$SSIM_{x,i,j}(x,y) = \frac{(2\bar{x}_{i,j}\bar{y}_{i,j} + C_1)(2cov_{xy,i,j} + C_2)}{(\bar{x}_{i,j}^2 + \bar{y}_{i,j}^2 + C_1)(\sigma_{x,i,j}^2 + \sigma_{y,i,j}^2 + C_2)} \tag{3.11}$$

SSIM ranges between -1 and 1 for each patch. A value of 1 indicates that two patches being compared are identical, with no degradation in structural information, luminance, or texture. A value close to 1 signifies that the patches are very similar in terms of structural integrity and visual appearance. In practice, a sliding Gaussian weighted kernel can be employed to compute the mean μ, variance σ, and covariance of the patches around the ij^{th} pixel position. Mean SSIM (MSSIM) value across an entire image can then be calculated. In this work, the implementation available in scikit-image was used[1].

3.2.3 Blur Metric

While MSE and SSIM are intrinsically designed to compare two images, Frederique et al.[57] designed a no-reference perceptual blur metric that operates on a scale of 0 to 1, with zero signifying the worst possible blur quality and one representing the best. This metric is based on the idea that the level of blurriness in an image can be assessed by comparing the differences in pixel intensity before and after applying a low-pass filter. The application of this filter involves calculating the variations in intensity between neighbouring pixels in the original image and then comparing these variations after the filter has been applied. By examining these discrepancies, it was possible to evaluate the degree of blurriness in the original image. A significant difference between the original and filtered

[1] https://scikit-image.org/docs/stable/api/skimage.metrics.html#skimage.metrics.structural_similarity

images indicates that the former was sharp, while a less pronounced difference suggests that it was somewhat blurred. This measure can be utilized as a blur quality metric or as a means of enhancing picture sharpness or fine-tuning artifact corrections to prevent the emergence of a blur. In this work, the implementation available in scikit-image was used[2].

3.2.4 limitations

Although SSIM and PSNR are commonly used to evaluate image quality, their limitations must be considered. The limitations of MSE (and PSNR) are shown by Wang et al. [326, 327] where they indicated that identical MSE scores can correspond to vastly different visual qualities.In contrast, SSIM provides a more accurate representation of differences in visual quality. This is because SSIM evaluates the structural distortion, whereas MSE does not. The human visual system is highly sensitive to structural distortions and can readily compensate for non-structural distortions. However, SSIM faces the same challenge as the PSNR in this regard. Kotevski et al.[161] presented a study in which SSIM scores were similar but visual quality differed significantly. Dosselmann et al.[78] suggested that there may exist a direct association between SSIM and MSE. Empirical evidence was provided by correlating the MSE with the SSIM for a variety of images. Hore et al.[7] discovered a simple analytical relationship between PSNR and SSIM, suggesting that PSNR values can be predicted from SSIM, and vice versa. Although the SSIM can produce useful results in some instances, it can also yield counterintuitive outcomes in many other cases. However, Pambrun et al.[227] found that local SSIM estimates are not stable for medical images with fragments of high or low average brightness. In addition, the luminance and variance terms of the SSIM are dice coefficients that depend on the absolute values of the input parameters, which can cause the model to become unstable when the input parameters are close to zero. Reibman et al.[247] found that popular quality metrics, such as PSNR and SSIM, were not applicable for SR tasks.

3.3 COBRAI

To assess the presence of residual artifacts and compare reconstructions using quantitative metrics, the COrrelartion-Based Residual Artefact Index (COBRAI) is proposed as an average local measure of residual artifacts. The process starts with the computation of a RM as the difference between a reference or ground-truth image x and a degraded or predicted image y. The similarity between the RM and the ground-truth image should then reflect the level of information still present in the residual map and consequently absent from the reconstruction.

To construct this index, similar to SSIM, COBRAI utilized a patch-based approach and will be evaluated here on 2D images. For each patch in an image, the correlation coefficient between the RM of the current reconstruction and the ground truth is computed. A COBRA map is produced by computing the local correlation in the window. The COBRAI is then taken as the mean absolute value of the correlation coefficients of all patches within a relevant mask.

Let us consider the mathematical formulation of COBRAI. Given a reconstructed image y and its corresponding ground truth x, recall equation 3.2 is used for the error calculation:

[2]https://scikit-image.org/docs/stable/api/skimage.measure.html#skimage.measure.blur_effect

$$RM = x - y \tag{3.12}$$

The patch-wise correlation coefficient, $c_{i,j}$, between RM and x for a patch of size $(2n+1) \times (2n+1)$ centered at pixel (i,j) is given by:

$$c_{i,j} = \frac{\sum_{k=i-n}^{i+n} \sum_{l=j-n}^{j+n} (r_{k,l} - \bar{r}_{i,j})(x_{k,l} - \bar{x}_{i,j})}{\sqrt{\sum_{k=i-n}^{i+n} \sum_{l=j-n}^{j+n} (r_{k,l} - \bar{r}_{i,j})^2 + \epsilon} \sqrt{\sum_{k=i-n}^{i+n} \sum_{l=j-n}^{j+n} (x_{k,l} - \bar{x}_{i,j})^2 + \epsilon}} \tag{3.13}$$

Where:

- n is the size of patch at location i, j
- $r_{k,l}$ represents the pixel values of the RM at position (k,l).
- $x_{k,l}$ denotes the pixel values of the ground truth image at position (k,l).
- $\bar{r}_{i,j} = \frac{1}{(2n+1)^2} \sum_{k=i-n}^{i+n} \sum_{l=j-n}^{j+n} r_{k,l}$ and $\bar{x}_{i,j} = \frac{1}{(2n+1)^2} \sum_{k=i-n}^{i+n} \sum_{l=j-n}^{j+n} x_{k,l}$ are the mean pixel values of the RM and ground truth on the patch centered at (i, j), respectively.
- ϵ is a very small value added to avoid division by zero (typically 10^{-6}).

This computation gave rise to an image that we will refer to as the COBRA map. Except the small number added to avoid division by 0, $c_{i,j}$ is almost equivalent to the structural part of SSIM applied to the image x and to the residual: $c_{i,j} \approx s_{i,j}(x, x-y)$. Note in the above computation, it is assumed that the values outside the image boundaries are set to 0 (for $i, j < 1$, $i > M$, $j > N$).

The COBRAI is then computed as the mean absolute value of these correlation coefficients over all patches (mean of the COBRA map):

$$COBRAI = \frac{1}{M \times N} \sum_{i=1}^{M} \sum_{j=1}^{N} |c_{i,j}| \tag{3.14}$$

Where $M \times N$ is the total number of patches. To avoid boundary effects its computation can also be restricted to valid patches, or to a selected mask.

The COBRAI scores range between 0 to 1. Having a residual at 0 everywhere leads to $COBRAI = 0$. When the residual is identical to the ground-truth image x, $COBRAI = 1$. An adequate reconstruction is thus expected to have a small COBRAI, such as when unstructured noise in the RM is obtained. This occurs when the reconstruction accurately captures the details, resulting in a low correlation with the ground truth. Therefore, when comparing several reconstructions with COBRAI, a lower score indicates a lower average correlation between local patches in the corresponding RM as compared to the ground truth. Conversely, a higher COBRAI score suggests that the corresponding RM has more structural information that is either not present in the ground truth (hallucinations) or that details that are present in the ground truth are not reflected in the degraded or predicted image.

3.3.1 Metrics evaluation for various degradations

This section illustrates and evaluates the performance of COBRAI in quantifying image quality using the sheep-logan phantom dataset with varying degrees of noise, blurring, translation, and rotation. Subsequently, the analysis will be extended to natural images and MRI reconstruction using data from the fasMRI dataset.

3.3.1.1 Shepp-Logan phantom with various degradations

In order to provide an idea of the variation of the various image metrics presented above, the Shepp-Logan phantom was used. Implementation was done using Python. Starting from the ground-truth image, we investigated the effectiveness of COBRAI and other metrics under various image degradations:

- contrast stretching (using a scaling factor)
- Gaussian noise addition (with increasing variance levels)
- Gaussian blurring (i.e., smoothing with a Gaussian filter kernel with increasing variance levels)
- salt-and-pepper noise insertion,
- spatial translation (with increasing shifts)
- rotations (with increasing angles)

Results are illustrated in Figure 3.1 on which the metrics are reported. As expected, COBRAI and SSIM are insensitive to scaling, while pSNR is. This feature is a consequence of the construction of the metrics. COBRAI is relatively insensitive to noise (both Gaussian and salt and pepper) while SSIM is. On the other hand, COBRAI varies a lot with blurring, while the latter does not affect SSIM much.

Figures 3.2 a and b help to interpret COBRAI sensitivity to several degradation types for a range of degradation levels.

- Upon examining Figure 3.2a, which illustrates the cobra's sensitivity to blur, it becomes apparent that the cobra exhibits strong performance in the range of blur strengths from 0-0.2, as evidenced by the favourable results shown by other metrics. This suggests that the cobra can detect blur in this range as compared to other metrics.
- As depicted in Figure 3.2b, it is evident that the cobra demonstrates insensitivity to noise. This is because the noise has been artificially introduced, and the RM will not exhibit any discernible structured artifacts.
- COBRAI is more sensitive to both translation and rotation as compared to the other metrics, as indicated by a steeper transition for smaller degradation levels (Figure 3.2c and d).

3.3.1.2 Subtle natural image degradations

In this section, we present some examples of natural images that have been degraded in such a way that they are indistinguishable to the naked eye. Specifically, a kernel, also known as a filter or mask, is selected. This kernel employs either averaging or smoothing. By utilizing convolution, a kernel of size 3×3 is applied to the image, resulting in each pixel being replaced by a weighted average of its

surrounding pixels. Inspired by edge detection and smoothing filters, the filters are set to 0 for the first and third columns, and 1/3 for the second column. As depicted in Figure 3.3, raccoon, moon, and checkerboard images were chosen.

The SSIM value is high, indicating that the degraded image is indistinguishable. However, upon examining the error maps, it appears that there are structured artifacts associated with large COBRAI values. This further indicates that COBRAI provides different information as compared to SSIM, more sensitive to smoothing and missing details in these situations.

3.3.1.3 COBRAI and SSIM in the context of MRI Reconstruction

To assess the presence of residual artifacts in MR reconstruction COBRAI is proposed. The process is illustrated in Fig. 3.4. rRAKI and crRAKI will be discussed in detail in the coming chapter. A sample slice from FLAIR acquisition in the fastMRI dataset is shown.

To generate the ground-truth image, the fully sampled k-space was Fourier-tranformed and coil-combined using sum-of-squares. The three reconstruction strategies of retrospectively undersampled k-spaces with an acceleration factor of four are evaluated. Residual maps were calculated for GRAPPA, rRAKI, and crRAKI by using coil-combined reconstructed magnitude images. Patch-based covariance maps and COrrelation-Based Residual Artifact (COBRA) maps were also computed.

The COBRA map highlights the presence of artifacts in the rRAKI reconstruction, whereas the RM and COBRA maps enable a direct comparison of the performance of GRAPPA, rRAKI, and crRAKI in terms of accuracy, artifact suppression, and image characteristic preservation, as measured by COBRAI. In this example, GRAPPA and crRAKI had lower COBRAI values, indicating fewer reconstruction artifacts.

Fig. 3.5 displays a second example taken from in-house T2W-spin-echo data. Similarly, k-space was undersampled by a factor of four. Here, the SSIM map is also displayed. In this example, while SSIM is not different between the reconstructions, the residual map displays more structures for rRAKI reconstruction, which is highlighted on the covariance map as well as on the COBRA map. A lower COBRAI synthetizes this by a single number.

3.4 Discussion and Conclusion

Evaluating the quality of MRI images is crucial because it has a direct impact on interpretation and subsequent clinical decisions. We examine various traditional reference-based metrics. We observed that these metrics do not always equate to the desired image quality, as measured by metrics such as SSIM/PSNR, which can be similar even when traditional metrics do not accurately reflect the difference between the original and degraded images. For example, subtle distortions in natural images may not be visible to the human eye, and the SSIM/PSNR cannot distinguish between degraded and original images. Such distortions frequently occur in various reconstruction methods, where, even though traditional metrics may show improvement, there may be hallucinations or blurriness. In this study, we show that COBRAI is effective in identifying these degradations because the correlation is directly computed with the residual maps. After assessing the characteristics of COBRAI under various degradation scenarios, we determined that it exhibits a high sensitivity to blurring compared to noise. Consequently, we suggest integrating this metric with other conventional

metrics to gain a more comprehensive perspective. Furthermore, in the next chapter, COBRAI is used to evaluate the performance of various scan-specific reconstruction methods along with traditional metrics.

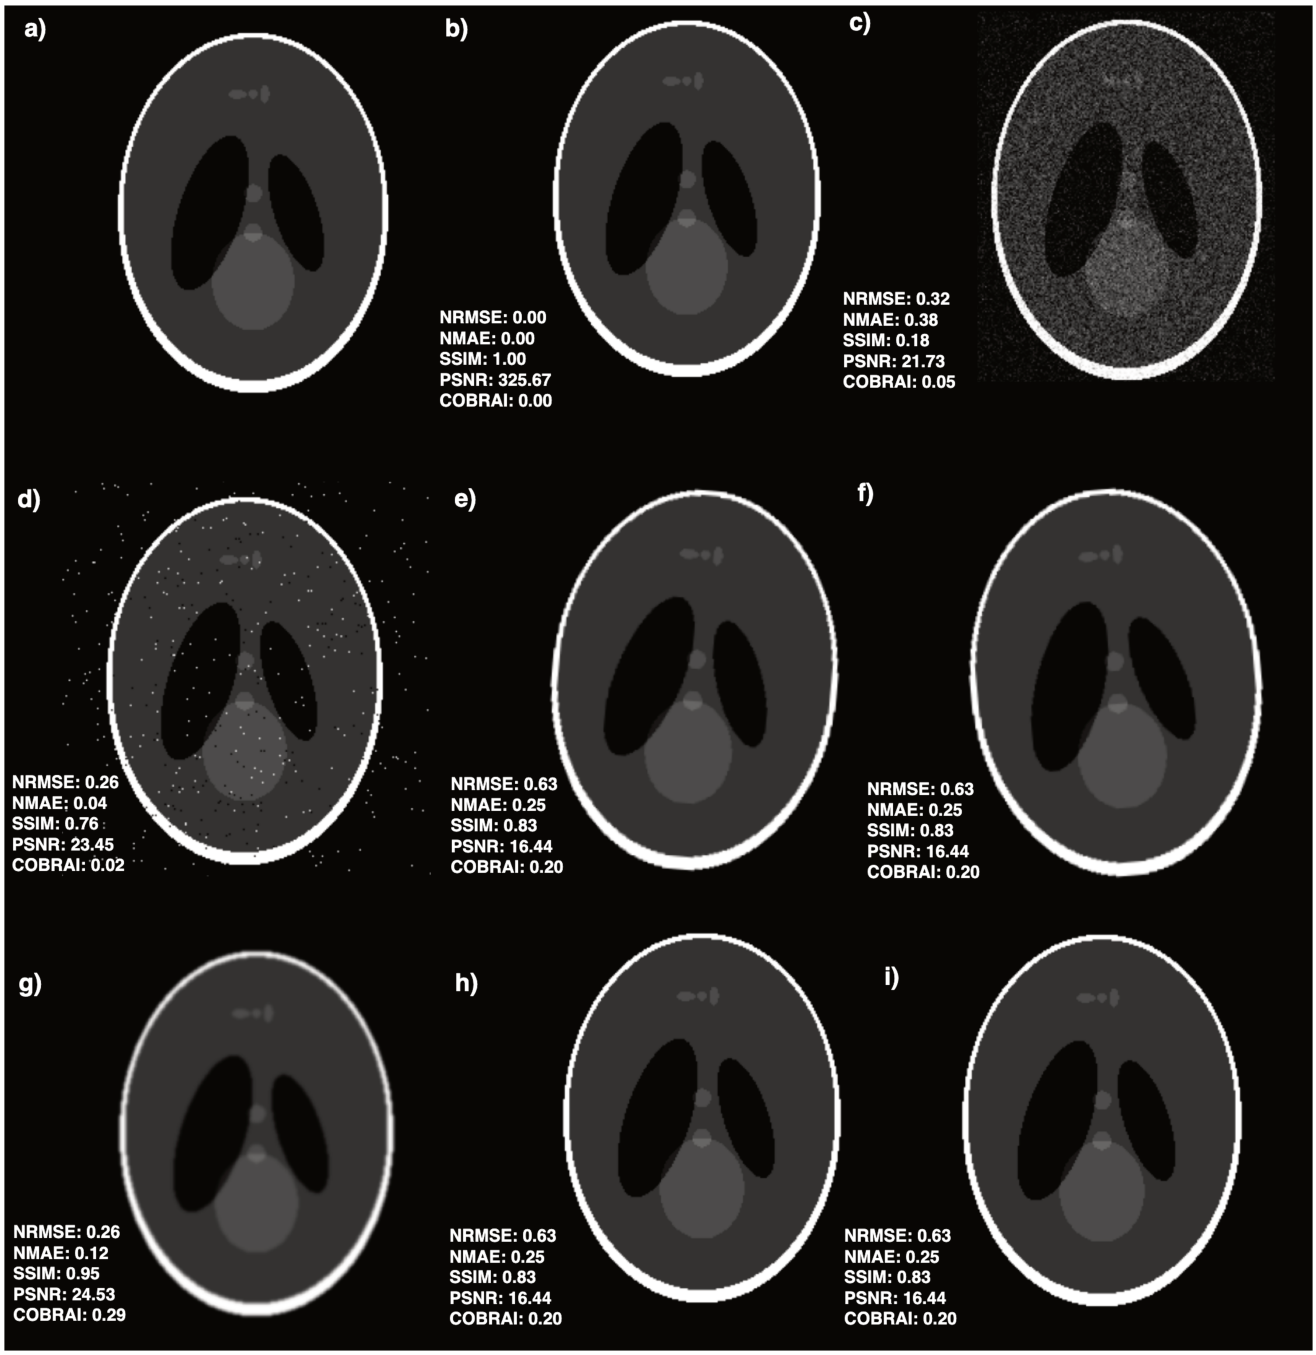

Figure 3.1: COBRAI performance on the Shepp-Logan phantom for various degradations a) Original Image: The unaltered reference image, b) Contrast Stretching: Enhances image contrast by stretching the intensity range, c) Gaussian noise: Random noise is added following a Gaussian distribution, d) Salt-and-pepper noise: Random noise is introduced as white and black pixels, e) Clockwise rotation: The image was rotated clockwise by a specified angle, f) Anticlockwise rotation: The image is rotated anticlockwise at a specified angle, g) Blur: Apply Gaussian blur to the image, h) Spatial shift left: The image is shifted to the left by a specified number of pixels, i) Spatial shift right: The image is shifted to the right by a specified number of pixels, evaluated using metrics NRMSE, NMAE, SSIM, PSNR, COBRAI.

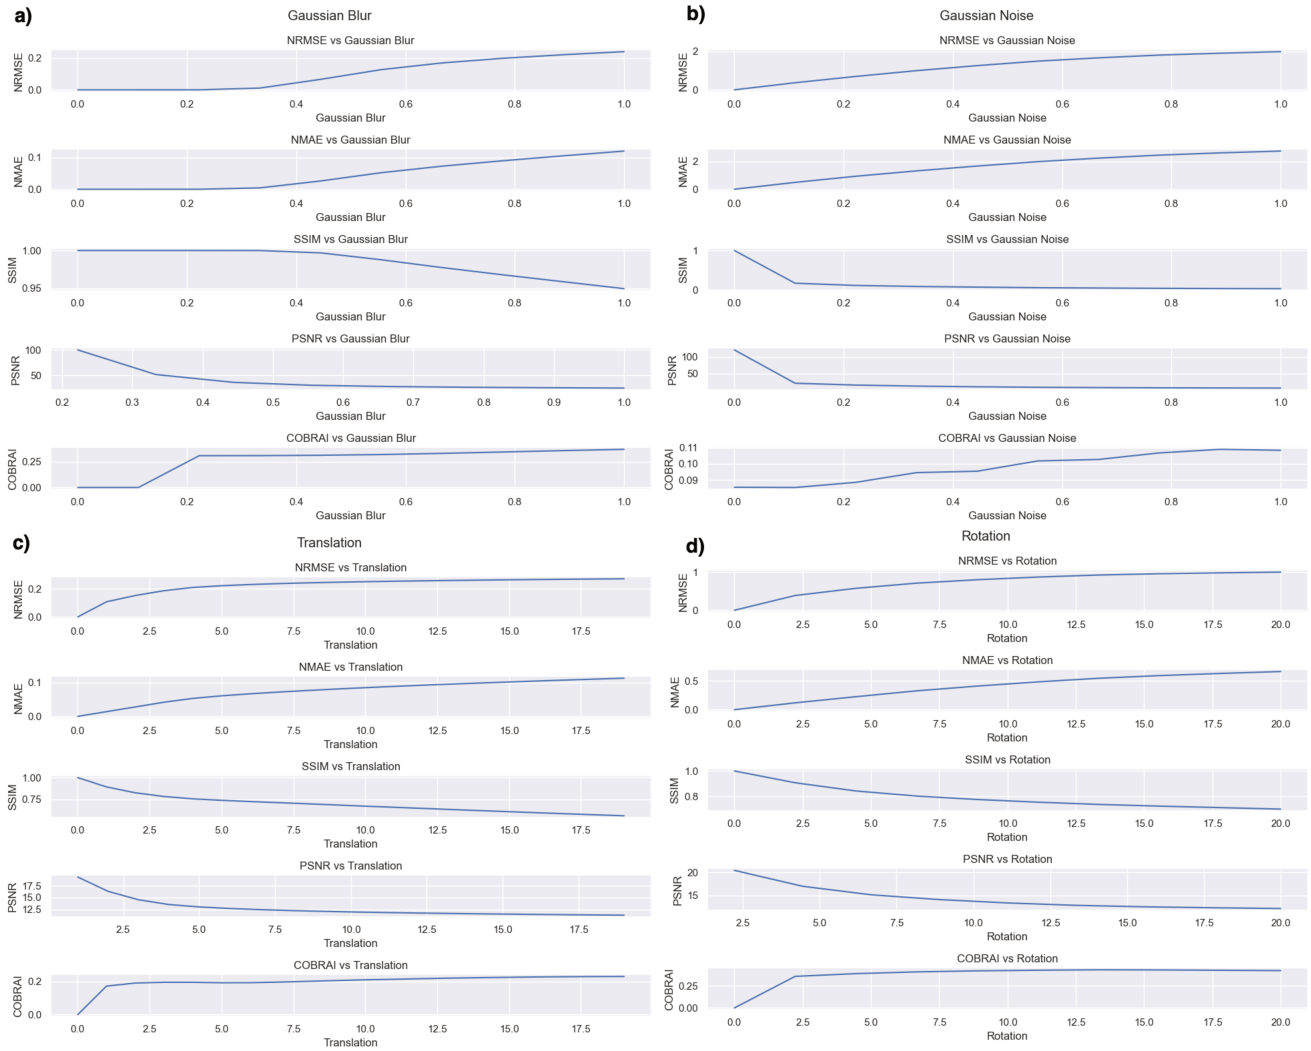

Figure 3.2: Metrics performance on degraded images with a) Gaussian Blur, COBRAI demonstrates strong performance in detecting blur in the range of 0-0.2, b) Gaussian Noise, COBRAI remains insensitive to noise, c) Translation, d) Rotation. Compared to other metrics, the COBRAI is more sensitive to both translation and rotation, as indicated by a steeper transition for smaller degradation levels.

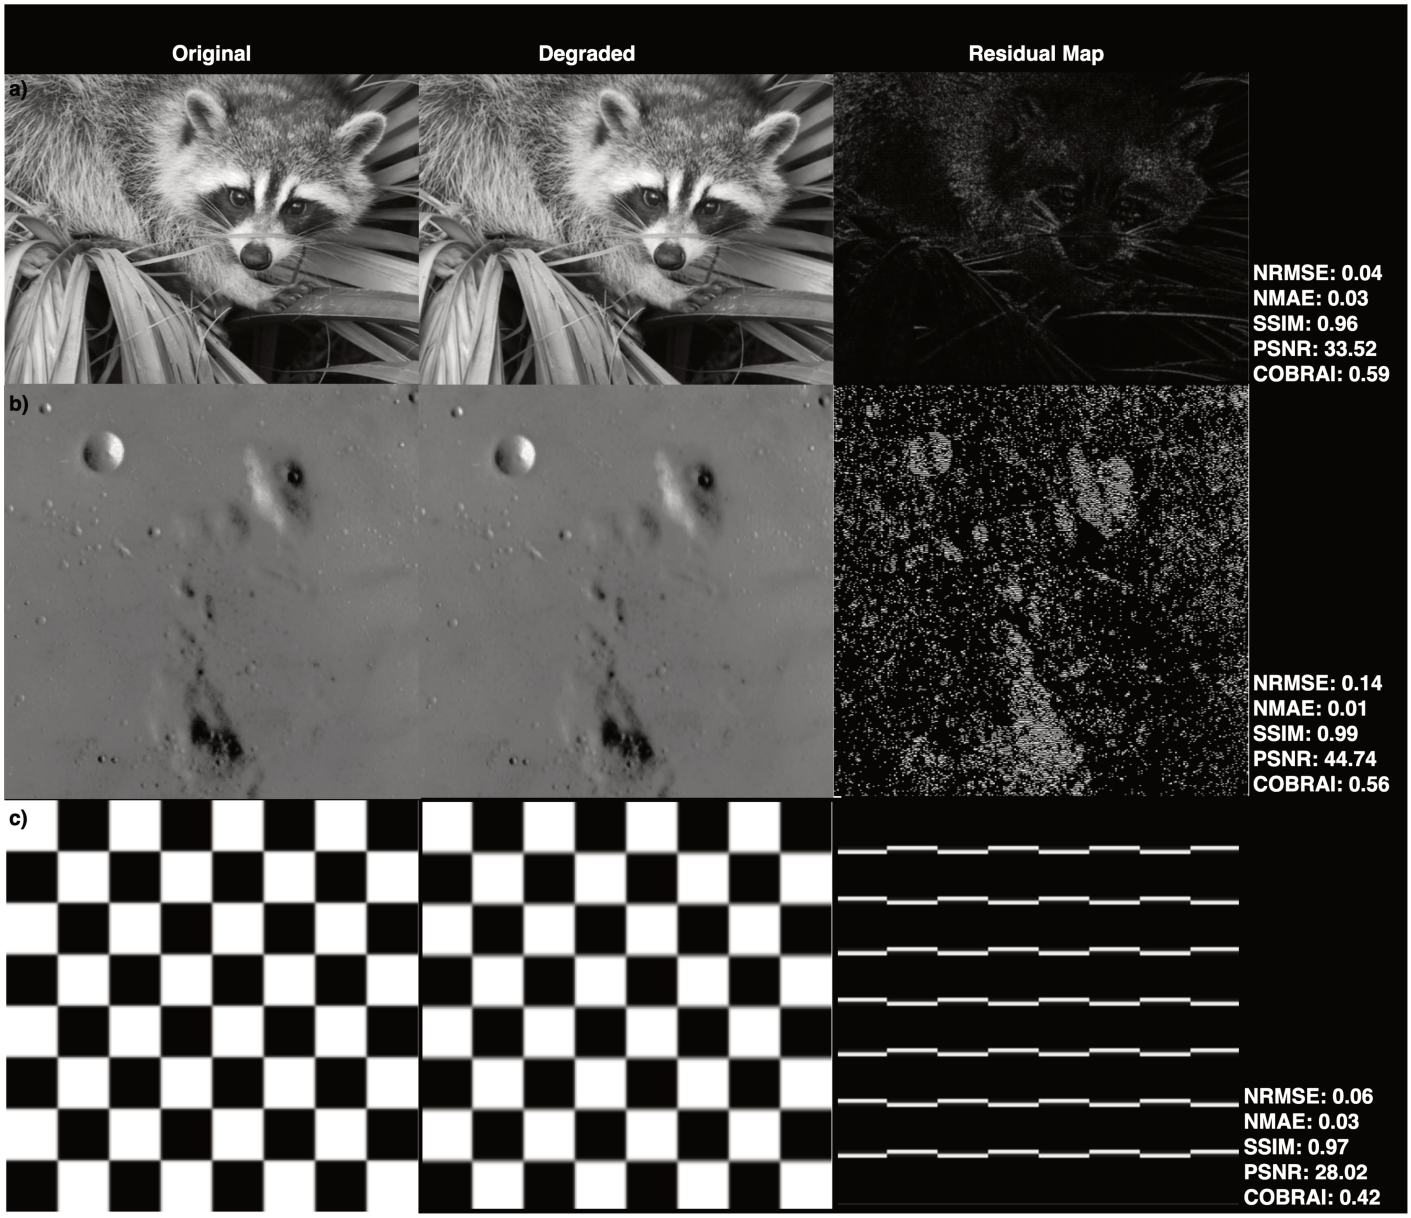

Figure 3.3: Degraded images with smoothing. The SSIM value is high, but error maps reveal structured artifacts. This suggests that COBRAI provides different information and is more sensitive to smoothing and missing details.

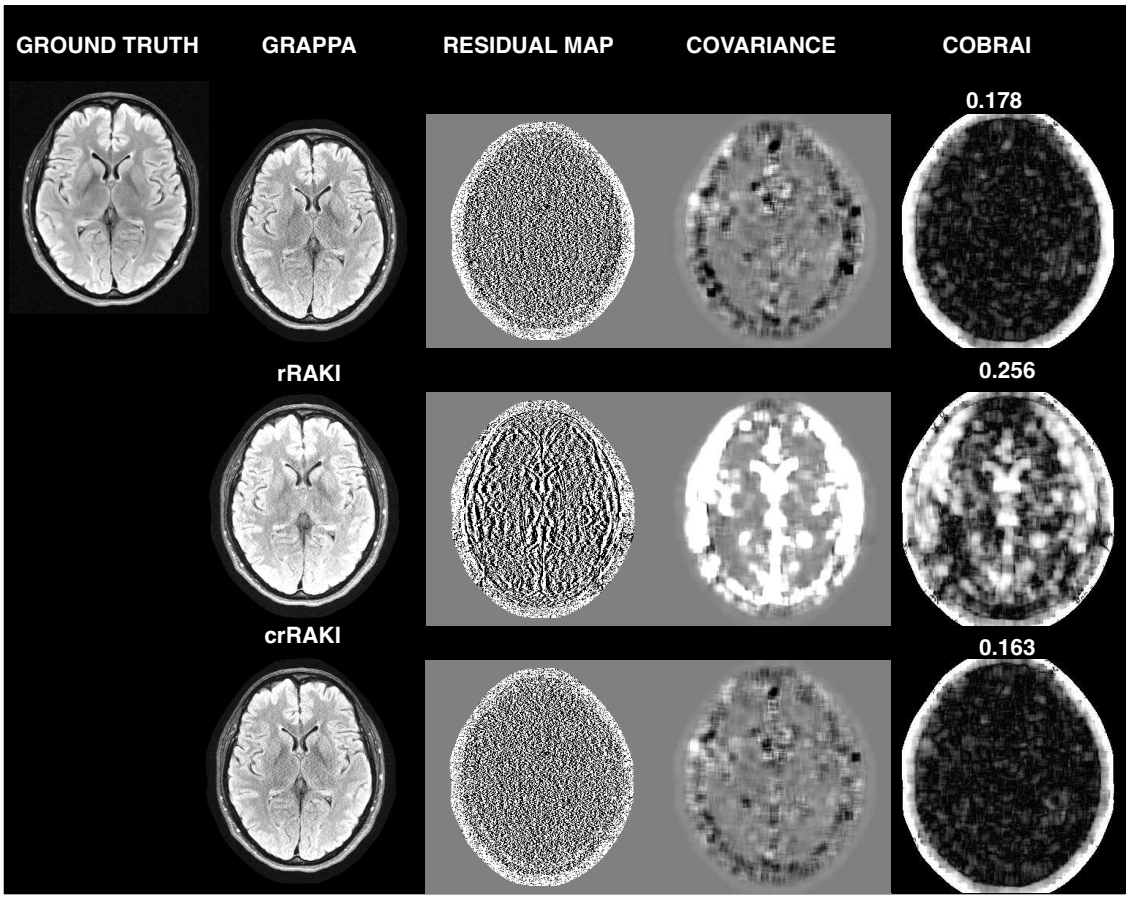

Figure 3.4: Illustration and explanation of COBRAI for providing a quantitative assessment of artifacts to compare accelerated reconstructions. An exemplary slice from a FLAIR acquisition from the fastMRI dataset is shown. The ground-truth corresponds to the coil-combined magnitude image from the fully-sampled k-space from which a mask is computed encompassing the brain. From the coil-combined reconstructed magnitude images (GRAPPA, rRAKI, and crRAKI), residual maps are computed (the masked residual map RM is shown). Patch-based covariance maps and COrrelation-Based Residual Artifact (COBRA, i.e. a local correlation coefficient) maps are computed. Among the various average covariance-based metrics, we report the mean absolute correlation coefficient over the mask as the COBRA Index (COBRAI). The COBRA map highlights without ambiguity the presence of artifacts in the rRAKI reconstruction here, which is synthesized by a single number (COBRAI). The RM and COBRA maps enable to directly see the limitations of GRAPPA, rRAKI, and crRAKI, shedding light on their performance in terms of accuracy, artifact suppression, and preservation of image characteristics with COBRAI. In this example, GRAPPA and crRAKI have smaller COBRAI, i.e. fewer reconstruction artifacts.

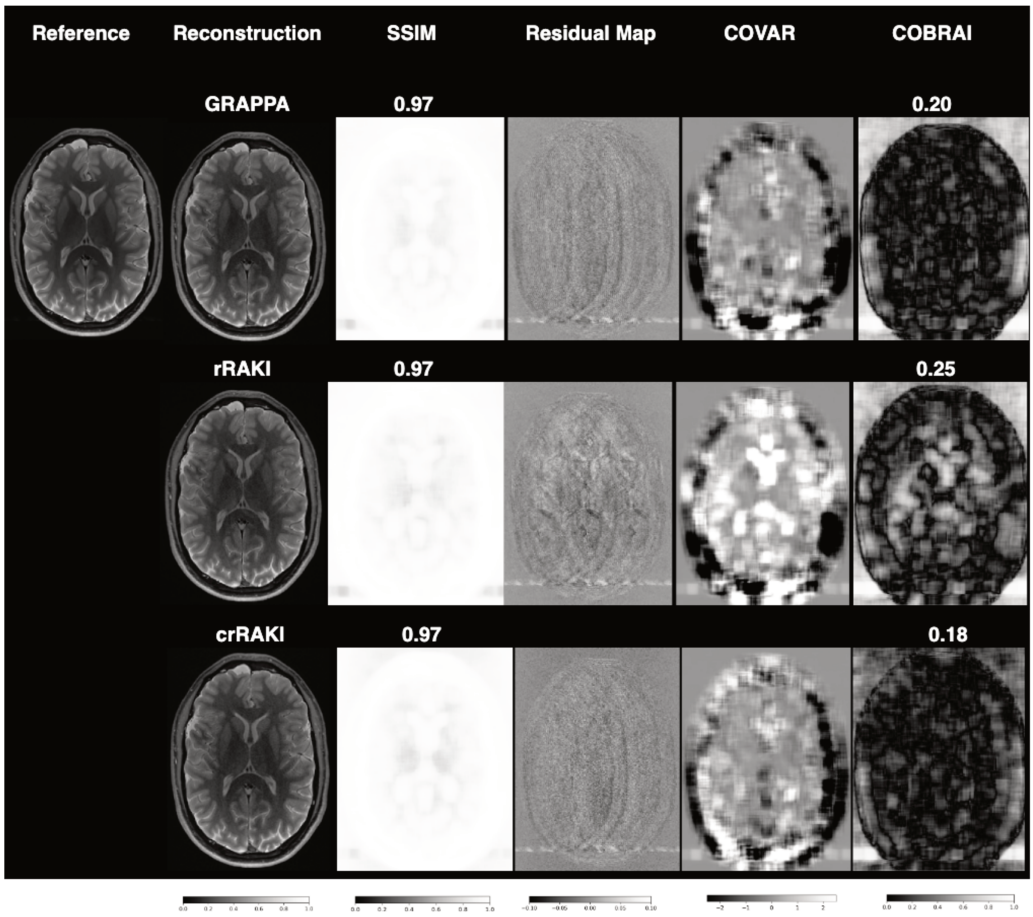

Figure 3.5: Illustration and explanation of COBRAI for providing a quantitative assessment of artifacts to compare accelerated reconstructions. An exemplary slice from an in-house dataset T2W is shown. The ground-truth corresponds to the coil-combined magnitude image from the fully-sampled k-space from which a mask is computed encompassing the brain. From the coil-combined reconstructed magnitude images (GRAPPA, rRAKI, and crRAKI), residual maps are computed (the masked residual map RM is shown). Patch-based ssim maps, covariance maps and COrrelation-Based Residual Artifact (COBRA, i.e. a local correlation coefficient) maps are computed. Among the various average covariance-based metrics, we report the mean absolute correlation coefficient over the mask as the COBRA Index (COBRAI). The COBRA map highlights without ambiguity the presence of artifacts in the rRAKI reconstruction here, which is synthesized by a single number (COBRAI). The RM and COBRA maps enable to directly see the limitations of GRAPPA, rRAKI, and crRAKI, shedding light on their performance in terms of accuracy, artifact suppression, and preservation of image characteristics with COBRAI. In this example, GRAPPA and crRAKI have smaller COBRAI, i.e. fewer reconstruction artifacts.

4 Focus on scan-specific convolutional models for MRI reconstruction

Summary

4.1 Introduction . . . 75
4.2 Method and Material . . . 76
4.2.1 Train-Validation Partition (TVP) . . . 76
4.2.2 Models . . . 77
4.2.2.1 Real and complex layers . . . 77
4.2.2.2 Residual RAKI and complex residual RAKI . . . 77
4.2.2.3 Grid-search: K-fold cross validation . . . 78
4.2.3 Implementation details . . . 79
4.2.3.1 Code and computing resources . . . 79
4.2.3.2 Model training . . . 80
4.2.3.3 Network architecture hyperparameters . . . 80
4.2.3.4 Statistics . . . 81
4.2.3.5 fastMRI and separated-ACS mode data . . . 81
4.3 Results . . . 82
4.3.1 Grid-search in 2D and separated ACS mode . . . 83
4.3.2 2D results . . . 83
4.3.2.1 fastMRI retrospective undersampling (integrated ACS mode) . . . 83
4.3.2.2 T2 Weighted in separated ACS mode . . . 89
4.3.2.3 2D multi-echo GRE in separated mode . . . 95
4.3.2.4 Effect of the number of ACS lines . . . 95
4.3.3 3D RESULTS in separated mode . . . 100
4.3.3.1 3D MPRAGE . . . 100
4.3.3.2 3D Multi-Echo GRE . . . 100

4.4 Discussion . . . 104
4.4.1 Main findings . . . 104
4.4.2 Non-linearity . . . 104
4.4.3 Real vs complex . . . 105
4.4.4 TVP and K-fold cross-validation . . . 106
4.4.5 Kernel sizes and depths . . . 106
4.4.6 ACS and Timing . . . 106
4.4.7 COBRAI . . . 107
4.4.8 Improvements and practical implications . . . 107
4.5 Conclusions . . . 108

4.1 Introduction

In the previous chapters, it was reviewed that deep learning techniques have been developed to address MRI reconstruction tasks in a variety of ways, including supervised learning[217, 235, 285, 286, 375], GAN-based methods[60, 186, 194, 195, 276, 318], and self-supervised methods[62, 143, 341, 342]. Although supervised learning methods can produce impressive results with higher acceleration rates and improved image quality, they require a large amount of training data and may not generalize well outside of the specific application on which they were trained. This is particularly relevant for MRI because images are generated under a wide range of experimental conditions, such as multiple anatomic territories, user-defined scan parameters, or various sequence types. If the training database does not include sufficient examples of all scan situations, especially those with pathologies of interest, it can negatively impact the accuracy of diagnosis by introducing incorrect information in the reconstructed images or by removing fine details[152, 217]. Additionally, deep learning-based MRI models do not generalize across datasets collected with different coils, which is an important consideration for clinical adoption[20].

To address the challenges of DL-MRI in reconstructing high-quality images for PI from different acquisition parameters, anatomical regions, and receiver coil configurations, scan-specific methods in which training data are acquired during the scan have been developed[5, 11, 193, 341, 356]. For example, Robust Artificial neural networks for K-space Interpolation (RAKI)[5] train neural networks on the so-called ACS[98] to estimate unacquired lines in k-space in a scan-specific manner. However, most prior RAKI-based studies, such as those based on the fastMRI dataset[153, 353], consider integrated-mode ACS. In the integrated mode, the ACS is acquired along with the undersampled k-space and shares the same contrast. ACS can also be acquired in the separated mode prior to the undersampled k-space. As opposed to the former, the latter acquisition method incurs only a marginal increase in the total scan time as shorter repetition times can be used. The counterpart is that the magnitude and phase usually differ between the ACS and the undersampled k-space, in particular due to different echo times and acquisition bandwidths, such that scan-specific models may fail to generalize well, leading to reduced image quality. This may be particularly true in applications in which the phase is of utmost importance for field mapping, phase-contrast velocity mapping[301], fat-water separation[68], or chemical shift imaging[216].

Additionally, although complex convolution layers have been introduced in the DL community[299], the majority of MRI reconstruction networks, including RAKI, separate the real and imaginary components of complex numbers into separate real-valued channels. Recent studies suggest that deep neural networks based on complex-valued convolutions[54, 66, 338] have high representational power and accuracy, which allows them to be applied to MRI reconstruction, including magnitude- and phase-based applications based on highly undersampled data. Thus, using real-valued versus complex-valued implementations may also have consequences for model generalizability. The recently proposed iRAKI[66] explored the use of complex-valued networks in k-space-based reconstruction in 2D imaging. It is important to emphasize that the optimization of various hyperparameters was performed empirically. In addition, an extension to 3D was not performed.

Notably, the existing literature lacks comprehensive studies that investigate both optimal hyperparameters and model architecture for scan-specific k-space reconstruction tasks, including the influence of nonlinearities and real/complex implementations. While prior research has explored hyperparameter optimization in general machine-learning contexts, their specific application to the domain of scan-specific k-space reconstruction remains limited. Additionally, reliable image quality evaluation is a major challenge in evaluating the quality of MRI reconstructions. Commonly reported metrics, such as NRMSE, NMAE, PSNR, and SSIM, often fail to adequately account for local variations and fine details, leading to suboptimal assessments[170, 208].

To investigate these issues, we evaluated rRAKI[356] and its complex counterpart crRAKI, and compared it with GRAPPA. Our work aims to conduct a thorough investigation into the impact of hyperparameters on the reconstruction performance and to evaluate the generalizability with regard to the ACS mode. To achieve this, we employed a grid search strategy combined with K-fold cross-validation and systematically explored a range of hyperparameter values to identify optimal configurations that maximize the accuracy and effectiveness of the models and guide hyperparameter choice. As such, we leverage a TVP scheme with early stopping to avoid or reduce overfitting effects. Focusing on implications for future research and clinical transfer, quantitative and qualitative comparisons between the different models were performed in various brain imaging situations, using different ACS modes, including 2D and 3D, various contrasts (T1W, T2W, T2*W, FLAIR), as well as phase-based applications such as field mapping with high acceleration rates.

4.2 Method and Material

4.2.1 Train-Validation Partition (TVP)

All models considered in this work are learned in a scan-specific mode, meaning that the models are trained on reference (ACS) data acquired during the scan and then used to infer the full k-space.

The ACS correspond to a fully-sampled *small* k-space, i.e. a 4D tensor of complex numbers:

$$\begin{aligned} S_f^{acs} = \big\{ & S^{acs}(i,j,k,c) \in \mathbb{C}, \\ & i,j,k,c \in [\![1, N_x^{acs}]\!] \times [\![1, N_y^{acs}]\!] \\ & \times [\![1, N_z^{acs}]\!] \times [\![1, N_c]\!] \big\} \end{aligned} \tag{4.1}$$

where $N_x^{acs} \leq N_x$ is the number of sampled points along the readout direction for the ACS, $N_y^{acs} \leq N_y$,

the number of sampled points in the first phase encoding direction, and $N_z^{acs} \leq N_z$ in the second phase encoding direction.

To train the models $f_{\theta_{m,n}}$, the fully-sampled ACS k-space is split into training and validation sets. (Recall from chapter2) Figure 2.7-(a) illustrates how a fully sampled k-space (the ACS in our case) may be used to learn a reconstruction model for an acceleration rate $R_y = 3$ on the simpler 2D case (generalization to 3D is straightforward).

Assume the index of the acquired *undersampled k-space* lines are $T_R(1) \equiv \{1,4,7,...\}$ in the phase encoding direction. Then one has to build models able to infer lines with index in $T_R(2) \equiv \{2,5,8,...\}$ and $T_R(3) \equiv \{3,6,9,...\}$. It is obvious that reconstructing lines in $T_R(2)$ from lines in $T_R(1)$ is the same task as reconstructing lines in $T_R(3)$ from lines in $T_R(2)$ but that reconstructing lines in $T_R(2)$ from lines in $T_R(1)$ is different from reconstructing lines in $T_R(3)$ from the lines in $T_R(1)$. Hence there are $(R-1)$ reconstruction subtasks for each of which one learns a reconstruction model. For instance, a first model is learnt to reconstruct the first lines missing right after the observed ones (Task 1). It is learnt both to infer lines in $T_R(2)$ from $T_R(1)$ and to infer lines in $T_R(3)$ from lines in $T_R(2)$... Actually, there are R such sets of pairs (input lines, output lines). Since it is common practice in machine learning to distinguish between training data and validation data, which are not used to learn the model but are used to estimate their prediction ability and to avoid overfitting, we chose to use $(R-1)$ sets of pairs for learning a model and the last one for validation. For instance for Task 1 in Figure 2.7-(a) the last pair of input ($T_R(3)$) and output lines ($T_R(4)$) is left for validation.

4.2.2 Models

4.2.2.1 Real and complex layers

The deployed model architectures incorporate real- and complex-valued convolutions without any bias. In addition, nonlinear real and complex activation layers were used. This has already been discussed in the previous chapter, so we will briefly recap the activation functions implemented include the LeakyReLU (for real tensors) and its complex counterpart, $\mathbb{C}LeakyReLU$.

$$LeakyReLU(x) = \begin{cases} x, & \text{if } x \geq 0, \\ cx, & \text{otherwise,} \end{cases} \tag{4.2}$$

$$\mathbb{C}LeakyReLU(x) = LeakyReLU(\mathrm{Re}(x)) + i \times LeakyReLU(\mathrm{Im}(x)) \tag{4.3}$$

The hyperparameter $c \in [\![0,1]\!]$ can be used to scale the amount of non-linearity between a pure ReLU ($c = 0$) up to pure linearity ($c = 1$).

4.2.2.2 Residual RAKI and complex residual RAKI

RAKI was improved using residual connections[114]. Residual connections enable to design networks with many layers without encountering the vanishing gradient problem. These connections allow the network to learn a "shortcut" between the input and output of the layer. The rRAKI approach implements a linear residual connection part (referred to as GRAPPA short connection), together with a RAKI section, using real-valued convolutions.

The GRAPPA short connection and RAKI section are explicitly separated and jointly trained. The training is done via a weighted MSE loss function consisting of two terms: the first term calculates the loss between the rRAKI predicted and actual output (model loss), and the second term consists of the loss between the GRAPPA short-connection prediction and the actual output (short-connection loss). A weighting factor, $\lambda = 1$, is introduced to balance the two terms in the total loss function.

Here, to evaluate the impact of the short-connection loss term and the effect of non-linearities, an ablation study is conducted by setting λ to zero and by varying the value of c in LeakyReLU. See the section 4.2.3.3 for details.

Then, we investigate the performances of a complex implementation of rRAKI, namely crRAKI, and extend its feasibility to 3D. Figure 4.2 illustrates these architectures which are a combination of the complex RAKI (cRAKI) section and of the GRAPPA short connection (rRAKI architecture can be found in Figure 4.1).

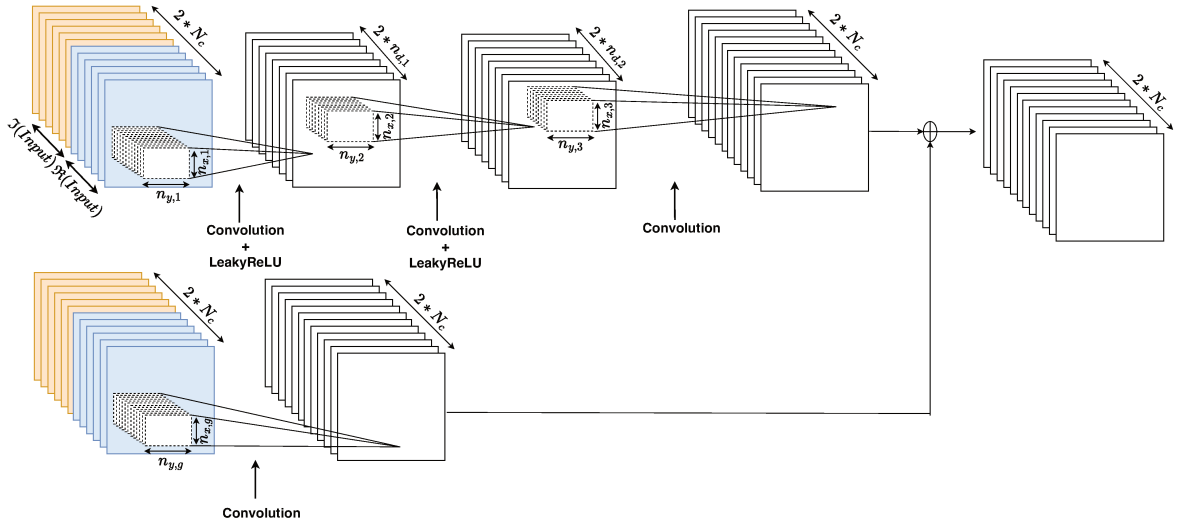

Figure 4.1: rRAKI architecture with its real-valued implementation concatenating the real and imaginary parts.

4.2.2.3 Grid-search: K-fold cross validation

A grid search was performed in conjunction with a K-fold cross-validation. K-fold cross-validation involves dividing the available data into K equally sized subsets, or "folds," and then training and validating the model K times. The advantage of employing K-fold cross-validation is its ability to provide a more robust evaluation of a model's performance, particularly in terms of its generalization capabilities. In the context of the k-space reconstruction task, it is important to note that the value of K is equal to R, reflecting the alignment between the number of folds used for cross-validation from the R datasets discussed in Section 4.2.1, meaning that R-folds can be generated from this dataset. In each round of training and validation, one of the R folds was used as the validation set, whereas the remaining R-1 folds were used as the training set. The model was trained on the training set and its performance was evaluated on the validation set. This process was repeated R times, and each fold was used once as the validation set. The results from each round of validation were averaged to obtain an overall estimate of the model's performance. The evaluation metric for model performance

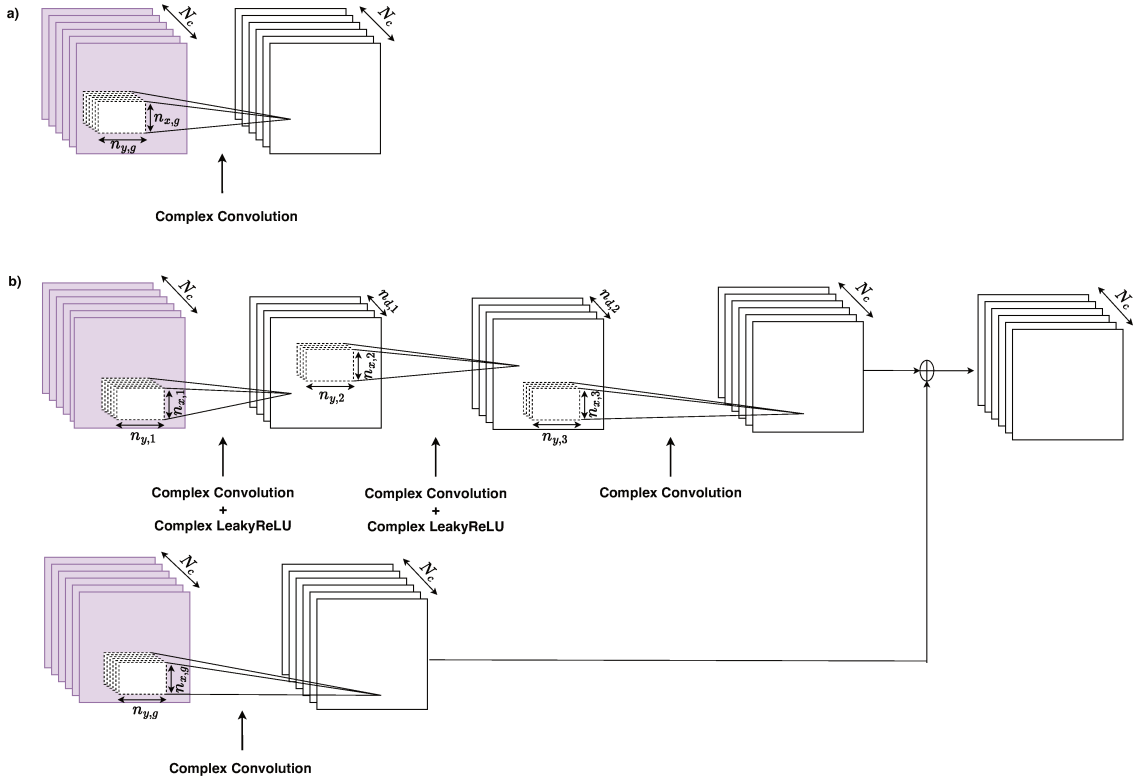

Figure 4.2: Network architectures a) GRAPPA b) crRAKI (shown here for the 2D case). These architectures are defined with few hyperparameters (see text for explanations). Both RAKI and crRAKI use two paths to process a single input multi-coil *undersampled k-space*, which is drawn twice here to improve clarity, a GRAPPA short connection and a cRAKI section, whose outputs are summed to produce the predicted output.

was the MSE in the k-space. By averaging the results from each round of validation, we obtained a more accurate representation of the model's overall performance and its potential for generalization. This approach is typically more reliable than simply evaluating the model on a single training/testing split, because it provides a more robust measure of how well the model is likely to perform on new, unseen data. Details are provided on the grid search parameter selection and implementation in section 4.2.3.

4.2.3 Implementation details

4.2.3.1 Code and computing resources

The study employed Python 3.7 as the software platform for development. Pytorch 1.9, a deep learning framework, was utilized as the core tool for building and training the model. To ensure efficient computation, the training process was performed on a GPU-powered slurm cluster equipped with an NVIDIA GeForce RTX 2080 Ti GPU.

4.2.3.2 Model training

The reference ACS data were first normalized by computing the 2-norm of all samples and coils and subsequently dividing each data point by this value. The reconstruction models were trained using a supervised learning approach, in which the MSE was employed as the objective function. The MSE was implemented in either the real or complex domain depending on the specific context. A maximum of 1000 training iterations, referred to as epochs, were conducted using the ADAM optimizer with a learning rate of 10^{-3}. An early stopping strategy was implemented to prevent overfitting. The performance of the model was evaluated using a validation set at each epoch. Training was terminated when the performance on the validation set began to deteriorate, specifically when the MSE loss on the validation set increased for more than 20 consecutive epochs. The model parameters corresponding to the epoch with the lowest MSE loss in the validation set were retained in the final reconstruction. PyTorch default weight initialization is used, where the weights of a layer are random values drawn from a normal distribution with zero mean and standard deviation of $\sqrt{\frac{2}{n_{\text{in}}}}$, where n_{in} is the number of input units in the layer.

4.2.3.3 Network architecture hyperparameters

Here, the LeakyReLU nonlinearity parameter (c, see section 4.2.2.1) and kernel sizes and depths (see Figure 4.2) must be chosen. To this end, an initial grid search[136] was conducted to study the effect on nonlinearities by tuning c, building on prior knowledge from the RAKI /rRAKI[5, 356] in which 2D models were used. Subsequently, using the best c value, the best architecture was found using a grid search for each method.

Non-linearity study for rRAKI and crRAKI

Using T2W spin-echo data (separated ACS mode), we evaluated the effect of adding this mixed loss and determined the optimal value for c. The architecture kernel and depth sizes were identical to those of the original rRAKI architecture and the model was trained using K-fold cross-validation for $c \in [\![0, 0.01, 0.1, 0.2, 0.3, 0.5, 0.7, 0.9, 1]\!]$ and $\lambda \in [\![0, 1]\!]$. The value of c that produces the smallest MSE is selected as the optimal choice. To validate the significance of the optimal c value, one-tailed Student's t-tests were performed to compare its performance with that of the models that incorporate other values of c.

Kernel sizes and depths

Using the optimal c obtained above, optimization of the architecture was performed considering for GRAPPA in 2D the following kernel sizes: $[n_{x,g}, n_{y,g}]$ = [3,2], [5,2], [5,4], [5,5], [7,2], [7,4], [9,2], [9,4].

To choose rRAKI and crRAKI hyperparameters in 2D, all the combinations of the following were tested:

- first layer size $[n_{x,1}, n_{y,1}]$=[5,2] and [5,4],
- second layer size set to $[n_{x,2}, n_{y,2}]$=[1,1],
- layer depths $2 \times n_{d,1}$ and $2 \times n_{d,2}$ = 32, 16 and 8,

- third layer size [$n_{x,3}$, $n_{y,3}$] = [3,2] and [3,4],
- residual layer size [$n_{x,g}$, $n_{y,g}$] = [5,2], [1,1],

leading to 72 different architectures for both rRAKI and crRAKI.

rRAKI and crRAKI models using these hyperparameter sets were trained for all slices and folds of the T2W spin-echo ACS data. This leads to a validation loss distribution for each hyperparameter set. A one-tailed Student's t-test was then performed to test the hypothesis that the architecture kernel size and depth provide a significantly smaller MSE, on average, because this R is significantly better than the other hyperparameter sets. The architecture hyperparameters for which this hypothesis can be rejected are then considered valid architecture choices ($p<0.05$).

4.2.3.4 Statistics

When evaluating the performance prospectively from the ACS (before the trained model was applied to new data), the method with the lowest mean k-space MSE over all slices was determined and considered as the best one for a given reconstruction strategy. A one-tailed Student's t-test was then performed assuming unequal variance to determine if other strategies statistically produced a larger k-space MSE on average. To evaluate the performance retrospectively in terms of image metrics (when reference images reconstructed from full k-spaces were available), the method providing the best results on average was compared to the others using a one-tail paired Student's t-test (consistent with earlier work[356]) after a normality test (Shapiro-Wilk). For all tests, alpha-risks of 0.05 were considered significant.

4.2.3.5 fastMRI and separated-ACS mode data

This thesis evolves around data from fastMRI[353] and in-house data. Reconstruction was performed on the fastMRI dataset, which encompasses integrated-mode ACS and in-house separated-mode ACS scans. For the latter, experiments were performed on two healthy volunteers at 3T (Vida; Siemens, Erlangen, Germany). Written informed consent was obtained and the study protocol was approved by the ethics committee. Acquisitions were performed on the brain using T2W, 2D/3D multi-echo gradient echo (GRE), and 3D MPRAGE sequences. The relevant sequence parameters are listed in Table 4.1. Raw k-space data were converted into an ismrm raw data format[130] and saved for offline processing.

fastMRI data

Twenty volumes were randomly selected for each of the four different brain contrasts available, namely FLAIR, T1W, T2W, and T1POST, resulting in approximately 300 slices for each contrast from the fastMRI dataset. In order to simulate undersampled acquisitions, full k-spaces were retrospectively undersampled for R=4 keeping 24 and 40 ACS lines.

T2-weighted spin-echo

Interleaved multi-slice axial T2W dual-echo (TE_1/TE_2=11/90 ms) 2D spin-echo brain images were acquired with a 64-channel head coil (N_c=52 selected). Field-of-View (FOV) was 250x187.5 mm and

Table 4.1: Main acquisition parameters for the reference data (ACS), as well as for the full and undersampled k-spaces. Note that the scan time for the ACS is significantly smaller than the undersampled k-space acquisition time.

Sequence	**T2W Spin Echo**	**2D Multi echo GRE**	**3D MPRAGE**	**3D Multi Echo GRE**
ACS size	64x40(x49)	64x64(x24)	64x40x40	64x64x64
ACS scan time (s)	8.2	5	5	13
Acquisition Matrix (Full) ($N_x \times N_y \times N_z$)	256x192(x49)	256x205(x24)	256x256x224	256x208x160
Scan Time (Full, min)	4.1	2.7	9.8	19.4
$R = R_y \times R_z$	Full, 2, 3, 4, 5, 6	Full, 2, 3, 4, 5, 6	Full, 2x1,3x1	2x1,3x1,2x2,3x2,3x3,4x3

Abbreviations: R, acceleration factor acquired in prospective mode.
ACS was always acquired in a separated mode ($TR_{ACS} \sim 3ms$, fast low flip angle gradient echo scan). For 2D, ($\times N_z$) stands for the number of slices.

the slice thickness 3 mm. Repetition and inversion times were respectively 9020 and 900 ms, and flip angle 150°. Fully-sampled k-space (4.1 minutes), as well as undersampled ones with R=2 to 6 were acquired (see Table 4.1).

3D MPRAGE

T1W 3D-MPRAGE brain images were obtained with a 64-channel head coil (N_c = 64 selected) using the following parameters: FOV = 240x240x200 mm, TR/TE = 2300/2.32ms, flip angle=8°, inversion time = 900 ms. The fully-sampled along with R=2x1,3x1 were also acquired. In addition, the fully sampled acquisition was also retrospectively undersampled to R = 2x1,3x1,2x2,3x2,3x3,4x3.

2D/3D multi-echo GRE

A first setup was used to acquire interleaved multi-slice (24 slices) 2D GRE multi-echo (6 echos, TE1/dTE=3.2/5.23 ms) transverse brain acquisitions with FOV=256x256 mm, slice thickness 3 mm, 80% phase encoding resolution, TR=802 ms, T_{acq}=2.4 min, flip angle = 70°, with a 20-channel coil array (N_c=16 selected). Acquisitions with acceleration factors R=2 were performed with 24 ACS lines in separated mode. In order to study the effect of ACS size, a second setup was used to acquire a fully sampled 2D GRE multi-echo (6 echoes) with a 64-channel coil array (N_c = 52 selected) with 64 ACS lines. Finally, multi-echo (6 echoes, TE1/dTE=4.2/5.25 ms) 3D GRE brain acquisitions with FOV=256x256x160 mm, 80% phase encoding resolution, TR=35 ms, flip angle = 15°, N_c = 64, were acquired. For the latter, the acquisition time for a full k-space acquisition was too long (19.4 min), and only accelerated acquisitions were performed (see Table 4.1).

4.3 Results

Following the ablation study conducted on separated-mode rRAKI using the kernel sizes and depth used in the literature, we deduce the nonlinearity parameters λ and c that provide the best results. Using the latter, the optimal kernel sizes and depths were determined using a grid-search strategy. The effect of nonlinearity on crRAKI was subsequently reported. Utilizing the optimized architectures,

the performance of the models and the influence of nonlinearity on the fastMRI brain dataset (in integrated mode) are then presented. Thereafter, the performances of the optimized GRAPPA, rRAKI, and crRAKI models for T2W and 2D GRE imaging (in separated mode) were analyzed. Additionally, a comprehensive analysis highlighting the impact of ACS on 2D multi-echo GRE sequences was presented. Finally, an extension to 3D MPRAGE and 3D GRE was qualitatively reported.

4.3.1 Grid-search in 2D and separated ACS mode

Detrimental effects of non-linearities

The ablation study (Fig 4.3, Fig 4.4 and Tables 4.2 and 4.3) in the separated-mode indicates that the short-connection loss plays a significant role in the overall performance of rRAKI and crRAKI, respectively, as removing this loss resulted in a notable decrease in the reconstruction performance and visual quality for $c = 0$ (pure ReLU).

Upon visual inspection, artifactual dots were observed for both rRAKI and crRAKI with $c \leq 0.5$. The qualitative performance was confirmed by the quantitative performance evaluated using KSPACE-MSE from k-fold cross validation, NRMSE, NMAE, SSIM, Blur metrics, and COBRAI on the final reconstruction. The KSPACE-MSE, NRMSE, and NMAE metrics demonstrated the best performance when rRAKI was linear, exhibiting superior reconstruction quality.

Notably, the KSPACE-MSE was minimal for $\lambda = 0$ and $c = 1$ for both rRAKI and crRAKI, without a significant difference compared to $\lambda = 1$ and $c = 1$ for rRAKI. These equivalent optimal conditions also exhibited a smaller COBRAI. This shows that the KSPACE-MSE obtained from K-fold cross-validation can be used for hyperparameter selection. This implies that the linear variant of rRAKI/crRAKI performed better and that there was no need to add a short-connection loss during training. In the following, we report the results of training all models with $\lambda = 0$ and $c = 1$ for both rRAKI and crRAKI for the separated mode.

Kernels size and depth

The grid-search led to a list of valid architecture hyperparameters for GRAPPA, rRAKI and crRAKI (Table 4.4). Based on this list, GRAPPA and rRAKI sizes were chosen to be the same as in earlier RAKI-based works[5, 356]. For crRAKI, the residual layer size did not significantly influence the performance; therefore, the smallest size was selected for this layer to reduce the total number of model parameters. The selected parameters are listed in Table 4.5. To extend the models in 3D, similar parameters were chosen, additionally setting $n_{z,l}$ to the value $n_{y,l}$ for $l = 1, 3$, and g(GRAPPA short connection) and setting the layer depth $n_{d,2}^{opt}$ to 16 to reduce the computational overhead. The next section reports the results obtained with these fixed hyperparameters.

4.3.2 2D results

4.3.2.1 fastMRI retrospective undersampling (integrated ACS mode)

Overall, the results obtained with our implementation are consistent with earlier studies comparing the performances of rRAKI and GRAPPA(TVP) with ACS 40 in most cases and for all reported metrics

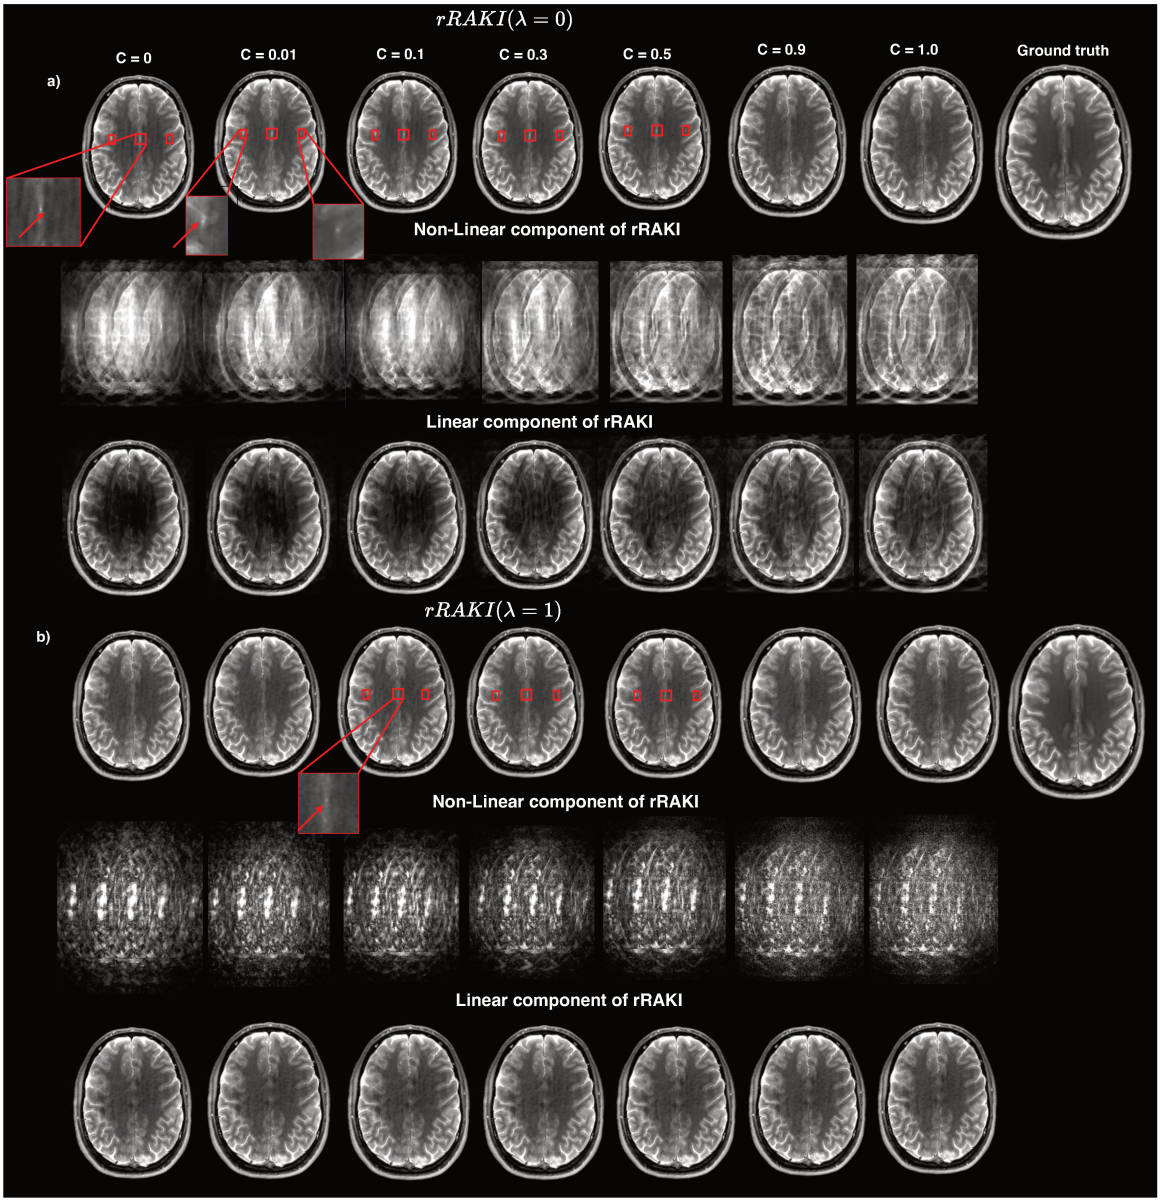

Figure 4.3: Effect of non-linearities on rRAKI with LeakyReLU activation function by varying the coefficient parameter c, where a) corresponds to $\lambda = 0$ and b) corresponds to $\lambda = 1$. The first row is the reconstructed image estimate, the second raw is an image generated from the RAKI section only (F in [356]), the third row is an image generated from the short-connection section of the model (G in [356]). These results are consistent with what was reported in this latter work, with an already meaningful estimate coming from the GRAPPA short-connection of the model. However, artefactual dots are seen (red boxes), which have been magnified for $c = 0$ and $c = 0.01$ at $\lambda = 0$ and $c = 0.1$ at $\lambda = 1$. These artefactual dots are noticeable visually for up to $c < 0.9$ both for $\lambda = 0$ and $\lambda = 1$. There is a decrease in the intensity of these artifacts when $\lambda = 1$, but they still persist in the final reconstruction. Overall, the optimal performance is achieved when $c = 1$ and $\lambda = 1$, indicating the importance of using a linear RAKI section, regardless of the chosen loss function.

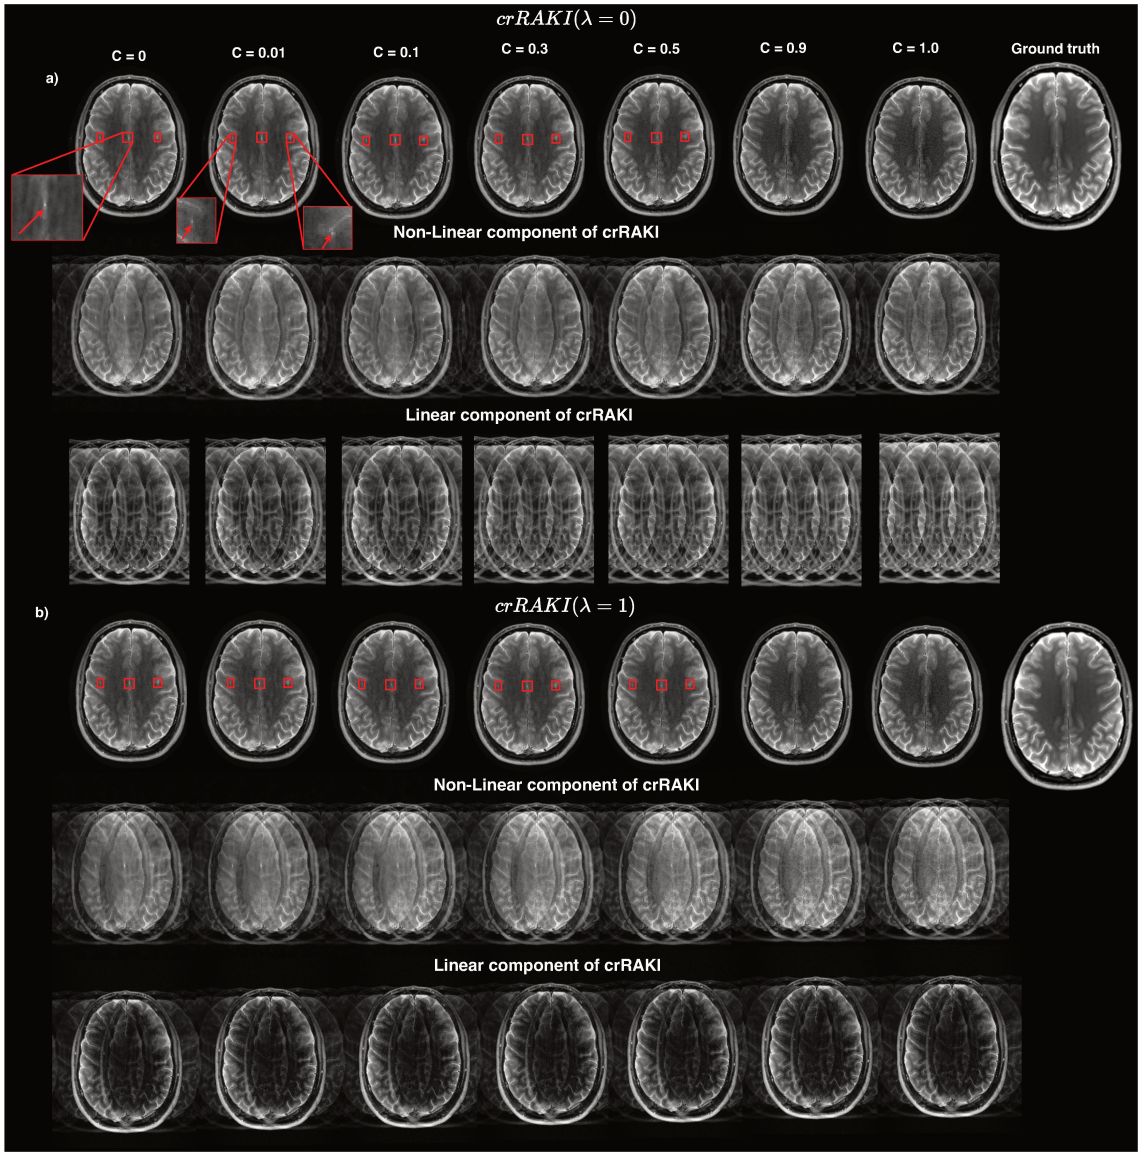

Figure 4.4: Effect of non-linearities on crRAKI with LeakyReLU activation function by varying the coefficient parameter c, where a) corresponds to $\lambda = 0$ and b) corresponds to $\lambda = 1$. The first row is the reconstructed image estimate, the second raw is an image generated from the complex RAKI section only (F in [356]), the third row is an image generated from the short-connection section of the model (G in [356]). Artefactual dots are seen (red boxes), which have been magnified for $c = 0$ and $c = 0.01$ at $\lambda = 0$ and $c = 0.1$ at $\lambda = 1$. These artefactual dots are noticeable visually for up to $c < 0.9$ both for $\lambda = 0$ and $\lambda = 1$. There is a decrease in the intensity of these artifacts when $\lambda = 1$, but they still persist in the final reconstruction. Overall, the optimal performance is achieved when $c = 1$ and $\lambda = 0$.

Table 4.2: Effect of non-linearities on rRAKI with LeakyReLU activation function by varying the parameter c between 0 and 1, with $\lambda = 0$ or 1. K-FOLD cross-validation was performed to determine the optimal parameters based on K-SPACE MSE, which corresponds to the model loss on the validation set (mean over all folds). The model loss decreases with increasing c values for both λ conditions. For a given c, the MSE loss was lower when $\lambda = 1$ compared to $\lambda = 0$ when $c < 1$, indicating that rRAKI performed better with the additional short-connection loss term. However, the best performance was observed when $c = 1$ (linear model) and $\lambda = 0$. The final reconstruction performance of the trained model was evaluated using NRMSE, Normalized Mean Absolute Error (NMAE) and PSNR, yielding the best results for $\lambda = 0$ and $c = 1$. Metrics are better for $c = 1$, signifying linear models are the best. We conclude that non-linearity degrades the performances and that an additional loss term is then not needed. The best-performing results are denoted in **bold**, and metrics without significatively different mean are indicated in *italics*.

λ	C	KSPACE-MSE (1e-7)	NRMSE	NMAE	SSIM	PSNR	Blur Metrics	COBRAI
REF		**K-FOLD**					0.340 ± 0.065	
	0.000	4.818 ± 5.103	0.125 ± 0.049	0.114 ± 0.042	0.943 ± 0.026	34.368 ± 2.309	0.338 ± 0.064	0.365 ± 0.092
	0.010	4.750 ± 5.037	0.124 ± 0.049	0.113 ± 0.043	0.943 ± 0.026	34.398 ± 2.364	0.338 ± 0.065	0.364 ± 0.094
	0.100	4.342 ± 4.657	0.121 ± 0.048	0.110 ± 0.042	0.945 ± 0.026	34.719 ± 2.334	0.338 ± 0.064	0.354 ± 0.091
	0.200	3.710 ± 3.911	0.115 ± 0.046	0.106 ± 0.040	0.946 ± 0.026	35.158 ± 2.297	0.337 ± 0.065	0.340 ± 0.086
0	0.300	3.165 ± 3.300	0.111 ± 0.044	0.102 ± 0.039	0.948 ± 0.025	35.501 ± 2.232	0.337 ± 0.065	0.327 ± 0.081
	0.500	2.272 ± 2.302	0.106 ± 0.041	0.098 ± 0.037	$\mathbf{0.949 \pm 0.025}$	35.905 ± 2.192	0.336 ± 0.065	0.306 ± 0.072
	0.700	1.476 ± 1.325	0.104 ± 0.039	0.097 ± 0.037	0.948 ± 0.026	36.069 ± 2.148	0.335 ± 0.065	0.291 ± 0.068
	0.900	0.767 ± 0.486	0.100 ± 0.039	0.095 ± 0.038	0.947 ± 0.028	36.419 ± 2.223	0.333 ± 0.066	0.264 ± 0.056
	1.000	$\mathbf{0.602 \pm 0.357}$	$\mathbf{0.099 \pm 0.040}$	$\mathbf{0.094 \pm 0.039}$	0.947 ± 0.029	$\mathbf{36.556 \pm 2.254}$	0.333 ± 0.066	0.250 ± 0.050
	0.000	1.718 ± 1.338	0.113 ± 0.043	0.107 ± 0.042	0.940 ± 0.030	35.426 ± 2.215	0.333 ± 0.065	0.294 ± 0.071
	0.010	1.679 ± 1.240	0.112 ± 0.043	0.106 ± 0.042	0.940 ± 0.030	35.452 ± 2.238	0.333 ± 0.065	0.294 ± 0.071
	0.100	1.663 ± 1.328	0.110 ± 0.042	0.104 ± 0.042	0.942 ± 0.029	35.647 ± 2.224	0.333 ± 0.065	0.288 ± 0.069
	0.200	1.596 ± 1.228	0.108 ± 0.043	0.102 ± 0.042	0.943 ± 0.030	35.829 ± 2.237	0.333 ± 0.065	0.282 ± 0.067
1	0.300	1.487 ± 1.159	0.107 ± 0.043	0.101 ± 0.041	0.944 ± 0.029	35.952 ± 2.181	0.333 ± 0.066	0.278 ± 0.064
	0.500	1.275 ± 0.923	0.103 ± 0.041	0.097 ± 0.04	0.946 ± 0.029	36.229 ± 2.225	0.333 ± 0.066	0.271 ± 0.062
	0.700	1.027 ± 0.710	0.101 ± 0.041	0.096 ± 0.04	0.946 ± 0.029	36.439 ± 2.214	0.332 ± 0.066	0.263 ± 0.058
	0.900	0.703 ± 0.419	0.100 ± 0.041	0.095 ± 0.041	0.945 ± 0.030	$\mathit{36.537 \pm 2.244}$	0.331 ± 0.066	0.251 ± 0.053
	1.000	$\mathit{0.622 \pm 0.363}$	0.101 ± 0.043	0.097 ± 0.043	0.944 ± 0.032	36.451 ± 2.314	$\mathbf{0.331} \pm \mathit{0.066}$	$\mathbf{0.248 \pm 0.051}$

Standard deviations are reported over all slices.

Table 4.3: Effect of non-linearities on crRAKI with LeakyReLU activation function by varying the parameter c between 0 and 1, with $\lambda = 0$ or 1. K-FOLD cross-validation was performed to determine the optimal parameters based on K-SPACE MSE, which corresponds to the model loss on the validation set (mean over all folds). The model loss decreases with increasing c values for both λ conditions. For a given c, the MSE loss was lower when $\lambda = 1$ compared to $\lambda = 0$ when $c < 1$, indicating that crRAKI performed better with the additional short-connection loss term. However, the best performance was observed when $c = 1$ (linear model) and $\lambda = 0$. The final reconstruction performance of the trained model was evaluated using NRMSE, MAE and PSNR, yielding the best results for $\lambda = 0$ and $c = 1$. Metrics are better for $c = 1$, signifying linear models are the best. We conclude that non-linearity degrades the performances and that an additional loss term is then not needed. The best-performing results are denoted in **bold**, and metrics without significatively different mean are indicated in *italics*.

λ	C	KSPACE-MSE (1e-7)	NRMSE	NMAE	SSIM	PSNR	Blur Metrics	COBRAI
REF		**K-FOLD**					0.340 ± 0.065	
	0.000	7.003 ± 2.630	0.120 ± 0.045	0.112 ± 0.040	0.948 ± 0.022	34.572 ± 2.229	0.339 ± 0.068	0.400 ± 0.091
	0.010	7.155 ± 2.865	0.105 ± 0.037	0.088 ± 0.034	0.972 ± 0.026	35.815 ± 2.027	0.311 ± 0.068	0.370 ± 0.109
	0.100	5.872 ± 2.217	0.096 ± 0.035	0.082 ± 0.033	0.974 ± 0.026	36.568 ± 2.144	0.310 ± 0.068	0.349 ± 0.106
	0.200	4.840 ± 1.768	0.089 ± 0.033	0.076 ± 0.033	0.975 ± 0.027	37.392 ± 2.130	0.310 ± 0.067	0.325 ± 0.104
0	0.300	4.239 ± 1.558	0.083 ± 0.032	0.071 ± 0.033	0.975 ± 0.028	38.009 ± 2.298	0.310 ± 0.067	0.305 ± 0.104
	0.500	2.735 ± 0.970	0.075 ± 0.031	0.065 ± 0.033	**0.976 ± 0.029**	38.982 ± 2.421	0.309 ± 0.067	0.263 ± 0.100
	0.700	1.876 ± 0.633	**0.071 ± 0.030**	**0.063 ± 0.033**	0.975 ± 0.029	39.555 ± 2.556	0.308 ± 0.067	0.221 ± 0.098
	0.900	1.210 ± 0.415	0.071 ± 0.031	0.063 ± 0.033	0.973 ± 0.030	**39.558 ± 2.649**	0.307 ± 0.066	0.193 ± 0.098
	1.000	**1.105 ± 0.410**	0.072 ± 0.031	0.064 ± 0.033	0.972 ± 0.030	39.437 ± 2.672	**0.307 ± 0.067**	**0.188 ± 0.098**
	0.000	5.256 ± 2.027	0.114 ± 0.043	0.112 ± 0.042	0.938 ± 0.026	35.148 ± 2.071	0.330 ± 0.062	0.341 ± 0.083
	0.010	4.954 ± 1.842	0.092 ± 0.030	0.079 ± 0.025	0.972 ± 0.016	36.971 ± 1.922	0.308 ± 0.066	0.295 ± 0.078
	0.100	4.190 ± 1.586	0.087 ± 0.028	0.075 ± 0.025	0.973 ± 0.017	37.485 ± 1.885	0.307 ± 0.066	0.279 ± 0.075
	0.200	3.468 ± 1.193	0.101 ± 0.038	0.099 ± 0.038	0.945 ± 0.025	36.306 ± 2.001	0.330 ± 0.063	0.299 ± 0.069
1	0.300	2.881 ± 0.951	0.096 ± 0.036	0.095 ± 0.036	0.947 ± 0.025	36.725 ± 1.982	0.330 ± 0.063	0.280 ± 0.064
	0.500	2.123 ± 0.653	0.089 ± 0.034	0.088 ± 0.034	0.950 ± 0.025	37.387 ± 2.007	0.329 ± 0.065	0.245 ± 0.052
	0.700	1.578 ± 0.466	0.087 ± 0.033	0.085 ± 0.034	0.950 ± 0.026	37.707 ± 2.105	0.329 ± 0.065	0.214 ± 0.043
	0.900	1.318 ± 0.397	0.088 ± 0.034	0.086 ± 0.035	0.947 ± 0.027	37.643 ± 2.194	0.327 ± 0.066	0.193 ± 0.037
	1.000	1.280 ± 0.395	0.088 ± 0.035	0.087 ± 0.035	0.946 ± 0.028	37.569 ± 2.233	0.327 ± 0.066	0.190 ± 0.036

Standard deviations are reported over all slices.

Table 4.4: List of valid models obtained through a grid search. In order to identify the most suitable set of hyperparameters, a grid search based on the K-fold cross-validation was performed and based on the MSE metric in k-space the potential list of architecture was obtained. To determine whether the performance differences observed across different hyperparameters were statistically significant, we conducted a one-tailed Student's t-test, which enabled us to compare the average performance metric among various hyperparameters. The architecture was optimized for GRAPPA in 2D, using kernel sizes of $[n_{x,g}, n_{y,g}]$ = [3,2], [5,2], [5,4], [5,5], [7,2], [7,4], [9,2], and [9,4]. In 2D, to determine the hyperparameters for both rRAKI and crRAKI, all possible combinations were tested for the following: first layer size $[n_{x,1}, n_{y,1}]$ = [5,2] and [5,4], second layer size set to $[n_{x,2}, n_{y,2}]$ = [1,1], layer depths $2 \times n_{d,1}$ and $2 \times n_{d,2}$ = 32, 16 and 8, third layer size $[n_{x,3}, n_{y,3}]$ = [3,2] and [3,4], and residual layer size $[n_{x,g}, n_{y,g}]$ = [5,2] and [1,1]. This led to a total of 72 different architectures for both rRAKI and crRAKI. Only non-significantly different cases are reported. In the column labelled 'R', bold indicates the best solutions for a given R. As can seen, several kernel sizes can be selected for GRAPPA, the best being [5, 4] which is suitable for all acceleration rates and is a commonly-used kernel size. For rRAKI, several hyperparameters are possible, and the ones used (in bold) correspond to a situation suitable for several acceleration rates. For crRAKI, as can be seen, several kernel sizes and depths can be chosen, and we selected one with a smaller kernel size in the residual layer which is suitable for several acceleration rates.

Model	**KERNEL SIZE**	**KERNEL DEPTH**	**R**
GRAPPA	[3, 4]	-	**2**, **3**, 4, **6**
	[5, 4]	-	2, 3, **4**, **5**, 6
	[7, 4]	-	6
rRAKI	[5 2], [1 1], [3 2], [1 1]	32, 32	2, 3, 4, 5
	[5 2], [1 1], [3 2], [5 2]	8, 8	2, **3**, 4, 6
	[5 2], [1 1], [3 2], [5 2]	8, 16	**2**, 3, 4, 5, 6
	[5 2], [1 1], [3 2], [5 2]	8, 32	2, 3, 4, 5
	[5 2], [1 1], [3 2], [5 2]	16, 8	2, 3, 4, 6
	[5 2], [1 1], [3 2], [5 2]	16, 16	2, 3, 4, 5, 6
	[5 2], [1 1], [3 2], [5 2]	16, 32	2, 3, 4, 5
	[5 2], [1 1], [3 2], [5 2]	**32, 8**	2, 3, 4
	[5 2], [1 1], [3 2], [5 2]	32, 16	2, 3, 4, 5, 6
	[5 2], [1 1], [3 2], [5 2]	32, 32	2, 3, **4**, **5**
crRAKI	**[5 2], [1 1], [3 2], [1 1]**	**16, 32**	2, 3, 4, 5
	[5 2], [1 1], [3 2], [1 1]	32, 32	2, 3, 4, 5
	[5 2], [1 1], [3 2], [5 2]	8, 8	2, 3, 4, 5, 6
	[5 2], [1 1], [3 2], [5 2]	8, 16	2, 3, 4, 5
	[5 2], [1 1], [3 2], [5 2]	8, 32	**2**, **3**, **4**, **5**
	[5 2], [1 1], [3 2], [5 2]	16, 8	2, 3, 4, 5, 6
	[5 2], [1 1], [3 2], [5 2]	16, 16	2, 3, 4, 5
	[5 2], [1 1], [3 2], [5 2]	16, 32	2, 3, 4, 5
	[5 2], [1 1], [3 2], [5 2]	32, 8	2, 3, 4, 5, **6**
	[5 2], [1 1], [3 2], [5 2]	32, 16	2, 3, 4, 5, 6
	[5 2], [1 1], [3 2], [5 2]	32, 32	2, 3, 4, 5

Table 4.5: Model hyperparameter choice.

Model	**2D** $[n_{x,l}, n_{y,l}, n_d^{in}, n_d^{out}]$	**3D** $[n_{x,l}, n_{y,l}, n_{z,l}, n_d^{in}, n_d^{out}]$
GRAPPA (complex)	[5, 4, N_c, N_c]×2	[5, 4, 4, N_c, N_c]×2
rRAKI (real)		
layer 1	[5, 2, 2 × N_c, 32]	[5, 2, 2, 2 × N_c, 32]
layer 2	[1, 1, 32, 8]	[1, 1, 1, 32, 8]
layer 3	[3, 2, 8, 2 × N_c]	[3, 2, 2, 8, 2 × N_c]
residual layer	[5, 2, 2 × N_c, 2 × N_c]	[5, 2, 2, 2 × N_c, 2 × N_c]
crRAKI (complex)		
layer 1	[5, 2, N_c, 16]×2	[5, 2, 2, N_c, 16]×2
layer 2	[1, 1, 16, 32]×2	[1, 1, 1, 16, 16]×2
layer 3	[3, 2, 32, N_c]×2	[3, 2, 2, 16, N_c]×2
residual layer	[1, 1, N_c, N_c]×2	[1, 1, 1, N_c, N_c]×2

Kernel sizes are reported for each layer and model. N_c is the number of coils. n_d^{in} and n_d^{out} stand for input and output depth, respectively. For models implementing complex convolutions, two kernels (2×) of the same size are used.

that were used[356]. This provides confidence in our TVP scheme with an early stopping and reconstruction implementation. GRAPPA(TVP) sometimes outperformed GRAPPA(TR) with Tikhonov regularization in the case of FLAIR and T1W. Figures 4.5, 4.6, 4.7, and 4.8 provide a qualitative evaluation. We evaluated linear (L) rRAKI/crRAKI against nonlinear (NL) counterparts in conjunction with GRAPPA. Here, the linear model is defined with the parameters $\lambda = 0$ and $c = 1$. The nonlinear model employed $\lambda = 1$ and $c = 0$. Specifically, crRAKI(L) exhibited superior visual quality compared with GRAPPA, crRAKI(NL), and rRAKI(L/NL) for ACS = 24, except for T1W. Upon examining the residual maps, distinct patterns were observed across different methods. GRAPPA, rRAKI(L), and crRAKI(L) demonstrated higher errors in the central region, whereas rRAKI(NL) exhibited structured residuals, and crRAKI(NL) exhibited artifactual dots. For ACS = 40, rRAKI(NL) produces high-fidelity reconstructions for T1W data, visually surpassing the performance of GRAPPA and crRAKI. This observation was corroborated by the presence of noticeable noise in the central region for the GRAPPA and crRAKI(L) reconstructions. However, in the case of FLAIR and T2W data, rRAKI(NL) yielded reconstructions with noticeable blurriness and artifactual dots in the center. A similar observation was made in crRAKI(NL). Additionally, for T1POST, rRAKI(NL) has limitations in capturing fine details that are more visible in the reconstructions than those obtained using GRAPPA and crRAKI(L). Tables 4.6 and 4.7 provide a quantitative evaluation of the fastMRI data in the same line. Conversely, for ACS = 40, rRAKI demonstrated statistically significant improvements over GRAPPA and crRAKI in FLAIR and T1W. Overall, crRAKI(L) consistently outperformed COBRAI across different ACS levels and contrasts, indicating lower artifact levels in the fastMRI dataset.

4.3.2.2 T2 Weighted in separated ACS mode

Figure 4.9 shows a comparison of the quantitative performance for different acceleration rates. The comparison was performed under the same experimental conditions in terms of the number of ACS lines as suggested by Zhang et al.[356]. To determine the significance of the results, a one-tailed

Table 4.6: Mean and standard deviations (over slices) for NRMSE, NMAE, SSIM, PSNR, Blur metrics and COBRAI for the tested reconstruction methods for FLAIR, T1W, T1POST and T2W from the fastMRI dataset at R = 4 for ACS=40. The method with the best metric **in bold** is compared with the other methods. Metrics without statistical significance (p > 0.05 - 5% level) are indicated in *italics*. While rRAKI has better image quality metrics (NRMSE, NMAE, SSIM, PSNR) in the case of T1W and FLAIR images, this is not the case for T1POST and T2W.

TYPE	METHOD	NRMSE	NMAE	SSIM	PSNR	BLUR METRICS	COBRAI
FLAIR	GRAPPA(TR)	0.093 ± 0.035	0.086 ± 0.033	0.903 ± 0.054	33.723 ± 3.312	0.338 ± 0.033	0.171 ± 0.063
	GRAPPA	0.089 ± 0.026	0.062 ± 0.024	0.900 ± 0.055	33.489 ± 3.298	**0.335 ± 0.033**	0.166 ± 0.064
	crRAKI(L)	0.093 ± 0.036	0.085 ± 0.034	0.904 ± 0.054	33.756 ± 3.339	0.337 ± 0.033	**0.162 ± 0.065**
	crRAKI(NL)	0.088 ± 0.026	0.080 ± 0.024	0.907 ± 0.052	33.986 ± 2.784	0.349 ± 0.033	0.179 ± 0.059
	rRAKI(L)	0.095 ± 0.036	0.086 ± 0.034	0.899 ± 0.054	33.392 ± 3.164	0.339 ± 0.03	0.185 ± 0.072
	rRAKI(NL)	**0.082 ± 0.025**	**0.057 ± 0.023**	**0.912 ± 0.050**	**34.020 ± 2.981**	0.356 ± 0.034	0.209 ± 0.072
T1W	GRAPPA(TR)	0.098 ± 0.041	0.096 ± 0.038	0.884 ± 0.071	35.714 ± 3.449	0.399 ± 0.079	0.189 ± 0.067
	GRAPPA	0.097 ± 0.037	0.079 ± 0.034	0.884 ± 0.072	35.573 ± 3.419	**0.399 ± 0.079**	0.190 ± 0.068
	crRAKI(L)	0.098 ± 0.041	0.095 ± 0.038	0.885 ± 0.073	35.779 ± 3.540	0.400 ± 0.080	**0.183 ± 0.071**
	crRAKI(NL)	0.077 ± 0.034	0.076 ± 0.032	0.926 ± 0.051	37.722 ± 2.817	0.426 ± 0.074	0.205 ± 0.080
	rRAKI(L)	0.088 ± 0.038	0.087 ± 0.036	0.902 ± 0.062	36.570 ± 3.228	0.413 ± 0.077	0.198 ± 0.071
	rRAKI(NL)	**0.075 ± 0.03**	**0.062 ± 0.027**	**0.923 ± 0.051**	**37.563 ± 2.898**	0.430 ± 0.078	0.216 ± 0.074
T1POST	GRAPPA(TR)	**0.056 ± 0.031**	0.053 ± 0.028	0.939 ± 0.058	39.185 ± 3.505	0.453 ± 0.094	0.155 ± 0.058
	GRAPPA	0.061 ± 0.030	0.048 ± 0.030	0.934 ± 0.060	38.402 ± 3.428	**0.445 ± 0.090**	0.140 ± 0.059
	crRAKI(L)	0.058 ± 0.031	0.054 ± 0.029	0.938 ± 0.06	**39.020 ± 3.639**	0.449 ± 0.093	**0.131 ± 0.058**
	crRAKI(NL)	0.065 ± 0.027	0.061 ± 0.025	0.93 ± 0.058	37.82 ± 2.851	0.442 ± 0.082	0.137 ± 0.064
	rRAKI(L)	0.060 ± 0.031	0.056 ± 0.029	0.938 ± 0.058	38.641 ± 3.405	0.448 ± 0.089	0.146 ± 0.063
	rRAKI(NL)	0.059 ± 0.03	**0.046 ± 0.029**	**0.940 ± 0.056**	38.648 ± 3.079	0.454 ± 0.089	0.151 ± 0.064
T2W	GRAPPA(TR)	0.060 ± 0.029	0.060 ± 0.028	0.958 ± 0.042	39.879 ± 3.250	0.394 ± 0.043	0.137 ± 0.047
	GRAPPA	0.063 ± 0.025	**0.051 ± 0.025**	0.956 ± 0.042	39.180 ± 3.078	**0.389 ± 0.042**	0.122 ± 0.050
	crRAKI(L)	**0.060 ± 0.029**	0.059 ± 0.028	**0.959 ± 0.043**	**39.897 ± 3.385**	0.392 ± 0.043	**0.117 ± 0.048**
	crRAKI(NL)	0.073 ± 0.025	0.071 ± 0.024	0.947 ± 0.043	38.011 ± 2.567	0.388 ± 0.04	0.123 ± 0.052
	rRAKI(L)	0.065 ± 0.029	0.064 ± 0.027	0.955 ± 0.043	39.107 ± 3.22	0.392 ± 0.042	0.125 ± 0.049
	rRAKI(NL)	0.063 ± 0.025	0.052 ± 0.026	0.957 ± 0.042	39.089 ± 2.988	0.396 ± 0.042	0.132 ± 0.050

Abbreviations: GRAPPA: GRAPPA with TVP, GRAPPA(TR): GRAPPA with Tikhonov Regularization, L: Linear ($\lambda = 0, c = 1$), NL: Non-linear ($\lambda = 1, c = 0$)

Table 4.7: Mean and standard deviations (over slices) for NRMSE, NMAE, SSIM, PSNR, Blur metrics and COBRAI for the tested reconstruction methods for FLAIR, T1W, T1POST and T2W from the fastMRI dataset at R = 4 for ACS=24. The method with the best metric **in bold** is compared with the other methods. Metrics without statistical significance (p > 0.05 - 5% level) are indicated in *italics*. While rRAKI has better image quality metrics (NRMSE, NMAE, SSIM, PSNR) in the case of T1W and FLAIR images, this is not the case for T1POST and T2W.

TYPE	METHOD	NRMSE	NMAE	SSIM	PSNR	BLUR METRICS	COBRAI
FLAIR	**Reference**					0.309 ± 0.046	
	GRAPPA(TR)	0.100 ± 0.037	0.092 ± 0.035	0.893 ± 0.062	33.171 ± 3.298	0.336 ± 0.034	0.179 ± 0.062
	GRAPPA	0.102 ± 0.026	0.070 ± 0.026	0.886 ± 0.062	32.423 ± 3.309	**0.331 ± 0.033**	0.181 ± 0.066
	crRAKI(L)	0.099 ± 0.038	0.090 ± 0.036	**0.895 ± 0.059**	**33.221 ± 3.353**	0.333 ± 0.033	**0.160 ± 0.066**
	crRAKI(NL)	0.106 ± 0.03	0.096 ± 0.028	0.882 ± 0.061	32.372 ± 3.018	0.339 ± 0.032	0.208 ± 0.079
	rRAKI(L)	0.103 ± 0.035	0.093 ± 0.034	0.892 ± 0.057	32.600 ± 3.137	0.340 ± 0.030	0.208 ± 0.076
	rRAKI(NL)	**0.097 ± 0.027**	**0.067 ± 0.026**	*0.895 ± 0.058*	32.557 ± 3.07	0.351 ± 0.032	0.237 ± 0.085
T1W	**Reference**					0.397 ± 0.066	
	GRAPPA(TR)	0.110 ± 0.043	0.107 ± 0.039	0.870 ± 0.077	34.770 ± 3.354	*0.393 ± 0.078*	0.196 ± 0.066
	GRAPPA	0.118 ± 0.042	0.094 ± 0.04	0.872 ± 0.077	34.016 ± 3.359	0.401 ± 0.077	0.217 ± 0.073
	crRAKI(L)	0.106 ± 0.044	0.102 ± 0.041	0.877 ± 0.075	35.137 ± 3.501	**0.393 ± 0.079**	**0.185 ± 0.076**
	crRAKI(NL)	**0.089 ± 0.033**	0.088 ± 0.031	**0.908 ± 0.054**	**36.249 ± 2.430**	0.419 ± 0.070	0.211 ± 0.075
	rRAKI(L)	0.104 ± 0.044	0.101 ± 0.04	0.882 ± 0.072	35.221 ± 3.418	0.405 ± 0.08	0.219 ± 0.078
	rRAKI(NL)	0.093 ± 0.034	**0.075 ± 0.032**	0.901 ± 0.061	35.888 ± 2.885	0.421 ± 0.076	0.229 ± 0.076
T1POST	**Reference**					0.412 ± 0.068	
	GRAPPA(TR)	**0.060 ± 0.031**	0.057 ± 0.029	**0.936 ± 0.062**	**38.612 ± 3.483**	0.453 ± 0.094	0.173 ± 0.062
	GRAPPA	0.082 ± 0.033	0.061 ± 0.032	0.922 ± 0.065	36.107 ± 3.581	0.445 ± 0.090	0.165 ± 0.062
	crRAKI(L)	0.062 ± 0.032	0.057 ± 0.030	0.933 ± 0.063	38.509 ± 3.580	**0.444 ± 0.093**	**0.126 ± 0.058**
	crRAKI(NL)	0.082 ± 0.026	0.076 ± 0.023	0.909 ± 0.062	35.814 ± 2.641	0.430 ± 0.075	0.164 ± 0.074
	rRAKI(L)	0.066 ± 0.031	0.061 ± 0.029	0.933 ± 0.061	37.785 ± 3.398	0.450 ± 0.088	0.165 ± 0.069
	rRAKI(NL)	0.072 ± 0.032	**0.055 ± 0.032**	0.928 ± 0.061	37.042 ± 3.102	0.450 ± 0.086	0.169 ± 0.074
T2W	**Reference**					0.365 ± 0.053	
	GRAPPA(TR)	**0.062 ± 0.030**	0.062 ± 0.028	0.956 ± 0.044	**39.531 ± 3.226**	0.393 ± 0.043	0.145 ± 0.049
	GRAPPA	0.072 ± 0.026	**0.058 ± 0.027**	0.952 ± 0.045	38.058 ± 3.027	**0.387 ± 0.042**	0.140 ± 0.053
	crRAKI(L)	0.064 ± 0.030	0.062 ± 0.029	**0.956 ± 0.045**	39.390 ± 3.332	0.389 ± 0.043	**0.110 ± 0.049**
	crRAKI(NL)	0.092 ± 0.025	0.090 ± 0.024	0.930 ± 0.045	35.925 ± 2.567	0.38 ± 0.038	0.155 ± 0.066
	rRAKI(L)	0.072 ± 0.03	0.069 ± 0.028	0.951 ± 0.045	38.188 ± 3.286	0.392 ± 0.043	0.142 ± 0.055
	rRAKI(NL)	0.077 ± 0.029	0.061 ± 0.030	0.950 ± 0.045	37.518 ± 3.273	0.394 ± 0.042	0.154 ± 0.060

Abbreviations: TR: Traditional (GRAPPA with Tikhonov regularization), L: Linear ($\lambda = 0, c = 1$), NL: Non-linear ($\lambda = 1, c = 0$)

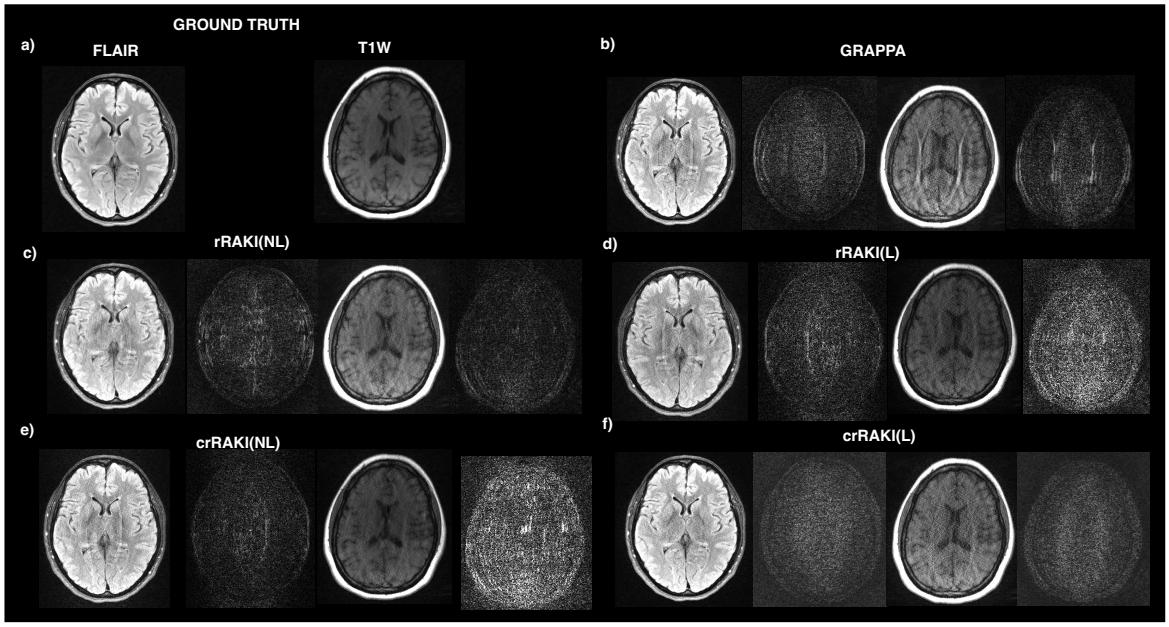

Figure 4.5: Comparative Evaluation of Reconstruction Methods a) Ground Truth b) GRAPPA, c) rRAKI(NL) and d) rRAKI(L) e) ccrRAKI(NL) f) crRAKI(L) for FLAIR and T1W Images with ACS = 24 from the fastMRI Dataset. The crRAKI method demonstrates superior capability in reconstructing high-quality images with enhanced fidelity and accuracy when compared to the GRAPPA and rRAKI. Undersampling artifacts can be observed in both rRAKI and GRAPPA corroborated by the absolute residual maps.

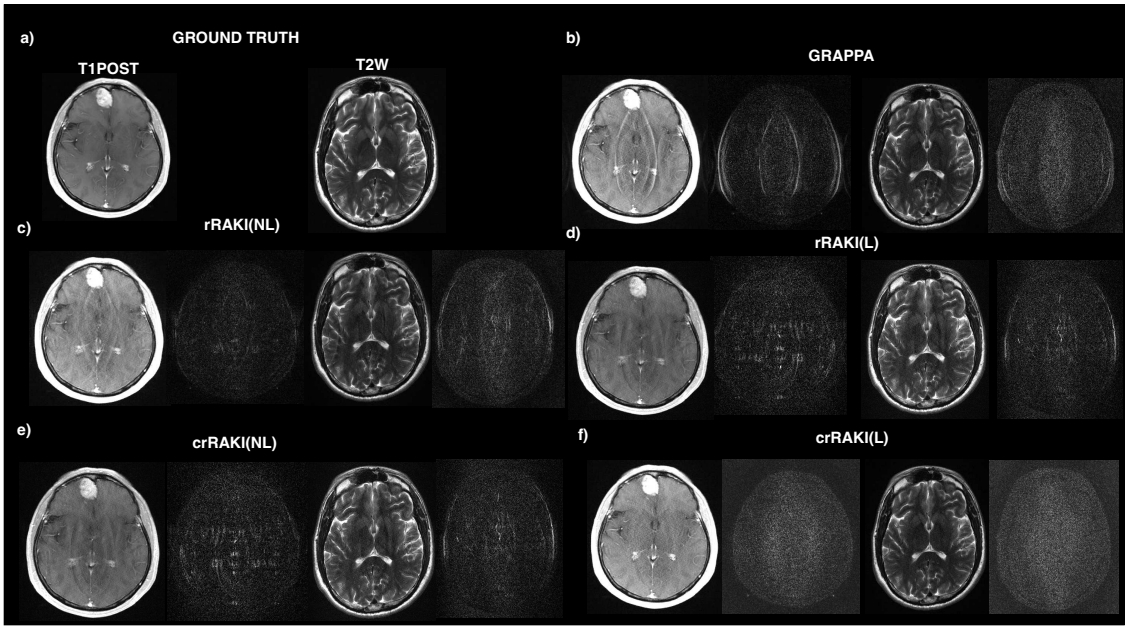

Figure 4.6: Comparative Evaluation of Reconstruction Methods a) Ground Truth b) GRAPPA, c) rRAKI(NL) and d) rRAKI(L) e) crRAKI(NL) f) crRAKI(L) for T1POST, and T2W Images with ACS = 24 from the fastMRI dataset. The crRAKI method demonstrates superior capability in reconstructing high-quality images with enhanced fidelity and accuracy when compared to the GRAPPA and rRAKI. Undersampling artifacts can be observed in both rRAKI and GRAPPA corroborated by the absolute residual maps.

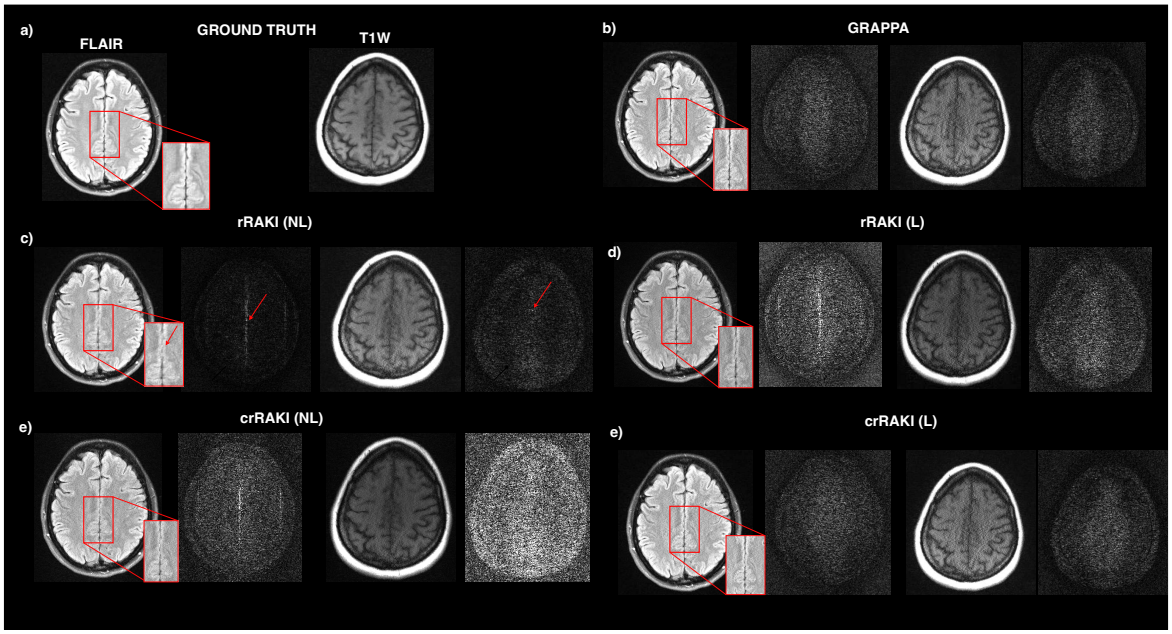

Figure 4.7: Comparative Evaluation of Reconstruction Methods a) Ground Truth b) GRAPPA, c) rRAKI(NL) and d) rRAKI(L) e) crRAKI(NL) f) crRAKI(L) for FLAIR, and T1W Images with ACS = 40 from the fastMRI dataset (best viewed in zoom). The reconstructions obtained with rRAKI demonstrate superior performance for T1-weighted data compared to GRAPPA and crRAKI.

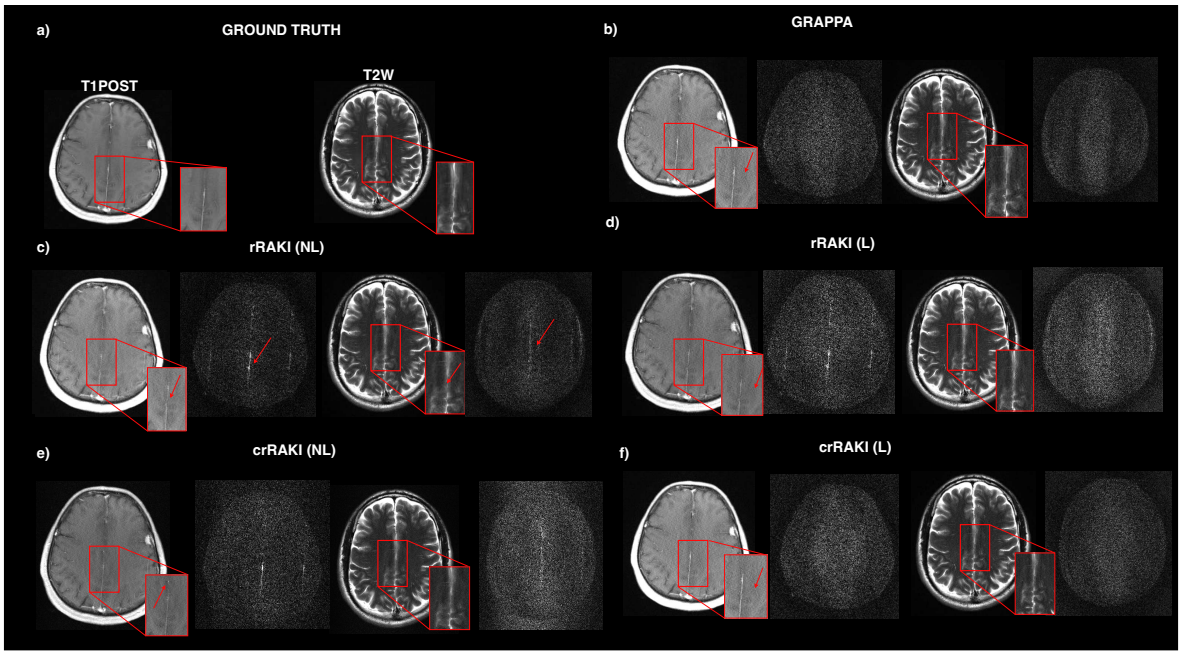

Figure 4.8: Comparative Evaluation of Reconstruction Methods: a) Ground Truth, b) GRAPPA, c) rRAKI(NL), and d) rRAKI(L) e) crRAKI(NL) f) crRAKI(L) for T1POST, and T2W Images with ACS = 40 from the fastMRI dataset (best viewed in zoom). The reconstructions obtained with rRAKI demonstrated superior performance for T1-weighted data compared with GRAPPA and crRAKI.

paired Student's t-test was applied, with the hypothesis that crRAKI is superior to the other methods. The results showed that crRAKI outperformed GRAPPA and rRAKI in all metrics except blur metrics at R=3,4,5,6, the exception being at GRAPPA at R=2. crRAKI demonstrates significant superiority over both rRAKI and GRAPPA in terms of NRMSE, NMAE, PSNR, and COBRAI, particularly at acceleration rates of 4, 5, and 6.

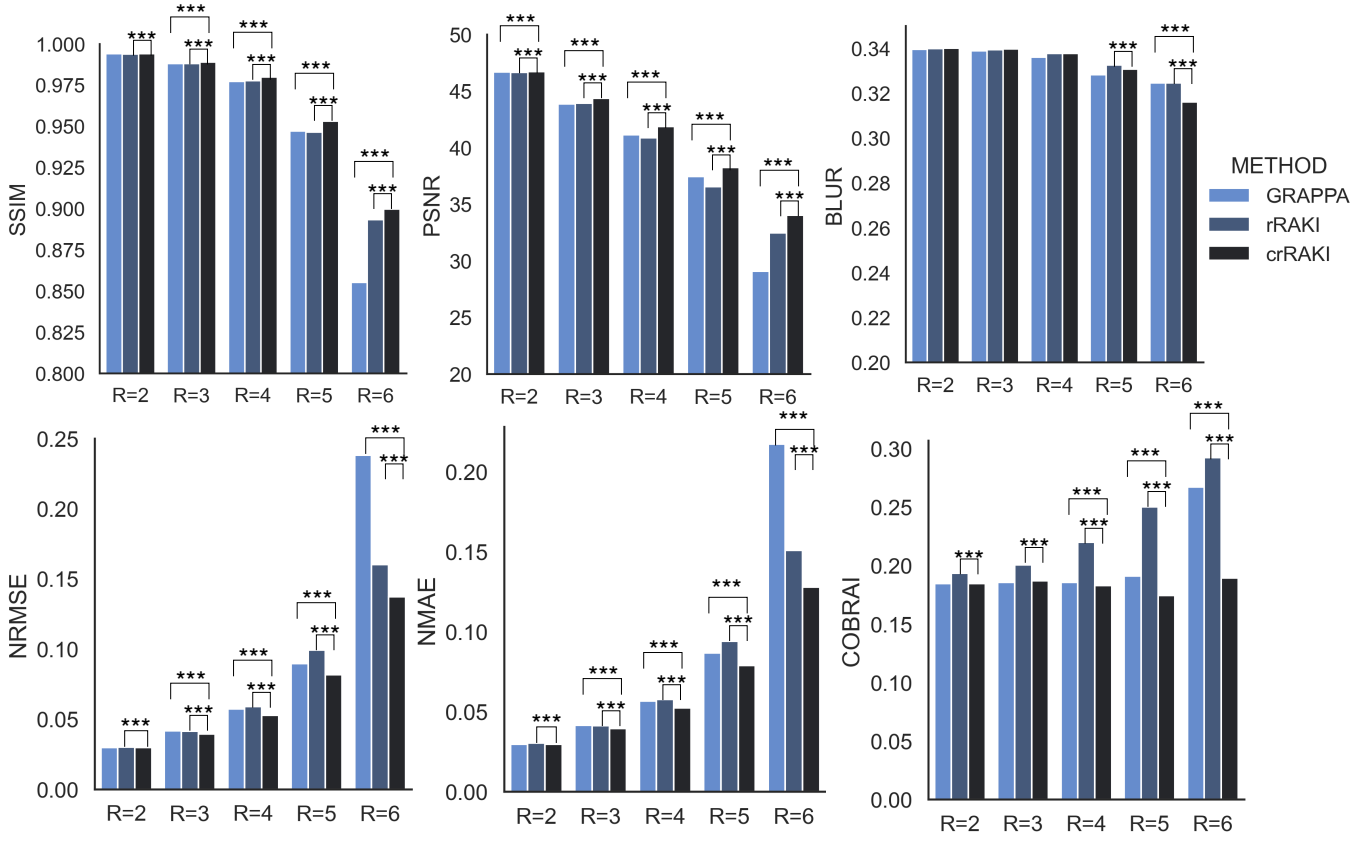

Figure 4.9: Quantitative evaluation of the reconstruction performance of GRAPPA, rRAKI, and crRAKI in terms of SSIM, PSNR, NRMSE, NMAE, Blur and COBRAI at acceleration rates of R = 2, 3, 4, 5, and 6. The values are the mean metric over all 49 slices for the 2 echoes of data from the T2W scan on one volunteer, using the sum-of-squares coil combination. The figure displays the results of the comparison and accompanying p-values (*, **, ***) calculated from a one-tailed paired t-test. A small p-value (typically less than 0.05) indicates strong evidence against the null hypothesis (that the methods have equal performance) and supports the alternative hypothesis (that crRAKI has better performance than the other methods) The symbols *, **, and *** in statistical t-tests are used to indicate the level of statistical significance of the results. * signifies $p < 0.05$ (5% level), ** signifies $p < 0.01$ (1% level), and *** signifies $p < 0.001$ (0.1% level).

We observed that for R=2,3, both rRAKI and crRAKI were able to reconstruct images that matched GRAPPA reconstruction quality. At R=4,5, we start to see undersampling artifacts on the GRAPPA and rRAKI reconstructions, which is indicated by the residual maps as shown in Figure 4.10. In contrast, the crRAKI method has fewer artifacts, suggesting that it can better preserve the global features of the images. At R=6, the GRAPPA method significantly degraded the reconstruction, resulting in images of poor quality and low fidelity. The rRAKI method produces noisy reconstructions, which make it

difficult to see the details of the images. Overall, the crRAKI method can produce reconstructions with acceptable visual quality up to R=5. Even in noisier reconstructions at R=6, the overall features of the image are preserved, making the images produced by this method visually acceptable even at this high value of R.

4.3.2.3 2D multi-echo GRE in separated mode

Figure 4.11 compares the performance of the three different methods for phase-mapping reconstructions with $N_c = 16$. It shows the reconstructed magnitude and phase maps for two different values of R (2 and 4) and for the first and last echoes, with a total of 24 ACS. The results show that the rRAKI method has some limited undersampling artifacts at R=2; however, these artifacts are more pronounced for both GRAPPA and rRAKI at R=4. In contrast, the crRAKI method has almost no undersampling artifacts in either the magnitude or phase images at R=4. The filtered-field maps also support this finding, showing artifact-free images for crRAKI at R=4, whereas GRAPPA and rRAKI still have significant undersampling artifacts. This indicates that the crRAKI method is capable of performing more accurate phase-mapping reconstructions at higher values of R, such as R=4, in these 2D acquisition settings. It is important to note that there exists a difference in contrast information between the ACS and the undersampled image scan. A noticeable degradation was observed in the rRAKI reconstruction of the last echo in Figure 4.11, indicating a limited generalization capability that may result from learning the specific ACS phase not adapted to different echo times. Similar artifacts were observed for $N_c = 52$ for R=4 (Figure 4.5).

4.3.2.4 Effect of the number of ACS lines

To study the impact of ACS size, the number of ACS lines was varied from 24 to 64 at R = 4,5,6, and GRAPPA, rRAKI, and crRAKI performances were compared using 2D multi-echo GRE data with $N_c = 52$. The quantitative evaluation is presented in Figure 4.12 with the corresponding p-values. The results showed that the performance of all methods improved as the number of ACS lines increased. Notably, crRAKI exhibited good performance even with a lower number of ACS lines, indicating its efficiency in improving the image quality with a reduced number of ACS lines, which is consistent with the reduced number of parameters for this model. These results also suggest that increasing the number of ACS lines can enhance the performance of GRAPPA.

We also performed a similar qualitative evaluation as above based on a number of artifacts from undersampling and sharpness of the images. The filtered phase was also studied during the evaluation. A qualitative comparison (Figures 4.13, 4.14 and 4.15) is provided in the supplementary material. At R=4, there was a significant improvement in the visual quality of GRAPPA when the number of ACS lines increased. Notably, rRAKI also showed improvements; however, the reconstructed images still contained undersampling artifacts. It can be noted that at R=5, crRAKI reconstruction achieves acceptable visual quality with the remaining undersampling artifacts, whereas the reconstruction from other methods suffers from degradation.

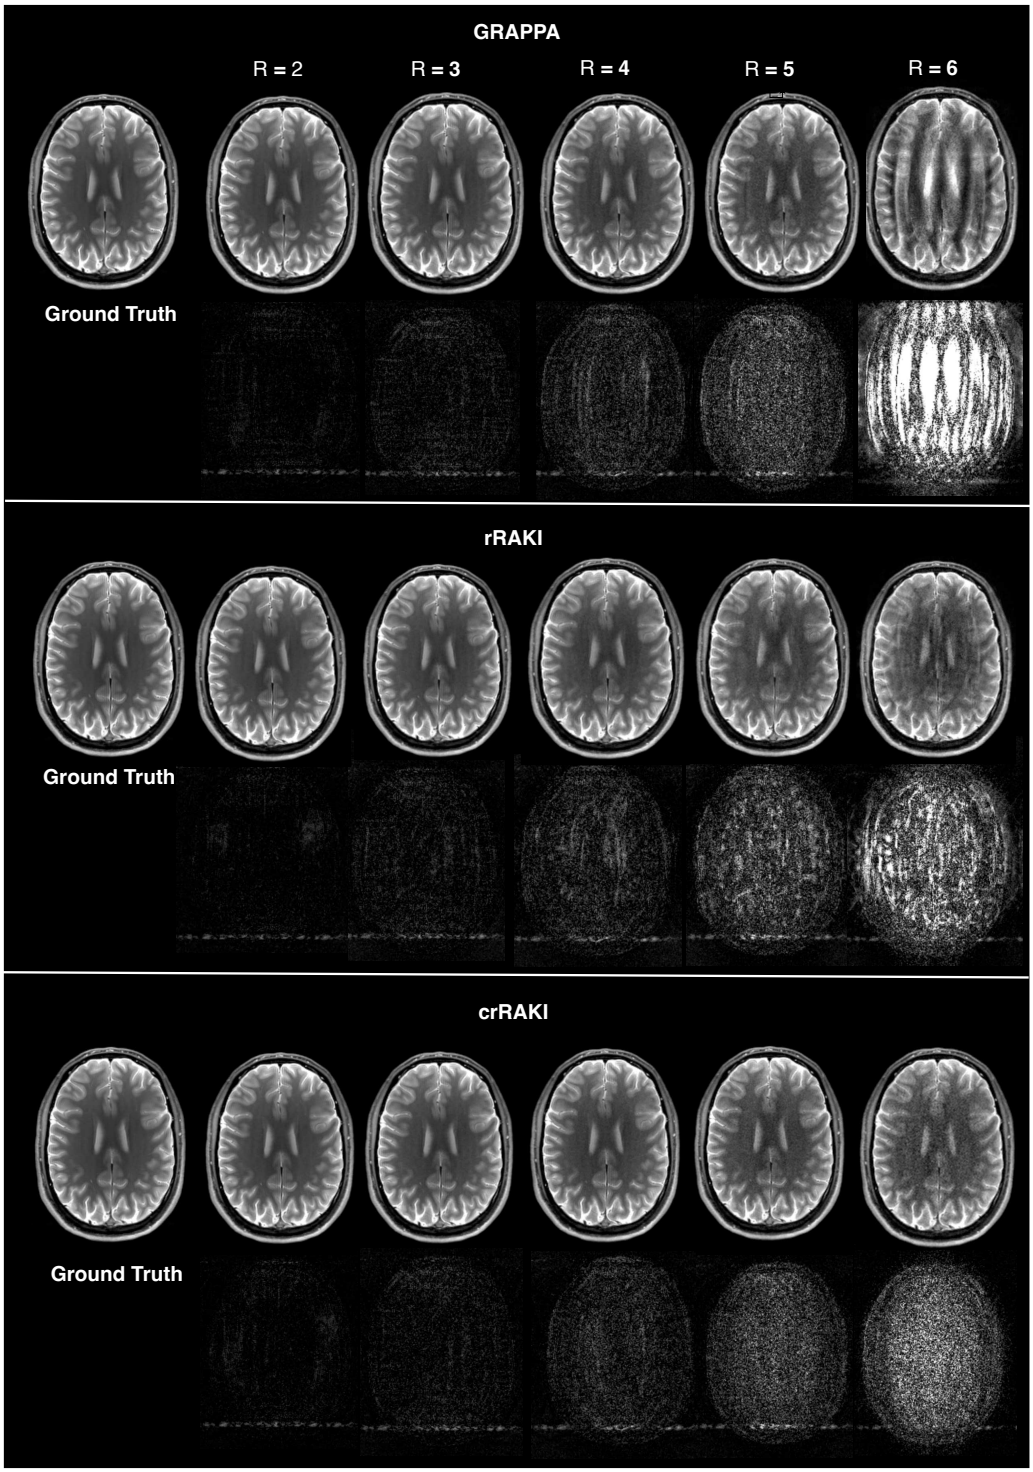

Figure 4.10: A qualitative comparison of GRAPPA, rRAKI, and crRAKI was presented in terms of their ability to reconstruct images at various acceleration rates (R = 2, 3, 4, 5, 6). Residual maps and fully sampled images were used as the ground truths for this comparison. Residual maps are represented in grayscale, with the brightness of each pixel indicating the magnitude of the absolute difference at the same window level. By examining the residual map, we can assess the quality of the images by examining the presence of artifacts and comparing the residual levels. At R=4,5 structured artifacts started to appear on GRAPPA and rRAKI reconstructions, whereas crRAKI reduced the number of artifacts. At R=6, GRAPPA methods have degraded reconstruction, rRAKI has a noisier reconstruction, and crRAKI shows better preservation of the global feature, which makes it visually acceptable.

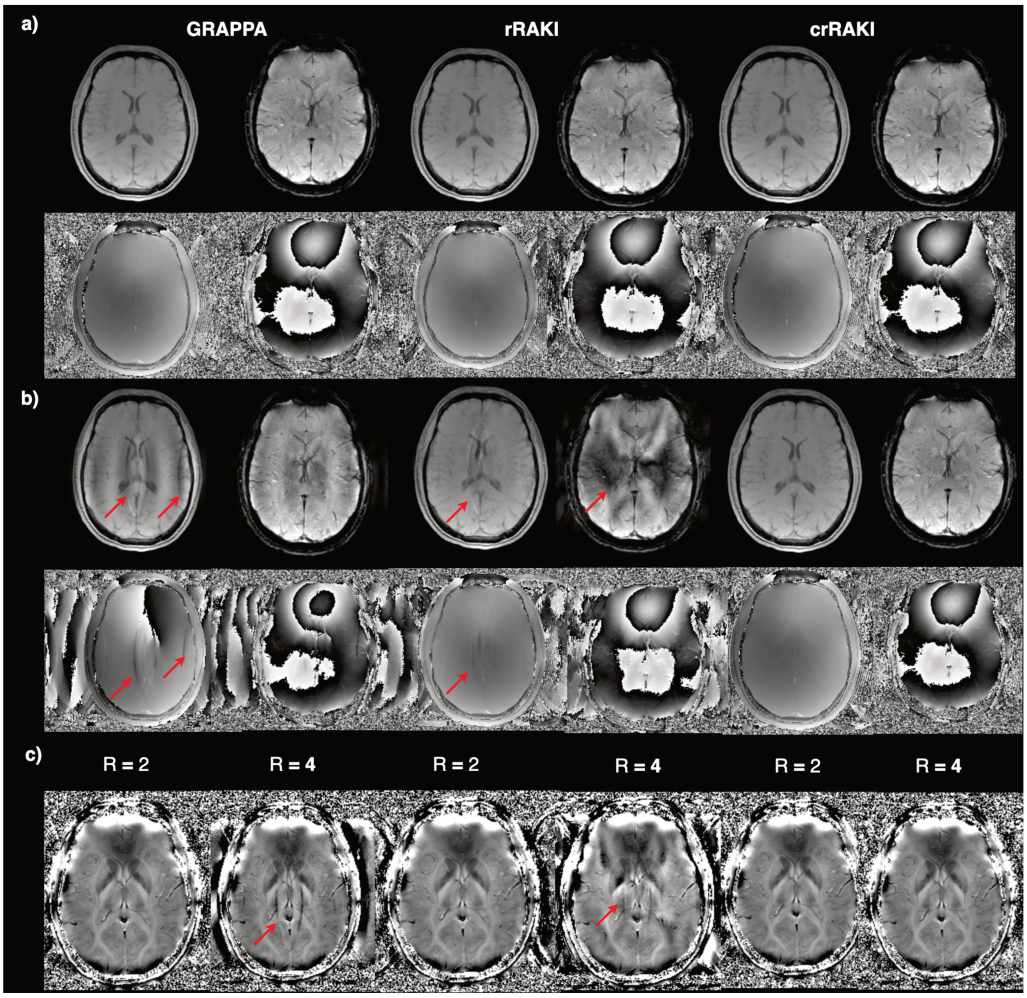

Figure 4.11: 2D GRE results with adaptive coil-combination[313] for $N_c = 16$ (a): Magnitude and phase at R=2 for fist and last echoes, (b): Magnitude and phase at R=4 for first and last echo, (c): Filtered-field maps for GRAPPA, rRAKI, crRAKI at R=2 and 4. The red arrows highlight the presence of undersampling artifacts. Notably, a contrast-loss artifact is observed in the rRAKI reconstructions for the last echo.

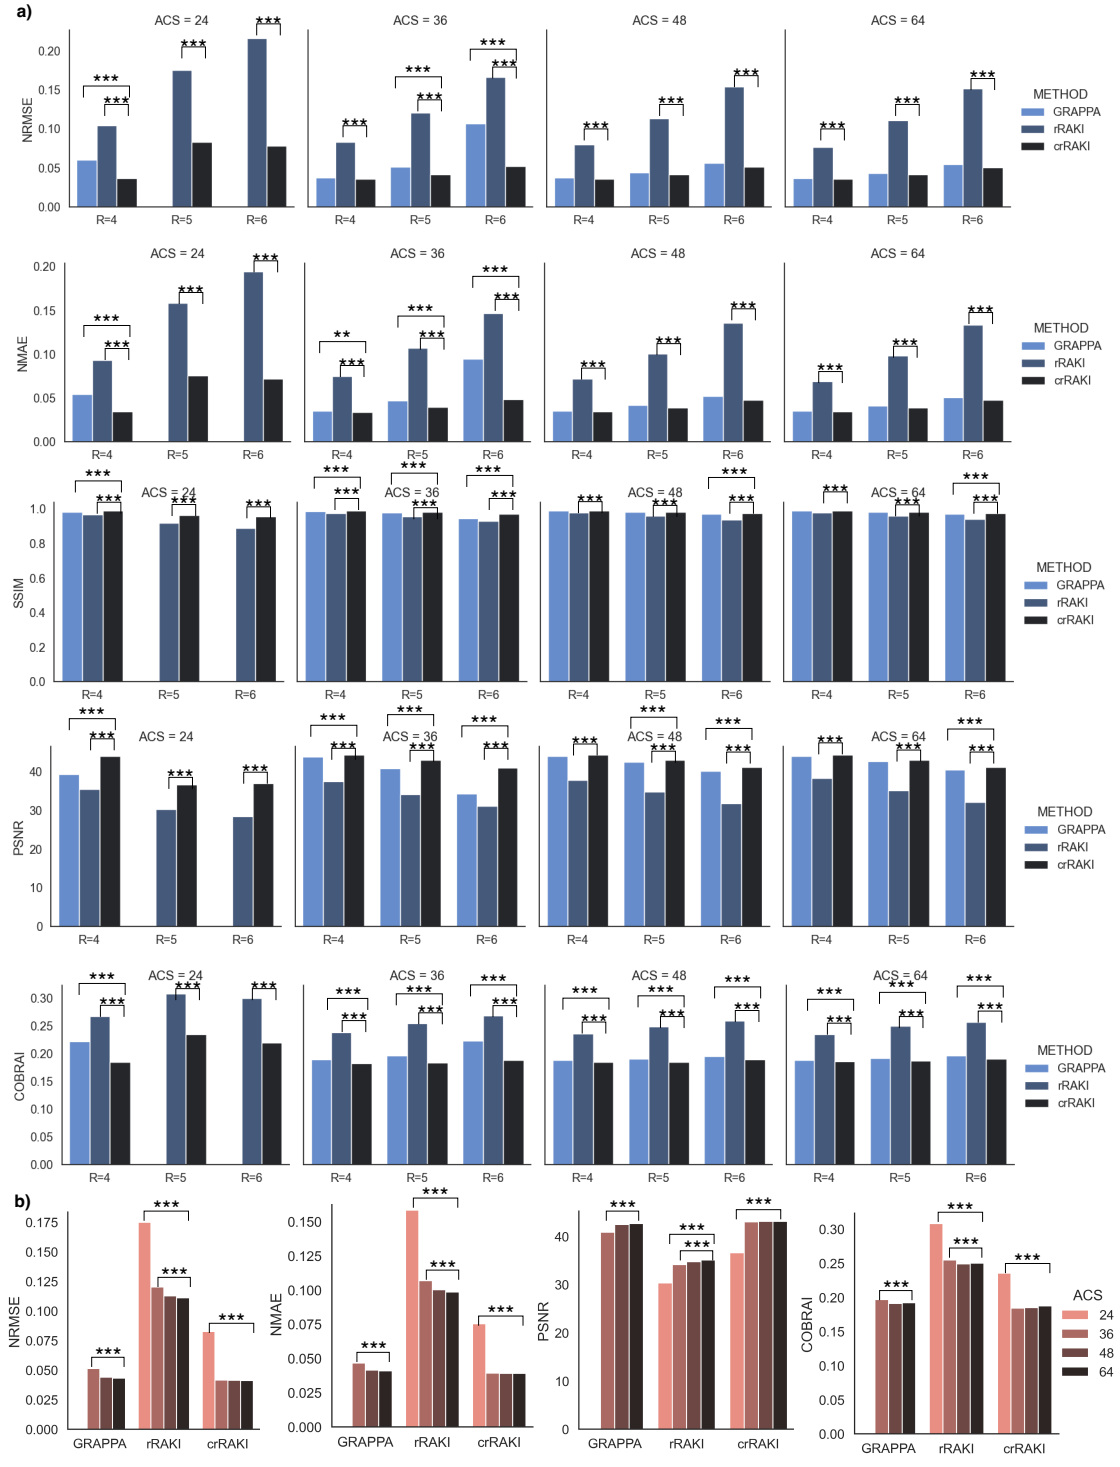

Figure 4.12: Quantitative comparison of various methods to study the effect of the number of ACS lines on the performance of the reconstruction methods (GRAPPA, rRAKI, and crRAKI) based on NRMSE, NMAE, SSIM, PSNR, and COBRAI. a) Performance for all metrics (blue) at R=4,5,6, One-tailed Student's t-test was conducted with the hypothesis that crRAKI would demonstrate superior performance across all values of R and ACS. b) Effect of varying ACS size on methods with R=5. One-tailed Student's t-test was performed to test the hypothesis that all methods would show better performance when the ACS size is 64. The relationship between the number of ACS lines and the performance of the methods was found to be correlated (*, **, and *** indicate p<0.05, 0.01 and 0.001, respectively).

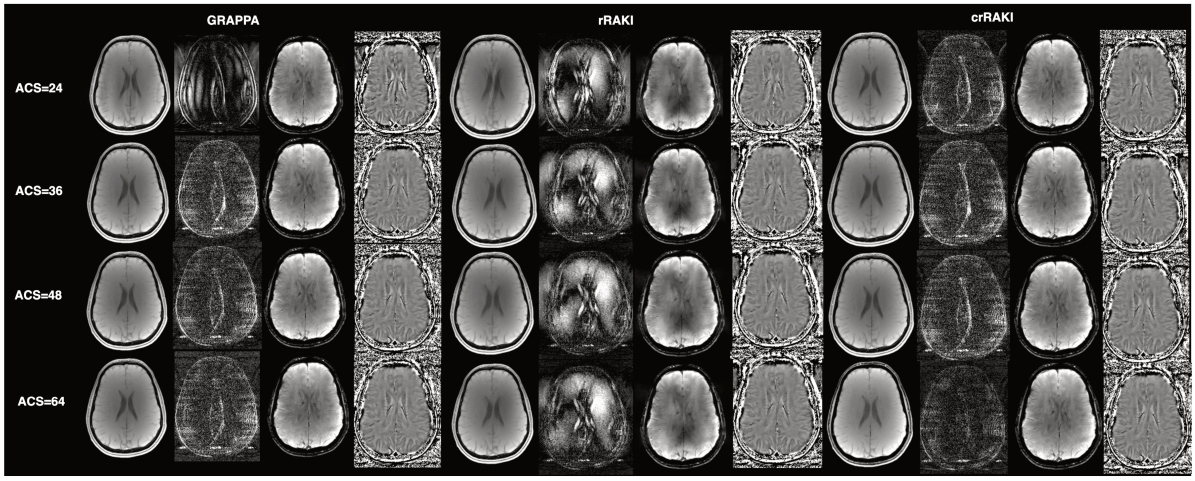

Figure 4.13: 2D GRE qualitative results for $N_c = 52$ at R=4 (a): Magnitude and residual maps for 1st echo, high-pass filtered field maps for GRAPPA, rRAKI, crRAKI. The window levels of the residual maps are the same.

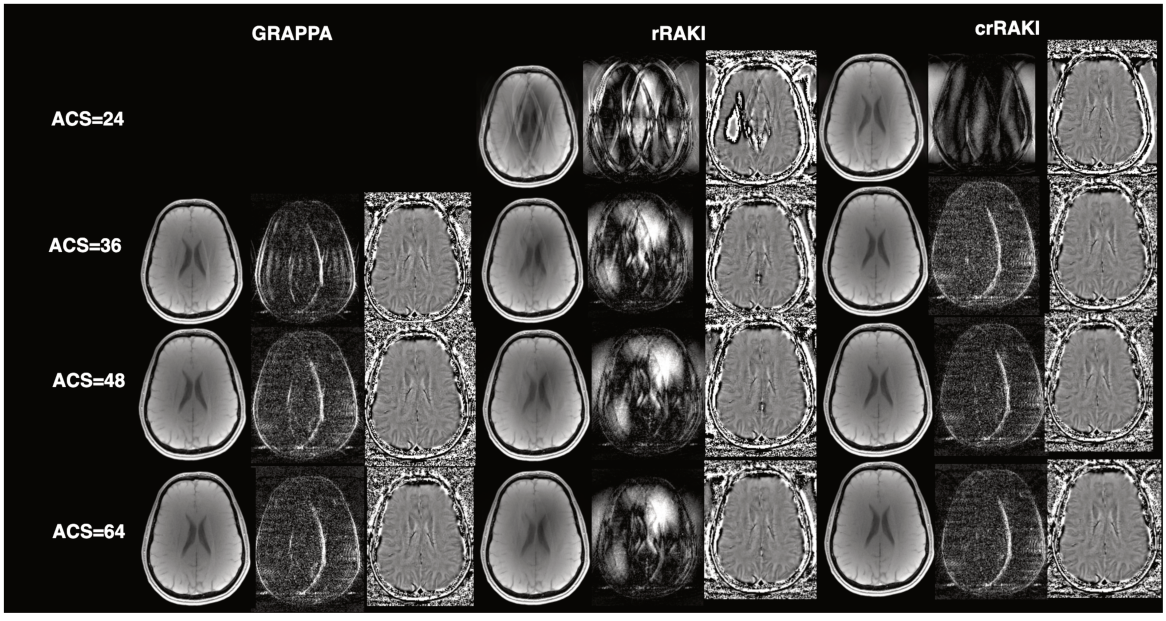

Figure 4.14: 2D GRE qualitative results for $N_c = 52$ at R=5 (a): Magnitude and error maps for 1st echo, high-pass filtered field maps for GRAPPA, rRAKI, crRAKI. Window levels of the residual maps are the same.

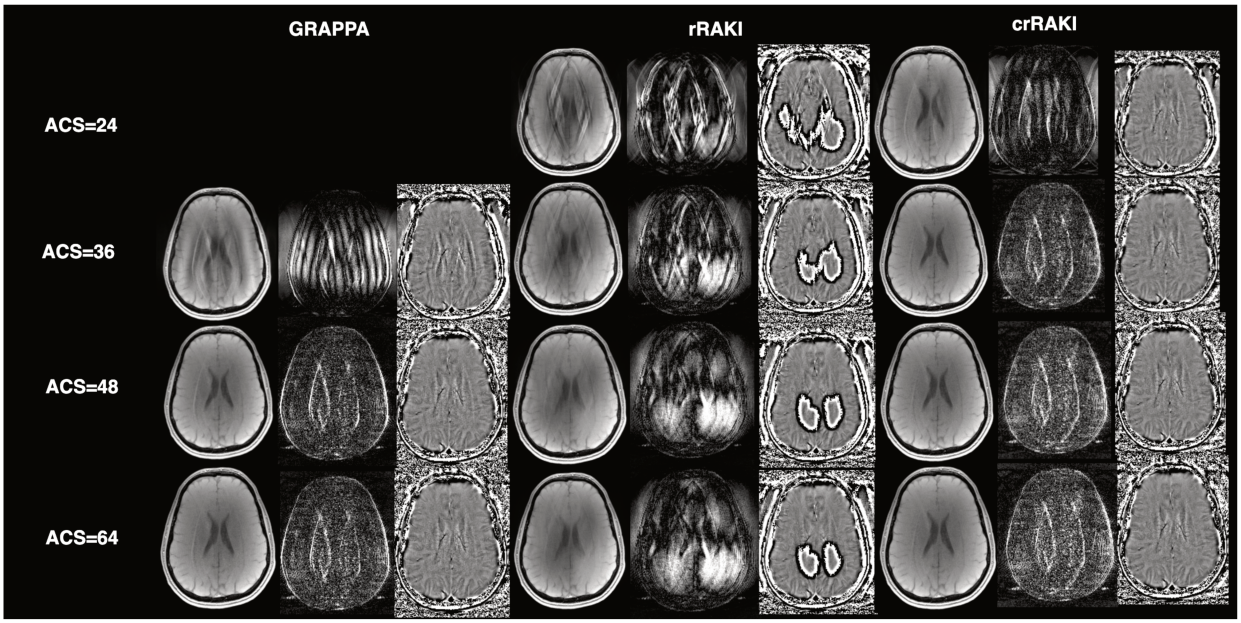

Figure 4.15: 2D GRE qualitative results for $N_c = 52$ at R=6 (a): Magnitude and error maps at for 1st echo, high-pass filtered field maps for GRAPPA, rRAKI, crRAKI. Window levels of the residual maps are the same.

4.3.3 3D RESULTS in separated mode

4.3.3.1 3D MPRAGE

The qualitative comparison of the images obtained from the MPRAGE data using GRAPPA and crRAKI 3D models trained on ACS of size 40x40 showed similar results with reduced artifacts and enhanced quality, supported by smaller residuals. A representative slice from a 3D volume was included at various retrospective acceleration rates: 2x1, 3x1, 3x2, 3x3, 4x2, and 4x3 as shown in Figure 4.16 and 4.17.

At lower acceleration rates (R=2x1, 3x1, 3x2 and 4x2), both crRAKI and GRAPPA produced reconstructed images of similar visual quality. However, as the acceleration rate increased (R=3x3 and 4x3), the GRAPPA reconstructions had higher residual values as compared to the crRAKI reconstructions. High acceleration rates are also associated with noisy images, as can be observed in the central brain regions. Images reconstructed from the prospective undersampled MPRAGE data showed consistent visual quality (not shown).

4.3.3.2 3D Multi-Echo GRE

We first examined the effect of ACS on 3D reconstructions in a manner similar to 2D (Figure 4.18), and the results consistently emphasized the improved performances of both GRAPPA and crRAKI with ACS size, as well as the higher quality of crRAKI images when training with small ACS sizes. For feasibility, Figure 4.19 shows the echo combined magnitude image, relaxation rate R_2^* map, and QSM for the GRAPPA and crRAKI reconstructions at R=4x3 with the largest ACS size available (64x64). Despite a visible reduction in the SNR, both methods can perform acceptable reconstructions for R= 4×3 and display similar paramagnetic deep grey matter structures. This suggests the use of a high

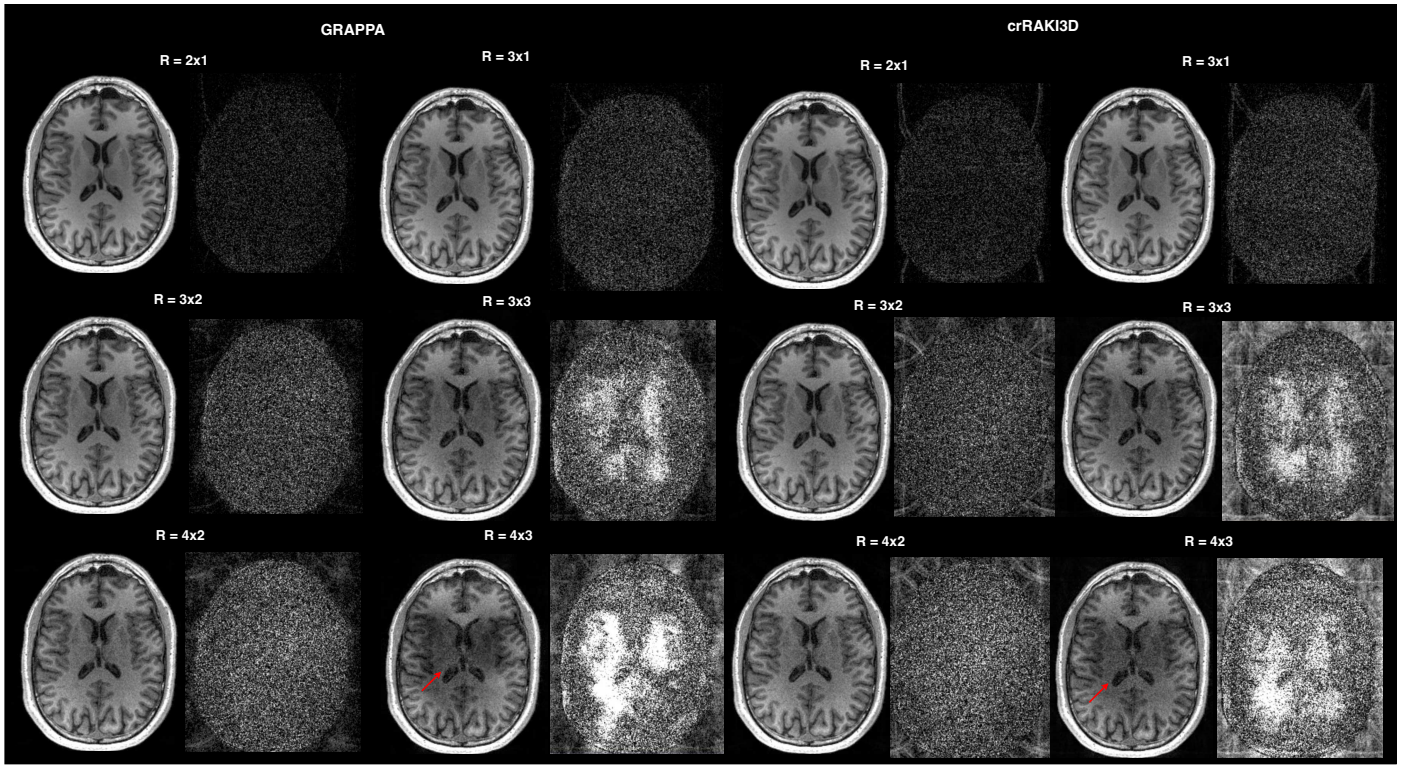

Figure 4.16: Qualitative comparison of the performance of GRAPPA, crRAKI for MPRAGE, at various retrospective acceleration rates (indicated above the images). Residual maps are presented at the same window level for comparison. The red arrows indicate contrast-loss artifacts.

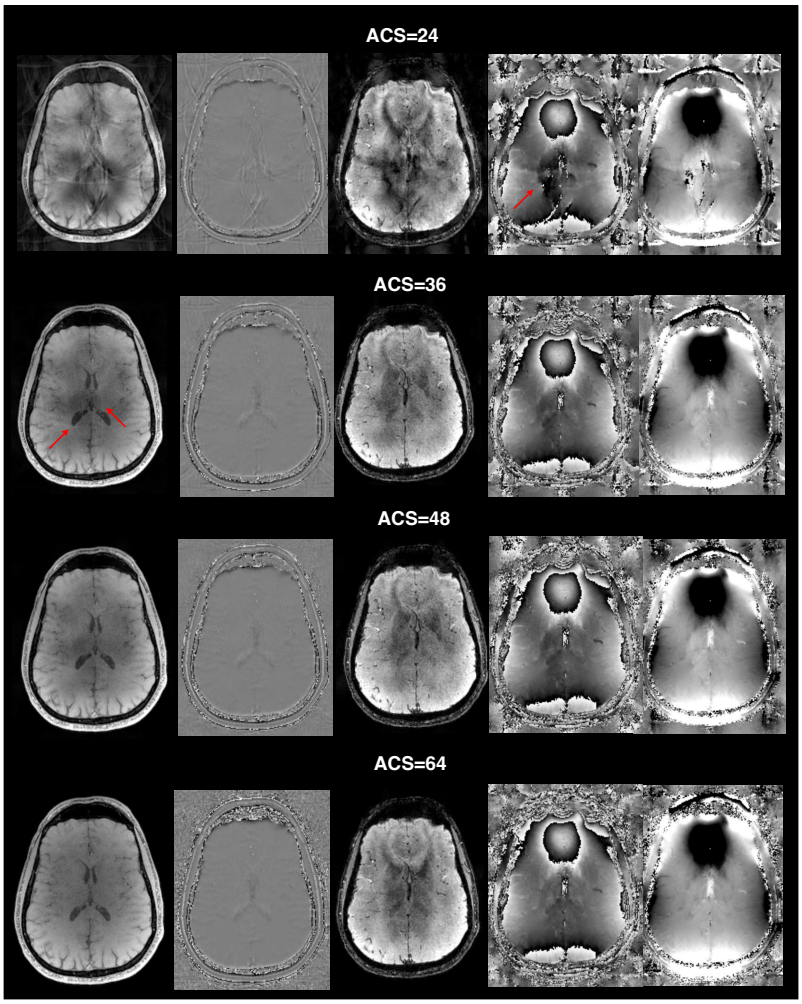

Figure 4.17: 3D GRE qualitative results demonstrating the impact of ACS on the magnitude, phase for first and last echoes, and B_0 maps at R=4x3 for GRAPPA with different ACS sizes. Coil-combination was performed using the low-pass filtered first echo image. The undersampling artifacts are denoted by red arrows.

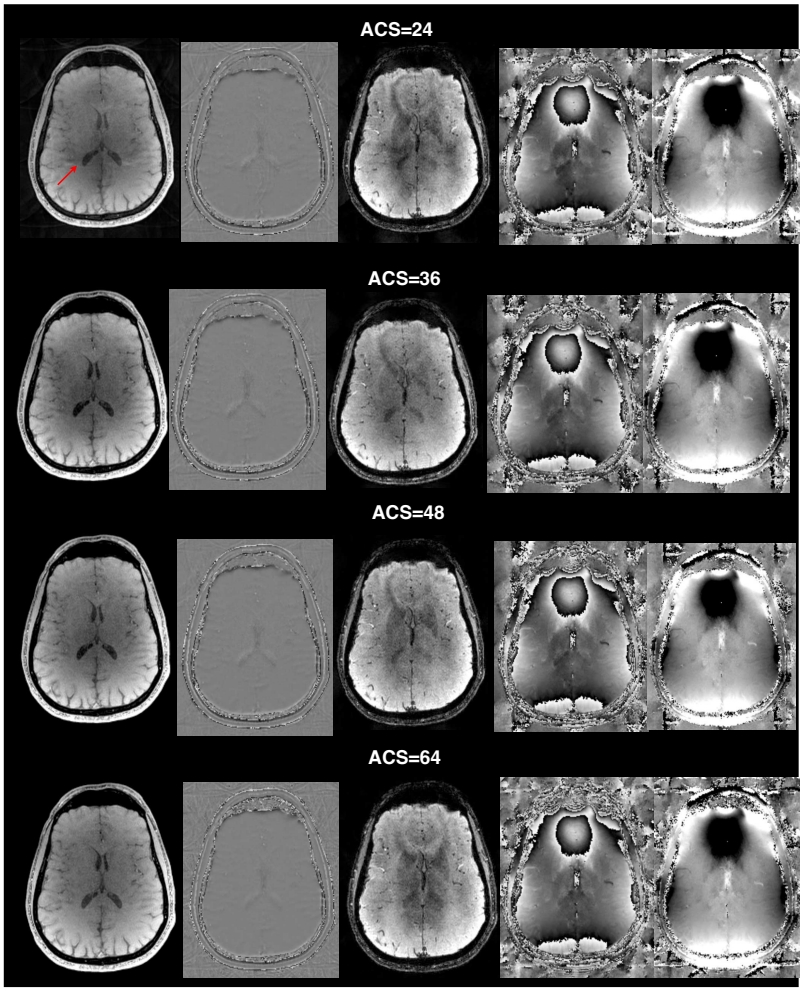

Figure 4.18: 3D GRE qualitative results demonstrating the impact of ACS on the magnitude, phase for first and last echoes, and B_0 maps at R=4x3 crRAKI with different ACS sizes. coil combination was performed using a low-pass filtered first-echo image. The undersampling artifacts are indicated by red arrows. The results indicate the requirement of a larger ACS size for GRAPPA with higher R, whereas crRAKI remains robust even with a limited number of ACS at high R.

acceleration rate with both GRAPPA with a sufficiently large ACS size and crRAKI for phase-based reconstruction derived from multi-echo GRE, such as QSM.

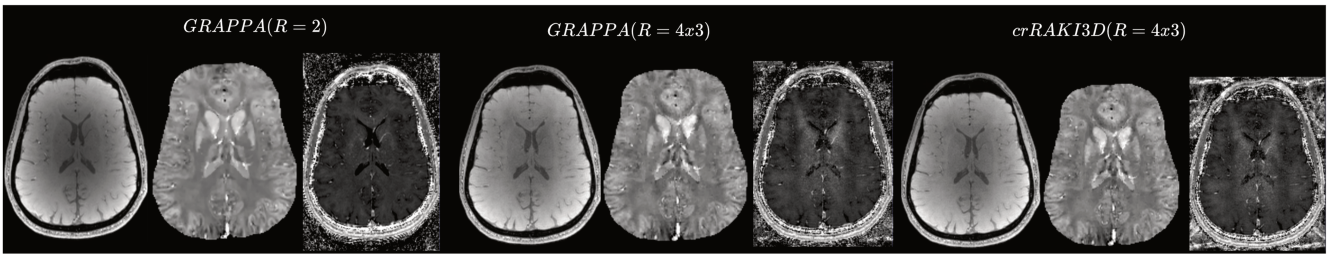

Figure 4.19: 3D multi-echo GRE magnitude image, relaxation rate R2* map, and QSM for GRAPPA and crRAKI at 4x3 acceleration compared to the default system reconstruction.

4.4 Discussion

4.4.1 Main findings

The main contribution of the study is the demonstration of the superiority of the crRAKI model as compared to GRAPPA and rRAKI with Cartesian sampling in PI for brain imaging with optimized hyperparameters and TVP, especially in separeted-mode ACS. This superiority is supported by enhanced quantitative metrics and the visual quality of magnitude and phase images in both 2D and 3D, as well as for various sequence types up to high acceleration rates. As such, COBRAI has been introduced to support visual quality. The grid-search strategy with TVP and K-fold cross-validation was implemented for hyperparameter selection, demonstrating that i) nonlinearity is detrimental, ii) MSE loss alone is sufficient for training, and iii) optimized architecture (kernel sizes, depths, linearity, model type) can be chosen solely based on ACS.

The presented results also emphasize the influence of the ACS size in all models, and, as one may expect, crRAKI, which has fewer parameters, needs less training data. Additionally, GRAPPA trained using the TVP strategy performs well up to acceleration rates higher than standard clinical practice, providing that a large number of ACS are acquired. The implication is that routine sequences such as axial T2W spin-echo and MPRAGE can be readily accelerated at a higher rate (at least up to 5 and 4x2, respectively), with an expected decrease in SNR, but no visible undersampling artifacts, when combined with the proposed DL reconstruction scheme. Finally, it was shown that crRAKI can also be used to accelerate phase-based acquisitions in 2D and 3D with enhanced phase map reconstruction compared to GRAPPA and rRAKI.

4.4.2 Non-linearity

Neural networks incorporating non-linearities, such as rRAKI, in fact have demonstrated their efficacy in learning intricate representations, making them well-suited for modelling data that require non-linear transformations when the goal is not on generalizability. However, it is important to acknowledge that these networks, typically characterized by a large number of parameters, are prone to overfitting when data availability is limited. One of the major findings of our study is that the

introduction of non-linearities can actually be detrimental for model generalization, particularly in separated ACS modes. Furthermore, when a significant number of parameters are present, there is a potential risk of network learning features from the ACS, such as contrast, which may not generalize effectively to multi-echo sequences. This is contrary to the general assumption that nonlinearities always improve reconstruction performance.

The performance of rRAKI is linked to the ability of the GRAPPA short-connection to provide a good estimate. If this linear term has a poor performance, it logically follows that rRAKI is affected, as observed in the ablation study.By increasing the degree of linearity with LeakyReLU[79] and the short-connexion loss, the rRAKI model becomes more efficient in modelling k-space, but effectively produce artifacts, indicating that trained models do not generalize well to new data with non-linearities. Similar observations have been reported in a crRAKI ablation study. These artifacts disappear with a linear model, leading to performances that are similar to those obtained using the single model loss term. From these findings, it can be concluded that activation functions such as ReLU should not be used for rRAKI in separated ACS acquisition mode. Even though the nonlinear rRAKI/crRAKI variant performance has better standard quantitative metrics than its linear variant in the integrated mode, COBRAI is higher, meaning that structural artifacts still remain. The use of a short-connection loss was not required in the training process, leading to a consistent use of the MSE loss in k-space for GRAPPA, rRAKI, and crRAKI linear models, avoiding the use of additional hyperparameters and reducing the computational requirements.

4.4.3 Real vs complex

Recently, there has been a growing interest in comparing the efficiency of complex-valued neural networks to real-valued networks with extensive analysis[54]. While iRAKI proposed a complex implementation of RAKI incorporating additional self-consistency, the foundational architecture may seem familiar here. However, our study on (c)rRAKI diverges by focusing solely on calibration consistency, that is, ACS, along with the use of a residual neural network. We demonstrate that, given an optimal choice of architecture, it is feasible to achieve satisfactory reconstructions in accelerated scenarios with a smaller ACS size. The use of data self-consistency in iRAKI may also be improved by using crRAKI estimates as an initialization.

Complex-valued linear models such as GRAPPA and crRAKI offer a simpler architecture devoid of non-linearities. These linear neural networks exhibit implicit regularization properties, constraining learned representations and mitigating the risk of overfitting. This intrinsic characteristic of implicit regularization endows linear models with greater robustness to contrast changes. However, rRAKI linearity is restricted to the real field, while GRAPPA and crRAKI linearity extends to the complex field. This may explain why rRAKI does not generalize well in the phase-based applications tested here, particularly for long echo times, an issue that has not been observed with GRAPPA and crRAKI. To reduce this effect, the concept of ACS augmentation has recently been proposed[66]. Indeed, to enforce a real-value neural network to be linear in a complex field, it may be sufficient to train the network using phase-shifted ACS, for example, multiplied by i, thus enforcing this property on average. However, it is a more straightforward strategy to enforce this algebraic property by design using complex convolutions with crRAKI, as pointed out earlier in the context of fat-water separation and phase-contrast velocity mapping[54] and that we confirm here for field-mapping in the brain.

4.4.4 TVP and K-fold cross-validation

In this work, we applied a train-validation partition scheme for training the tested models. While this concept is standard in deep learning[224], it has also been applied in similar work recently to regularize the training stage by early stopping[341]. Here, it is performed simply by partitioning the ACS. Indeed, early stopping involves monitoring the performance of the model on the validation set during training and stopping the training process once the model performance on the validation set starts to degrade. This helps to prevent overfitting and encourages the model to converge towards solutions with minimal norms. This strategy was applied consistently to regulate GRAPPA, rRAKI, and crRAKI for a fair comparison. As we train and validate using a single ACS acquisition, fully-sampled data is not needed to test the capacity to generalize. Our contribution is leveraging the use of K-fold cross-validation to assist in hyperparameter selection. Indeed, we demonstrate that the KSPACE-MSE loss alone (using the ACS validation sets) can predict image quality metrics that will be obtained after application to new data. This approach is a solution to the issue of heuristically determining the hyperparameters that were noted in an earlier study[66].

4.4.5 Kernel sizes and depths

Kernel size and depth were optimized in a specific 2D case and for a limited number of sizes and depth combinations. Interestingly, we confirmed the pertinence of the sizes used in earlier works for GRAPPA and rRAKI, as we obtained the architecture that they implemented[5, 98] in the list of valid ones. Additionally, we showed that several hyperparameter sets could also have been selected. For comparison with earlier studies, we used the same architectures as those previously studied for GRAPPA and rRAKI.

The number of parameters in neural networks plays a pivotal role in determining performance, influenced by factors such as the number of coils and the size of the ACS data. Striking an appropriate balance between underfitting and overfitting is critical and contingent upon the number of parameters and the available ACS data size. Additionally, while GRAPPA has demonstrated competitive performance in certain instances, it is hampered by limited flexibility in network design, rendering it less adaptable when faced with changes in the ACS size and number of receiver coils. Employing a K-fold cross-validation strategy aids in making informed decisions regarding the network architecture size. The chosen crRAKI model was more compact than the GRAPPA model. The residual section has $n_{x,g} = 1$, $n_{y,g} = 1$ (and $n_{z,g} = 1$ in 3D), and thus has a single k-space point for all coils. It is then equivalent to the coil-by-coil SMASH[211] implementation of the short-connection in parallel with a cRAKI section with three linear filters in cascade. In our experiments, we did not observe instability in learning crRAKI, such as divergence, which may be expected when learning cascade filters. To further reduce this risk, the first two layers can be regularized by imposing a unit norm to improve stability without any expected loss of performance.

4.4.6 ACS and Timing

In addition to evaluating the performance of different reconstruction methods, we also tested the effect of ACS lines on the quality of the reconstructed images at various acceleration rates for separated mode. As expected, all methods examined in our study benefited from the increase in ACS. However, the performance was saturated when the ACS lines exceeded 36. There is a significant impact on the

performance of GRAPPA in 3D, and increasing the number of ACS lines improves image quality and reduces aliasing artifacts. Additionally, we found that the crRAKI method demonstrated robustness and produced reliable results, even when utilizing a lower number of ACS lines.

On average, it takes between 15 and 30 seconds to train each task-specific model using 2D data and between 150 and 180 seconds using 3D data in our current implementation in which there is room for improvement in computation efficiency using, e.g., parallelization. These timings in 2D are consistent with the ones reported for RAKI, and significantly lower than iRAKI. Self-consistency is more demanding in computing resources than strategies based on using only the ACS, which is one advantage of the proposed approach compared to [66]. These limited training times render the technique attractive in the time-efficient separated ACS mode and feasible in a routine workflow without dead times as, being done on ACS only, training can be performed right after the ACS are acquired (at the beginning of the undersampled k-space) while the acquisition of the under-sampled lines is going on. The evaluation can then be rapidly performed using the trained network.

The practical consequences are that we recommend acquiring separated ACS with at least 36 lines in 2D, and 36x36 lines in 3D to reach the acceleration rates reported here, with a minor impact on total scan time. The other consequence is that to train larger models, possibly with non-linearity, or to reach an even larger acceleration rate, an effort should be made in the acquisition of ACS training data, such as with different scan conditions and sizes to improve generalizability.

4.4.7 COBRAI

In the previous chapter, we discussed COBRAI as the metrics to assess the image quality using a ground truth image as a reference. Existing metrics such as NRMSE, NMAE, PSNR, and SSIM are insufficient for detecting artifacts that are clearly visible in RM. In k-space reconstruction, supervised upsampling seeks to minimize the MSE between the simulated data and target pairs obtained from low-resolution ACS data. This optimization objective is advantageous because it appears to maximize widely accepted evaluation metrics such as PSNR and SSIM. However, it is essential to recognize that NRMSE, NMAE, PSNR, and SSIM have inherent limitations in capturing perceptually significant distinctions, particularly those related to contrast, because they rely on pixel-wise differences or structural similarity.

Although COBRAI metrics provide insights into performance, it should be noted that extreme corruption may not be fully captured. Therefore, a comprehensive analysis combining multiple metrics, such as NRMSE, NMAE, SSIM, and PSNR, is recommended for a more thorough evaluation. Due to its simplicity and capacity to summarize one type of reconstruction artifact, we recommend including COBRAI in future DL-MRI studies. COBRAI could also be included in composite cost functions when training DL-models. With large reference image databases, deep learning models may benefit from COBRAI to train models with reduced artifacts, opening new perspectives for DL-MRI.

4.4.8 Improvements and practical implications

The current choice of model parameters using grid-search has offered valuable insights, while we did not thoroughly examine the influence of the model's architecture. Indeed, it was done for 2D to determine the hyperparameters, and performances in terms of image metrics were verified for a larger R, in both 2D and 3D, for various sequence types and coil setups. However, there might exist different

hyperparameters may be more specific in these situations. This presents an opportunity to further optimize the performance of the model through a comprehensive exploration of its architecture by utilizing the presented hyperparameter optimization techniques.

To the best of our knowledge, this is the first feasibility study of scan-specific DL strategies with a focus on generalizability both for magnitude and phase-based applications in 2D and 3D. Using crRAKI, we have shown that 2D T2W spin-echo scans can be accelerated up to 5 without noticeable artifacts, whereas R=2 is often used in practice, reducing the total scan time to 1 min. For MP2RAGE and 3D Multi-echo GRE, a 4 x 2 acceleration rate can be achieved with reduced SNR but no noticeable artifact, bringing scan times down to 1.25 and 2 min, respectively. To address the issue of SNR loss expected when scan time is reduced, denoising approaches, including the use of image-domain DL approaches[19, 83], can be employed to enhance the image quality. Here, we have shown application on a regular Cartesian GRAPPA-like under-sampling pattern. Based on the formalization of the reconstruction task of a *full k-space* from an *undersampled k-space*, extension to other structured undersampling patterns, such as CAIPIRINHA[37] is straightforward.

4.5 Conclusions

We presented a complex-valued neural network that differs from prior studies by focusing exclusively on calibration consistency (ACS) and employing a residual neural network that can adapt to scanner specifications such as receiver coils. It is demonstrated that the crRAKI model is superior to GRAPPA and rRAKI, especially in the separated ACS mode, which is supported by enhanced quantitative metrics and visual quality of magnitude and phase images in both 2D and 3D for various sequence types, including those focusing on phase. Higher acceleration rates than those commonly used can be achieved using this method. Our study introduces a K-fold cross-validation strategy combined with a TVP scheme to fine-tune the network hyperparameters, such as kernel size, layer depth, and slope for leakyReLU, based solely on ACS data, without requiring fully sampled data. This level of customization in optimization is unparalleled in previous research, including iRAKI[66]. In the iRAKI paper, the authors specified the total number of hidden layers, assigned channel number for each layer, learning rate, and slope for leakyReLU through a heuristic search for hyperparameters since there were no existing methods for optimally tuning the network parameters. One of the major findings of our study is that the introduction of nonlinearities can be detrimental to model generalization, particularly in separated ACS modes. This is contrary to the general assumption that non-linearities always improve reconstruction performance, as we have shown situations here, when a low amount of training data is available as well as when training and evaluation data do not share the same contrast, in which linear models are performing better.

5 SANGRIA: one-Shot leArNinG super-ResolutIon with Adversarial training for accelerated Magnetic Resonance Imaging

Summary

5.1 Introduction . . . 109
5.2 Background and Related works . . . 109
5.2.1 One-shot learning . . . 109
5.2.2 Single-image patch-based generative models . . . 110
5.2.3 Data-driven regularization . . . 111
5.3 SANGRIA . . . 112
5.3.1 Model . . . 113
5.3.2 Model description . . . 113
5.4 Experiments . . . 114
5.4.1 Data and experimental setting . . . 114
5.4.2 Quantitative evaluation of the reconstruction quality . . . 115
5.4.2.1 Synthetic data . . . 115
5.4.2.2 In-vivo data . . . 115
5.4.3 Qualitative evaluation of reconstruction quality . . . 118
5.5 Discussion and Conclusion . . . 118

5.1 Introduction

This chapter delves into the extension of existing scan-specific reconstruction methods with a focus on tackling the super-resolution problem in k-space by incorporating a CNN-based approach and utilizing an adversarial strategy with a patch discriminator in the image space to achieve higher undersampling rates. Notably, both components of the proposed framework were trained within a one-shot learning setting, thereby enabling a rapid adaptation to a wide range of imaging scenarios.

We begin our discussion with a one-shot learning paradigm, state-of-the-art patch-based methods for single images and data-driven regularization. It was demonstrated using both simulated and in vivo brain experiments that this combined approach provides enhanced image quality for undersampling rates larger than those used in a clinical routine on multislice 2D T2W imaging sequences, making the approach readily applicable as an alternative reconstruction strategy for 2D parallel imaging, which could be further extended to accelerate 3D imaging and enhance reconstructions.

5.2 Background and Related works

5.2.1 One-shot learning

One-shot learning is a paradigm within machine learning that aims to enable models to learn from only a single example such as images, in contrast to traditional deep-learning approaches, which often require vast amounts of labeled data for training. This approach is of significant importance, particularly in scenarios where obtaining large datasets is not practical or costly. Siamese Neural Networks[156] represent an approach in one-shot learning, particularly in image-recognition tasks. These networks learn to compute similarity metrics by comparing pairs of datapoints. By utilizing the shared weights across two identical neural networks, Siamese networks can effectively learn a similarity function that discriminates between classes. Metric learning techniques are another class of methods that are used in one-shot learning[22]. These approaches aim to learn the distance metrics in the input space, enabling the model to differentiate between classes effectively. By optimizing the embedding space, metric learning methods facilitate more discriminative representations, thereby enhancing the model's ability to generalize from limited examples. Memory-augmented neural networks[147] leverage an external memory component to store past information and enhance the performance on new, similar examples. These networks can effectively retain information from limited training data and leverage it to make accurate predictions for unseen instances. By dynamically updating the memory content during training, memory-augmented networks can adapt to varying tasks and input distributions, making them well-suited for one-shot learning scenarios. Meta-learning[311] represents a promising direction in which a model learns from few-shot learning tasks. By exposing the model to a diverse set of tasks during meta-training, meta-learning algorithms can learn a good initialization or adaptation strategy for new unseen examples. Popular meta-learning algorithms, such as Model-Agnostic Meta-Learning (MAML)[87], have shown promising results in enabling models to adapt to new tasks with limited data quickly. In the next section, we discuss how patch-based networks can be effectively integrated into one-shot learning by using local descriptors to direct the learning process, thereby improving the model's ability to generalize from limited examples.

5.2.2 Single-image patch-based generative models

Modelling the internal structure of patches within a single natural image has been useful prior in various computer vision tasks. One-shot models currently available leverage this approach to generate diverse images with characteristics resembling those of the training image. The process of one-shot image synthesis involves creating multiple images from a single training example, with the aim of generating images that share similar characteristics with the training example, but with

different shapes. The synthesized images should be visually consistent and plausible.

Low-data regimes present challenges in training GANs due to the risk of overfitting, which can result in replication of the training data or poor synthesis quality. For example, few-shot image synthesis models such as FastGAN are prone to memorization issues when only one image or video is used as input. To address this issue, SinGAN[275] models utilize a pyramid of fully convolutional lightweight GANs to capture the distribution of patches on multiple scales. These models employ a cascade of multiscale patch GANs trained in multiple stages to learn the internal-patch-based distribution of an image. ExSinGAN[364] is learned not only from the internal patches of the given image, as in previous studies, but also from the external prior obtained by the GAN inversion technique. Granot et al.[97] proposed a Generative Patch Nearest Neighbors (GPNN) based on the SinGAN architecture but replaced its GAN module with a patch-nearest-neighbors component, resulting in a simpler generative patch-based algorithm that builds upon traditional methods. SIV-GAN[291], an unconditional generative model that can generate new scene compositions using a single training image or video clip. The model uses a two-branch discriminator architecture that separately evaluates the internal content and scene layout realism of generated images. KernelGAN[21] estimates the best SR kernel to preserve the distribution of patches across the scales of an LR image. Li et al.[173] studied a patch-based generative model for 3D natural scenes, which are often intricate and complex. Drawing inspiration from traditional patch-based image models, 3D scenes at the patch level using a single example were synthesised. However, these models struggle to learn high-level semantic properties of the scene and often result in the incoherent shuffling of image patches, which can lead to distorted objects and unrealistic layouts. Recently, Wang et al.[325] proposed patch-wise diffusion training (Patch Diffusion), which is a flexible training method that can be used with any UNet architecture, sampler, or noise schedule. Instead of training on the full image, training was performed on the image patches, which significantly reduced the computational burden per iteration. A pixel-level coordinate system was created to incorporate patch location information and encode it as additional coordinate channels.

Several studies have been conducted to exploit patch-based approaches for MRI reconstruction. Patch-based directional wavelets were proposed to reconstruct images from undersampled k-space data[245]. Yang et al.[345] present the Cross-contrasts CHannel Selection (CCHS) module, that effectively combines auxiliary- and target-contrast features on the fastMRI dataset by using multi-scale patch-based learning. Davy et al.[65] introduced training a neural network with nonlocal patch information using image redundancy for the denoising task. The main principle is to identify the most similar patches, collect the central values of these patches, and concatenate them to form additional layers.

An alternative approach is PatchGAN[131] that uses patches generated from a single image to take advantage of patch-based statistics in GAN discriminator. Unlike traditional discriminators that classify entire images as real or fake, PatchGAN classifies individual patches of the image, making it especially useful for tasks where the local image structure is crucial, such as image-to-image translation tasks. Wang et al.[318] proposed a GAN framework using a patch-discriminator for 3D single image super-resolution on T1-weighted brain MRI images. Tezcan et al.[297] proposed VAE-based priors which approximate the distributions of large image patches, termed Deep Density Prior (DDP). This model comprises a prior and a data consistency term, which encapsulates the encoding operation. PNCS[123] is a pixel-level nonlocal method-based compressed sensing undersampled MRI image reconstruction method. Block matching and row matching were performed on the reconstructed zero-filled image to obtain similar pixel groups and to establish a better sparse representation under

the patch-level non-local self-similarity prior. The separable Haar transform was applied to these pixel groups, with hard thresholding and Wiener filtering used to effectively remove any introduced artifacts.

5.2.3 Data-driven regularization

Data-driven regularization is particularly useful for addressing inverse problems. In contrast to traditional methods, this approach learns the regularization term from the data itself, which can be advantageous when the underlying physical process is not well understood, or when traditional methods may not be optimal. Data-driven regularization methods aim to learn a regulariser instead of manually crafting it. One advantage of these methods is that they automate the process of learning regularization directly from training data. Notable methods include adversarial regularization[192, 296], Tikhonov network[172], total deep variation[154], and regularization by denoising. As also seen in the previous chapter, DIP[300] use of neural networks via parametrization can intrinsically regularize. Solving inverse problems using a pretrained denoiser is also possible using PnP methods[4, 40, 146, 255]. PnP methods leverage pretrained denoising networks within the optimization framework. The denoiser acts as an implicit regulariser, guiding the reconstruction of solutions with less noise. These methods are similar to data-driven regularization but with additional constraints placed on the denoiser. Variational methods with learned priors involve the use of a neural network to represent the prior distribution (regularizer) that encodes the desirable properties of the solution.

It is generally easier to analyze learned regularization schemes using traditional functional analysis techniques[308]. However, these methods tend to underperform compared with end-to-end supervised approaches. Additionally, they involve solving a challenging, high-dimensional, potentially non-convex optimization problem that lacks a provably convergent algorithm, leading to slow reconstruction times. To address these limitations, a combination of end-to-end training with learned regularization was utilized, which involved supervised learning for the reconstruction task and adversarial training for the regularization network[218]. This approach allows joint optimization. Altekruger et al. [8]proposed a regularizer that incorporates a normalizing flow[155] trained on small image patches, which allows for explicit density evaluation. PatchNR consists of a function trained to approximate the distribution of patches, as the structure of small patches is simpler than that of whole images. With a large database of patches extracted from a few images, only a small number of training images is required. During the reconstruction, learned patchNR is used as a regularizer of the variational formulation to solve the inverse problem.

PatchGAN as a Regularizer: Lunz et al.[192] proposed an Adversarial Regularizer (AR) to address inverse problems using data-driven methods and a neural network as a regularization function. This variational framework incorporates knowledge of the forward operator and noise model into the reconstruction process, allowing for effective training on limited data. AR has been extended to learning patch-based regularizers (localARs)[238]. This network acts as a local regularizer and imparts prior knowledge of patch distribution to the entire image without size limitations. The CNN-based network was learned as a regularizer as a patch discriminator through WGAN. The PatchGAN discriminator fulfils a dual role, serving both as a data-driven regularizer and an image synthesis model. As a data-driven regularizer, it urges the generator to produce images that are indistinguishable from the real images to the discriminator. This training process fosters the generation of images with the same local structure as real images, which serves as a form of regularization for the GANs.

In conclusion, patch-based discriminators act as data-driven regularizations within GAN frameworks by compelling the model to capture the localized features of images. This approach forms the basis of SANGRIA and is defined in detail in the next section.

5.3 SANGRIA

Recall methods such as CS-like and SPIRiT, which enforce calibration and data consistency during reconstruction. Calibration consistency ensures that the estimated sensitivity maps accurately represent the actual sensitivity profiles of the receiver coils, whereas data consistency refers to the consistency between the reconstructed image and acquired k-space data. methods, such as iGRAPPA and iRAKI, integrate data consistency have been proposed. The introduction of data-driven regularization in the reconstruction process by introducing adversarial loss while simultaneously enforcing data consistency in a one-shot manner is the main concept behind SANGRIA. This approach has the advantage of requiring less data for training, including scan-specific data, and it uses a patch discriminator to enhance the reconstructed image quality.

In this work, we focus on uniform undersampled acquisitions. A convolutional neural network operating in the k-space domain is combined with patch-based adversarial learning in the image domain and trained for every new image. It has been demonstrated both in simulation and in vivo that this approach can achieve higher visual quality at higher acceleration rates than the current standard MRI reconstructions of undersampled 2D acquisitions of brain images. The adversarial discriminator improves the quality of the reconstructed images, although this does not translate into objective measurements and the generalization ability of the learned models. Our approach can be extended to further accelerate the acquisition of 3D MRI k-space data.

5.3.1 Model

We propose an MRI reconstruction model driven by an adversarial loss as shown in Fig 5.1. We use as SR model, G, state-of-the-art zero-shot reconstruction models such as GRAPPA and RAKI for uniform sampling. Furthermore, to enhance the quality of the reconstructed images, an adversarial discriminator that operates on patches in image space (after inverse Fourier transform and coil combination) is proposed. The adversarial model is trained from a single fully-sampled coil-combined image. This patch-based discriminative network, D, is trained to distinguish between patches generated from the fully sampled image as well as from the reconstructed image. It is a convolutional neural network that takes a patch as input and outputs a classification score.

5.3.2 Model description

Patch discriminator.

We introduce a discriminator to increase the quality of the upsampled image as it has been proposed for natural image super-resolution[170] and more recently for MRI images[318]. As the goal is to learn a super-resolution model for each slice, one can only rely on a patch-based adversarial discriminator, which has to be learned on a few data. To cope with this particular setting, we experimented with a

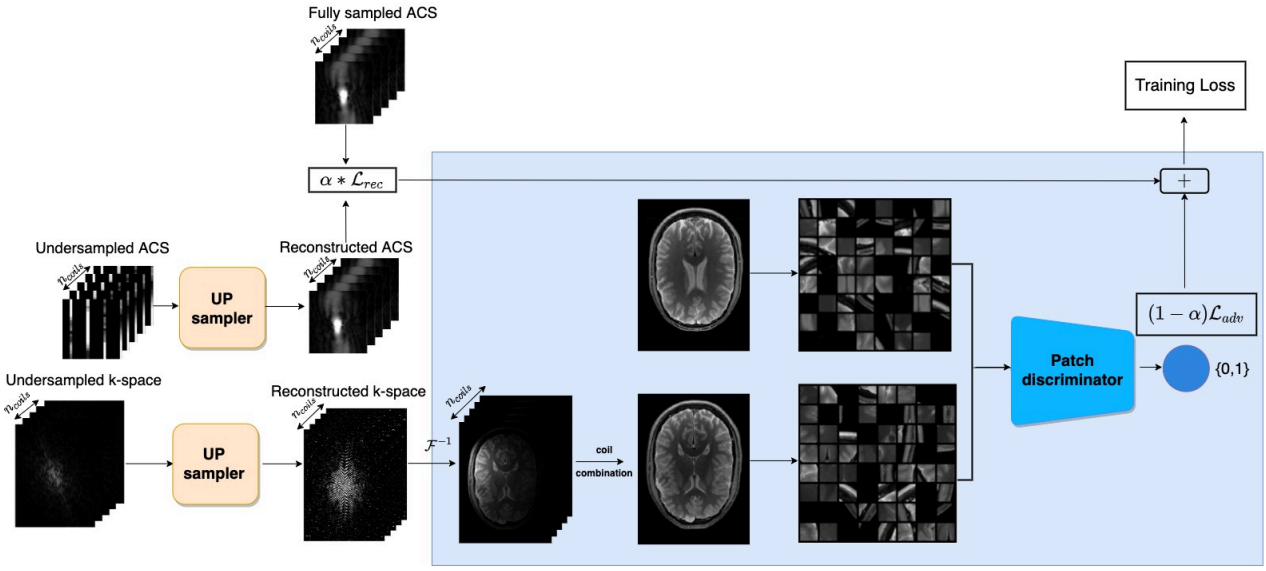

Figure 5.1: The architecture of the proposed one-shot learning Super-ResolutIon with Adversarial training, SANGRIA. It consists of a super-resolution model (upsampler) coupled with a patch-discriminator. Super-resolution models are either GRAPPA, RAKI or cRAKI.

shallow patch GAN discriminator with a limited receptive field. We choose it to accept input as 17x17 patch. The discriminator consists of three convolution blocks, each including a convolution layer, batch normalization and LeakyReLU, followed by a fully connected layer with a single decision output. Convolution layers have successively 1, 32 and 64 filters with size equal to 5×5. The discriminator takes the coil-combined image as input and outputs a decision score for each patch.

Training.

The full model is trained with a combined loss which is the sum of the reconstruction loss and of the adversarial loss.

$$\min_G \max_D \alpha \mathscr{L}_{rec}(G) + (1-\alpha)\mathscr{L}_{adv}(G,D) \tag{5.1}$$

where α is a regularization parameter between 0 and 1 that balances the two losses. The reconstruction loss is the mean squared error loss between the reconstructed ACS $\hat{y}$ and the fully-sampled ACS y, restricted to the training partition.

$$\mathscr{L}_{rec}(y,\hat{y}) = \frac{1}{N}\sum_{j=1}^{N}(y_j - \hat{y}_j)^2 \tag{5.2}$$

where N denotes the total number of k-space points. The adversarial loss $\mathscr{L}_{adv}$ penalizes the reconstructor for not generating realistic patches. discriminator is trained to classify correctly the patches coming from both the patches from fully sampled image and from the reconstructed image generated by G. The fully sampled image we call x_{ref} is the single fully sampled coil-combined (SoS) in the image space. Simultaneously G is trained to fool the discriminator and to generate a reconstructed k-space $G(S)$ (with S the multi-coils undersampled k-space), whose patches (after transformation in the image space) are classified as realistic by the discriminator.

$$\mathscr{L}_{adv}(G,D)(S) = \mathbb{E}_{p\sim x_{ref}}[\log D(p)] + \mathbb{E}_{p\sim i_R}[1-\log D(p)] \tag{5.3}$$

where i_R stands for the reconstructed image in the image space $i_R = \sqrt{\sum_{i=1}^{n_c}(\mathscr{F}^{-1}G(S_i))^2}$ from the inputs coils S, $p \sim x_{ref}$ means that p follows the patch distribution of x_{ref} and $p \sim i_R$ means that p follows the patch distribution of i_R.

5.4 Experiments

5.4.1 Data and experimental setting

Data.

We performed experiments on both synthetic brain images and in-vivo brain data acquired from a healthy volunteer. In-vivo data consists of various multi-slice T2-weighted 2D spin-echo brain acquisitions with in-plane resolution of 0.8 mm, matrix size 256x196 at 3 Tesla (Siemens Vida). Informed consent was obtained. Fully-sampled data (acquisition time T_{acq} of 4 minutes), as well as under-sampled data with acceleration factors ranging from R=2 to 6, were acquired, resulting in R times faster acquisitions. ACS of size (49,52,64,40) is acquired in the separated fast scan. Concerning synthetic data, brain volumes from the MNI152 standard-space T1-weighted average structural template were generated[88]. The template provides tissue probability maps of the following tissue classes: grey matter, white matter, cerebrospinal fluid, bone, skin and air. The MNI template is provided in the Statistical Parametric Mapping (SPM12) Matlab software. Synthetic brain volumes were simulated by setting a different signal level for each class. Coil sensitivities were simulated using Biot and Savart law assuming an $n_c = 8$ receiver-coil array with coils positioned regularly spaced around the brain on a cylindrical surface. From these simulated coil sensitivities and synthetic brain volumes, fully-sampled and under-sampled k-spaces with different acceleration factors R and noise levels, as well as SoS-image combinations were generated.

Experimental setting.

All the methods were trained for a maximum of 1000 epochs with early stopping. The reconstruction model and the discriminator are learned with learning rates that are grid searched on set by cross-validation (learning rates from 1e-4 to 1e-2). In any case, training ends when performance on the validation dataset starts to degrade. We used Adam optimizer with $\beta 1 = 0.9$, $\beta 2 = 0.999$.

5.4.2 Quantitative evaluation of the reconstruction quality

In the first series of experiments, we compare all the methods on data where train and test images correspond to the same acceleration rate and uniform undersampling mask. We provide results for simulated data and for in-vivo data for uniform undersampling.

5.4.2.1 Synthetic data

Table 5.1 shows the quantitative comparison of various methods for the simulated data. We observe that at $R = 4,5$ cRAKI has superior performance in terms of SSIM and PSNR as compared to GRAPPA and RAKI. When GRAPPA is coupled with an adversarial discriminator D, we observe improved performance for both $R = 4,5$ and for all the metrics. Besides the addition of the adversarial discriminator in RAKI and cRAKI modelling slightly degrades the performance in terms of SSIM/PSNR but sighlty improves with respect to NRMSE/NMAE both for $R = 4,5$.

Reconstructor	**R**	**NRMSE**	**NMAE**	**SSIM**	**PSNR**
GRAPPA	4x	0.058 ± 0.060	0.191 ± 0.088	0.943 ± 0.080	38.322 ± 4.463
RAKI	4x	0.066 ± 0.082	0.215 ± 0.084	0.971 ± 0.013	42.785 ± 6.845
cRAKI	4x	0.060 ± 0.085	0.164 ± 0.105	0.984 ± 0.009	44.599 ± 6.267
GRAPPA w/ discriminator	4x	0.050 ± 0.017	0.154 ± 0.028	0.982 ± 0.011	39.526 ± 1.606
RAKI w/ discriminator	4x	0.063 ± 0.064	0.216 ± 0.069	0.942 ± 0.041	37.440 ± 2.886
cRAKI w/ discriminator	4x	0.025 ± 0.015	0.112 ± 0.055	0.974 ± 0.029	44.081 ± 2.680
GRAPPA	5x	0.217 ± 0.239	0.378 ± 0.147	0.808 ± 0.189	28.174 ± 8.813
RAKI	5x	0.541 ± 0.138	0.523 ± 0.074	0.680 ± 0.096	23.893 ± 9.801
cRAKI	5x	0.083 ± 0.076	0.245 ± 0.079	0.966 ± 0.014	40.769 ± 7.360
GRAPPA w/ discriminator	5x	0.113 ± 0.032	0.380 ± 0.049	0.895 ± 0.023	30.551 ± 2.232
RAKI w/ discriminator	5x	0.181 ± 0.250	$0.368 \pm .159$	0.799 ± 0.221	30.700 ± 8.094
cRAKI w/ discriminator	5x	0.047 ± 0.032	0.188 ± 0.061	0.958 ± 0.031	39.100 ± 1.892

Table 5.1: Average metrics over 121 slices for the synthetic data for various reconstructors coupled (or not) with an adversarial patch-discriminator at R=4,5. We performed experiments on uniform undersampling.

5.4.2.2 In-vivo data

Table 5.2 compares the same quantitative metrics of various methods as above but for volunteer data, averaged over 49 slices, and with a higher acceleration rate of 5 and 6. A larger acceleration rate could be obtained in vivo. In this case, data were acquired with n_c=52 receiver coils (while the synthetic data had n_c=8), and the acquisition field-of-view was slightly larger than the brain dimension along the phase encoding. This could explain why the acceleration rate that can be achieved in vivo is larger than the one reached in simulation. Unlike the simulated data, image qualitative metrics such as SSIM/PSNR decreased when GRAPPA and RAKI is coupled with discriminator at both R=5,6. However, we observed an improvement in NRMSE/NMAE when RAKI was trained with the discriminator at both R=5,6. cRAKI and cRAKI with discriminator has similar quantitative metrics at R=5, but it degrades at R=6.

Reconstructor	**R**	**NRMSE**	**NMAE**	**SSIM**	**PSNR**
GRAPPA	5x	0.182 ± 0.046	0.545 ± 0.087	0.946 ± 0.016	37.176 ± 1.665
RAKI	5x	0.209 ± 0.045	0.514 ± 0.072	0.955 ± 0.114	39.276 ± 2.030
cRAKI	5x	0.158 ± 0.038	0.461 ± 0.066	0.898 ± 0.033	36.401 ± 1.835
GRAPPA w/ discriminator	5x	0.239 ± 0.082	0.725 ± 0.173	0.902 ± 0.035	35.013 ± 2.157
RAKI w/ discriminator	5x	0.164 ± 0.039	0.482 ± 0.069	0.920 ± 0.022	36.755 ± 1.668
cRAKI w/ discriminator	5x	0.158 ± 0.038	0.463 ± 0.066	0.901 ± 0.032	36.451 ± 1.809
GRAPPA	6x	0.354 ± 0.046	0.769 ± 0.074	0.844 ± 0.054	28.846 ± 2.539
RAKI	6x	0.265 ± 0.044	0.606 ± 0.063	0.915 ± 0.022	36.088 ± 2.099
cRAKI	6x	0.215 ± 0.046	0.576 ± 0.067	0.867 ± 0.032	33.099 ± 1.701
GRAPPA w/ discriminator	6x	0.427 ± 0.097	1.102 ± 0.190	0.732 ± 0.055	27.079 ± 1.982
RAKI w/ discriminator	6x	0.230 ± 0.055	0.647 ± 0.090	0.876 ± 0.032	33.182 ± 1.926
cRAKI w/ discriminator	6x	0.268 ± 0.072	0.703 ± 0.105	0.854 ± 0.058	32.361 ± 2.532

Table 5.2: Average metrics over 49 slices for the volunteer data for various reconstructors coupled (or not) with an adversarial patch-discriminator, discriminator at R=5,6. The MRI zero-shot methods such as GRAPPA, RAKI and cRAKI are trained with/without adversarial training. We performed experiments on uniform undersampling.

5.4.3 Qualitative evaluation of reconstruction quality

While prior results have not clearly demonstrated an advantage to using an adversarial discriminator with respect to objective evaluation measures, including image quality-based criteria such as SSIM/PSNR, we provide here qualitative results that demonstrate that these objective criteria do not directly relate to subjective image quality. It is reported in the computer vision literature[170] that perceptual quality may not correspond to the objective metrics.

The first series of images in figure 5.2 provides a qualitative comparison of the various methods. It shows the superior visual quality of reconstructed images when the methods are coupled with D. RAKI perform better in terms of the image-quality index such as SSIM/PSNR. However, we observed the images obtained from RAKI has more undersampling artefacts than GRAPPA and artefactual 'dots' around the centre region. When RAKI is coupled with D, artefactual 'dots' disappears and SSIM/PSNR is degraded. NRMSE/NMAE is lower in cRAKI and cRAKI with D, indicating that the model better represents the ACS data. However, in terms of image-based quantitative measured SSIM/PSNR degrades a little, although Figure 5.2 indicates cRAKI with discriminatorimproved visual quality both R=5,6. The cRAKI reconstruction has fewer undersampling artefacts than GRAPPA and the absence of artefactual 'dots' (red square) appeared in RAKI reconstruction. RAKI and cRAKI with discriminator both can perform reconstruction up to R=6. Although, we obtained noisy reconstructions at R=6 but the global features of the brain image were restored, while all other methods failed to reconstruct at this acceleration rate.

The second series of images in figure 5.3 illustrates how the learning changes when the tradeoff between the reconstruction loss and the adversarial loss is varied (through the α hyperparameter in Eq.5.1. Various values of α are used in the range of 1e-6 to 1. The quantitative measures are similar for all the α greater than 1e-3. Visually, however, the reconstruction quality improves for α = 0.1. When α = 1, only reconstruction loss dominates, and undersampling artefacts are visible.

5.5 Discussion and Conclusion

We proposed a one-shot learning approach for MRI reconstruction, which combines k-space-based reconstruction methods such as GRAPPA and RAKI with an adversarial patch discriminator operating in the image domain. Such integration has the potential for more accurate and efficient reconstruction, benefiting from the strengths of both the k-space-based methods and discriminators. Subjective analysis shows an improved quality reconstruction when an adversarial discriminator is introduced. Through adversarial loss, the reconstructed images are more realistic and detailed, focusing on preserving anatomical sharpness and fine image details. We observed that the images appeared better despite the decrease in SSIM/PSNR. This phenomenon is similar to that observed in the previous chapter, which leads to the similar conclusion that these traditional metrics may not fully capture the perceptual image quality. GANs, are designed to improve the perceptual realism of images, which often emphasizes texture and finer details that may not be reflected SSIM and PSNR scores. These metrics do not fully capture aspects of image quality such as texture and perceived sharpness.

However, there are notable drawbacks when adversarial loss is applied in a one-shot setting. One challenge is the potential generation of unrealistic details due to the lack of constraints in the k-space domain. In this study, we incorporate this constraint by adding a MSE in the k-space as a loss function during training to reduce such hallucinations. Furthermore, although adversarial loss

improves perceptual quality, it may not always guarantee fidelity to the original image, which is critical in medical imaging for accurate diagnosis. This can result in the accidental removal of "rare" pathological signals that are medically significant, even if the discriminator is able to remove artifacts.

In conclusion, although SSIM/PSNR are useful objective metrics, they have limitations in assessing perceptual image quality. The use of discriminators focuses on perceptual realism, which can result in images that are subjectively better for human observers despite lower SSIM/PSNR values. while an adversarial loss in MRI reconstruction offers significant improvements in image quality by enhancing texture and detail, it also introduces challenges, such as the potential for unrealistic detail generation and a trade-off between perceptual quality and image fidelity. These drawbacks are particularly pronounced in one-shot scenarios, in which the model must generalize well from a single example. Future research is needed to balance these aspects and to ensure that the reconstructed images are both visually and clinically reliable. A large-scale subjective evaluation could be performed in order to confirm the superiority of this adversarial strategy for MRI reconstruction.

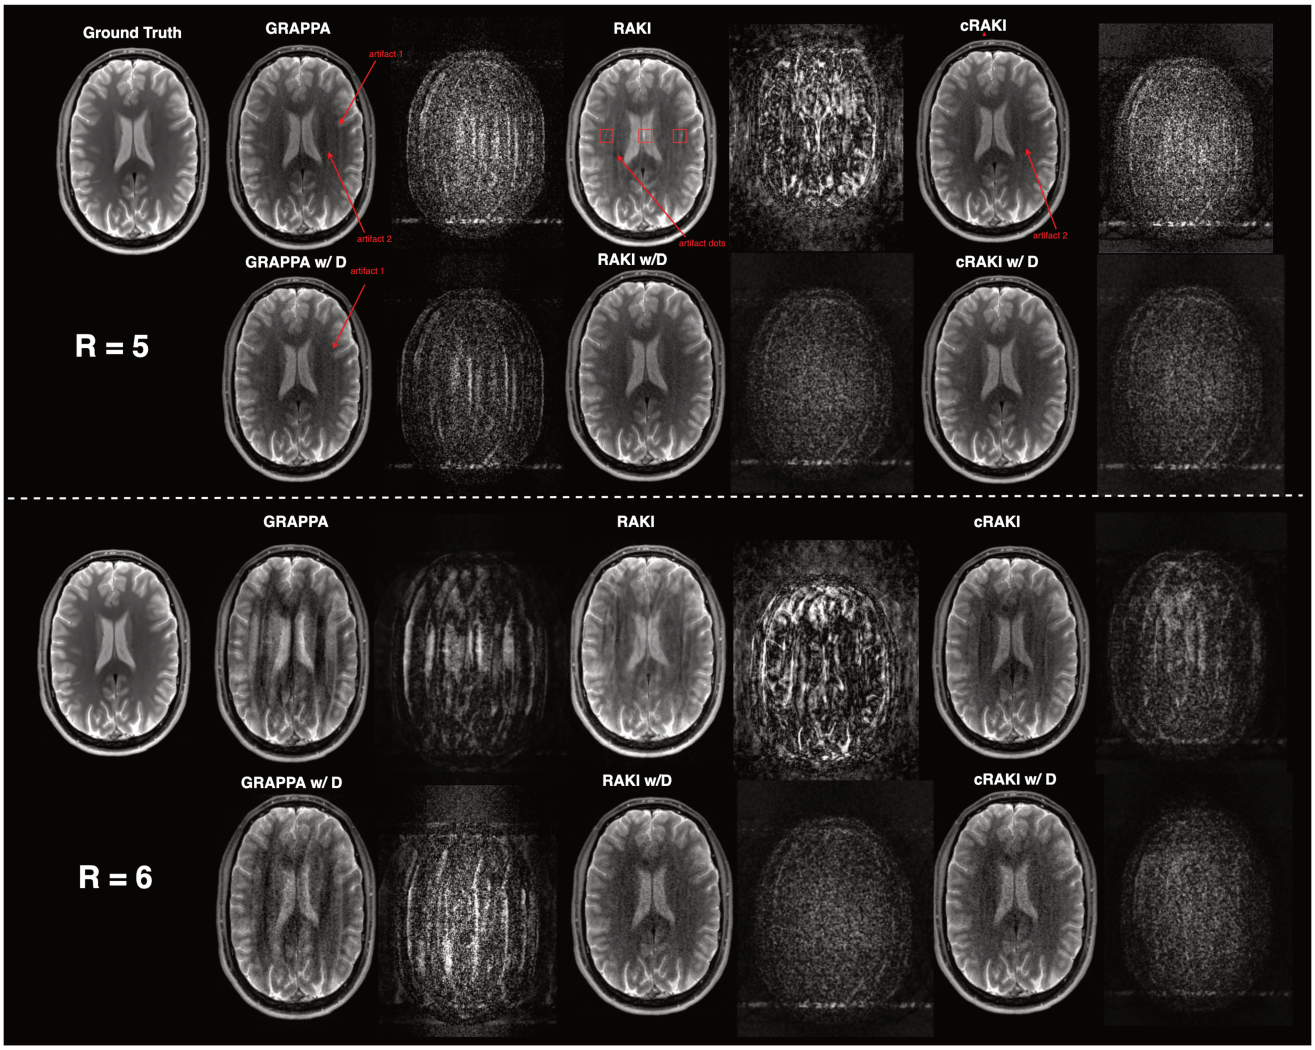

Figure 5.2: Reconstruction results on various slices where the test slice corresponds to the same contrast, acceleration rate and undersampling mask as the train slice for. The red arrows represent the aliasing artifacts from the undersampling. GRAPPA, GRAPPA with D, RAKI & cRAKI have the undersampling. RAKI reconstruction contains some artefactual dots at the centre represented in red squares. RAKI w/D do not display these artefactual dots. cRAKI with D and RAKI with D both have clinically acceptable image quality at R=5. At R=6, all the methods fail to reconstruct except RAKI and cRAKI with D.

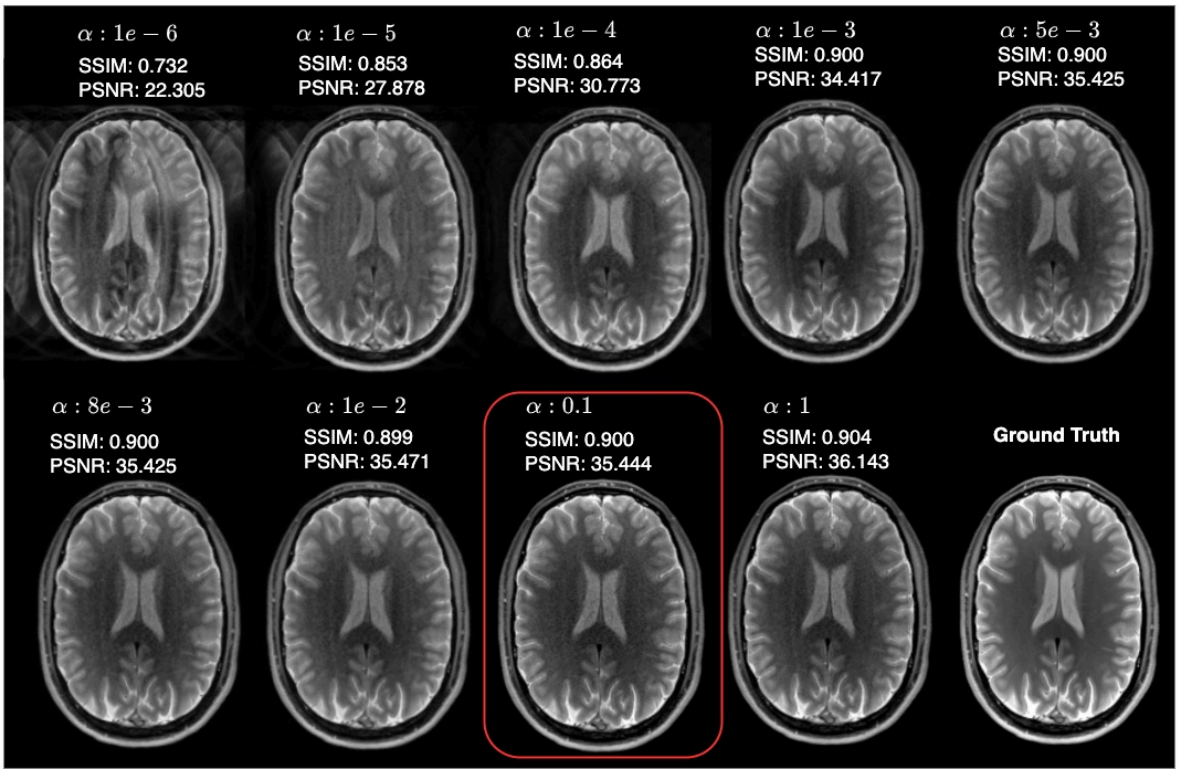

Figure 5.3: Effect of α in Eq.5.1, trade-off parameter between the reconstruction loss and the adversarial loss. The reconstructions are performed with cRAKI at R=5 and varying α from 1e-6 to 1. Generally, the lower the value of α, the smaller the contribution of reconstruction loss

Conclusion

Contributions and Limitations

Contributions

In this thesis, we focus on the reconstruction task in MRI using undersampled k-space and evaluating various scan-specific reconstruction methods. We considered the case of Cartesian sampling including T1W, T2W, several contrasts in multi-contrast 2D/3D gradient echo-type, MPRAGE, and acquisitions. We proposed a new image metric, called COBRAI, which evaluates the effectiveness of the methods in minimizing structured artifacts.

Then, we present a study on enhancing brain image quality in accelerated PI, focusing on time-efficient separated-mode ACS acquisition. Recently, k-space interpolation algorithms based on deep convolutional neural networks have demonstrated the potential for improved reconstruction of highly accelerated MRI scans compared to standard parallel imaging approaches. Often the focus is on the integrated-mode ACS acquisition in the previous studies. In this study, we present a DL-based formalism for the GRAPPA, which can be viewed as a single-layer complex-valued CNN. We compared the performances of rRAKI, its complex implementation (crRAKI]), and GRAPPA. In this study, we extended rRAKI by implementing complex convolutions as well as complex activation functions (complex leaky ReLU) and extended the method for accelerated 3D imaging with regular Cartesian undersampling in two directions. The results demonstrate that crRAKI is more effective in reducing structured artifacts than GRAPPA and rRAKI and that COBRAI is a useful tool for evaluating the performance of these methods across different sequence types.

We conducted a hyperparameter grid search to determine the optimal network parameters for rRAKI and crRAKI, including the fine-tuning of the nonlinearity parameter in the network activation function. We conducted experiments using various receiver coil arrays and different sequence types for 2D and 3D imaging. The resulting crRAKI images for both magnitude- and phase-sensitive applications were compared to those of standard parallel imaging (GRAPPA) and rRAKI. This study was of interest for several reasons. First, although complex-valued neural networks have been previously investigated for MR image reconstruction, we demonstrate the benefits of complex-valued convolutions and activation functions for k-space-based reconstruction in phase-sensitive applications. Additionally, we employed a training–validation split approach for model training to avoid overfitting and allow for early stopping of the learning process. Furthermore, the results of the hyperparameter search for crRAKI provide insights into the choice of reconstruction parameters. Finally, it was found that crRAKI performs robustly, even with a relatively small number of ACS lines. We achieved an acceleration rate of 6x in 2D and 12x in 3D. Another interesting observation is that nonlinearities are detrimental when separated-mode ACS is used. This reduces the generalization performance in situations where ACS size is smaller and the acceleration rates are high.

Additionally, we used a patch-based adversarial training approach called SANGRIA for MRI reconstruction using PI. We concluded that the performance drop is similar to crRAKI, where adding a discriminator decreases the SSIM/PSNR. However, in the subjective quality evaluation part, the discriminator helps reduce the number of artifacts in the reconstruction (and thus increases the subjective quality). Indeed, a discriminator does not really help improve quality (in the sense of distortion), but rather improves the plausibility of an image.

Limitations

Calibration data vs motion: We assumed that the acquired calibration data were perfect, indicating that the patient did not move. Motion artifacts often result in blurring, or ghosting. Although the acquisition of ACS was rapid, motion is still possible. It is important to note that all scan-specific methods rely heavily on the calibration data. This can significantly impact image quality since the training may significantly get affected by degraded ACS.

Limitation of COBRAI: COBRAI is intended to recognize structural artifacts in reconstruction by calculating the correlation between fully-sampled ground truth images and residual maps. However, it is unable to perform this task in cases of severe corruption, such as noise, where it is difficult to discern structures. COBRAI is not sensitive to noise and cannot detect extreme corruption in the absence of discernible structures. As a result, it cannot be considered a universal metric. Other reference-based metrics, such as SSIM and PSNR, should be combined with COBRAI to address this limitation.

Acceleration rate is limited: By reducing the number of k-space samples, it is expected that SNR would decrease. The amount of SNR is influenced by the number of receiver coils, magnetic field strength, and pulse sequences. In our study, when acs=24 and coils=16, we achieved a rate of 4. However, for coil =52, the rate increased to 6. To overcome this limitation, data-driven approaches such as those discussed in hybrid methods can be employed. This strategy combines the adaptability of scan-specific methods with pattern extraction from big dataset.

Runtime: We suggest determining the optimal hyperparameters without the need for fully-sampled ground-truth data. Instead, we propose using only calibration data. To achieve this, we utilized a K-fold cross-validation along with the TVP scheme. Notably, if the search parameter grid is large, it may lead to a computationally expensive grid-search process. Therefore, it is necessary to parallelize such approaches in order to optimize efficiency. It is also worth mentioning that the training adversarial models is lengthy.

Hallucinatory effects of discriminator: The findings from SANGRIA appear reasonable because a discriminator typically aims to reconstruct plausible images and minimize artifacts. However, if the discriminator is able to remove artifacts, it may inadvertently eliminate "rare" structural details which is significant clinically. In future studies, we intend to conduct additional subjective evaluations. Nonetheless, it is important to perform a pathology detection task assessment to ensure that the method does not diminish the likelihood of disease detection. With the availability of datasets such as fastMRI+ [369], where data have been annotated by experts for clinical pathology, it is possible to develop such pipelines. These pipelines may facilitate the fast and robust detection of pathology, which is either hallucinated or removed by reconstruction models. Furthermore, the feasibility of using COBRAI in such pipelines can be studied. Chung et al. [53] proposed the formalization of

pathology-detection assessment tasks as an object detection task. YOLOv5[1] was used. The training data for the YOLOv5 model consists of images from fastMRI+. The model weights were fine-tuned using an open-source pre-trained model.

Perspectives

Incorporating scan-specific methods with neural networks: In scan-specific settings, an image-based error correction is not practical, given that ACS is low resolution. The errors in the image space of the ACS region do not match the reconstruction artifacts found in the higher-resolution reconstruction. Thus, k-space-based techniques are more suitable in a scan-specific setting because they are better equipped to handle specific errors that arise during scanning. An effective image-based super-resolution task may require a database of accelerated reconstructions and the corresponding fully sampled images. Therefore, a natural extension is to combine scan-specific and data-driven techniques. The examples previously mentioned such as GRAPPANET[286] and CSGM[135], leverage the learned representation from the data and adapt to the individual patient. This is due to the fact that the patient's position in the scanner and anatomy can affect coil sensitivity, potentially resulting in a scenario where the dataset used for training does not have an equivalent example.

In this study, we noticed that linear methods such as GRAPPA and crRAKI led to reconstructions with residual maps that exhibited a Gaussian-like profile. This outcome implies that the reconstructions were noisy, and the denoisers could be trained to remove this noise. Currently, various denoising methods are available. Supervised learning has been suggested for MRI denoising, but they are not feasible for clinical use because of the challenges of obtaining paired high- and low-SNR image pairs across various anatomies, field strengths, vendors, and clinical scenarios, which leads to a significant decline in model performance, as demonstrated by Darestani et al.[61]. Therefore, unsupervised denoising techniques are required for this purpose. One such approach is DDM2[336], an unsupervised denoising technique for diffusion MRI that utilizes DM. However, denoisers tend to eliminate the clinically relevant fine details in the process of beautifying it. To overcome this challenge, one proposal is using COBRAI as a loss function that can balance the details eliminated from the reconstruction by integrating it with traditional metrics, such as MSE or SSIM for future work. Another potential strategy could be to train end-to-end these linear methods (scan-specific methods) and denoisers on a dataset. Then, fine-tuning the scan-specific layer each time new data becomes available.

Extension of crRAKI to other regular patterns: Another popular regular sampling pattern employed in the MRI scanner is Controlled Aliasing in Parallel Imaging Results in Higher Acceleration (CAIPIRINHA) for 3D breath-holding abdominal and 2D imaging [36, 37]. In contrast to SENSE and GRAPPA, which can be performed in 3D acquisition mode using acceleration in two phase-encoding directions simultaneously, CAIPIRINHA uses a unique k-space sampling pattern to reduce pixel aliasing and overlap on the reconstructed images. To achieve this, the acquired points in k-space are shifted from one another by applying additional offsets in the phase-encoding gradient direction. There may exist other "locally-regular" sampling patterns. The key idea is determining an invertible reordering transformation of the locally-regular patterns to grappa-like uniform patterns. Then integrate this transformation function with crRAKI. Wave-CAIPI[30] can achieve an acceleration rate of 9-12x in 3D while maintaining a low level of artifacts and negligible g-factor penalties. This simplistic approach is

[1] https://github.com/ultralytics/yolov5

particularly interesting for accelerating high-resolution volumetric data. In future studies, identify how combining the locally-regular patterns with crRAKI would scale in terms of acceleration rate.

COBRAI as a loss function/Pathology-driven sampling: The use of restoration and degradation models is crucial in the fields of image and signal processing because they allow for the analysis, simulation, and manipulation of processes that affect images and signals, including blurring and noise. By simulating the degradation, it is possible to evaluate image-processing algorithms under various conditions. Restoration methods aim to improve the quality of degraded images by removing noise, reducing blurring, and enhancing the details etc. However, by using metrics such as MSE/SSIM for training, such restoration models can result in over-smooth reconstructions[257]. Therefore, integrating COBRAI into the training process can enhance performance by striking a balance between beautifying and preserving the relevant details. COBRAI is naturally differentiable and easy to integrate, thus making it a simple and effective training metric. It can also be used as a loss function for optimizing the sampling mask and is a natural extension of optimizing k-space undersampling to improve the detection and localization of target pathologies in reconstructed images without explicitly using annotated datasets for pathology.

Multi-contrast MRI k-space and image-space SR: Clinical studies have several limitations when using multiparametric MRI, including low signal-to-noise ratios and long scan times. These limitations can lead to difficulties in obtaining a dataset for multiparametric maps, which can affect generalizability. In addition, super-resolution models trained on a dataset may introduce biases that could lead to potential misdiagnosis. Therefore, scan-specific models are clinically relevant to overcome these limitations and improve diagnosis. However, multiparametric isotropic 3D scans are not always feasible because of time constraints, motion, and patient conditions. SR can provide an alternative method for obtaining isotropic 3D scans from anisotropic 2D scans. To address this, McGinnis et al. proposed a multicontrast super-resolution framework that requires only low-resolution MR scans of different sequences and views for supervision[209]. An implicit neural representation (INR) jointly learns two different contrasts of complementary views and exchanges anatomical information among them. Mutual Information (MI) is used as a metric. Another potential approach may be to directly acquire highly undersampled data in 3D and train the models in k-space. As previously observed, an acceleration rate of 12x was achieved using crraki for phase mapping in 3D. This is because crRAKI learning is restricted to uniform undersampling and a single slice in 2D or a single volume in 3D. The acceleration rate can be further improved by leveraging other regular patterns or mutual information between various contrasts in both 2D and 3D.

Bibliography

[1] Richard G. Abramson, Pei-Fang Su, and Yu Shyr. "Quantitative Metrics in Clinical Radiology Reporting: A Snapshot Perspective from a Single Mixed Academic-Community Practice". In: *Magnetic Resonance Imaging*. Quantitative Imaging in Cancer 30.9 (Nov. 1, 2012), pp. 1357–1366. ISSN: 0730-725X. DOI: 10.1016/j.mri.2012.04.018. URL: https://www.sciencedirect.com/science/article/pii/S0730725X12002020 (visited on 03/21/2024) (cit. on p. 26).

[2] Hemant K. Aggarwal and Mathews Jacob. "JOINT OPTIMIZATION OF SAMPLING PATTERN AND PRIORS IN MODEL BASED DEEP LEARNING". In: *Proceedings. IEEE International Symposium on Biomedical Imaging* 2020 (Apr. 2020), pp. 926–929. ISSN: 1945-7928. DOI: 10.1109/isbi45749.2020.9098639. pmid: 33584975 (cit. on p. 55).

[3] Hemant Kumar Aggarwal and Mathews Jacob. "J-MoDL: Joint Model-Based Deep Learning for Optimized Sampling and Reconstruction". In: *IEEE Journal of Selected Topics in Signal Processing* 14.6 (Oct. 2020), pp. 1151–1162. ISSN: 1941-0484. DOI: 10.1109/JSTSP.2020.3004094. URL: https://ieeexplore.ieee.org/document/9122388 (visited on 04/04/2024) (cit. on pp. 52, 55).

[4] Rizwan Ahmad, Charles A. Bouman, Gregery T. Buzzard, et al. "Plug and Play Methods for Magnetic Resonance Imaging (Long Version)". In: *IEEE Signal Processing Magazine* 37.1 (Jan. 2020), pp. 105–116. ISSN: 1053-5888, 1558-0792. DOI: 10.1109/MSP.2019.2949470. arXiv: 1903.08616 [cs]. URL: http://arxiv.org/abs/1903.08616 (visited on 04/10/2024) (cit. on pp. 53, 111).

[5] Mehmet Akçakaya, Steen Moeller, Sebastian Weingärtner, et al. "Scan-Specific Robust Artificial-Neural-Networks for k-Space Interpolation (RAKI) Reconstruction: Database-free Deep Learning for Fast Imaging". In: *Magnetic Resonance in Medicine* 81.1 (2019), pp. 439–453. ISSN: 1522-2594. DOI: 10.1002/mrm.27420. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27420 (visited on 10/05/2020) (cit. on pp. 47, 75, 80, 83, 106).

[6] Mehmet Akçakaya, Maxine Tang, and Reza Nezafat. "Cardiac Magnetic Resonance Imaging Physics". In: *Cardiovascular Magnetic Resonance Imaging*. Ed. by Raymond Y. Kwong, Michael Jerosch-Herold, and Bobak Heydari. New York, NY: Springer, 2019, pp. 1–16. ISBN: 978-1-4939-8841-9. DOI: 10.1007/978-1-4939-8841-9_1. URL: https://doi.org/10.1007/978-1-4939-8841-9_1 (visited on 06/13/2024) (cit. on p. 43).

[7] Horé Alain and Djemel Ziou. *Image Quality Metrics: PSNR vs. SSIM*. URL: https://ieeexplore.ieee.org/document/5596999 (visited on 06/18/2024) (cit. on p. 66).

[8] Fabian Altekrüger, Alexander Denker, Paul Hagemann, et al. *PatchNR: Learning from Very Few Images by Patch Normalizing Flow Regularization*. Nov. 21, 2022. arXiv: 2205.12021 [cs, eess, math]. URL: http://arxiv.org/abs/2205.12021 (visited on 01/24/2023). preprint (cit. on p. 112).

[9] Hifz Aniq and Robert Campbell. “Chapter 13 - Magnetic Resonance Imaging”. In: *Pain Management (Second Edition)*. Ed. by Steven D. Waldman. Philadelphia: W.B. Saunders, Jan. 1, 2011, pp. 106–116. ISBN: 978-1-4377-0721-2. DOI: 10.1016/B978-1-4377-0721-2.00013-1. URL: https://www.sciencedirect.com/science/article/pii/B9781437707212000131 (visited on 03/20/2024) (cit. on p. 15).

[10] Saeed Anwar and Nick Barnes. “Densely Residual Laplacian Super-Resolution”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44.3 (Mar. 2022), pp. 1192–1204. ISSN: 1939-3539. DOI: 10.1109/TPAMI.2020.3021088. URL: https://ieeexplore.ieee.org/document/9185010 (visited on 04/02/2024) (cit. on p. 42).

[11] Yamin Arefeen, Onur Beker, Jaejin Cho, et al. “Scan-Specific Artifact Reduction in k-Space (SPARK) Neural Networks Synergize with Physics-Based Reconstruction to Accelerate MRI”. In: *Magnetic Resonance in Medicine* 87.2 (2022), pp. 764–780. ISSN: 1522-2594. DOI: 10.1002/mrm.29036. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29036 (visited on 10/06/2022) (cit. on pp. 47, 75).

[12] Martin Arjovsky, Soumith Chintala, and Léon Bottou. “Wasserstein Generative Adversarial Networks”. In: *Proceedings of the 34th International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, July 17, 2017, pp. 214–223. URL: https://proceedings.mlr.press/v70/arjovsky17a.html (visited on 04/02/2024) (cit. on p. 42).

[13] Martin Arjovsky, Amar Shah, and Yoshua Bengio. *Unitary Evolution Recurrent Neural Networks*. May 25, 2016. DOI: 10.48550/arXiv.1511.06464. arXiv: 1511.06464 [cs, stat]. URL: http://arxiv.org/abs/1511.06464 (visited on 03/31/2024). preprint (cit. on p. 36).

[14] Cagla D. Bahadir, Alan Q. Wang, Adrian V. Dalca, et al. “Deep-Learning-Based Optimization of the Under-Sampling Pattern in MRI”. In: *IEEE Transactions on Computational Imaging* 6 (2020), pp. 1139–1152. ISSN: 2333-9403. DOI: 10.1109/TCI.2020.3006727. URL: https://ieeexplore.ieee.org/document/9133281 (visited on 04/04/2024) (cit. on p. 55).

[15] Tim Bakker, Matthew Muckley, Adriana Romero-Soriano, et al. *On Learning Adaptive Acquisition Policies for Undersampled Multi-Coil MRI Reconstruction*. Mar. 30, 2022. DOI: 10.48550/arXiv.2203.16392. arXiv: 2203.16392 [cs, eess]. URL: http://arxiv.org/abs/2203.16392 (visited on 12/31/2022). preprint (cit. on pp. 55, 56).

[16] Tim Bakker, Herke van Hoof, and Max Welling. *Experimental Design for MRI by Greedy Policy Search*. Dec. 15, 2020. DOI: 10.48550/arXiv.2010.16262. arXiv: 2010.16262 [cs]. URL: http://arxiv.org/abs/2010.16262 (visited on 04/04/2024). preprint (cit. on p. 55).

[17] R. Bakshi, S. Ariyaratana, R. H. Benedict, et al. “Fluid-Attenuated Inversion Recovery Magnetic Resonance Imaging Detects Cortical and Juxtacortical Multiple Sclerosis Lesions”. In: *Archives of Neurology* 58.5 (May 2001), pp. 742–748. ISSN: 0003-9942. DOI: 10.1001/archneur.58.5.742. pmid: 11346369 (cit. on p. 25).

[18] Chaity Banerjee, Tathagata Mukherjee, and Eduardo Pasiliao. “An Empirical Study on Generalizations of the ReLU Activation Function”. In: *Proceedings of the 2019 ACM Southeast Conference*. ACM SE ’19. New York, NY, USA: Association for Computing Machinery, Apr. 18, 2019, pp. 164–167. ISBN: 978-1-4503-6251-1. DOI: 10.1145/3299815.3314450. URL: https://doi.org/10.1145/3299815.3314450 (visited on 01/23/2023) (cit. on p. 47).

[19] Joshua Batson and Loic Royer. *Noise2Self: Blind Denoising by Self-Supervision*. June 8, 2019. DOI: 10.48550/arXiv.1901.11365. arXiv: 1901.11365 [cs, stat]. URL: http://arxiv.org/abs/1901.11365 (visited on 04/23/2023). preprint (cit. on p. 108).

[20] Youssef Beauferris, Jonas Teuwen, Dimitrios Karkalousos, et al. “Multi-Coil MRI Reconstruction Challenge—Assessing Brain MRI Reconstruction Models and Their Generalizability to Varying Coil Configurations”. In: *Frontiers in Neuroscience* 16 (July 6, 2022), p. 919186. ISSN: 1662-4548. DOI: 10.3389/fnins.2022.919186. pmid: 35873808. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9298878/ (visited on 12/12/2022) (cit. on pp. 45, 75).

[21] Sefi Bell-Kligler, Assaf Shocher, and Michal Irani. *Blind Super-Resolution Kernel Estimation Using an Internal-GAN*. Jan. 7, 2020. DOI: 10.48550/arXiv.1909.06581. arXiv: 1909.06581 [cs]. URL: http://arxiv.org/abs/1909.06581 (visited on 04/13/2024). preprint (cit. on p. 110).

[22] Aurélien Bellet, Amaury Habrard, and Marc Sebban. *A Survey on Metric Learning for Feature Vectors and Structured Data*. Feb. 12, 2014. DOI: 10.48550/arXiv.1306.6709. arXiv: 1306.6709 [cs, stat]. URL: http://arxiv.org/abs/1306.6709 (visited on 04/16/2024). preprint (cit. on p. 110).

[23] Noam Ben-Eliezer, Daniel K. Sodickson, and Kai Tobias Block. “Rapid and Accurate T2 Mapping from Multi–Spin-Echo Data Using Bloch-simulation-based Reconstruction”. In: *Magnetic Resonance in Medicine* 73.2 (2015), pp. 809–817. ISSN: 1522-2594. DOI: 10.1002/mrm.25156. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.25156 (visited on 04/07/2024) (cit. on p. 58).

[24] Yoshua Bengio. *Springtime for AI: The Rise of Deep Learning*. Scientific American. June 1, 2016. URL: https://www.scientificamerican.com/article/springtime-for-ai-the-rise-of-deep-learning/ (visited on 03/23/2024) (cit. on p. 29).

[25] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. *Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation*. Aug. 15, 2013. DOI: 10.48550/arXiv.1308.3432. arXiv: 1308.3432 [cs]. URL: http://arxiv.org/abs/1308.3432 (visited on 04/04/2024). preprint (cit. on p. 55).

[26] Inbal Beracha, Amir Seginer, and Assaf Tal. “Adaptive Model-Based Magnetic Resonance”. In: *Magnetic Resonance in Medicine* 90.3 (2023), pp. 839–851. ISSN: 1522-2594. DOI: 10.1002/mrm.29688. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29688 (visited on 04/04/2024) (cit. on p. 56).

[27] Kamal Berahmand, Fatemeh Daneshfar, Elaheh Sadat Salehi, et al. “Autoencoders and Their Applications in Machine Learning: A Survey”. In: *Artificial Intelligence Review* 57.2 (Feb. 3, 2024), p. 28. ISSN: 1573-7462. DOI: 10.1007/s10462-023-10662-6. URL: https://doi.org/10.1007/s10462-023-10662-6 (visited on 03/23/2024) (cit. on p. 35).

[28] Matt A. Bernstein, Kevin F. King, and Xiaohong Joe Zhou. *Handbook of MRI Pulse Sequences*. Elsevier Inc., Sept. 7, 2004. ISBN: 978-0-12-092861-3. DOI: 10.1016/B978-0-12-092861-3.X5000-6. URL: http://www.scopus.com/inward/record.url?scp=84941965835&partnerID=8YFLogxK (visited on 03/20/2024) (cit. on p. 13).

[29] Wanyu Bian, Albert Jang, and Fang Liu. “Improving Quantitative MRI Using Self-Supervised Deep Learning with Model Reinforcement: Demonstration for Rapid T1 Mapping”. In: *Magnetic Resonance in Medicine* n/a.n/a (). ISSN: 1522-2594. DOI: 10.1002/mrm.30045. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.30045 (visited on 04/08/2024) (cit. on p. 59).

[30] Berkin Bilgic, Borjan A. Gagoski, Stephen F. Cauley, et al. “Wave-CAIPI for Highly Accelerated 3D Imaging”. In: *Magnetic Resonance in Medicine* 73.6 (June 2015), pp. 2152–2162. ISSN: 1522-2594. DOI: 10.1002/mrm.25347. pmid: 24986223 (cit. on p. 124).

[31] Chris M. Bishop. “Training with Noise Is Equivalent to Tikhonov Regularization”. In: *Neural Computation* 7.1 (Jan. 1995), pp. 108–116. ISSN: 0899-7667. DOI: 10.1162/neco.1995.7.1.108. URL: https://ieeexplore.ieee.org/abstract/document/6796505 (visited on 03/23/2024) (cit. on p. 33).

[32] F. Bloch, W. W. Hansen, and Martin Packard. “Nuclear Induction”. In: *Physical Review* 69.3-4 (Feb. 1, 1946), pp. 127–127. DOI: 10.1103/PhysRev.69.127. URL: https://link.aps.org/doi/10.1103/PhysRev.69.127 (visited on 02/12/2024) (cit. on pp. 13, 17).

[33] Ashish Bora, Ajil Jalal, Eric Price, et al. *Compressed Sensing Using Generative Models*. Mar. 9, 2017. DOI: 10.48550/arXiv.1703.03208. arXiv: 1703.03208 [cs, math, stat]. URL: http://arxiv.org/abs/1703.03208 (visited on 04/04/2024). preprint (cit. on pp. 54, 56).

[34] H. Bourlard and Y. Kamp. “Auto-Association by Multilayer Perceptrons and Singular Value Decomposition”. In: *Biological Cybernetics* 59.4 (Sept. 1, 1988), pp. 291–294. ISSN: 1432-0770. DOI: 10.1007/BF00332918. URL: https://doi.org/10.1007/BF00332918 (visited on 03/23/2024) (cit. on p. 34).

[35] M. Brant-Zawadzki, G. D. Gillan, and W. R. Nitz. “MP RAGE: A Three-Dimensional, T1-weighted, Gradient-Echo Sequence–Initial Experience in the Brain”. In: *Radiology* 182.3 (Mar. 1992), pp. 769–775. ISSN: 0033-8419. DOI: 10.1148/radiology.182.3.1535892. pmid: 1535892 (cit. on p. 25).

[36] Felix A. Breuer, Martin Blaimer, Robin M. Heidemann, et al. “Controlled Aliasing in Parallel Imaging Results in Higher Acceleration (CAIPIRINHA) for Multi-Slice Imaging”. In: *Magnetic Resonance in Medicine* 53.3 (Mar. 2005), pp. 684–691. ISSN: 0740-3194. DOI: 10.1002/mrm.20401. pmid: 15723404 (cit. on p. 124).

[37] Felix A. Breuer, Martin Blaimer, Matthias F. Mueller, et al. “Controlled Aliasing in Volumetric Parallel Imaging (2D CAIPIRINHA)”. In: *Magnetic Resonance in Medicine* 55.3 (2006), pp. 549–556. ISSN: 1522-2594. DOI: 10.1002/mrm.20787. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.20787 (visited on 03/21/2023) (cit. on pp. 108, 124).

[38] Andrew Brock, Jeff Donahue, and Karen Simonyan. *Large Scale GAN Training for High Fidelity Natural Image Synthesis*. Feb. 25, 2019. DOI: 10.48550/arXiv.1809.11096. arXiv: 1809.11096 [cs, stat]. URL: http://arxiv.org/abs/1809.11096 (visited on 03/23/2024). preprint (cit. on p. 36).

[39] Wilhelm Burger and Mark James Burge. *Principles of Digital Image Processing: Core Algorithms*. Undergraduate Topics in Computer Science. London: Springer, 2009. ISBN: 978-1-84800-194-7 978-1-84800-195-4. DOI: 10.1007/978-1-84800-195-4. URL: https://link.springer.com/10.1007/978-1-84800-195-4 (visited on 06/05/2024) (cit. on p. 42).

[40] Gregery T. Buzzard, Stanley H. Chan, Suhas Sreehari, et al. *Plug-and-Play Unplugged: Optimization Free Reconstruction Using Consensus Equilibrium*. June 13, 2018. DOI: 10.48550/arXiv.1705.08983. arXiv: 1705.08983 [cs, math]. URL: http://arxiv.org/abs/1705.08983 (visited on 04/10/2024). preprint (cit. on pp. 53, 111).

[41] Stuart Cantrill. “A Shift in Expectations”. In: *Nature Physics* 4.1 (Feb. 2008), S11–S11. ISSN: 1745-2481. DOI: 10.1038/nphys865. URL: https://www.nature.com/articles/milespin10 (visited on 03/20/2024) (cit. on p. 13).

[42] Y. Cao and D. N. Levin. “Feature-Recognizing MRI”. In: *Magnetic Resonance in Medicine* 30.3 (Sept. 1993), pp. 305–317. ISSN: 0740-3194. DOI: 10.1002/mrm.1910300306. pmid: 8412601 (cit. on p. 54).

[43] Y. Cao and D. N. Levin. “Using Prior Knowledge of Human Anatomy to Constrain MR Image Acquisition and Reconstruction: Half k-Space and Full k-Space Techniques”. In: *Magnetic Resonance Imaging* 15.6 (1997), pp. 669–677. ISSN: 0730-725X. DOI: 10.1016/s0730-725x(97)00027-1. pmid: 9285806 (cit. on p. 45).

[44] Stephen F. Cauley, Kawin Setsompop, Dan Ma, et al. “Fast Group Matching for MR Fingerprinting Reconstruction”. In: *Magnetic Resonance in Medicine* 74.2 (Aug. 2015), pp. 523–528. ISSN: 1522-2594. DOI: 10.1002/mrm.25439. pmid: 25168690 (cit. on p. 58).

[45] Stanley H. Chan. *Performance Analysis of Plug-and-Play ADMM: A Graph Signal Processing Perspective*. May 17, 2019. DOI: 10.48550/arXiv.1809.00020. arXiv: 1809.00020 [cs, eess]. URL: http://arxiv.org/abs/1809.00020 (visited on 04/10/2024). preprint (cit. on p. 53).

[46] Qi Chen, Fang Liu, Huiyu Duan, et al. “MRIQA: Subjective Method and Objective Model for Magnetic Resonance Image Quality Assessment”. In: *2022 IEEE International Conference on Visual Communications and Image Processing (VCIP)*. 2022 IEEE International Conference on Visual Communications and Image Processing (VCIP). Dec. 2022, pp. 1–5. DOI: 10.1109/VCIP56404.2022.10008885. URL: https://ieeexplore.ieee.org/document/10008885 (visited on 06/23/2024) (cit. on p. 61).

[47] Xi Chen, Wenchuan Wu, and Mark Chiew. “Locally Structured Low-Rank MR Image Reconstruction Using Submatrix Constraints”. In: *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*. 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI). Mar. 2022, pp. 1–4. DOI: 10.1109/ISBI52829.2022.9761692. URL: https://ieeexplore.ieee.org/document/9761692 (visited on 04/09/2024) (cit. on p. 52).

[48] Xiangyu Chen, Xintao Wang, Jiantao Zhou, et al. “Activating More Pixels in Image Super-Resolution Transformer”. In: *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). June 2023, pp. 22367–22377. DOI: 10.1109/CVPR52729.2023.02142. URL: https://ieeexplore.ieee.org/document/10204527 (visited on 04/02/2024) (cit. on p. 42).

[49] Yuhua Chen, Feng Shi, Anthony G. Christodoulou, et al. “Efficient and Accurate MRI Super-Resolution Using a Generative Adversarial Network and 3D Multi-level Densely Connected Network”. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*. Ed. by Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, et al. Cham: Springer International Publishing, 2018, pp. 91–99. ISBN: 978-3-030-00928-1. DOI: 10.1007/978-3-030-00928-1_11 (cit. on p. 43).

[50] Li Sze Chow and Raveendran Paramesran. “Review of Medical Image Quality Assessment”. In: *Biomedical Signal Processing and Control* 27 (May 1, 2016), pp. 145–154. ISSN: 1746-8094. DOI: 10.1016/j.bspc.2016.02.006. URL: https://www.sciencedirect.com/science/article/pii/S1746809416300180 (visited on 04/17/2024) (cit. on p. 62).

[51] Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, et al. “Diffusion Posterior Sampling for General Noisy Inverse Problems”. In: The Eleventh International Conference on Learning Representations. Sept. 29, 2022. URL: https://openreview.net/forum?id=OnD9zGAGT0k (visited on 04/04/2024) (cit. on p. 57).

[52] Hyungjin Chung, Eun Sun Lee, and Jong Chul Ye. “MR Image Denoising and Super-Resolution Using Regularized Reverse Diffusion”. In: *IEEE Transactions on Medical Imaging* 42.4 (Apr. 2023), pp. 922–934. ISSN: 1558-254X. DOI: 10.1109/TMI.2022.3220681. URL: https://ieeexplore.ieee.org/document/9941138 (visited on 04/02/2024) (cit. on p. 43).

[53] Hyungjin Chung and Jong Chul Ye. *Score-Based Diffusion Models for Accelerated MRI*. July 16, 2022. DOI: 10.48550/arXiv.2110.05243. arXiv: 2110.05243 [cs, eess]. URL: http://arxiv.org/abs/2110.05243 (visited on 05/16/2024). preprint (cit. on p. 123).

[54] Elizabeth Cole, Joseph Cheng, John Pauly, et al. “Analysis of Deep Complex-Valued Convolutional Neural Networks for MRI Reconstruction and Phase-Focused Applications”. In: *Magnetic Resonance in Medicine* 86.2 (Aug. 2021), pp. 1093–1109. ISSN: 1522-2594. DOI: 10.1002/mrm.28733. pmid: 33724507 (cit. on pp. 35, 76, 105).

[55] Elizabeth K. Cole, Frank Ong, Shreyas S. Vasanawala, et al. “Fast Unsupervised MRI Reconstruction Without Fully-Sampled Ground Truth Data Using Generative Adversarial Networks”. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, pp. 3988–3997. URL: https://openaccess.thecvf.com/content/ICCV2021W/LCI/html/Cole_Fast_Unsupervised_MRI_Reconstruction_Without_Fully-Sampled_Ground_Truth_Data_Using_ICCVW_2021_paper.html (visited on 04/08/2024) (cit. on p. 53).

[56] Jade Copet, Felix Kreuk, Itai Gat, et al. *Simple and Controllable Music Generation*. Jan. 29, 2024. DOI: 10.48550/arXiv.2306.05284. arXiv: 2306.05284 [cs, eess]. URL: http://arxiv.org/abs/2306.05284 (visited on 03/23/2024). preprint (cit. on p. 36).

[57] Frédérique Crété-Roffet, Thierry Dolmiere, Patricia Ladret, et al. “The Blur Effect: Perception and Estimation with a New No-Reference Perceptual Blur Metric”. In: *SPIE Electronic Imaging Symposium Conf Human Vision and Electronic Imaging*. Vol. XII. San Jose, United States, Jan. 2007, EI 6492–16. URL: https://hal.science/hal-00232709 (visited on 04/16/2024) (cit. on p. 65).

[58] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, et al. “Diffusion Models in Vision: A Survey”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45.9 (Sept. 2023), pp. 10850–10869. ISSN: 1939-3539. DOI: 10.1109/TPAMI.2023.3261988. URL: https://ieeexplore.ieee.org/abstract/document/10081412 (visited on 04/02/2024) (cit. on p. 40).

[59] G. Cybenko. “Approximation by Superpositions of a Sigmoidal Function”. In: *Mathematics of Control, Signals and Systems* 2.4 (Dec. 1, 1989), pp. 303–314. ISSN: 1435-568X. DOI: 10.1007/BF02551274. URL: https://doi.org/10.1007/BF02551274 (visited on 03/23/2024) (cit. on pp. 29, 30).

[60] Salman U.H. Dar, Mahmut Yurt, Mohammad Shahdloo, et al. “Prior-Guided Image Reconstruction for Accelerated Multi-Contrast MRI via Generative Adversarial Networks”. In: *IEEE Journal of Selected Topics in Signal Processing* 14.6 (Oct. 2020), pp. 1072–1087. ISSN: 1941-0484. DOI: 10.1109/JSTSP.2020.3001737 (cit. on pp. 53, 75).

[61] Mohammad Zalbagi Darestani, Akshay S. Chaudhari, and Reinhard Heckel. "Measuring Robustness in Deep Learning Based Compressive Sensing". In: *Proceedings of the 38th International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, July 1, 2021, pp. 2433–2444. URL: https://proceedings.mlr.press/v139/darestani21a.html (visited on 04/10/2024) (cit. on pp. 54, 59, 124).

[62] Mohammad Zalbagi Darestani and Reinhard Heckel. "Accelerated MRI With Un-Trained Neural Networks". In: *IEEE Transactions on Computational Imaging* 7 (2021), pp. 724–733. ISSN: 2333-9403. DOI: 10.1109/TCI.2021.3097596 (cit. on pp. 51, 75).

[63] Mohammad Zalbagi Darestani, Jiayu Liu, and Reinhard Heckel. "Test-Time Training Can Close the Natural Distribution Shift Performance Gap in Deep Learning Based Compressed Sensing". In: *Proceedings of the 39th International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, June 28, 2022, pp. 4754–4776. URL: https://proceedings.mlr.press/v162/darestani22a.html (visited on 04/03/2024) (cit. on p. 54).

[64] Mohammad Zalbagi Darestani, Vishwesh Nath, Wenqi Li, et al. "IR-FRestormer: Iterative Refinement With Fourier-Based Restormer for Accelerated MRI Reconstruction". In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. 2024, pp. 7655–7664. URL: https://openaccess.thecvf.com/content/WACV2024/html/Darestani_IR-FRestormer_Iterative_Refinement_With_Fourier-Based_Restormer_for_Accelerated_MRI_Reconstruction_WACV_2024_paper.html (visited on 04/03/2024) (cit. on p. 53).

[65] Axel Davy, Thibaud Ehret, Jean-Michel Morel, et al. "A Non-Local CNN for Video Denoising". In: *2019 IEEE International Conference on Image Processing (ICIP)*. 2019 IEEE International Conference on Image Processing (ICIP). Sept. 2019, pp. 2409–2413. DOI: 10.1109/ICIP.2019.8803314. URL: https://ieeexplore.ieee.org/document/8803314 (visited on 04/13/2024) (cit. on p. 111).

[66] Peter Dawood, Felix Breuer, Jannik Stebani, et al. "Iterative Training of Robust K-Space Interpolation Networks for Improved Image Reconstruction with Limited Scan Specific Training Samples". In: *Magnetic Resonance in Medicine* 89.2 (2023), pp. 812–827. ISSN: 1522-2594. DOI: 10.1002/mrm.29482. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29482 (visited on 01/31/2023) (cit. on pp. 35, 51, 76, 105–108).

[67] 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". In: *Scientific Reports* 12.1 (Apr. 16, 2022), p. 6362. ISSN: 2045-2322. DOI: 10.1038/s41598-022-10298-6. URL: https://www.nature.com/articles/s41598-022-10298-6 (visited on 03/24/2024) (cit. on p. 41).

[68] Emmanuelle M. Delfaut, Javier Beltran, Glyn Johnson, et al. "Fat Suppression in MR Imaging: Techniques and Pitfalls". In: *RadioGraphics* 19.2 (Mar. 1999), pp. 373–382. ISSN: 0271-5333. DOI: 10.1148/radiographics.19.2.g99mr03373. URL: https://pubs.rsna.org/doi/abs/10.1148/radiographics.19.2.g99mr03373 (visited on 01/13/2023) (cit. on p. 75).

[69] Ludovic de Rochefort, Tian Liu, Bryan Kressler, et al. "Quantitative Susceptibility Map Reconstruction from MR Phase Data Using Bayesian Regularization: Validation and Application to Brain Imaging". In: *Magnetic Resonance in Medicine* 63.1 (Jan. 2010), pp. 194–206. ISSN: 1522-2594. DOI: 10.1002/mrm.22187. pmid: 19953507 (cit. on p. 27).

[70] Arjun D. Desai, Batu M. Ozturkler, Christopher M. Sandino, et al. “Noise2Recon: Enabling SNR-robust MRI Reconstruction with Semi-Supervised and Self-Supervised Learning”. In: *Magnetic Resonance in Medicine* 90.5 (2023), pp. 2052–2070. ISSN: 1522-2594. DOI: 10.1002/mrm.29759. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29759 (visited on 04/08/2024) (cit. on p. 53).

[71] Prafulla Dhariwal, Heewoo Jun, Christine Payne, et al. *Jukebox: A Generative Model for Music*. Apr. 30, 2020. DOI: 10.48550/arXiv.2005.00341. arXiv: 2005.00341 [cs, eess, stat]. URL: http://arxiv.org/abs/2005.00341 (visited on 03/23/2024). preprint (cit. on p. 36).

[72] Katharine A. Dishner, Bala McRae-Posani, Arka Bhowmik, et al. “A Survey of Publicly Available MRI Datasets for Potential Use in Artificial Intelligence Research”. In: *Journal of magnetic resonance imaging: JMRI* 59.2 (Feb. 2024), pp. 450–480. ISSN: 1522-2586. DOI: 10.1002/jmri.29101. pmid: 37888298 (cit. on p. 52).

[73] Mariya Doneva, Peter Börnert, Holger Eggers, et al. “Compressed Sensing Reconstruction for Magnetic Resonance Parameter Mapping”. In: *Magnetic Resonance in Medicine* 64.4 (Oct. 2010), pp. 1114–1120. ISSN: 1522-2594. DOI: 10.1002/mrm.22483. pmid: 20564599 (cit. on p. 58).

[74] Chao Dong, Chen Change Loy, Kaiming He, et al. *Image Super-Resolution Using Deep Convolutional Networks*. July 31, 2015. DOI: 10.48550/arXiv.1501.00092. arXiv: 1501.00092 [cs]. URL: http://arxiv.org/abs/1501.00092 (visited on 04/02/2024). preprint (cit. on p. 42).

[75] Chao Dong, Chen Change Loy, and Xiaoou Tang. “Accelerating the Super-Resolution Convolutional Neural Network”. In: *Computer Vision – ECCV 2016*. Ed. by Bastian Leibe, Jiri Matas, Nicu Sebe, et al. Cham: Springer International Publishing, 2016, pp. 391–407. ISBN: 978-3-319-46475-6. DOI: 10.1007/978-3-319-46475-6_25 (cit. on p. 42).

[76] David L. Donoho. “For Most Large Underdetermined Systems of Linear Equations the Minimal 1-Norm Solution Is Also the Sparsest Solution”. In: *Communications on Pure and Applied Mathematics* 59.6 (2006), pp. 797–829. ISSN: 1097-0312. DOI: 10.1002/cpa.20132. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.20132 (visited on 03/20/2024) (cit. on pp. 14, 15).

[77] Simon J. Doran. “Chapter 21 Quantitative Magnetic Resonance Imaging: Applications and Estimation of Errors”. In: *Data Handling in Science and Technology*. Ed. by D. N. Rutledge. Vol. 18. Signal Treatment and Signal Analysis in NMR. Elsevier, Jan. 1, 1996, pp. 452–488. DOI: 10.1016/S0922-3487(96)80058-X. URL: https://www.sciencedirect.com/science/article/pii/S092234879680058X (visited on 03/21/2024) (cit. on p. 26).

[78] Richard Dosselmann and Xue Dong Yang. “A Comprehensive Assessment of the Structural Similarity Index”. In: *Signal, Image and Video Processing* 5.1 (Mar. 1, 2011), pp. 81–91. ISSN: 1863-1711. DOI: 10.1007/s11760-009-0144-1. URL: https://doi.org/10.1007/s11760-009-0144-1 (visited on 04/17/2024) (cit. on p. 66).

[79] Arun Kumar Dubey and Vanita Jain. “Comparative Study of Convolution Neural Network’s Relu and Leaky-Relu Activation Functions”. In: *Applications of Computing, Automation and Wireless Systems in Electrical Engineering*. Ed. by Sukumar Mishra, Yog Raj Sood, and Anuradha Tomar. Lecture Notes in Electrical Engineering. Singapore: Springer, 2019, pp. 873–880. ISBN: 9789811367724. DOI: 10.1007/978-981-13-6772-4_76 (cit. on p. 105).

[80] Serge O. Dumoulin, Alessio Fracasso, Wietske van der Zwaag, et al. “Ultra-High Field MRI: Advancing Systems Neuroscience towards Mesoscopic Human Brain Function”. In: *NeuroImage*. Neuroimaging with Ultra-high Field MRI: Present and Future 168 (Mar. 1, 2018), pp. 345–357. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2017.01.028. URL: https://www.sciencedirect.com/science/article/pii/S1053811917300289 (visited on 03/21/2024) (cit. on p. 14).

[81] Bradley Efron. “Tweedie’s Formula and Selection Bias”. In: *Journal of the American Statistical Association* 106.496 (2011), pp. 1602–1614. ISSN: 0162-1459. DOI: 10.1198/jasa.2011.tm11181. pmid: 22505788 (cit. on p. 57).

[82] Taejoon Eo, Yohan Jun, Taeseong Kim, et al. “KIKI-net: Cross-Domain Convolutional Neural Networks for Reconstructing Undersampled Magnetic Resonance Images”. In: *Magnetic Resonance in Medicine* 80.5 (Nov. 2018), pp. 2188–2201. ISSN: 1522-2594. DOI: 10.1002/mrm.27201. pmid: 29624729 (cit. on p. 53).

[83] Shreyas Fadnavis, Joshua Batson, and Eleftherios Garyfallidis. “Patch2Self: Denoising Diffusion MRI with Self-Supervised Learning”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, et al. Vol. 33. Curran Associates, Inc., 2020, pp. 16293–16303. URL: https://proceedings.neurips.cc/paper/2020/file/bc047286b224b7bfa73d4cb02de1238d-Paper.pdf (cit. on p. 108).

[84] David A. Feinberg, Alexander J. S. Beckett, An T. Vu, et al. “Next-Generation MRI Scanner Designed for Ultra-High-Resolution Human Brain Imaging at 7 Tesla”. In: *Nature Methods* 20.12 (Dec. 2023), pp. 2048–2057. ISSN: 1548-7105. DOI: 10.1038/s41592-023-02068-7. URL: https://www.nature.com/articles/s41592-023-02068-7 (visited on 03/20/2024) (cit. on p. 14).

[85] Li Feng, Dan Ma, and Fang Liu. “Rapid MR Relaxometry Using Deep Learning: An Overview of Current Techniques and Emerging Trends”. In: *NMR in Biomedicine* 35.4 (2022), e4416. ISSN: 1099-1492. DOI: 10.1002/nbm.4416. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/nbm.4416 (visited on 05/02/2022) (cit. on p. 58).

[86] Ruimin Feng, Qing Wu, Yuyao Zhang, et al. “A Scan-Specific Unsupervised Method for Parallel MRI Reconstruction Via Implicit Neural Representation”. In: *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*. 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). Apr. 2023, pp. 1–5. DOI: 10.1109/ISBI53787.2023.10230813. URL: https://ieeexplore.ieee.org/abstract/document/10230813 (visited on 04/03/2024) (cit. on p. 51).

[87] Chelsea Finn, Pieter Abbeel, and Sergey Levine. *Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks*. July 18, 2017. DOI: 10.48550/arXiv.1703.03400. arXiv: 1703.03400 [cs]. URL: http://arxiv.org/abs/1703.03400 (visited on 04/16/2024). preprint (cit. on p. 110).

[88] VS Fonov, AC Evans, RC McKinstry, et al. “Unbiased Nonlinear Average Age-Appropriate Brain Templates from Birth to Adulthood”. In: *NeuroImage*. Organization for Human Brain Mapping 2009 Annual Meeting 47 (July 1, 2009), S102. ISSN: 1053-8119. DOI: 10.1016/S1053-8119(09)70884-5. URL: https://www.sciencedirect.com/science/article/pii/S1053811909708845 (visited on 04/16/2024) (cit. on p. 115).

[89] Zhiyang Fu, Sagar Mandava, Mahesh B. Keerthivasan, et al. “A Multi-Scale Residual Network for Accelerated Radial MR Parameter Mapping”. In: *Magnetic Resonance Imaging* 73 (Nov. 1, 2020), pp. 152–162. ISSN: 0730-725X. DOI: 10.1016/j.mri.2020.08.013. URL: https://www.sciencedirect.com/science/article/pii/S0730725X20301521 (visited on 04/08/2024) (cit. on p. 58).

[90] Y. Gao and S.J. Reeves. “Optimal K-Space Sampling in MRSI for Images with a Limited Region of Support”. In: *IEEE Transactions on Medical Imaging* 19.12 (Dec. 2000), pp. 1168–1178. ISSN: 1558-254X. DOI: 10.1109/42.897808. URL: https://ieeexplore.ieee.org/document/897808 (visited on 04/04/2024) (cit. on p. 54).

[91] Saurabh Garg, Sivaraman Balakrishnan, J. Zico Kolter, et al. *RATT: Leveraging Unlabeled Data to Guarantee Generalization*. Nov. 6, 2021. DOI: 10.48550/arXiv.2105.00303. arXiv: 2105.00303 [cs, stat]. URL: http://arxiv.org/abs/2105.00303 (visited on 03/23/2024). preprint (cit. on p. 32).

[92] Federico Girosi, Michael Jones, and Tomaso Poggio. “Regularization Theory and Neural Networks Architectures”. In: *Neural Computation* 7.2 (Mar. 1995), pp. 219–269. ISSN: 0899-7667. DOI: 10.1162/neco.1995.7.2.219. URL: https://ieeexplore.ieee.org/document/6796297 (visited on 03/23/2024) (cit. on p. 32).

[93] James W. Goldfarb, Jason Craft, and J. Jane Cao. “Water–Fat Separation and Parameter Mapping in Cardiac MRI via Deep Learning with a Convolutional Neural Network”. In: *Journal of Magnetic Resonance Imaging* 50.2 (2019), pp. 655–665. ISSN: 1522-2586. DOI: 10.1002/jmri.26658. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.26658 (visited on 03/30/2024) (cit. on p. 35).

[94] Fabrício Guimarães Gonçalves, Suraj D. Serai, and Giulio Zuccoli. “Synthetic Brain MRI: Review of Current Concepts and Future Directions”. In: *Topics in Magnetic Resonance Imaging* 27.6 (Dec. 2018), p. 387. ISSN: 1536-1004. DOI: 10.1097/RMR.0000000000000189. URL: https://journals.lww.com/topicsinmri/fulltext/2018/12000/synthetic_brain_mri__review_of_current_concepts.1.aspx (visited on 03/13/2024) (cit. on p. 57).

[95] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, et al. *Generative Adversarial Networks*. June 10, 2014. DOI: 10.48550/arXiv.1406.2661. arXiv: 1406.2661 [cs, stat]. URL: http://arxiv.org/abs/1406.2661 (visited on 03/23/2024). preprint (cit. on p. 36).

[96] Hans Van Gorp, Iris Huijben, Bastiaan S. Veeling, et al. “Active Deep Probabilistic Subsampling”. In: *Proceedings of the 38th International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, July 1, 2021, pp. 10509–10518. URL: https://proceedings.mlr.press/v139/van-gorp21a.html (visited on 04/04/2024) (cit. on p. 55).

[97] Niv Granot, Ben Feinstein, Assaf Shocher, et al. *Drop the GAN: In Defense of Patches Nearest Neighbors as Single Image Generative Models*. Aug. 24, 2021. DOI: 10.48550/arXiv.2103.15545. arXiv: 2103.15545 [cs]. URL: http://arxiv.org/abs/2103.15545 (visited on 04/17/2024). preprint (cit. on p. 110).

[98] Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, et al. “Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA)”. In: *Magnetic Resonance in Medicine* 47.6 (June 2002), pp. 1202–1210. ISSN: 0740-3194. DOI: 10.1002/mrm.10171. pmid: 12111967 (cit. on pp. 14, 15, 23, 46, 75, 106).

[99] Nitzan Guberman. *On Complex Valued Convolutional Neural Networks*. Feb. 29, 2016. DOI: 10.48550/arXiv.1602.09046. arXiv: 1602.09046 [cs]. URL: http://arxiv.org/abs/1602.09046 (visited on 03/31/2024). preprint (cit. on p. 36).

[100] Vikas Gulani and Nicole Seiberlich. "Quantitative MRI: Rationale and Challenges". In: *Advances in Magnetic Resonance Technology and Applications*. Ed. by Nicole Seiberlich, Vikas Gulani, Fernando Calamante, et al. Vol. 1. Quantitative Magnetic Resonance Imaging. Academic Press, Jan. 1, 2020, pp. xxxvii–li. DOI: 10.1016/B978-0-12-817057-1.00001-9. URL: https://www.sciencedirect.com/science/article/pii/B9780128170571000019 (visited on 03/21/2024) (cit. on p. 26).

[101] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, et al. "Improved Training of Wasserstein GANs". In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., Dec. 4, 2017, pp. 5769–5779. ISBN: 978-1-5108-6096-4 (cit. on p. 42).

[102] Justin P. Haldar. "Low-Rank Modeling of Local k-Space Neighborhoods (LORAKS) for Constrained MRI". In: *IEEE transactions on medical imaging* 33.3 (Mar. 2014), pp. 668–681. ISSN: 1558-254X. DOI: 10.1109/TMI.2013.2293974. pmid: 24595341 (cit. on p. 52).

[103] Justin P. Haldar and Daeun Kim. "OEDIPUS: An Experiment Design Framework for Sparsity-Constrained MRI". In: *IEEE transactions on medical imaging* 38.7 (July 2019), pp. 1545–1558. ISSN: 0278-0062. DOI: 10.1109/TMI.2019.2896180. pmid: 30716031. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6669033/ (visited on 04/04/2024) (cit. on p. 54).

[104] Justin P. Haldar and Jingwei Zhuo. "P-LORAKS: Low-rank Modeling of Local k-Space Neighborhoods with Parallel Imaging Data". In: *Magnetic Resonance in Medicine* 75.4 (Apr. 2016), pp. 1499–1514. ISSN: 1522-2594. DOI: 10.1002/mrm.25717. pmid: 25952136 (cit. on p. 52).

[105] Ahmet Mesrur Halefoglu and David Mark Yousem. "Susceptibility Weighted Imaging: Clinical Applications and Future Directions". In: *World Journal of Radiology* 10.4 (Apr. 28, 2018), pp. 30–45. ISSN: 1949-8470. DOI: 10.4329/wjr.v10.i4.30. pmid: 29849962. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5971274/ (visited on 03/21/2024) (cit. on p. 27).

[106] Jesse Hamilton, Dominique Franson, and Nicole Seiberlich. "Recent Advances in Parallel Imaging for MRI". In: *Progress in Nuclear Magnetic Resonance Spectroscopy* 101 (Aug. 1, 2017), pp. 71–95. ISSN: 0079-6565. DOI: 10.1016/j.pnmrs.2017.04.002. URL: https://www.sciencedirect.com/science/article/pii/S0079656517300031 (visited on 01/17/2024) (cit. on p. 46).

[107] Jesse I. Hamilton and Nicole Seiberlich. "Machine Learning for Rapid Magnetic Resonance Fingerprinting Tissue Property Quantification". In: *Proceedings of the IEEE* 108.1 (Jan. 2020), pp. 69–85. ISSN: 1558-2256. DOI: 10.1109/JPROC.2019.2936998. URL: https://ieeexplore.ieee.org/abstract/document/8831396 (visited on 03/20/2024) (cit. on p. 14).

[108] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, et al. "Learning a Variational Network for Reconstruction of Accelerated MRI Data". In: *Magnetic Resonance in Medicine* 79.6 (2018), pp. 3055–3071. ISSN: 1522-2594. DOI: 10.1002/mrm.26977. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26977 (visited on 10/09/2020) (cit. on pp. 35, 52).

[109] Yoseo Han, Leonard Sunwoo, and Jong Chul Ye. “K -Space Deep Learning for Accelerated MRI”. In: *IEEE Transactions on Medical Imaging* 39.2 (Feb. 2020), pp. 377–386. ISSN: 1558-254X. DOI: 10.1109/TMI.2019.2927101. URL: https://ieeexplore.ieee.org/document/8756028 (visited on 04/09/2024) (cit. on p. 52).

[110] Kazuyuki Hara, Daisuke Saito, and Hayaru Shouno. “Analysis of Function of Rectified Linear Unit Used in Deep Learning”. In: *2015 International Joint Conference on Neural Networks (IJCNN)*. 2015 International Joint Conference on Neural Networks (IJCNN). July 2015, pp. 1–8. DOI: 10.1109/IJCNN.2015.7280578 (cit. on p. 47).

[111] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. “Deep Back-Projection Networks for Super-Resolution”. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE Computer Society, June 1, 2018, pp. 1664–1673. ISBN: 978-1-5386-6420-9. DOI: 10.1109/CVPR.2018.00179. URL: https://www.computer.org/csdl/proceedings-article/cvpr/2018/642000b664/17D45WKWnHX (visited on 04/02/2024) (cit. on p. 42).

[112] Maximilian Harl, Marvin Herchenbach, Sven Kruschel, et al. *A Light in the Dark: Deep Learning Practices for Industrial Computer Vision*. Jan. 6, 2022. DOI: 10.48550/arXiv.2201.02028. arXiv: 2201.02028 [cs]. URL: http://arxiv.org/abs/2201.02028 (visited on 06/23/2024). preprint (cit. on p. 62).

[113] Inger Havsteen, Anders Ohlhues, Kristoffer H. Madsen, et al. “Are Movement Artifacts in Magnetic Resonance Imaging a Real Problem?—A Narrative Review”. In: *Frontiers in Neurology* 8 (May 30, 2017), p. 232. ISSN: 1664-2295. DOI: 10.3389/fneur.2017.00232. pmid: 28611728. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5447676/ (visited on 03/20/2024) (cit. on p. 14).

[114] Kaiming He, Xiangyu Zhang, Shaoqing Ren, et al. “Deep Residual Learning for Image Recognition”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Las Vegas, NV, USA: IEEE, June 2016, pp. 770–778. ISBN: 978-1-4673-8851-1. DOI: 10.1109/CVPR.2016.90. URL: http://ieeexplore.ieee.org/document/7780459/ (visited on 08/06/2020) (cit. on pp. 33, 34, 77).

[115] Rahel Heule, Jonas Bause, Orso Pusterla, et al. “Multi-Parametric Artificial Neural Network Fitting of Phase-Cycled Balanced Steady-State Free Precession Data”. In: *Magnetic Resonance in Medicine* 84.6 (Dec. 2020), pp. 2981–2993. ISSN: 1522-2594. DOI: 10.1002/mrm.28325. pmid: 32479661 (cit. on p. 58).

[116] Amir Heydari, Abbas Ahmadi, Tae Hyung Kim, et al. “Joint MAPLE: Accelerated Joint T1 and T2*\ \mathrmT_2^\ast \ Mapping with Scan-Specific Self-Supervised Networks”. In: *Magnetic Resonance in Medicine* 91.6 (2024), pp. 2294–2309. ISSN: 1522-2594. DOI: 10.1002/mrm.29989. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29989 (visited on 04/08/2024) (cit. on p. 59).

[117] Tom Hilbert, Tilman J. Sumpf, Elisabeth Weiland, et al. “Accelerated T2 Mapping Combining Parallel MRI and Model-Based Reconstruction: GRAPPATINI”. In: *Journal of Magnetic Resonance Imaging* 48.2 (2018), pp. 359–368. ISSN: 1522-2586. DOI: 10.1002/jmri.25972. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.25972 (visited on 04/07/2024) (cit. on p. 58).

[118] G. E. Hinton and R. R. Salakhutdinov. “Reducing the Dimensionality of Data with Neural Networks”. In: *Science* 313.5786 (July 28, 2006), pp. 504–507. DOI: 10.1126/science.1127647. URL: https://www.science.org/doi/10.1126/science.1127647 (visited on 03/23/2024) (cit. on p. 34).

[119] Jonathan Ho, Ajay Jain, and Pieter Abbeel. *Denoising Diffusion Probabilistic Models*. Dec. 16, 2020. DOI: 10.48550/arXiv.2006.11239. arXiv: 2006.11239 [cs, stat]. URL: http://arxiv.org/abs/2006.11239 (visited on 03/23/2024). preprint (cit. on p. 40).

[120] Guan Qiu Hong, Yuan Tao Wei, William A. W. Morley, et al. “Dual-Domain Accelerated MRI Reconstruction Using Transformers with Learning-Based Undersampling”. In: *Computerized Medical Imaging and Graphics* 106 (June 1, 2023), p. 102206. ISSN: 0895-6111. DOI: 10.1016/j.compmedimag.2023.102206. URL: https://www.sciencedirect.com/science/article/pii/S0895611123000241 (visited on 04/04/2024) (cit. on p. 55).

[121] Elisabeth Hoppe, Gregor Körzdörfer, Tobias Würfl, et al. “Deep Learning for Magnetic Resonance Fingerprinting: A New Approach for Predicting Quantitative Parameter Values from Time Series”. In: *Studies in Health Technology and Informatics* 243 (2017), pp. 202–206. ISSN: 1879-8365. pmid: 28883201 (cit. on p. 58).

[122] Seyed Amir Hossein Hosseini, Steen Moeller, Sebastian Weingärtner, et al. “ACCELERATED CORONARY MRI USING 3D SPIRIT-RAKI WITH SPARSITY REGULARIZATION”. In: *Proceedings. IEEE International Symposium on Biomedical Imaging* 2019 (Apr. 2019), pp. 1692–1695. ISSN: 1945-7928. DOI: 10.1109/ISBI.2019.8759459. pmid: 31893013 (cit. on p. 51).

[123] Hao Hou, Yuchen Shao, Yang Geng, et al. “PNCS: Pixel-Level Non-Local Method Based Compressed Sensing Undersampled MRI Image Reconstruction”. In: *IEEE Access* 11 (2023), pp. 42389–42402. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2023.3270900. URL: https://ieeexplore.ieee.org/document/10109723 (visited on 04/13/2024) (cit. on p. 111).

[124] Andrew Howard, Mark Sandler, Grace Chu, et al. “Searching for MobileNetV3”. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2019, pp. 1314–1324. URL: https://openaccess.thecvf.com/content_ICCV_2019/html/Howard_Searching_for_MobileNetV3_ICCV_2019_paper.html (visited on 03/23/2024) (cit. on p. 33).

[125] Chenxi Hu and Dana C. Peters. “SUPER: A Blockwise Curve-Fitting Method for Accelerating MR Parametric Mapping with Fast Reconstruction”. In: *Magnetic Resonance in Medicine* 81.6 (2019), pp. 3515–3529. ISSN: 1522-2594. DOI: 10.1002/mrm.27662. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27662 (visited on 04/07/2024) (cit. on p. 58).

[126] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, et al. “Densely Connected Convolutional Networks”. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). July 2017, pp. 2261–2269. DOI: 10.1109/CVPR.2017.243. URL: https://ieeexplore.ieee.org/document/8099726 (visited on 04/02/2024) (cit. on p. 43).

[127] Gao Huang, Zhuang Liu, Laurens van der Maaten, et al. *Densely Connected Convolutional Networks*. Jan. 28, 2018. DOI: 10.48550/arXiv.1608.06993. arXiv: 1608.06993 [cs]. URL: http://arxiv.org/abs/1608.06993 (visited on 03/23/2024). preprint (cit. on p. 33).

[128] Iris A.M. Huijben, Bastiaan S. Veeling, and Ruud J.G. van Sloun. “Learning Sampling and Model-Based Signal Recovery for Compressed Sensing MRI”. In: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). May 2020, pp. 8906–8910. DOI: 10.1109/ICASSP40776.2020.9053331. URL: https://ieeexplore.ieee.org/document/9053331 (visited on 04/04/2024) (cit. on p. 55).

[129] Chang Min Hyun, Hwa Pyung Kim, Sung Min Lee, et al. “Deep Learning for Undersampled MRI Reconstruction”. In: *Physics in Medicine & Biology* 63.13 (June 2018), p. 135007. ISSN: 0031-9155. DOI: 10.1088/1361-6560/aac71a. URL: https://dx.doi.org/10.1088/1361-6560/aac71a (visited on 03/24/2024) (cit. on p. 40).

[130] Souheil J. Inati, Joseph D. Naegele, Nicholas R. Zwart, et al. “ISMRM Raw Data Format: A Proposed Standard for MRI Raw Datasets”. In: *Magnetic Resonance in Medicine* 77.1 (2017), pp. 411–421. ISSN: 1522-2594. DOI: 10.1002/mrm.26089. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26089 (visited on 02/01/2023) (cit. on p. 81).

[131] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, et al. *Image-to-Image Translation with Conditional Adversarial Networks*. Nov. 26, 2018. DOI: 10.48550/arXiv.1611.07004. arXiv: 1611.07004 [cs]. URL: http://arxiv.org/abs/1611.07004 (visited on 03/23/2024). preprint (cit. on pp. 36, 42, 111).

[132] C. R. Jack. “Hippocampal T2 Relaxometry in Epilepsy: Past, Present, and Future”. In: *AJNR. American journal of neuroradiology* 17.10 (1996), pp. 1811–1814. ISSN: 0195-6108. pmid: 8933862 (cit. on p. 26).

[133] Ali Jahanian, Lucy Chai, and Phillip Isola. *On the "Steerability" of Generative Adversarial Networks*. Feb. 16, 2020. DOI: 10.48550/arXiv.1907.07171. arXiv: 1907.07171 [cs]. URL: http://arxiv.org/abs/1907.07171 (visited on 03/23/2024). preprint (cit. on p. 36).

[134] Geon-Ho Jahng, Ka-Loh Li, Leif Ostergaard, et al. “Perfusion Magnetic Resonance Imaging: A Comprehensive Update on Principles and Techniques”. In: *Korean Journal of Radiology* 15.5 (2014), pp. 554–577. ISSN: 1229-6929. DOI: 10.3348/kjr.2014.15.5.554. pmid: 25246817. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4170157/ (visited on 03/21/2024) (cit. on p. 26).

[135] Ajil Jalal, Marius Arvinte, Giannis Daras, et al. *Robust Compressed Sensing MRI with Deep Generative Priors*. Dec. 6, 2021. DOI: 10.48550/arXiv.2108.01368. arXiv: 2108.01368 [cs, math, stat]. URL: http://arxiv.org/abs/2108.01368 (visited on 04/26/2024). preprint (cit. on pp. 54, 124).

[136] Álvaro Barbero Jiménez, Jorge López Lázaro, and José R. Dorronsoro. “Finding Optimal Model Parameters by Discrete Grid Search”. In: *Innovations in Hybrid Intelligent Systems*. Ed. by Emilio Corchado, Juan M. Corchado, and Ajith Abraham. Advances in Soft Computing. Berlin, Heidelberg: Springer, 2007, pp. 120–127. ISBN: 978-3-540-74972-1. DOI: 10.1007/978-3-540-74972-1_17. URL: https://doi.org/10.1007/978-3-540-74972-1_17 (visited on 01/25/2023) (cit. on p. 80).

[137] Kyong Hwan Jin, Dongwook Lee, and Jong Chul Ye. “A General Framework for Compressed Sensing and Parallel MRI Using Annihilating Filter Based Low-Rank Hankel Matrix”. In: *IEEE Transactions on Computational Imaging* 2.4 (Dec. 2016), pp. 480–495. ISSN: 2333-9403. DOI: 10.1109/TCI.2016.2601296. URL: https://ieeexplore.ieee.org/document/7547372 (visited on 04/09/2024) (cit. on p. 52).

[138] Kyong Hwan Jin, Michael Unser, and Kwang Moo Yi. *Self-Supervised Deep Active Accelerated MRI*. Jan. 14, 2019. DOI: 10.48550/arXiv.1901.04547. arXiv: 1901.04547 [cs]. URL: http://arxiv.org/abs/1901.04547 (visited on 04/04/2024). preprint (cit. on p. 55).

[139] Kyong Hwan Jin and Jong Chul Ye. "Annihilating Filter-Based Low-Rank Hankel Matrix Approach for Image Inpainting". In: *IEEE Transactions on Image Processing* 24.11 (Nov. 2015), pp. 3498–3511. ISSN: 1941-0042. DOI: 10.1109/TIP.2015.2446943. URL: https://ieeexplore.ieee.org/document/7127011 (visited on 04/09/2024) (cit. on p. 52).

[140] Patricia M. Johnson, Geunu Jeong, Kerstin Hammernik, et al. "Evaluation of the Robustness of Learned MR Image Reconstruction to Systematic Deviations Between Training and Test Data for the Models from the fastMRI Challenge". In: *Machine Learning for Medical Image Reconstruction*. Ed. by Nandinee Haq, Patricia Johnson, Andreas Maier, et al. Cham: Springer International Publishing, 2021, pp. 25–34. ISBN: 978-3-030-88552-6. DOI: 10.1007/978-3-030-88552-6_3 (cit. on p. 54).

[141] Yohan Jun, Jaejin Cho, Xiaoqing Wang, et al. "SSL-QALAS: Self-Supervised Learning for Rapid Multiparameter Estimation in Quantitative MRI Using 3D-QALAS". In: *Magnetic Resonance in Medicine* 90.5 (2023), pp. 2019–2032. ISSN: 1522-2594. DOI: 10.1002/mrm.29786. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29786 (visited on 04/08/2024) (cit. on p. 59).

[142] Yohan Jun, Hyungseob Shin, Taejoon Eo, et al. "Deep Model-Based Magnetic Resonance Parameter Mapping Network (DOPAMINE) for Fast T1 Mapping Using Variable Flip Angle Method". In: *Medical Image Analysis* 70 (May 1, 2021), p. 102017. ISSN: 1361-8415. DOI: 10.1016/j.media.2021.102017. URL: https://www.sciencedirect.com/science/article/pii/S1361841521000633 (visited on 04/08/2024) (cit. on p. 59).

[143] Woojin Jung, Hyun-Soo Lee, Minkook Seo, et al. "MR-self Noise2Noise: Self-Supervised Deep Learning–Based Image Quality Improvement of Submillimeter Resolution 3D MR Images". In: *European Radiology* (Nov. 15, 2022). ISSN: 1432-1084. DOI: 10.1007/s00330-022-09243-y. URL: https://doi.org/10.1007/s00330-022-09243-y (visited on 12/12/2022) (cit. on p. 75).

[144] Hiroyuki Kabasawa. "MR Imaging in the 21st Century: Technical Innovation over the First Two Decades". In: *Magnetic Resonance in Medical Sciences* 21.1 (Apr. 16, 2021), pp. 71–82. ISSN: 1347-3182. DOI: 10.2463/mrms.rev.2021-0011. pmid: 33867419. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9199974/ (visited on 03/20/2024) (cit. on p. 13).

[145] Maya Kabkab, Pouya Samangouei, and Rama Chellappa. *Task-Aware Compressed Sensing with Generative Adversarial Networks*. Feb. 5, 2018. DOI: 10.48550/arXiv.1802.01284. arXiv: 1802.01284 [cs, stat]. URL: http://arxiv.org/abs/1802.01284 (visited on 04/04/2024). preprint (cit. on p. 56).

[146] Hasan H. Karaoglu and Ender M. Eksioglu. "A Plug-and-Play Deep Denoiser Prior Model for Accelerated MRI Reconstruction". In: *2022 45th International Conference on Telecommunications and Signal Processing (TSP)*. 2022 45th International Conference on Telecommunications and Signal Processing (TSP). July 2022, pp. 260–263. DOI: 10.1109/TSP55681.2022.9851376. URL: https://ieeexplore.ieee.org/document/9851376 (visited on 04/10/2024) (cit. on pp. 53, 111).

[147] Geethan Karunaratne, Manuel Schmuck, Manuel Le Gallo, et al. “Robust High-Dimensional Memory-Augmented Neural Networks”. In: *Nature Communications* 12.1 (Apr. 29, 2021), p. 2468. ISSN: 2041-1723. DOI: 10.1038/s41467-021-22364-0. URL: https://www.nature.com/articles/s41467-021-22364-0 (visited on 04/16/2024) (cit. on p. 110).

[148] Chan Kyo Kim, Byung Kwan Park, Jae Joon Han, et al. “Diffusion-Weighted Imaging of the Prostate at 3 T for Differentiation of Malignant and Benign Tissue in Transition and Peripheral Zones: Preliminary Results”. In: *Journal of Computer Assisted Tomography* 31.3 (2007), pp. 449–454. ISSN: 0363-8715. DOI: 10.1097/01.rct.0000243456.00437.59. pmid: 17538295 (cit. on p. 26).

[149] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. “Accurate Image Super-Resolution Using Very Deep Convolutional Networks”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). June 2016, pp. 1646–1654. DOI: 10.1109/CVPR.2016.182. URL: https://ieeexplore.ieee.org/document/7780551 (visited on 04/01/2024) (cit. on p. 42).

[150] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. “Deeply-Recursive Convolutional Network for Image Super-Resolution”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). June 2016, pp. 1637–1645. DOI: 10.1109/CVPR.2016.181. URL: https://ieeexplore.ieee.org/document/7780550/footnotes#footnotes (visited on 04/01/2024) (cit. on p. 42).

[151] Diederik P. Kingma and Max Welling. *Auto-Encoding Variational Bayes*. Dec. 10, 2022. DOI: 10.48550/arXiv.1312.6114. arXiv: 1312.6114 [cs, stat]. URL: http://arxiv.org/abs/1312.6114 (visited on 03/23/2024). preprint (cit. on pp. 36, 38).

[152] Florian Knoll, Tullie Murrell, Anuroop Sriram, et al. “Advancing Machine Learning for MR Image Reconstruction with an Open Competition: Overview of the 2019 fastMRI Challenge”. In: *Magnetic Resonance in Medicine* 84.6 (2020), pp. 3054–3070. ISSN: 1522-2594. DOI: 10.1002/mrm.28338. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28338 (visited on 12/12/2022) (cit. on pp. 45, 54, 75).

[153] Florian Knoll, Jure Zbontar, Anuroop Sriram, et al. “fastMRI: A Publicly Available Raw k-Space and DICOM Dataset of Knee Images for Accelerated MR Image Reconstruction Using Machine Learning”. In: *Radiology. Artificial intelligence* 2.1 (Jan. 29, 2020), e190007. ISSN: 2638-6100. DOI: 10.1148/ryai.2020190007. pmid: 32076662. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6996599/ (visited on 04/04/2024) (cit. on pp. 56, 75).

[154] Erich Kobler, Alexander Effland, Karl Kunisch, et al. *Total Deep Variation for Linear Inverse Problems*. Feb. 17, 2020. DOI: 10.48550/arXiv.2001.05005. arXiv: 2001.05005 [cs, math]. URL: http://arxiv.org/abs/2001.05005 (visited on 04/16/2024). preprint (cit. on p. 111).

[155] Ivan Kobyzev, Simon J. D. Prince, and Marcus A. Brubaker. “Normalizing Flows: An Introduction and Review of Current Methods”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43.11 (Nov. 1, 2021), pp. 3964–3979. ISSN: 0162-8828, 2160-9292, 1939-3539. DOI: 10.1109/TPAMI.2020.2992934. arXiv: 1908.09257 [cs, stat]. URL: http://arxiv.org/abs/1908.09257 (visited on 04/02/2024) (cit. on pp. 40, 112).

[156] Gregory R. Koch. “Siamese Neural Networks for One-Shot Image Recognition”. In: 2015. URL: https://www.semanticscholar.org/paper/Siamese-Neural-Networks-for-One-Shot-Image-Koch/f216444d4f2959b4520c61d20003fa30a199670a (visited on 04/16/2024) (cit. on p. 109).

[157] Shinya Kojima, Hiroyuki Shinohara, Takeyuki Hashimoto, et al. “Undersampling Patterns in K-Space for Compressed Sensing MRI Using Two-Dimensional Cartesian Sampling”. In: *Radiological Physics and Technology* 11.3 (Sept. 2018), pp. 303–319. ISSN: 1865-0341. DOI: 10.1007/s12194-018-0469-y. pmid: 30078080 (cit. on p. 40).

[158] Zhifeng Kong and Wei Ping. *On Fast Sampling of Diffusion Probabilistic Models*. June 23, 2021. DOI: 10.48550/arXiv.2106.00132. arXiv: 2106.00132 [cs]. URL: http://arxiv.org/abs/2106.00132 (visited on 04/02/2024). preprint (cit. on p. 40).

[159] Yilmaz Korkmaz, Tolga Cukur, and Vishal Patel. *Self-Supervised MRI Reconstruction with Unrolled Diffusion Models*. June 28, 2023. DOI: 10.48550/arXiv.2306.16654. arXiv: 2306.16654 [cs, eess]. URL: http://arxiv.org/abs/2306.16654 (visited on 04/08/2024). preprint (cit. on p. 53).

[160] Yilmaz Korkmaz, Salman U. H. Dar, Mahmut Yurt, et al. “Unsupervised MRI Reconstruction via Zero-Shot Learned Adversarial Transformers”. In: *IEEE Transactions on Medical Imaging* 41.7 (July 2022), pp. 1747–1763. ISSN: 1558-254X. DOI: 10.1109/TMI.2022.3147426. URL: https://ieeexplore.ieee.org/abstract/document/9695412 (visited on 04/03/2024) (cit. on p. 53).

[161] Zoran Kotevski and Pece Mitrevski. “Experimental Comparison of PSNR and SSIM Metrics for Video Quality Estimation”. In: *ICT Innovations 2009*. Ed. by Danco Davcev and Jorge Marx Gómez. Berlin, Heidelberg: Springer, 2010, pp. 357–366. ISBN: 978-3-642-10781-8. DOI: 10.1007/978-3-642-10781-8_37 (cit. on p. 65).

[162] Felix Kreuk, Gabriel Synnaeve, Adam Polyak, et al. *AudioGen: Textually Guided Audio Generation*. Mar. 5, 2023. DOI: 10.48550/arXiv.2209.15352. arXiv: 2209.15352 [cs, eess]. URL: http://arxiv.org/abs/2209.15352 (visited on 03/23/2024). preprint (cit. on p. 36).

[163] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. “ImageNet Classification with Deep Convolutional Neural Networks”. In: *Advances in Neural Information Processing Systems*. Vol. 25. Curran Associates, Inc., 2012. URL: https://papers.nips.cc/paper_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html (visited on 03/23/2024) (cit. on p. 33).

[164] Sergei Krylov and Vladimir Krylov. “Inverse Problem Solving Approach Using Deep Network Trained by GAN Simulated Data”. In: *Journal of Physics: Conference Series* 1727.1 (Jan. 2021), p. 012002. ISSN: 1742-6596. DOI: 10.1088/1742-6596/1727/1/012002. URL: https://dx.doi.org/10.1088/1742-6596/1727/1/012002 (visited on 04/04/2024) (cit. on p. 56).

[165] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, et al. “Deep Laplacian Pyramid Networks for Fast and Accurate Super-Resolution”. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). July 2017, pp. 5835–5843. DOI: 10.1109/CVPR.2017.618. URL: https://ieeexplore.ieee.org/document/8100101 (visited on 04/02/2024) (cit. on p. 42).

[166] Christian Langkammer, Ferdinand Schweser, Nikolaus Krebs, et al. “Quantitative Susceptibility Mapping (QSM) as a Means to Measure Brain Iron? A Post Mortem Validation Study”. In: *NeuroImage* 62.3 (Sept. 2012), pp. 1593–1599. ISSN: 1095-9572. DOI: 10.1016/j.neuroimage.2012.05.049. pmid: 22634862 (cit. on p. 26).

[167] P. C. Lauterbur. “Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance”. In: *Nature* 242.5394 (Mar. 1973), pp. 190–191. ISSN: 1476-4687. DOI: 10.1038/242190a0. URL: https://www.nature.com/articles/242190a0 (visited on 03/20/2024) (cit. on p. 13).

[168] Y. Lecun, L. Bottou, Y. Bengio, et al. “Gradient-Based Learning Applied to Document Recognition”. In: *Proceedings of the IEEE* 86.11 (Nov. 1998), pp. 2278–2324. ISSN: 1558-2256. DOI: 10.1109/5.726791. URL: https://ieeexplore.ieee.org/document/726791 (visited on 03/23/2024) (cit. on p. 33).

[169] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep Learning”. In: *Nature* 521.7553 (May 2015), pp. 436–444. ISSN: 1476-4687. DOI: 10.1038/nature14539. URL: https://www.nature.com/articles/nature14539 (visited on 01/27/2020) (cit. on p. 29).

[170] Christian Ledig, Lucas Theis, Ferenc Huszar, et al. “Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network”. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Honolulu, HI: IEEE, July 2017, pp. 105–114. ISBN: 978-1-5386-0457-1. DOI: 10.1109/CVPR.2017.19. URL: http://ieeexplore.ieee.org/document/8099502/ (visited on 04/08/2022) (cit. on pp. 42, 43, 76, 113, 118).

[171] Haoying Li, Yifan Yang, Meng Chang, et al. “SRDiff: Single Image Super-Resolution with Diffusion Probabilistic Models”. In: *Neurocomputing* 479 (Mar. 28, 2022), pp. 47–59. ISSN: 0925-2312. DOI: 10.1016/j.neucom.2022.01.029. URL: https://www.sciencedirect.com/science/article/pii/S0925231222000522 (visited on 04/02/2024) (cit. on p. 43).

[172] Housen Li, Johannes Schwab, Stephan Antholzer, et al. *NETT: Solving Inverse Problems with Deep Neural Networks*. Dec. 8, 2019. DOI: 10.48550/arXiv.1803.00092. arXiv: 1803.00092 [cs, math]. URL: http://arxiv.org/abs/1803.00092 (visited on 04/16/2024). preprint (cit. on p. 111).

[173] Weiyu Li, Xuelin Chen, Jue Wang, et al. “Patch-Based 3D Natural Scene Generation From a Single Example”. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2023, pp. 16762–16772. URL: https://openaccess.thecvf.com/content/CVPR2023/html/Li_Patch-Based_3D_Natural_Scene_Generation_From_a_Single_Example_CVPR_2023_paper.html (visited on 04/13/2024) (cit. on p. 110).

[174] Yuze Li, Yajie Wang, Haikun Qi, et al. “Deep Learning–Enhanced T1 Mapping with Spatial-Temporal and Physical Constraint”. In: *Magnetic Resonance in Medicine* 86.3 (2021), pp. 1647–1661. ISSN: 1522-2594. DOI: 10.1002/mrm.28793. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28793 (visited on 04/08/2024) (cit. on p. 57).

[175] Bee Lim, Sanghyun Son, Heewon Kim, et al. *Enhanced Deep Residual Networks for Single Image Super-Resolution*. July 10, 2017. DOI: 10.48550/arXiv.1707.02921. arXiv: 1707.02921 [cs]. URL: http://arxiv.org/abs/1707.02921 (visited on 04/02/2024). preprint (cit. on pp. 42, 43).

[176] Kang Lin and Reinhard Heckel. “Vision Transformers Enable Fast and Robust Accelerated MRI”. In: *Proceedings of The 5th International Conference on Medical Imaging with Deep Learning*. International Conference on Medical Imaging with Deep Learning. PMLR, Dec. 4, 2022, pp. 774–795. URL: https://proceedings.mlr.press/v172/lin22a.html (visited on 04/03/2024) (cit. on p. 53).

[177] Min Lin, Qiang Chen, and Shuicheng Yan. *Network In Network*. Mar. 4, 2014. DOI: 10.48550/arXiv.1312.4400. arXiv: 1312.4400 [cs]. URL: http://arxiv.org/abs/1312.4400 (visited on 03/23/2024). preprint (cit. on p. 33).

[178] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, et al. "A Survey on Deep Learning in Medical Image Analysis". In: *Medical Image Analysis* 42 (Dec. 1, 2017), pp. 60–88. ISSN: 1361-8415. DOI: 10.1016/j.media.2017.07.005. URL: https://www.sciencedirect.com/science/article/pii/S1361841517301135 (visited on 03/23/2024) (cit. on pp. 29, 40).

[179] Chang Liu, Xi Wu, Xi Yu, et al. "Fusing Multi-Scale Information in Convolution Network for MR Image Super-Resolution Reconstruction". In: *Biomedical Engineering Online* 17.1 (Aug. 25, 2018), p. 114. ISSN: 1475-925X. DOI: 10.1186/s12938-018-0546-9. pmid: 30144798 (cit. on p. 43).

[180] Chunlei Liu, Wei Li, Karen A. Tong, et al. "Susceptibility-Weighted Imaging and Quantitative Susceptibility Mapping in the Brain". In: *Journal of magnetic resonance imaging: JMRI* 42.1 (July 2015), pp. 23–41. ISSN: 1522-2586. DOI: 10.1002/jmri.24768. pmid: 25270052 (cit. on p. 25).

[181] Fang Liu, Li Feng, and Richard Kijowski. "MANTIS: Model-Augmented Neural neTwork with Incoherent k-Space Sampling for Efficient MR Parameter Mapping". In: *Magnetic Resonance in Medicine* 82.1 (July 2019), pp. 174–188. ISSN: 1522-2594. DOI: 10.1002/mrm.27707. pmid: 30860285 (cit. on p. 58).

[182] Fang Liu, Richard Kijowski, Georges El Fakhri, et al. "Magnetic Resonance Parameter Mapping Using Model-Guided Self-Supervised Deep Learning". In: *Magnetic Resonance in Medicine* 85.6 (2021), pp. 3211–3226. ISSN: 1522-2594. DOI: 10.1002/mrm.28659. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28659 (visited on 04/08/2024) (cit. on pp. 57, 59).

[183] Fang Liu, Richard Kijowski, Li Feng, et al. "High-Performance Rapid MR Parameter Mapping Using Model-Based Deep Adversarial Learning". In: *Magnetic Resonance Imaging* 74 (Dec. 1, 2020), pp. 152–160. ISSN: 0730-725X. DOI: 10.1016/j.mri.2020.09.021. URL: https://www.sciencedirect.com/science/article/pii/S0730725X2030477X (visited on 04/08/2024) (cit. on p. 58).

[184] Jing Liu, Tian Liu, Ludovic de Rochefort, et al. "Morphology Enabled Dipole Inversion for Quantitative Susceptibility Mapping Using Structural Consistency Between the Magnitude Image and the Susceptibility Map". In: *Neuroimage* 59.3 (Feb. 1, 2012), pp. 2560–2568. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2011.08.082. pmid: 21925276. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3254812/ (visited on 11/08/2022) (cit. on p. 27).

[185] Shufang Liu, Aurelien Bustin, Pauline Ferry, et al. "A Vectorized Levenberg-Marquardt Model Fitting Algorithm for Efficient Post-Processing of Cardiac T1 Mapping MRI". In: *Computers in Biology and Medicine* 96 (May 1, 2018), pp. 106–115. ISSN: 0010-4825. DOI: 10.1016/j.compbiomed.2018.03.009. URL: https://www.sciencedirect.com/science/article/pii/S001048251830060X (visited on 04/05/2024) (cit. on p. 57).

[186] Xianzhe Liu, Hongwei Du, Jinzhang Xu, et al. “DBGAN: A Dual-Branch Generative Adversarial Network for Undersampled MRI Reconstruction”. In: *Magnetic Resonance Imaging* 89 (June 1, 2022), pp. 77–91. ISSN: 0730-725X. DOI: 10.1016/j.mri.2022.03.003. URL: https://www.sciencedirect.com/science/article/pii/S0730725X2200042X (visited on 07/26/2022) (cit. on p. 75).

[187] Xiaohan Liu, Yanwei Pang, Ruiqi Jin, et al. “Dual-Domain Reconstruction Network with V-Net and K-Net for Fast MRI”. In: *Magnetic Resonance in Medicine* 88.6 (Dec. 2022), pp. 2694–2708. ISSN: 1522-2594. DOI: 10.1002/mrm.29400. pmid: 35942977 (cit. on p. 53).

[188] Li-xiong Liu and Yuan-quan Wang. “A Mean-Edge Structural Similarity for Image Quality Assessment”. In: *2009 Sixth International Conference on Fuzzy Systems and Knowledge Discovery*. 2009 Sixth International Conference on Fuzzy Systems and Knowledge Discovery. Vol. 5. Aug. 2009, pp. 311–315. DOI: 10.1109/FSKD.2009.610. URL: https://ieeexplore.ieee.org/document/5360608 (visited on 06/23/2024) (cit. on p. 61).

[189] Yiming Liu, Yanwei Pang, Ruiqi Jin, et al. “Reinforcement Learning and Transformer for Fast Magnetic Resonance Imaging Scan”. In: *IEEE Transactions on Emerging Topics in Computational Intelligence* (2024), pp. 1–14. ISSN: 2471-285X. DOI: 10.1109/TETCI.2024.3358180. URL: https://ieeexplore.ieee.org/abstract/document/10474060 (visited on 04/04/2024) (cit. on p. 56).

[190] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, et al. *A ConvNet for the 2020s*. Mar. 2, 2022. DOI: 10.48550/arXiv.2201.03545. arXiv: 2201.03545 [cs]. URL: http://arxiv.org/abs/2201.03545 (visited on 03/23/2024). preprint (cit. on p. 34).

[191] Cheng Lu, Yuhao Zhou, Fan Bao, et al. “DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps”. In: *Advances in Neural Information Processing Systems* 35 (Dec. 6, 2022), pp. 5775–5787. URL: https://proceedings.neurips.cc/paper_files/paper/2022/hash/260a14acce2a89dad36adc8eefe7c59e-Abstract-Conference.html (visited on 04/02/2024) (cit. on p. 40).

[192] Sebastian Lunz, Ozan Öktem, and Carola-Bibiane Schönlieb. “Adversarial Regularizers in Inverse Problems”. In: *Advances in Neural Information Processing Systems*. Vol. 31. Curran Associates, Inc., 2018. URL: https://proceedings.neurips.cc/paper_files/paper/2018/hash/d903e9608cfbf08910611e4346a0ba44-Abstract.html (visited on 03/22/2024) (cit. on pp. 111, 112).

[193] Michael Lustig and John M. Pauly. “SPIRiT: Iterative Self-Consistent Parallel Imaging Reconstruction from Arbitrary k-Space”. In: *Magnetic Resonance in Medicine* 64.2 (2010), pp. 457–471. ISSN: 1522-2594. DOI: 10.1002/mrm.22428. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.22428 (visited on 10/11/2021) (cit. on pp. 23, 50, 75).

[194] Jun Lv, Guangyuan Li, Xiangrong Tong, et al. “Transfer Learning Enhanced Generative Adversarial Networks for Multi-Channel MRI Reconstruction”. In: *Computers in Biology and Medicine* 134 (July 1, 2021), p. 104504. ISSN: 0010-4825. DOI: 10.1016/j.compbiomed.2021.104504. URL: https://www.sciencedirect.com/science/article/pii/S0010482521002985 (visited on 07/26/2022) (cit. on p. 75).

[195] Jun Lv, Chengyan Wang, and Guang Yang. “PIC-GAN: A Parallel Imaging Coupled Generative Adversarial Network for Accelerated Multi-Channel MRI Reconstruction”. In: *Diagnostics* 11.1 (1 Jan. 2021), p. 61. ISSN: 2075-4418. DOI: 10.3390/diagnostics11010061. URL: https://www.mdpi.com/2075-4418/11/1/61 (visited on 07/26/2022) (cit. on p. 75).

[196] Dan Ma, Vikas Gulani, Nicole Seiberlich, et al. “Magnetic Resonance Fingerprinting”. In: *Nature* 495.7440 (7440 Mar. 2013), pp. 187–192. ISSN: 1476-4687. DOI: 10.1038/nature11971. URL: https://www.nature.com/articles/nature11971 (visited on 11/17/2021) (cit. on pp. 14, 15, 57, 58).

[197] Ya-Jun Ma, Wentao Liu, Xuna Zhao, et al. “Improved Adaptive Reconstruction of Multichannel MR Images”. In: *Medical Physics* 42.2 (2015), pp. 637–644. ISSN: 2473-4209. DOI: 10.1118/1.4905163. URL: https://onlinelibrary.wiley.com/doi/abs/10.1118/1.4905163 (visited on 01/23/2023) (cit. on p. 46).

[198] Advait Maduskar, Aniket Ladukar, Shubhankar Gore, et al. “Music Generation Using Deep Generative Modelling”. In: *2020 International Conference on Convergence to Digital World - Quo Vadis (ICCDW)*. 2020 International Conference on Convergence to Digital World - Quo Vadis (ICCDW). Feb. 2020, pp. 1–4. DOI: 10.1109/ICCDW45521.2020.9318683. URL: https://ieeexplore.ieee.org/abstract/document/9318683 (visited on 03/23/2024) (cit. on p. 36).

[199] “Magnetic Field Inhomogeneity Effects and T2* Dephasing”. In: *Magnetic Resonance Imaging*. John Wiley & Sons, Ltd, 2014, pp. 569–617. ISBN: 978-1-118-63395-3. DOI: 10.1002/9781118633953.ch20. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118633953.ch20 (visited on 03/21/2024) (cit. on p. 25).

[200] Mahadevappa Mahesh. “The Essential Physics of Medical Imaging, Third Edition.” In: *Medical Physics* 40.7 (2013), p. 077301. ISSN: 2473-4209. DOI: 10.1118/1.4811156. URL: https://onlinelibrary.wiley.com/doi/abs/10.1118/1.4811156 (visited on 03/21/2024) (cit. on p. 24).

[201] Nancy M. Major, Mark W. Anderson, Clyde A. Helms, et al. “1 - Basic Principles of Musculoskeletal MRI”. In: *Musculoskeletal MRI (Third Edition)*. Ed. by Nancy M. Major, Mark W. Anderson, Clyde A. Helms, et al. St. Louis (MO): Elsevier, Jan. 1, 2020, pp. 1–22. ISBN: 978-0-323-41560-6. DOI: 10.1016/B978-0-323-41560-6.00001-9. URL: https://www.sciencedirect.com/science/article/pii/B9780323415606000019 (visited on 03/21/2024) (cit. on p. 24).

[202] José V. Manjón, Pierrick Coupé, Antonio Buades, et al. “Non-Local MRI Upsampling”. In: *Medical Image Analysis* 14.6 (Dec. 1, 2010), pp. 784–792. ISSN: 1361-8415. DOI: 10.1016/j.media.2010.05.010. URL: https://www.sciencedirect.com/science/article/pii/S1361841510000630 (visited on 04/02/2024) (cit. on p. 43).

[203] P. Mansfield. “Multi-Planar Image Formation Using NMR Spin Echoes”. In: *Journal of Physics C: Solid State Physics* 10.3 (Feb. 1977), p. L55. ISSN: 0022-3719. DOI: 10.1088/0022-3719/10/3/004. URL: https://dx.doi.org/10.1088/0022-3719/10/3/004 (visited on 03/20/2024) (cit. on p. 13).

[204] A. J. Marchand, E. Hitti, F. Monge, et al. “MRI Quantification of Diffusion and Perfusion in Bone Marrow by Intravoxel Incoherent Motion (IVIM) and Non-Negative Least Square (NNLS) Analysis”. In: *Magnetic Resonance Imaging* 32.9 (Nov. 2014), pp. 1091–1096. ISSN: 1873-5894. DOI: 10.1016/j.mri.2014.07.009. pmid: 25093628 (cit. on p. 57).

[205] M. Markl, S. Schnell, C. Wu, et al. “Advanced Flow MRI: Emerging Techniques and Applications”. In: *Clinical radiology* 71.8 (Aug. 2016), pp. 779–795. ISSN: 0009-9260. DOI: 10.1016/j.crad.2016.01.011. pmid: 26944696. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4930408/ (visited on 03/21/2024) (cit. on p. 26).

[206] Michael Markl, Alex Frydrychowicz, Sebastian Kozerke, et al. “4D Flow MRI”. In: *Journal of Magnetic Resonance Imaging* 36.5 (2012), pp. 1015–1036. ISSN: 1522-2586. DOI: 10.1002/jmri.23632. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.23632 (visited on 03/30/2024) (cit. on p. 35).

[207] Kenneth Marro, Randolph Otto, Orpheus Kolokythas, et al. “A Simulation-Based Comparison of Two Methods for Determining Relaxation Rates from Relaxometry Images”. In: *Magnetic Resonance Imaging* 29.4 (May 2011), pp. 497–506. ISSN: 1873-5894. DOI: 10.1016/j.mri.2010.11.005. pmid: 21333480 (cit. on p. 58).

[208] Allister Mason, James Rioux, Sharon E. Clarke, et al. “Comparison of Objective Image Quality Metrics to Expert Radiologists’ Scoring of Diagnostic Quality of MR Images”. In: *IEEE Transactions on Medical Imaging* 39.4 (Apr. 2020), pp. 1064–1072. ISSN: 1558-254X. DOI: 10.1109/TMI.2019.2930338 (cit. on p. 76).

[209] Julian McGinnis, Suprosanna Shit, Hongwei Bran Li, et al. “Single-Subject Multi-contrast MRI Super-resolution via Implicit Neural Representations”. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. Ed. by Hayit Greenspan, Anant Madabhushi, Parvin Mousavi, et al. Cham: Springer Nature Switzerland, 2023, pp. 173–183. ISBN: 978-3-031-43993-3. DOI: 10.1007/978-3-031-43993-3_17 (cit. on p. 125).

[210] Debra F. McGivney, Eric Pierre, Dan Ma, et al. “SVD Compression for Magnetic Resonance Fingerprinting in the Time Domain”. In: *IEEE transactions on medical imaging* 33.12 (Dec. 2014), pp. 2311–2322. ISSN: 1558-254X. DOI: 10.1109/TMI.2014.2337321. pmid: 25029380 (cit. on p. 58).

[211] Charles A. McKenzie, Michael A. Ohliger, Ernest N. Yeh, et al. “Coil-by-Coil Image Reconstruction with SMASH”. In: *Magnetic Resonance in Medicine* 46.3 (2001), pp. 619–623. ISSN: 1522-2594. DOI: 10.1002/mrm.1236. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1236 (visited on 11/21/2023) (cit. on p. 106).

[212] Ziyu Meng, Rong Guo, Yudu Li, et al. “Accelerating T2 Mapping of the Brain by Integrating Deep Learning Priors with Low-Rank and Sparse Modeling”. In: *Magnetic Resonance in Medicine* 85.3 (2021), pp. 1455–1467. ISSN: 1522-2594. DOI: 10.1002/mrm.28526. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28526 (visited on 04/08/2024) (cit. on p. 57).

[213] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, et al. *NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis*. Aug. 3, 2020. DOI: 10.48550/arXiv.2003.08934. arXiv: 2003.08934 [cs]. URL: http://arxiv.org/abs/2003.08934 (visited on 04/03/2024). preprint (cit. on p. 51).

[214] S. Millman, P. Kusch, and I. I. Rabi. “The Nuclear Magnetic Moment of ${\mathrm{N}}^{14}$”. In: *Physical Review* 54.11 (Dec. 1, 1938), pp. 968–968. DOI: 10.1103/PhysRev.54.968. URL: https://link.aps.org/doi/10.1103/PhysRev.54.968 (visited on 02/12/2024) (cit. on p. 13).

[215] Marvin Minsky and Seymour A. Papert. *Perceptrons: An Introduction to Computational Geometry*. The MIT Press, Sept. 22, 2017. ISBN: 978-0-262-34393-0. DOI: 10.7551/mitpress/11301.001.0001. URL: https://direct.mit.edu/books/monograph/3132/PerceptronsAn-Introduction-to-Computational (visited on 03/23/2024) (cit. on p. 29).

[216] D G Mitchell, M Crovello, T Matteucci, et al. “Benign Adrenocortical Masses: Diagnosis with Chemical Shift MR Imaging.” In: *Radiology* 185.2 (Nov. 1992), pp. 345–351. ISSN: 0033-8419. DOI: 10.1148/radiology.185.2.1410337. URL: https://pubs.rsna.org/doi/abs/10.1148/radiology.185.2.1410337 (visited on 01/13/2023) (cit. on p. 75).

[217] Matthew J. Muckley, Bruno Riemenschneider, Alireza Radmanesh, et al. “Results of the 2020 fastMRI Challenge for Machine Learning MR Image Reconstruction”. In: *IEEE transactions on medical imaging* 40.9 (Sept. 2021), pp. 2306–2317. ISSN: 0278-0062. DOI: 10.1109/TMI.2021.3075856. pmid: 33929957. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8428775/ (visited on 04/10/2022) (cit. on pp. 35, 45, 52, 62, 75).

[218] Subhadip Mukherjee, Marcello Carioni, Ozan Öktem, et al. *End-to-End Reconstruction Meets Data-Driven Regularization for Inverse Problems*. June 7, 2021. DOI: 10.48550/arXiv.2106.03538. arXiv: 2106.03538 [cs, math]. URL: http://arxiv.org/abs/2106.03538 (visited on 04/16/2024). preprint (cit. on p. 112).

[219] Gustav Müller-Franzes, Teresa Nolte, Malin Ciba, et al. “Fast, Accurate, and Robust T2 Mapping of Articular Cartilage by Neural Networks”. In: *Diagnostics* 12.3 (3 Mar. 2022), p. 688. ISSN: 2075-4418. DOI: 10.3390/diagnostics12030688. URL: https://www.mdpi.com/2075-4418/12/3/688 (visited on 04/08/2024) (cit. on p. 59).

[220] Jim Nilsson and Tomas Akenine-Möller. *Understanding SSIM*. June 29, 2020. DOI: 10.48550/arXiv.2006.13846. arXiv: 2006.13846 [cs, eess]. URL: http://arxiv.org/abs/2006.13846 (visited on 04/16/2024). preprint (cit. on pp. 61, 64).

[221] Marco Nittscher, Michael Falk Lameter, Riccardo Barbano, et al. “SVD-DIP: Overcoming the Overfitting Problem in DIP-based CT Reconstruction”. In: Medical Imaging with Deep Learning. Apr. 4, 2023. URL: https://openreview.net/forum?id=ivC7VP2mof (visited on 04/03/2024) (cit. on p. 51).

[222] Tomoko Okuda, Yukunori Korogi, Yoshinori Shigematsu, et al. “Brain Lesions: When Should Fluid-attenuated Inversion-Recovery Sequences Be Used in MR Evaluation?” In: *Radiology* 212.3 (Sept. 1999), pp. 793–798. ISSN: 0033-8419. DOI: 10.1148/radiology.212.3.r99se07793. URL: https://pubs.rsna.org/doi/abs/10.1148/radiology.212.3.r99se07793 (visited on 03/21/2024) (cit. on p. 25).

[223] Benoît Oriol and Alexandre Miot. *On Some Theoretical Limitations of Generative Adversarial Networks*. Oct. 21, 2021. DOI: 10.48550/arXiv.2110.10915. arXiv: 2110.10915 [cs]. URL: http://arxiv.org/abs/2110.10915 (visited on 04/02/2024). preprint (cit. on p. 40).

[224] Samet Oymak, Mingchen Li, and Mahdi Soltanolkotabi. “Generalization Guarantees for Neural Architecture Search with Train-Validation Split”. In: *Proceedings of the 38th International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, July 1, 2021, pp. 8291–8301. URL: https://proceedings.mlr.press/v139/oymak21a.html (visited on 01/25/2023) (cit. on p. 106).

[225] Christopher C. Paige and Michael A. Saunders. “LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares”. In: *ACM Transactions on Mathematical Software* 8.1 (Mar. 1, 1982), pp. 43–71. ISSN: 0098-3500. DOI: 10.1145/355984.355989. URL: https://dl.acm.org/doi/10.1145/355984.355989 (visited on 04/09/2024) (cit. on p. 51).

[226] Myeongsuk Pak and Sanghoon Kim. “A Review of Deep Learning in Image Recognition”. In: *2017 4th International Conference on Computer Applications and Information Processing Technology (CAIPT)*. 2017 4th International Conference on Computer Applications and Information Processing Technology (CAIPT). Aug. 2017, pp. 1–3. DOI: 10.1109/CAIPT.2017.8320684. URL: https://ieeexplore.ieee.org/abstract/document/8320684 (visited on 03/23/2024) (cit. on p. 29).

[227] Jean-François Pambrun and Rita Noumeir. “Limitations of the SSIM Quality Metric in the Context of Diagnostic Imaging”. In: *2015 IEEE International Conference on Image Processing (ICIP)*. 2015 IEEE International Conference on Image Processing (ICIP). Sept. 2015, pp. 2960–2963. DOI: 10.1109/ICIP.2015.7351345. URL: https://ieeexplore.ieee.org/document/7351345 (visited on 04/17/2024) (cit. on p. 66).

[228] Paul M. Parizel, Luc van den Hauwe, Frank De Belder, et al. “Magnetic Resonance Imaging of the Brain”. In: *Clinical MR Imaging*. Ed. by Peter Reimer, Paul M. Parizel, James F. M. Meaney, et al. Berlin, Heidelberg: Springer, 2010, pp. 107–195. ISBN: 978-3-540-74504-4. DOI: 10.1007/978-3-540-74504-4_2. URL: https://doi.org/10.1007/978-3-540-74504-4_2 (visited on 03/20/2024) (cit. on p. 15).

[229] Hong Peng, Chen Jiang, Jing Cheng, et al. “One-Shot Generative Prior in Hankel-k-Space for Parallel Imaging Reconstruction”. In: *IEEE Transactions on Medical Imaging* 42.11 (Nov. 2023), pp. 3420–3435. ISSN: 1558-254X. DOI: 10.1109/TMI.2023.3288219. URL: https://ieeexplore.ieee.org/document/10158730 (visited on 04/09/2024) (cit. on p. 52).

[230] Xi Peng, Leslie Ying, Yuanyuan Liu, et al. “Accelerated Exponential Parameterization of T2 Relaxation with Model-Driven Low Rank and Sparsity Priors (MORASA)”. In: *Magnetic Resonance in Medicine* 76.6 (2016), pp. 1865–1878. ISSN: 1522-2594. DOI: 10.1002/mrm.26083. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26083 (visited on 04/08/2024) (cit. on p. 57).

[231] Chi-Hieu Pham, Aurélien Ducournau, Ronan Fablet, et al. “Brain MRI Super-Resolution Using Deep 3D Convolutional Networks”. In: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*. 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017). Apr. 2017, pp. 197–200. DOI: 10.1109/ISBI.2017.7950500. URL: https://ieeexplore.ieee.org/document/7950500 (visited on 04/02/2024) (cit. on p. 43).

[232] Luis Pineda, Sumana Basu, Adriana Romero, et al. “Active MR K-Space Sampling with Reinforcement Learning”. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. Ed. by Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, et al. Cham: Springer International Publishing, 2020, pp. 23–33. ISBN: 978-3-030-59713-9. DOI: 10.1007/978-3-030-59713-9_3 (cit. on pp. 55, 56).

[233] Esben Plenge, Dirk H. J. Poot, Monique Bernsen, et al. “Super-Resolution Methods in MRI: Can They Improve the Trade-off between Resolution, Signal-to-Noise Ratio, and Acquisition Time?” In: *Magnetic Resonance in Medicine* 68.6 (2012), pp. 1983–1993. ISSN: 1522-2594. DOI: 10.1002/mrm.24187. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24187 (visited on 03/24/2024) (cit. on p. 41).

[234] R. Pohmann, M. von Kienlin, and A. Haase. “Theoretical Evaluation and Comparison of Fast Chemical Shift Imaging Methods”. In: *Journal of Magnetic Resonance (San Diego, Calif.: 1997)* 129.2 (Dec. 1997), pp. 145–160. ISSN: 1090-7807. DOI: 10.1006/jmre.1997.1245. pmid: 9441879 (cit. on p. 35).

[235] Daniel Polak, Stephen Cauley, Berkin Bilgic, et al. “Joint Multi-Contrast Variational Network Reconstruction (jVN) with Application to Rapid 2D and 3D Imaging”. In: *Magnetic Resonance in Medicine* 84.3 (2020), pp. 1456–1469. ISSN: 1522-2594. DOI: 10.1002/mrm.28219. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28219 (visited on 10/01/2020) (cit. on p. 75).

[236] Dirk H. J. Poot and Stefan Klein. “Detecting Statistically Significant Differences in Quantitative MRI Experiments, Applied to Diffusion Tensor Imaging”. In: *IEEE transactions on medical imaging* 34.5 (May 2015), pp. 1164–1176. ISSN: 1558-254X. DOI: 10.1109/TMI.2014.2380830. pmid: 25532168 (cit. on p. 58).

[237] W. G. Proctor and F. C. Yu. “The Dependence of a Nuclear Magnetic Resonance Frequency upon Chemical Compound”. In: *Physical Review* 77.5 (Mar. 1, 1950), pp. 717–717. DOI: 10.1103/PhysRev.77.717. URL: https://link.aps.org/doi/10.1103/PhysRev.77.717 (visited on 03/20/2024) (cit. on p. 13).

[238] Jean Prost, Antoine Houdard, Andrés Almansa, et al. *Learning Local Regularization for Variational Image Restoration*. Feb. 11, 2021. DOI: 10.48550/arXiv.2102.06155. arXiv: 2102.06155 [cs, eess]. URL: http://arxiv.org/abs/2102.06155 (visited on 04/17/2024). preprint (cit. on p. 112).

[239] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, et al. “SENSE: Sensitivity Encoding for Fast MRI”. In: *Magnetic Resonance in Medicine* 42.5 (Nov. 1999), pp. 952–962. ISSN: 0740-3194. pmid: 10542355 (cit. on pp. 14, 15, 23).

[240] *Pulse Methods in 1D & 2D Liquid-Phase NMR - 1st Edition | Elsevier Shop*. URL: https://shop.elsevier.com/books/pulse-methods-in-1d-and-2d-liquid-phase-nmr/brey/978-0-08-091703-0 (visited on 03/21/2024) (cit. on p. 26).

[241] E. M. Purcell, H. C. Torrey, and R. V. Pound. “Resonance Absorption by Nuclear Magnetic Moments in a Solid”. In: *Physical Review* 69.1-2 (Jan. 1, 1946), pp. 37–38. DOI: 10.1103/PhysRev.69.37. URL: https://link.aps.org/doi/10.1103/PhysRev.69.37 (visited on 02/12/2024) (cit. on p. 13).

[242] Chen Qin, Jo Schlemper, Jose Caballero, et al. “Convolutional Recurrent Neural Networks for Dynamic MR Image Reconstruction”. In: *IEEE Transactions on Medical Imaging* 38.1 (Jan. 2019), pp. 280–290. ISSN: 1558-254X. DOI: 10.1109/TMI.2018.2863670. URL: https://ieeexplore.ieee.org/document/8425639 (visited on 04/04/2024) (cit. on p. 52).

[243] Shihan Qiu, Yuhua Chen, Sen Ma, et al. “Multiparametric Mapping in the Brain from Conventional Contrast-Weighted Images Using Deep Learning”. In: *Magnetic Resonance in Medicine* 87.1 (2022), pp. 488–495. ISSN: 1522-2594. DOI: 10.1002/mrm.28962. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28962 (visited on 04/08/2024) (cit. on p. 57).

[244] Peng Qu, Gary X. Shen, Chunsheng Wang, et al. “Tailored Utilization of Acquired K-Space Points for GRAPPA Reconstruction”. In: *Journal of Magnetic Resonance (San Diego, Calif.: 1997)* 174.1 (May 2005), pp. 60–67. ISSN: 1090-7807. DOI: 10.1016/j.jmr.2005.01.015. pmid: 15809173 (cit. on p. 47).

[245] Xiaobo Qu, Di Guo, Bende Ning, et al. “Undersampled MRI Reconstruction with Patch-Based Directional Wavelets”. In: *Magnetic Resonance Imaging* 30.7 (Sept. 1, 2012), pp. 964–977. ISSN: 0730-725X. DOI: 10.1016/j.mri.2012.02.019. URL: https://www.sciencedirect.com/science/article/pii/S0730725X12000604 (visited on 04/13/2024) (cit. on p. 111).

[246] Xiaobo Qu, Yingkun Hou, Fan Lam, et al. “Magnetic Resonance Image Reconstruction from Undersampled Measurements Using a Patch-Based Nonlocal Operator”. In: *Medical Image Analysis*. Sparse Methods for Signal Reconstruction and Medical Image Analysis 18.6 (Aug. 1, 2014), pp. 843–856. ISSN: 1361-8415. DOI: 10.1016/j.media.2013.09.007. URL: https://www.sciencedirect.com/science/article/pii/S1361841513001448 (visited on 04/02/2024) (cit. on p. 43).

[247] *Quality Assessment for Super-Resolution Image Enhancement*. URL: https://ieeexplore.ieee.org/document/4106955 (visited on 06/18/2024) (cit. on p. 66).

[248] Alec Radford, Jeff Wu, R. Child, et al. “Language Models Are Unsupervised Multitask Learners”. In: 2019. URL: https://www.semanticscholar.org/paper/Language-Models-are-Unsupervised-Multitask-Learners-Radford-Wu/9405cc0d6169988371b2755e573cc28650d14df (visited on 03/23/2024) (cit. on p. 36).

[249] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, et al. *Designing Network Design Spaces*. Mar. 30, 2020. DOI: 10.48550/arXiv.2003.13678. arXiv: 2003.13678 [cs]. URL: http://arxiv.org/abs/2003.13678 (visited on 03/23/2024). preprint (cit. on p. 34).

[250] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, et al. *Hierarchical Text-Conditional Image Generation with CLIP Latents*. Apr. 12, 2022. DOI: 10.48550/arXiv.2204.06125. arXiv: 2204.06125 [cs]. URL: http://arxiv.org/abs/2204.06125 (visited on 03/23/2024). preprint (cit. on p. 36).

[251] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, et al. *Zero-Shot Text-to-Image Generation*. Feb. 26, 2021. DOI: 10.48550/arXiv.2102.12092. arXiv: 2102.12092 [cs]. URL: http://arxiv.org/abs/2102.12092 (visited on 04/02/2024). preprint (cit. on p. 43).

[252] Zaccharie Ramzi, Philippe Ciuciu, and Jean-Luc Starck. “Benchmarking Deep Nets MRI Reconstruction Models on the Fastmri Publicly Available Dataset”. In: *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). Apr. 2020, pp. 1441–1445. DOI: 10.1109/ISBI45749.2020.9098335. URL: https://ieeexplore.ieee.org/document/9098335 (visited on 04/04/2024) (cit. on p. 52).

[253] Zaccharie Ramzi, Philippe Ciuciu, and Jean-Luc Starck. *XPDNet for MRI Reconstruction: An Application to the 2020 fastMRI Challenge*. July 7, 2021. DOI: 10.48550/arXiv.2010.07290. arXiv: 2010.07290 [physics, stat]. URL: http://arxiv.org/abs/2010.07290 (visited on 04/10/2024). preprint (cit. on p. 53).

[254] Garvesh Raskutti, Martin J. Wainwright, and Bin Yu. *Early Stopping and Non-Parametric Regression: An Optimal Data-Dependent Stopping Rule*. June 15, 2013. DOI: 10.48550/arXiv.1306.3574. arXiv: 1306.3574 [stat]. URL: http://arxiv.org/abs/1306.3574 (visited on 03/23/2024). preprint (cit. on p. 32).

[255] Arash Rasti-Meymandi, Aboozar Ghaffari, and Emad Fatemizadeh. “Plug and Play Augmented HQS: Convergence Analysis and Its Application in MRI Reconstruction”. In: *Neurocomputing* 518 (Jan. 21, 2023), pp. 1–14. ISSN: 0925-2312. DOI: 10.1016/j.neucom.2022.10.061. URL: https://www.sciencedirect.com/science/article/pii/S0925231222013406 (visited on 04/10/2024) (cit. on pp. 53, 111).

[256] Sriram Ravula, Brett Levac, Ajil Jalal, et al. *Optimizing Sampling Patterns for Compressed Sensing MRI with Diffusion Generative Models*. June 5, 2023. DOI: 10.48550/arXiv.2306.03284. arXiv: 2306.03284 [cs, eess]. URL: http://arxiv.org/abs/2306.03284 (visited on 04/03/2024). preprint (cit. on p. 56).

[257] Artem Razumov, Oleg Y. Rogov, and Dmitry V. Dylov. "Optimal MRI Undersampling Patterns for Pathology Localization". In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*. Ed. by Linwei Wang, Qi Dou, P. Thomas Fletcher, et al. Lecture Notes in Computer Science. Cham: Springer Nature Switzerland, 2022, pp. 768–779. ISBN: 978-3-031-16446-0. DOI: 10.1007/978-3-031-16446-0_73 (cit. on pp. 55, 125).

[258] Danilo Jimenez Rezende and Shakir Mohamed. *Variational Inference with Normalizing Flows*. June 14, 2016. DOI: 10.48550/arXiv.1505.05770. arXiv: 1505.05770 [cs, stat]. URL: http://arxiv.org/abs/1505.05770 (visited on 03/23/2024). preprint (cit. on p. 36).

[259] Salah Rifai, Yoshua Bengio, Yann Dauphin, et al. *A Generative Process for Sampling Contractive Auto-Encoders*. June 27, 2012. DOI: 10.48550/arXiv.1206.6434. arXiv: 1206.6434 [cs, stat]. URL: http://arxiv.org/abs/1206.6434 (visited on 03/23/2024). preprint (cit. on p. 35).

[260] Herbert Robbins and Sutton Monro. "A Stochastic Approximation Method". In: *The Annals of Mathematical Statistics* 22.3 (1951), pp. 400–407. ISSN: 0003-4851. JSTOR: 2236626. URL: https://www.jstor.org/stable/2236626 (visited on 03/23/2024) (cit. on p. 32).

[261] Simon Robinson and Jorge Jovicich. "B0 Mapping with Multi-Channel RF Coils at High Field". In: *Magnetic Resonance in Medicine* 66.4 (Oct. 2011), pp. 976–988. ISSN: 1522-2594. DOI: 10.1002/mrm.22879. pmid: 21608027 (cit. on p. 26).

[262] P. B. Roemer, W. A. Edelstein, C. E. Hayes, et al. "The NMR Phased Array". In: *Magnetic Resonance in Medicine* 16.2 (Nov. 1990), pp. 192–225. ISSN: 0740-3194. DOI: 10.1002/mrm.1910160203. pmid: 2266841 (cit. on p. 22).

[263] Robin Rombach, Andreas Blattmann, Dominik Lorenz, et al. *High-Resolution Image Synthesis with Latent Diffusion Models*. Apr. 13, 2022. DOI: 10.48550/arXiv.2112.10752. arXiv: 2112.10752 [cs]. URL: http://arxiv.org/abs/2112.10752 (visited on 04/02/2024). preprint (cit. on p. 43).

[264] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. "Learning Representations by Back-Propagating Errors". In: *Nature* 323.6088 (Oct. 1986), pp. 533–536. ISSN: 1476-4687. DOI: 10.1038/323533a0. URL: https://www.nature.com/articles/323533a0 (visited on 03/23/2024) (cit. on p. 30).

[265] Chitwan Saharia, Jonathan Ho, William Chan, et al. *Image Super-Resolution via Iterative Refinement*. June 30, 2021. DOI: 10.48550/arXiv.2104.07636. arXiv: 2104.07636 [cs, eess]. URL: http://arxiv.org/abs/2104.07636 (visited on 04/02/2024). preprint (cit. on p. 43).

[266] Irina Sanchez and Veronica Vilaplana. "Brain MRI Super-Resolution Using 3D Generative Adversarial Networks". In: Medical Imaging with Deep Learning. July 11, 2022. URL: https://openreview.net/forum?id=rJevSbniM (visited on 04/02/2024) (cit. on p. 43).

[267] Thomas Sanchez. "Learning to Sample in Cartesian MRI". May 18, 2022. DOI: 10.5075/epfl-thesis-9981. arXiv: 2312.04327 [cs, eess]. URL: http://arxiv.org/abs/2312.04327 (visited on 05/21/2024) (cit. on p. 56).

[268] Thomas Sanchez, Igor Krawczuk, Zhaodong Sun, et al. “Closed Loop Deep Bayesian Inversion: Uncertainty Driven Acquisition for Fast MRI”. In: (Sept. 25, 2019). URL: https://openreview.net/forum?id=BJlPOlBKDB (visited on 04/04/2024) (cit. on pp. 55, 56).

[269] Christopher M. Sandino, Joseph Y. Cheng, Feiyu Chen, et al. “Compressed Sensing: From Research to Clinical Practice with Deep Neural Networks”. In: *IEEE signal processing magazine* 37.1 (Jan. 2020), pp. 111–127. ISSN: 1053-5888. DOI: 10.1109/MSP.2019.2950433. pmid: 33192036. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7664163/ (visited on 03/20/2024) (cit. on p. 14).

[270] Umme Sara, Morium Akter, and Mohammad Shorif Uddin. “Image Quality Assessment through FSIM, SSIM, MSE and PSNR—A Comparative Study”. In: *Journal of Computer and Communications* 7.3 (3 Mar. 4, 2019), pp. 8–18. DOI: 10.4236/jcc.2019.73002. URL: https://www.scirp.org/journal/paperinformation?paperid=90911 (visited on 06/23/2024) (cit. on p. 61).

[271] Jo Schlemper, Jose Caballero, Joseph V. Hajnal, et al. “A Deep Cascade of Convolutional Neural Networks for MR Image Reconstruction”. In: *Information Processing in Medical Imaging*. Ed. by Marc Niethammer, Martin Styner, Stephen Aylward, et al. Cham: Springer International Publishing, 2017, pp. 647–658. ISBN: 978-3-319-59050-9. DOI: 10.1007/978-3-319-59050-9_51 (cit. on p. 52).

[272] Ferdinand Schweser, Andreas Deistung, Berengar W. Lehr, et al. “Differentiation between Diamagnetic and Paramagnetic Cerebral Lesions Based on Magnetic Susceptibility Mapping”. In: *Medical Physics* 37.10 (Oct. 2010), pp. 5165–5178. ISSN: 0094-2405. DOI: 10.1118/1.3481505. pmid: 21089750 (cit. on p. 26).

[273] Matthias Seeger, Hannes Nickisch, Rolf Pohmann, et al. “Optimization of K-Space Trajectories for Compressed Sensing by Bayesian Experimental Design”. In: *Magnetic Resonance in Medicine* 63.1 (2010), pp. 116–126. ISSN: 1522-2594. DOI: 10.1002/mrm.22180. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.22180 (visited on 04/04/2024) (cit. on p. 54).

[274] Viraj Shah and Chinmay Hegde. *Solving Linear Inverse Problems Using GAN Priors: An Algorithm with Provable Guarantees*. Feb. 23, 2018. DOI: 10.48550/arXiv.1802.08406. arXiv: 1802.08406 [cs, stat]. URL: http://arxiv.org/abs/1802.08406 (visited on 04/04/2024). preprint (cit. on p. 56).

[275] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. “SinGAN: Learning a Generative Model From a Single Natural Image”. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2019, pp. 4570–4580. URL: https://openaccess.thecvf.com/content_ICCV_2019/html/Shaham_SinGAN_Learning_a_Generative_Model_From_a_Single_Natural_Image_ICCV_2019_paper.html (visited on 04/13/2024) (cit. on p. 110).

[276] Roy Shaul, Itamar David, Ohad Shitrit, et al. “Subsampled Brain MRI Reconstruction by Generative Adversarial Neural Networks”. In: *Medical Image Analysis* 65 (Oct. 1, 2020), p. 101747. ISSN: 1361-8415. DOI: 10.1016/j.media.2020.101747. URL: https://www.sciencedirect.com/science/article/pii/S1361841520301110 (visited on 07/27/2022) (cit. on p. 75).

[277] Feng Shi, Jian Cheng, Li Wang, et al. “LRTV: MR Image Super-Resolution With Low-Rank and Total Variation Regularizations”. In: *IEEE Transactions on Medical Imaging* 34.12 (Dec. 2015), pp. 2459–2466. ISSN: 1558-254X. DOI: 10.1109/TMI.2015.2437894. URL: https://ieeexplore.ieee.org/document/7113897 (visited on 04/02/2024) (cit. on p. 43).

[278] Jun Shi, Qingping Liu, Chaofeng Wang, et al. “Super-Resolution Reconstruction of MR Image with a Novel Residual Learning Network Algorithm”. In: *Physics in Medicine and Biology* 63.8 (Apr. 19, 2018), p. 085011. ISSN: 1361-6560. DOI: 10.1088/1361-6560/aab9e9. pmid: 29583134 (cit. on p. 43).

[279] Peter J. Shin, Peder E. Z. Larson, Michael A. Ohliger, et al. “Calibrationless Parallel Imaging Reconstruction Based on Structured Low-Rank Matrix Completion”. In: *Magnetic Resonance in Medicine* 72.4 (Oct. 2014), pp. 959–970. ISSN: 1522-2594. DOI: 10.1002/mrm.24997. pmid: 24248734 (cit. on p. 51).

[280] Karen Simonyan and Andrew Zisserman. “Very Deep Convolutional Networks for Large-Scale Image Recognition”. Apr. 10, 2015. arXiv: 1409.1556 [cs]. URL: http://arxiv.org/abs/1409.1556 (visited on 08/06/2020) (cit. on p. 33).

[281] Seohee So, Hyun Wook Park, Byungjai Kim, et al. “BUDA-MESMERISE: Rapid Acquisition and Unsupervised Parameter Estimation for T1, T2, M0, B0, and B1 Maps”. In: *Magnetic Resonance in Medicine* 88.1 (2022), pp. 292–308. ISSN: 1522-2594. DOI: 10.1002/mrm.29228. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29228 (visited on 04/08/2024) (cit. on p. 58).

[282] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, et al. *Deep Unsupervised Learning Using Nonequilibrium Thermodynamics*. Nov. 18, 2015. DOI: 10.48550/arXiv.1503.03585. arXiv: 1503.03585 [cond-mat, q-bio, stat]. URL: http://arxiv.org/abs/1503.03585 (visited on 03/23/2024). preprint (cit. on pp. 36, 40).

[283] Yang Song and Stefano Ermon. *Generative Modeling by Estimating Gradients of the Data Distribution*. Oct. 10, 2020. DOI: 10.48550/arXiv.1907.05600. arXiv: 1907.05600 [cs, stat]. URL: http://arxiv.org/abs/1907.05600 (visited on 03/23/2024). preprint (cit. on p. 40).

[284] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, et al. *Score-Based Generative Modeling through Stochastic Differential Equations*. Feb. 10, 2021. DOI: 10.48550/arXiv.2011.13456. arXiv: 2011.13456 [cs, stat]. URL: http://arxiv.org/abs/2011.13456 (visited on 04/02/2024). preprint (cit. on p. 40).

[285] Anuroop Sriram, Jure Zbontar, Tullie Murrell, et al. “End-to-End Variational Networks for Accelerated MRI Reconstruction”. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. Ed. by Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, et al. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2020, pp. 64–73. ISBN: 978-3-030-59713-9. DOI: 10.1007/978-3-030-59713-9_7 (cit. on pp. 35, 52, 53, 56, 75).

[286] Anuroop Sriram, Jure Zbontar, Tullie Murrell, et al. “GrappaNet: Combining Parallel Imaging With Deep Learning for Multi-Coil MRI Reconstruction”. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). June 2020, pp. 14303–14310. DOI: 10.1109/CVPR42600.2020.01432 (cit. on pp. 35, 54, 75, 124).

[287] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, et al. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958. ISSN: 1533-7928. URL: http://jmlr.org/papers/v15/srivastava14a.html (visited on 03/23/2024) (cit. on p. 33).

[288] Carsten Stüber, David Pitt, and Yi Wang. “Iron in Multiple Sclerosis and Its Noninvasive Imaging with Quantitative Susceptibility Mapping”. In: *International Journal of Molecular Sciences* 17.1 (Jan. 14, 2016), p. 100. ISSN: 1422-0067. DOI: 10.3390/ijms17010100. pmid: 26784172 (cit. on p. 27).

[289] Tilman J. Sumpf, Martin Uecker, Susann Boretius, et al. “Model-Based Nonlinear Inverse Reconstruction for T2 Mapping Using Highly Undersampled Spin-Echo MRI”. In: *Journal of magnetic resonance imaging: JMRI* 34.2 (Aug. 2011), pp. 420–428. ISSN: 1522-2586. DOI: 10.1002/jmri.22634. pmid: 21780234 (cit. on p. 58).

[290] Zhaodong Sun, Thomas Sanchez, Fabian Latorre, et al. “Solving Inverse Problems with Hybrid Deep Image Priors: The Challenge of Preventing Overfitting”. In: *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. June 6, 2021, pp. 8103–8107. DOI: 10.1109/ICASSP39728.2021.9414879. arXiv: 2011.01748 [cs, eess]. URL: http://arxiv.org/abs/2011.01748 (visited on 04/03/2024) (cit. on p. 51).

[291] Vadim Sushko, Dan Zhang, Juergen Gall, et al. *Generating Novel Scene Compositions from Single Images and Videos*. Dec. 13, 2023. DOI: 10.48550/arXiv.2103.13389. arXiv: 2103.13389 [cs]. URL: http://arxiv.org/abs/2103.13389 (visited on 04/13/2024). preprint (cit. on p. 110).

[292] Christian Szegedy, Wei Liu, Yangqing Jia, et al. *Going Deeper with Convolutions*. Sept. 16, 2014. DOI: 10.48550/arXiv.1409.4842. arXiv: 1409.4842 [cs]. URL: http://arxiv.org/abs/1409.4842 (visited on 03/23/2024). preprint (cit. on p. 33).

[293] Ying Tai, Jian Yang, and Xiaoming Liu. “Image Super-Resolution via Deep Recursive Residual Network”. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). July 2017, pp. 2790–2798. DOI: 10.1109/CVPR.2017.298. URL: https://ieeexplore.ieee.org/document/8099781 (visited on 04/02/2024) (cit. on p. 42).

[294] Ying Tai, Jian Yang, Xiaoming Liu, et al. *MemNet: A Persistent Memory Network for Image Restoration*. Aug. 7, 2017. DOI: 10.48550/arXiv.1708.02209. arXiv: 1708.02209 [cs]. URL: http://arxiv.org/abs/1708.02209 (visited on 04/02/2024). preprint (cit. on p. 42).

[295] L. N. Tanenbaum, A. J. Tsiouris, A. N. Johnson, et al. “Synthetic MRI for Clinical Neuroimaging: Results of the Magnetic Resonance Image Compilation (MAGiC) Prospective, Multicenter, Multireader Trial”. In: *AJNR. American journal of neuroradiology* 38.6 (June 2017), pp. 1103–1110. ISSN: 1936-959X. DOI: 10.3174/ajnr.A5227. pmid: 28450439 (cit. on p. 57).

[296] Dávid Terjék. “Adversarial Lipschitz Regularization”. In: International Conference on Learning Representations. Sept. 25, 2019. URL: https://openreview.net/forum?id=Bke_DertPB (visited on 04/16/2024) (cit. on p. 111).

[297] Kerem C. Tezcan, Christian F. Baumgartner, Roger Luechinger, et al. “MR Image Reconstruction Using Deep Density Priors”. In: *IEEE Transactions on Medical Imaging* 38.7 (July 2019), pp. 1633–1642. ISSN: 1558-254X. DOI: 10.1109/TMI.2018.2887072. URL: https://ieeexplore.ieee.org/abstract/document/8579232 (visited on 04/12/2024) (cit. on pp. 54, 111).

[298] Thibaut Théate and Damien Ernst. “An Application of Deep Reinforcement Learning to Algorithmic Trading”. In: *Expert Systems with Applications* 173 (July 1, 2021), p. 114632. ISSN: 0957-4174. DOI: 10.1016/j.eswa.2021.114632. URL: https://www.sciencedirect.com/science/article/pii/S0957417421000737 (visited on 03/23/2024) (cit. on p. 29).

[299] Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, et al. *Deep Complex Networks*. Feb. 25, 2018. DOI: 10.48550/arXiv.1705.09792. arXiv: 1705.09792 [cs]. URL: http://arxiv.org/abs/1705.09792 (visited on 03/30/2024). preprint (cit. on pp. 35, 36, 75).

[300] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. "Deep Image Prior". In: *International Journal of Computer Vision* 128.7 (July 1, 2020), pp. 1867–1888. ISSN: 1573-1405. DOI: 10.1007/s11263-020-01303-4. URL: https://doi.org/10.1007/s11263-020-01303-4 (visited on 04/03/2024) (cit. on pp. 51, 111).

[301] S. R. Underwood, D. N. Firmin, R. H. Klipstein, et al. "Magnetic Resonance Velocity Mapping: Clinical Application of a New Technique." In: *Heart* 57.5 (May 1, 1987), pp. 404–412. ISSN: 1355-6037, 1468-201X. DOI: 10.1136/hrt.57.5.404. pmid: 3496109. URL: https://heart.bmj.com/content/57/5/404 (visited on 01/13/2023) (cit. on p. 75).

[302] Behroze Vachha and Susie Y. Huang. "MRI with Ultrahigh Field Strength and High-Performance Gradients: Challenges and Opportunities for Clinical Neuroimaging at 7 T and Beyond". In: *European Radiology Experimental* 5.1 (Aug. 26, 2021), p. 35. ISSN: 2509-9280. DOI: 10.1186/s41747-021-00216-2. URL: https://doi.org/10.1186/s41747-021-00216-2 (visited on 03/21/2024) (cit. on p. 14).

[303] Arash Vahdat, Karsten Kreis, and Jan Kautz. "Score-Based Generative Modeling in Latent Space". In: Advances in Neural Information Processing Systems. Nov. 9, 2021. URL: https://openreview.net/forum?id=P9TYG0j-wtG (visited on 04/02/2024) (cit. on p. 40).

[304] R. Van de Walle, I. Lemahieu, and E. Achten. "Magnetic Resonance Imaging and the Reduction of Motion Artifacts: Review of the Principles". In: *Technology and Health Care: Official Journal of the European Society for Engineering and Medicine* 5.6 (Dec. 1997), pp. 419–435. ISSN: 0928-7329. pmid: 9696161 (cit. on p. 14).

[305] Eric Van Reeth, Ivan W. K. Tham, Cher Heng Tan, et al. "Super-Resolution in Magnetic Resonance Imaging: A Review". In: *Concepts in Magnetic Resonance Part A* 40A.6 (2012), pp. 306–325. ISSN: 1552-5023. DOI: 10.1002/cmr.a.21249. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cmr.a.21249 (visited on 03/24/2024) (cit. on p. 41).

[306] Willem Van Valenberg, Stefan Klein, Frans M. Vos, et al. "An Efficient Method for Multi-Parameter Mapping in Quantitative MRI Using B-Spline Interpolation". In: *IEEE Transactions on Medical Imaging* 39.5 (2020), pp. 1681–1689. ISSN: 0278-0062. DOI: 10.1109/TMI.2019.2954751. URL: http://www.scopus.com/inward/record.url?scp=85084960923&partnerID=8YFLogxK (visited on 04/07/2024) (cit. on p. 58).

[307] Divya Varadarajan, Katherine L. Bouman, Andre van der Kouwe, et al. *Unsupervised Learning of MRI Tissue Properties Using MRI Physics Models*. July 6, 2021. DOI: 10.48550/arXiv.2107.02704. arXiv: 2107.02704 [cs, eess, q-bio]. URL: http://arxiv.org/abs/2107.02704 (visited on 04/08/2024). preprint (cit. on p. 59).

[308] *Variational Methods in Imaging*. URL: https://link.springer.com/book/10.1007/978-0-387-69277-7 (visited on 04/16/2024) (cit. on p. 112).

[309] Ashish Vaswani, Noam Shazeer, Niki Parmar, et al. "Attention Is All You Need". In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc., 2017. URL: https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html (visited on 03/23/2024) (cit. on p. 29).

[310] Singanallur V. Venkatakrishnan, Charles A. Bouman, and Brendt Wohlberg. “Plug-and-Play Priors for Model Based Reconstruction”. In: *2013 IEEE Global Conference on Signal and Information Processing*. 2013 IEEE Global Conference on Signal and Information Processing. Dec. 2013, pp. 945–948. DOI: 10.1109/GlobalSIP.2013.6737048. URL: https://ieeexplore.ieee.org/document/6737048 (visited on 04/10/2024) (cit. on p. 53).

[311] Ricardo Vilalta and Youssef Drissi. “A Perspective View and Survey of Meta-Learning”. In: *Artificial Intelligence Review* 18.2 (June 1, 2002), pp. 77–95. ISSN: 1573-7462. DOI: 10.1023/A:1019956318069. URL: https://doi.org/10.1023/A:1019956318069 (visited on 04/16/2024) (cit. on p. 110).

[312] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, et al. “Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion”. In: *The Journal of Machine Learning Research* 11 (Dec. 1, 2010), pp. 3371–3408. ISSN: 1532-4435 (cit. on p. 35).

[313] David O. Walsh, Arthur F. Gmitro, and Michael W. Marcellin. “Adaptive Reconstruction of Phased Array MR Imagery”. In: *Magnetic Resonance in Medicine* 43.5 (2000), pp. 682–690. ISSN: 1522-2594. DOI: 10.1002/(SICI)1522-2594(200005)43:5<682::AID-MRM10>3.0.CO;2-G. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291522-2594%28200005%2943%3A5%3C682%3A%3AAID-MRM10%3E3.0.CO%3B2-G (visited on 01/23/2023) (cit. on pp. 46, 97).

[314] Alan Q. Wang, Adrian V. Dalca, and Mert R. Sabuncu. “Neural Network-Based Reconstruction in Compressed Sensing MRI Without Fully-Sampled Training Data”. In: *Machine Learning for Medical Image Reconstruction*. Ed. by Farah Deeba, Patricia Johnson, Tobias Würfl, et al. Cham: Springer International Publishing, 2020, pp. 27–37. ISBN: 978-3-030-61598-7. DOI: 10.1007/978-3-030-61598-7_3 (cit. on p. 53).

[315] Fuyixue Wang, Zijing Dong, Timothy G. Reese, et al. “Echo Planar Time-Resolved Imaging (EPTI)”. In: *Magnetic Resonance in Medicine* 81.6 (June 2019), pp. 3599–3615. ISSN: 1522-2594. DOI: 10.1002/mrm.27673. pmid: 30714198 (cit. on p. 57).

[316] Guanhua Wang, Tianrui Luo, Jon-Fredrik Nielsen, et al. “B-Spline Parameterized Joint Optimization of Reconstruction and K-Space Trajectories (BJORK) for Accelerated 2D MRI”. In: *IEEE Transactions on Medical Imaging* 41.9 (Sept. 2022), pp. 2318–2330. ISSN: 1558-254X. DOI: 10.1109/TMI.2022.3161875. URL: https://ieeexplore.ieee.org/document/9740248 (visited on 04/04/2024) (cit. on p. 55).

[317] Hengkang Wang, Taihui Li, Zhong Zhuang, et al. “Early Stopping for Deep Image Prior”. In: (Sept. 29, 2022). URL: https://openreview.net/forum?id=JIl_kij_aov (visited on 04/03/2024) (cit. on p. 51).

[318] Jiancong Wang, Yuhua Chen, Yifan Wu, et al. “Enhanced Generative Adversarial Network for 3D Brain MRI Super-Resolution”. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. 2020, pp. 3627–3636. URL: https://openaccess.thecvf.com/content_WACV_2020/html/Wang_Enhanced_generative_adversarial_network_for_3D_brain_MRI_super-resolution_WACV_2020_paper.html (visited on 03/24/2024) (cit. on pp. 41, 43, 75, 111, 113).

[319] Jiechao Wang, Qinqin Yang, Qizhi Yang, et al. “Joint Optimization of Cartesian Sampling Patterns and Reconstruction for Single-contrast and Multi-contrast Fast Magnetic Resonance Imaging”. In: *Computer Methods and Programs in Biomedicine* 226 (Nov. 1, 2022), p. 107150. ISSN: 0169-2607. DOI: 10.1016/j.cmpb.2022.107150. URL: https://www.sciencedirect.com/science/article/pii/S0169260722005314 (visited on 04/04/2024) (cit. on p. 55).

[320] Jueqi Wang, Jacob Levman, Walter Hugo Lopez Pinaya, et al. “InverseSR: 3D Brain MRI Super-Resolution Using a Latent Diffusion Model”. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. Ed. by Hayit Greenspan, Anant Madabhushi, Parvin Mousavi, et al. Cham: Springer Nature Switzerland, 2023, pp. 438–447. ISBN: 978-3-031-43999-5. DOI: 10.1007/978-3-031-43999-5_42 (cit. on p. 43).

[321] Shanshan Wang, Zhenghang Su, Leslie Ying, et al. “Accelerating Magnetic Resonance Imaging via Deep Learning”. In: *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*. 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI). Apr. 2016, pp. 514–517. DOI: 10.1109/ISBI.2016.7493320. URL: https://ieeexplore.ieee.org/document/7493320 (visited on 04/04/2024) (cit. on p. 52).

[322] Xintao Wang, Ke Yu, Shixiang Wu, et al. “ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks”. In: *Computer Vision – ECCV 2018 Workshops*. Ed. by Laura Leal-Taixé and Stefan Roth. Cham: Springer International Publishing, 2019, pp. 63–79. ISBN: 978-3-030-11021-5. DOI: 10.1007/978-3-030-11021-5_5 (cit. on p. 42).

[323] Yi Wang and Tian Liu. “Quantitative Susceptibility Mapping (QSM): Decoding MRI Data for a Tissue Magnetic Biomarker”. In: *Magnetic Resonance in Medicine* 73.1 (2015), pp. 82–101. ISSN: 1522-2594. DOI: 10.1002/mrm.25358. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.25358 (visited on 03/08/2024) (cit. on p. 35).

[324] Yun-Heng Wang, Jiaqing Qiao, Jun-Bao Li, et al. “Sparse Representation-Based MRI Super-Resolution Reconstruction”. In: *Measurement* 47 (Jan. 1, 2014), pp. 946–953. ISSN: 0263-2241. DOI: 10.1016/j.measurement.2013.10.026. URL: https://www.sciencedirect.com/science/article/pii/S0263224113005253 (visited on 04/02/2024) (cit. on p. 43).

[325] Zhendong Wang, Yifan Jiang, Huangjie Zheng, et al. “Patch Diffusion: Faster and More Data-Efficient Training of Diffusion Models”. In: Thirty-Seventh Conference on Neural Information Processing Systems. Nov. 2, 2023. URL: https://openreview.net/forum?id=iv2sTQtbst&referrer=%5Bthe%20profile%20of%20Pengcheng%20He%5D(%2Fprofile%3Fid%3D~Pengcheng_He2) (visited on 04/09/2024) (cit. on p. 111).

[326] Zhou Wang, A.C. Bovik, H.R. Sheikh, et al. “Image Quality Assessment: From Error Visibility to Structural Similarity”. In: *IEEE Transactions on Image Processing* 13.4 (Apr. 2004), pp. 600–612. ISSN: 1941-0042. DOI: 10.1109/TIP.2003.819861 (cit. on pp. 64, 65).

[327] Zhou Wang and Alan C. Bovik. “Mean Squared Error: Love It or Leave It? A New Look at Signal Fidelity Measures”. In: *IEEE Signal Processing Magazine* 26.1 (Jan. 2009), pp. 98–117. ISSN: 1558-0792. DOI: 10.1109/MSP.2008.930649. URL: https://ieeexplore.ieee.org/document/4775883 (visited on 04/16/2024) (cit. on pp. 61, 63, 65).

[328] Kaixuan Wei, Angelica Aviles-Rivero, Jingwei Liang, et al. *TFPnP: Tuning-free Plug-and-Play Proximal Algorithm with Applications to Inverse Imaging Problems*. Sept. 18, 2021. DOI: 10.48550/arXiv.2012.05703. arXiv: 2012.05703 [cs, eess]. URL: http://arxiv.org/abs/2012.05703 (visited on 04/10/2024). preprint (cit. on p. 53).

[329] Tomer Weiss, Sanketh Vedula, Ortal Senouf, et al. *Joint Learning of Cartesian Undersampling and Reconstruction for Accelerated MRI*. Apr. 5, 2020. DOI: 10.48550/arXiv.1905.09324. arXiv: 1905.09324 [physics]. URL: http://arxiv.org/abs/1905.09324 (visited on 04/04/2024). preprint (cit. on p. 55).

[330] Felix A. Wichmann and Robert Geirhos. *Are Deep Neural Networks Adequate Behavioural Models of Human Visual Perception?* May 26, 2023. DOI: 10.48550/arXiv.2305.17023. arXiv: 2305.17023 [cs, q-bio]. URL: http://arxiv.org/abs/2305.17023 (visited on 06/23/2024). preprint (cit. on p. 62).

[331] M. L. Wood and R. M. Henkelman. "MR Image Artifacts from Periodic Motion". In: *Medical Physics* 12.2 (1985), pp. 143–151. ISSN: 0094-2405. DOI: 10.1118/1.595782. pmid: 4000069 (cit. on p. 14).

[332] Bichen Wu, Alvin Wan, Xiangyu Yue, et al. "Shift: A Zero FLOP, Zero Parameter Alternative to Spatial Convolutions". In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. June 2018, pp. 9127–9135. DOI: 10.1109/CVPR.2018.00951. URL: https://ieeexplore.ieee.org/document/8579049 (visited on 03/23/2024) (cit. on p. 33).

[333] Yan Wu, Yajun Ma, Jiang Du, et al. "Accelerating Quantitative MR Imaging with the Incorporation of $B1$ Compensation Using Deep Learning". In: *Magnetic Resonance Imaging* 72 (Oct. 1, 2020), pp. 78–86. ISSN: 0730-725X. DOI: 10.1016/j.mri.2020.06.011. URL: https://www.sciencedirect.com/science/article/pii/S0730725X20300217 (visited on 04/08/2024) (cit. on p. 58).

[334] Zhanxiong Wu, Xuanheng Chen, Sangma Xie, et al. "Super-Resolution of Brain MRI Images Based on Denoising Diffusion Probabilistic Model". In: *Biomedical Signal Processing and Control* 85 (Aug. 1, 2023), p. 104901. ISSN: 1746-8094. DOI: 10.1016/j.bspc.2023.104901. URL: https://www.sciencedirect.com/science/article/pii/S1746809423003348 (visited on 03/24/2024) (cit. on pp. 41, 43).

[335] Lei Xiang, Yong Chen, Weitang Chang, et al. "Ultra-Fast T2-Weighted MR Reconstruction Using Complementary T1-Weighted Information". In: *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention* 11070 (Sept. 2018), pp. 215–223. DOI: 10.1007/978-3-030-00928-1_25. pmid: 30906934. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6430217/ (visited on 05/27/2024) (cit. on p. 53).

[336] Tiange Xiang, Mahmut Yurt, Ali B. Syed, et al. *DDM^2: Self-Supervised Diffusion MRI Denoising with Generative Diffusion Models*. Feb. 6, 2023. DOI: 10.48550/arXiv.2302.03018. arXiv: 2302.03018 [cs, eess]. URL: http://arxiv.org/abs/2302.03018 (visited on 06/21/2024). preprint (cit. on p. 124).

[337] Heng Xiao, Xin Wang, Jun Wang, et al. "Single Image Super-Resolution with Denoising Diffusion GANS". In: *Scientific Reports* 14.1 (Feb. 21, 2024), p. 4272. ISSN: 2045-2322. DOI: 10.1038/s41598-024-52370-3. URL: https://www.nature.com/articles/s41598-024-52370-3 (visited on 04/02/2024) (cit. on p. 43).

[338] Linfang Xiao, Yilong Liu, Zheyuan Yi, et al. “Partial Fourier Reconstruction of Complex MR Images Using Complex-Valued Convolutional Neural Networks”. In: *Magnetic Resonance in Medicine* 87.2 (2022), pp. 999–1014. ISSN: 1522-2594. DOI: 10.1002/mrm.29033. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29033 (visited on 01/13/2023) (cit. on p. 76).

[339] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. *Tackling the Generative Learning Trilemma with Denoising Diffusion GANs*. Apr. 4, 2022. DOI: 10.48550/arXiv.2112.07804. arXiv: 2112.07804 [cs, stat]. URL: http://arxiv.org/abs/2112.07804 (visited on 04/02/2024). preprint (cit. on p. 40).

[340] Saining Xie, Ross Girshick, Piotr Dollár, et al. *Aggregated Residual Transformations for Deep Neural Networks*. Apr. 10, 2017. DOI: 10.48550/arXiv.1611.05431. arXiv: 1611.05431 [cs]. URL: http://arxiv.org/abs/1611.05431 (visited on 03/23/2024). preprint (cit. on p. 33).

[341] Burhaneddin Yaman, Seyed Amir Hossein Hosseini, and Mehmet Akcakaya. “Zero-Shot Self-Supervised Learning for MRI Reconstruction”. In: International Conference on Learning Representations. Mar. 9, 2022. URL: https://openreview.net/forum?id=085y6YPaYjP (visited on 07/18/2022) (cit. on pp. 54, 75, 106).

[342] Burhaneddin Yaman, Seyed Amir Hossein Hosseini, Steen Moeller, et al. “Self-Supervised Physics-Based Deep Learning MRI Reconstruction Without Fully-Sampled Data”. In: *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). Apr. 2020, pp. 921–925. DOI: 10.1109/ISBI45749.2020.9098514 (cit. on pp. 51, 53, 75).

[343] Ling Yang, Zhilong Zhang, Yang Song, et al. “Diffusion Models: A Comprehensive Survey of Methods and Applications”. In: *ACM Computing Surveys* 56.4 (Nov. 9, 2023), 105:1–105:39. ISSN: 0360-0300. DOI: 10.1145/3626235. URL: https://doi.org/10.1145/3626235 (visited on 04/02/2024) (cit. on p. 40).

[344] Mingrui Yang, Dan Ma, Yun Jiang, et al. “Low Rank Approximation Methods for MR Fingerprinting with Large Scale Dictionaries”. In: *Magnetic Resonance in Medicine* 79.4 (Apr. 2018), pp. 2392–2400. ISSN: 1522-2594. DOI: 10.1002/mrm.26867. pmid: 28804918 (cit. on p. 58).

[345] Zehua Yang, Chengliang Liu, and Yong Xu. “Multi-Contrast MRI Reconstruction via Multi-scale Patched-based Cross-contrast Channel Selection”. In: *Journal of Physics: Conference Series* 2171.1 (Jan. 2022), p. 012034. ISSN: 1742-6596. DOI: 10.1088/1742-6596/2171/1/012034. URL: https://dx.doi.org/10.1088/1742-6596/2171/1/012034 (visited on 04/13/2024) (cit. on p. 111).

[346] Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. “On Early Stopping in Gradient Descent Learning”. In: *Constructive Approximation* 26.2 (Aug. 1, 2007), pp. 289–315. ISSN: 1432-0940. DOI: 10.1007/s00365-006-0663-2. URL: https://doi.org/10.1007/s00365-006-0663-2 (visited on 01/25/2023) (cit. on p. 32).

[347] Vasily L. Yarnykh. “Actual Flip-Angle Imaging in the Pulsed Steady State: A Method for Rapid Three-Dimensional Mapping of the Transmitted Radiofrequency Field”. In: *Magnetic Resonance in Medicine* 57.1 (Jan. 2007), pp. 192–200. ISSN: 0740-3194. DOI: 10.1002/mrm.21120. pmid: 17191242 (cit. on p. 26).

[348] Ali Pour Yazdanpanah, Onur Afacan, and Simon K. Warfield. *Deep Plug-and-Play Prior for Parallel MRI Reconstruction*. Sept. 6, 2019. DOI: 10.48550/arXiv.1909.00089. arXiv: 1909.00089 [cs, eess]. URL: http://arxiv.org/abs/1909.00089 (visited on 04/10/2024). preprint (cit. on p. 53).

[349] Tianwei Yin, Zihui Wu, He Sun, et al. *End-to-End Sequential Sampling and Reconstruction for MRI*. July 16, 2022. DOI: 10.48550/arXiv.2105.06460. arXiv: 2105.06460 [cs, eess]. URL: http://arxiv.org/abs/2105.06460 (visited on 04/04/2024). preprint (cit. on pp. 55, 56).

[350] Takeshi Yokoo, Won C. Bae, Gavin Hamilton, et al. "A Quantitative Approach to Sequence and Image Weighting". In: *Journal of Computer Assisted Tomography* 34.3 (2010), pp. 317–331. ISSN: 1532-3145. DOI: 10.1097/RCT.0b013e3181d3449a. pmid: 20498530 (cit. on p. 26).

[351] Jaejun Yoo, Kyong Hwan Jin, Harshit Gupta, et al. "Time-Dependent Deep Image Prior for Dynamic MRI". In: *IEEE Transactions on Medical Imaging* 40.12 (Dec. 2021), pp. 3337–3348. ISSN: 1558-254X. DOI: 10.1109/TMI.2021.3084288. URL: https://ieeexplore.ieee.org/document/9442767 (visited on 04/03/2024) (cit. on p. 51).

[352] Tengfei Yuan, Jie Yang, Jieru Chi, et al. "A Cross-Domain Complex Convolution Neural Network for Undersampled Magnetic Resonance Image Reconstruction". In: *Magnetic Resonance Imaging* 108 (May 2024), pp. 86–97. ISSN: 1873-5894. DOI: 10.1016/j.mri.2024.02.004. pmid: 38331053 (cit. on p. 53).

[353] Jure Zbontar, Florian Knoll, Anuroop Sriram, et al. *fastMRI: An Open Dataset and Benchmarks for Accelerated MRI*. Dec. 11, 2019. DOI: 10.48550/arXiv.1811.08839. arXiv: 1811.08839 [physics, stat]. URL: http://arxiv.org/abs/1811.08839 (visited on 06/15/2023). preprint (cit. on pp. 75, 81).

[354] Kun Zeng, Hong Zheng, Congbo Cai, et al. "Simultaneous Single- and Multi-Contrast Super-Resolution for Brain MRI Images Based on a Convolutional Neural Network". In: *Computers in Biology and Medicine* 99 (Aug. 1, 2018), pp. 133–141. ISSN: 0010-4825. DOI: 10.1016/j.compbiomed.2018.06.010. URL: https://www.sciencedirect.com/science/article/pii/S0010482518301586 (visited on 04/02/2024) (cit. on p. 43).

[355] Chi Zhang, Seyed Amir Hossein Hosseini, Sebastian Weingärtner, et al. "Optimized Fast GPU Implementation of Robust Artificial-Neural-Networks for k-Space Interpolation (RAKI) Reconstruction". In: *PLOS ONE* 14.10 (Oct. 23, 2019), e0223315. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0223315. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0223315 (visited on 01/23/2023) (cit. on p. 47).

[356] Chi Zhang, Steen Moeller, Omer Burak Demirel, et al. "Residual RAKI: A Hybrid Linear and Non-Linear Approach for Scan-Specific k-Space Deep Learning". In: *NeuroImage* 256 (Aug. 1, 2022), p. 119248. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2022.119248. URL: https://www.sciencedirect.com/science/article/pii/S105381192200372X (visited on 06/30/2022) (cit. on pp. 75, 76, 80, 81, 83–85, 89).

[357] Di Zhang, Jiazhong He, Yun Zhao, et al. "MR Image Super-Resolution Reconstruction Using Sparse Representation, Nonlocal Similarity and Sparse Derivative Prior". In: *Computers in Biology and Medicine* 58 (Mar. 1, 2015), pp. 130–145. ISSN: 0010-4825. DOI: 10.1016/j.compbiomed.2014.12.023. URL: https://www.sciencedirect.com/science/article/pii/S0010482515000025 (visited on 04/02/2024) (cit. on p. 43).

[358] Jinwei Zhang, Hang Zhang, Alan Wang, et al. “Extending LOUPE for K-Space Under-Sampling Pattern Optimization in Multi-coil MRI”. In: *Machine Learning for Medical Image Reconstruction*. Ed. by Farah Deeba, Patricia Johnson, Tobias Würfl, et al. Cham: Springer International Publishing, 2020, pp. 91–101. ISBN: 978-3-030-61598-7. DOI: 10.1007/978-3-030-61598-7_9 (cit. on p. 55).

[359] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. *Adding Conditional Control to Text-to-Image Diffusion Models*. Nov. 26, 2023. DOI: 10.48550/arXiv.2302.05543. arXiv: 2302.05543 [cs]. URL: http://arxiv.org/abs/2302.05543 (visited on 03/23/2024). preprint (cit. on p. 36).

[360] Richard Zhang, Phillip Isola, Alexei A. Efros, et al. *The Unreasonable Effectiveness of Deep Features as a Perceptual Metric*. Apr. 10, 2018. DOI: 10.48550/arXiv.1801.03924. arXiv: 1801.03924 [cs]. URL: http://arxiv.org/abs/1801.03924 (visited on 06/23/2024). preprint (cit. on p. 61).

[361] Tao Zhang, John M. Pauly, and Ives R. Levesque. “Accelerating Parameter Mapping with a Locally Low Rank Constraint”. In: *Magnetic Resonance in Medicine* 73.2 (2015), pp. 655–661. ISSN: 1522-2594. DOI: 10.1002/mrm.25161. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.25161 (visited on 04/07/2024) (cit. on p. 58).

[362] Xinlin Zhang, Di Guo, Yiman Huang, et al. “Image Reconstruction with Low-Rankness and Self-Consistency of k-Space Data in Parallel MRI”. In: *Medical Image Analysis* 63 (July 1, 2020), p. 101687. ISSN: 1361-8415. DOI: 10.1016/j.media.2020.101687. URL: https://www.sciencedirect.com/science/article/pii/S1361841520300530 (visited on 09/20/2021) (cit. on p. 52).

[363] Yulun Zhang, Yapeng Tian, Yu Kong, et al. “Residual Dense Network for Image Super-Resolution”. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. June 2018, pp. 2472–2481. DOI: 10.1109/CVPR.2018.00262. URL: https://ieeexplore.ieee.org/document/8578360 (visited on 04/01/2024) (cit. on p. 42).

[364] ZiCheng Zhang, CongYing Han, and TianDe Guo. *ExSinGAN: Learning an Explainable Generative Model from a Single Image*. Jan. 5, 2022. DOI: 10.48550/arXiv.2105.07350. arXiv: 2105.07350 [cs]. URL: http://arxiv.org/abs/2105.07350 (visited on 04/13/2024). preprint (cit. on p. 110).

[365] Zizhao Zhang, Adriana Romero, Matthew J. Muckley, et al. “Reducing Uncertainty in Undersampled MRI Reconstruction With Active Acquisition”. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). June 2019, pp. 2049–2053. DOI: 10.1109/CVPR.2019.00215. URL: https://ieeexplore.ieee.org/document/8953216 (visited on 04/04/2024) (cit. on pp. 55, 56).

[366] Bo Zhao, Wenmiao Lu, T. Kevin Hitchens, et al. “Accelerated MR Parameter Mapping with Low-Rank and Sparsity Constraints”. In: *Magnetic Resonance in Medicine* 74.2 (2015), pp. 489–498. ISSN: 1522-2594. DOI: 10.1002/mrm.25421. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.25421 (visited on 04/08/2024) (cit. on p. 57).

[367] Can Zhao, Aaron Carass, Blake E. Dewey, et al. “Self Super-Resolution for Magnetic Resonance Images Using Deep Networks”. In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). Apr. 2018, pp. 365–368. DOI: 10.1109/ISBI.2018.8363594. URL: https://ieeexplore.ieee.org/document/8363594 (visited on 04/02/2024) (cit. on p. 43).

[368] Lin Zhao, Xiao Chen, Eric Z. Chen, et al. *JoJoNet: Joint-contrast and Joint-sampling-and-reconstruction Network for Multi-contrast MRI*. Oct. 26, 2022. DOI: 10.48550/arXiv.2210.12548. arXiv: 2210.12548 [cs, eess]. URL: http://arxiv.org/abs/2210.12548 (visited on 04/04/2024). preprint (cit. on p. 55).

[369] Ruiyang Zhao, Burhaneddin Yaman, Yuxin Zhang, et al. “fastMRI+, Clinical Pathology Annotations for Knee and Brain Fully Sampled Magnetic Resonance Imaging Data”. In: *Scientific Data* 9.1 (Apr. 5, 2022), p. 152. ISSN: 2052-4463. DOI: 10.1038/s41597-022-01255-z. URL: https://www.nature.com/articles/s41597-022-01255-z (visited on 05/14/2024) (cit. on p. 123).

[370] Tiejun Zhao and Xiaoping Hu. “Iterative GRAPPA (iGRAPPA) for Improved Parallel Imaging Reconstruction”. In: *Magnetic Resonance in Medicine* 59.4 (2008), pp. 903–907. ISSN: 1522-2594. DOI: 10.1002/mrm.21370. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21370 (visited on 02/21/2024) (cit. on p. 51).

[371] Xiaole Zhao, Yulun Zhang, Tao Zhang, et al. “Channel Splitting Network for Single MR Image Super-Resolution”. In: *IEEE Transactions on Image Processing* 28.11 (Nov. 2019), pp. 5649–5662. ISSN: 1941-0042. DOI: 10.1109/TIP.2019.2921882. URL: https://ieeexplore.ieee.org/document/8736987 (visited on 04/02/2024) (cit. on p. 43).

[372] Hongkai Zheng, Weili Nie, Arash Vahdat, et al. “Fast Sampling of Diffusion Models via Operator Learning”. In: *Proceedings of the 40th International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, July 3, 2023, pp. 42390–42402. URL: https://proceedings.mlr.press/v202/zheng23d.html (visited on 04/02/2024) (cit. on p. 40).

[373] Bo Zhou, Neel Dey, Jo Schlemper, et al. *DSFormer: A Dual-domain Self-supervised Transformer for Accelerated Multi-contrast MRI Reconstruction*. Aug. 16, 2022. DOI: 10.48550/arXiv.2201.10776. arXiv: 2201.10776 [cs, eess]. URL: http://arxiv.org/abs/2201.10776 (visited on 04/26/2024). preprint (cit. on p. 54).

[374] Bo Zhou, Jo Schlemper, Neel Dey, et al. “Dual-Domain Self-Supervised Learning for Accelerated Non-Cartesian MRI Reconstruction”. In: *Medical Image Analysis* 81 (Oct. 1, 2022), p. 102538. ISSN: 1361-8415. DOI: 10.1016/j.media.2022.102538. URL: https://www.sciencedirect.com/science/article/pii/S1361841522001852 (visited on 04/26/2024) (cit. on p. 53).

[375] Bo Zhu, Jeremiah Z. Liu, Stephen F. Cauley, et al. “Image Reconstruction by Domain-Transform Manifold Learning”. In: *Nature* 555.7697 (7697 Mar. 2018), pp. 487–492. ISSN: 1476-4687. DOI: 10.1038/nature25988. URL: https://www.nature.com/articles/nature25988 (visited on 10/12/2020) (cit. on pp. 15, 35, 75).

[376] Marcelo Victor Wust Zibetti, Florian Knoll, and Ravinder R. Regatte. “Alternating Learning Approach for Variational Networks and Undersampling Pattern in Parallel MRI Applications”. In: *IEEE Transactions on Computational Imaging* 8 (2022), pp. 449–461. ISSN: 2333-9403. DOI: 10.1109/TCI.2022.3176129. URL: https://ieeexplore.ieee.org/document/9779482 (visited on 04/04/2024) (cit. on p. 55).

[377] Yukai Zou and Ikbeom Jang. *Engineering AI Tools for Systematic and Scalable Quality Assessment in Magnetic Resonance Imaging*. Dec. 2, 2021. DOI: 10.48550/arXiv.2112.01629. arXiv: 2112.01629 [cs, eess]. URL: http://arxiv.org/abs/2112.01629 (visited on 06/23/2024). preprint (cit. on p. 61).

List of Figures

1.1 Emerging Trends in MR Image Reconstruction . . . 15
1.2 Gradient-echo sequence . . . 20
1.3 Spin-echo sequence . . . 22
1.4 Parallel Imaging . . . 23
1.5 Different MRI contrasts. . . . 24

2.1 MLP. . . . 31
2.2 CNN. . . . 34
2.3 Residual Block. . . . 35
2.4 GAN. . . . 37
2.5 VAE. . . . 38
2.6 Research axes of MRI reconstruction. . . . 41
2.7 Illustration of training strategy in the 2D case for an acceleration factor $R_y = 3$ . . . 48
2.8 Illustration of the a) GRAPPA, b) RAKI, and c) cRAKI models. . . . 49

3.1 COBRAI performance on the Shepp-Logan phantom for various degradations. . . . 70
3.2 Metrics performance on degraded images . . . 71
3.3 Degraded images with smoothing. . . . 72
3.4 COBRAI illustration . . . 73
3.5 COBRAI illustration 2 . . . 74

4.1 rRAKI architecture . . . 78
4.2 Network architectures . . . 79
4.3 Effect of non-linearities on rRAKI . . . 84
4.4 Effect of non-linearities on crRAKI with LeakyReLU activation . . . 85
4.5 Qualitative Evaluation on fastMRI@ACS=24 (FLAIR and T1W) . . . 92
4.6 Qualitative Evaluation on fastMRI@ACS=24 (T1POST and T2W) . . . 92
4.7 Qualitative Evaluation on fastMRI@ACS=40 (FLAIR and T1W) . . . 93
4.8 Qualitative Evaluation on fastMRI@ACS=40 (T1POST and T2W) . . . 93
4.9 Quantitative on T2W for R = (2,3,4,5,6) for ACS=24 at separated mode . . . 94
4.10 Qualitative evaluation on T2W for R = (2,3,4,5,6) for ACS=40 at separated mode . . . 96
4.11 Qualitative evaluation on 2D GRE for $N_c = 16$ at R=2,4. . . . 97
4.12 Quantitative comparison of various methods to study the effect of the number of ACS 98
4.13 Qualitative Evaluation on 2D GRE for $N_c = 52$ at R=4 . . . 99
4.14 Qualitative Evaluation on 2D GRE for $N_c = 52$ at R=5 . . . 99
4.15 Qualitative Evaluation on 2D GRE for $N_c = 52$ at R=6 . . . 100
4.16 Qualitative evaluation on MPRAGE . . . 101
4.17 Qualitative evaluation on 3D GRE for ACS comparison . . . 102
4.18 Qualitative evaluation on 3D GRE . . . 103
4.19 Magnitude, R2*, QSM image at R=4x3 for crRAKI . . . 104

5.1 SANGRIA architecture. 113
5.2 Qualitative Evaluation of T2W images at R=5,6. 120
5.3 Effect of α. 121

List of Tables

2.1 Activation Functions 33
2.2 Complex Activation Functions 36

4.1 Main acquisition parameters for ACS 82
4.2 Effect of non-linearities on rRAKI with LeakyReLU activation 86
4.3 Effect of non-linearities on crRAKI with LeakyReLU activation 87
4.4 List of valid models 88
4.5 Model hyperparameter choice. 89
4.6 Quantitative Evaluation on fastMRI dataset at R = 4 for ACS=40 90
4.7 Quantitative Evaluation on fastMRI dataset at R = 4 for ACS=24 91

5.1 Quantitaive evaluation on synthetic data. 116
5.2 Quantitative evaluation on T2W images. 117

Acronyms list

ACS

Autocalibration Signal. 11, 12, 46, 47, 50, 54, 75–77, 80, 81, 83, 89–93, 95, 98, 100, 102–108, 122–124

ADC

Analog-to-Digital Converter. 20

ANN

Artificial Neural Network. 58

CAIPIRINHA

Controlled Aliasing in Parallel Imaging Results in Higher Acceleration. 124

cGAN

Conditional Generative Adversarial Networks. 37, 40, 43, 53, 56

CNN

Convolutional Neural Network. 11, 12, 29, 33–35, 42, 43, 47, 51–53, 55, 56, 58, 59, 109, 112, 122

COBRAI

COrreletion-Based Residual Artifact Index. 11, 12, 62, 66–69, 71, 98, 107, 122–125

cRAKI

complex RAKI. 78, 79, 106, 113, 115, 117, 118, 120, 121

crRAKI

complex residual RAKI. 11, 68, 76, 78–83, 88, 89, 92–101, 103–108, 122, 124, 125

CS

Compressive Sensing. 50, 52, 54, 58

CSGM

Compressed Sensing using Generative Models. 54, 56, 124

CT

Computational Tomography. 10, 64

cVAE

Conditional Variational Autoencoders. 39

DDPM

Diffusion Denoising Probabilistic Models. 40, 43

DDQN

Double Deep Q-Network. 56

DIP

Deep Image Prior. 51, 111

DL

Deep learning. 11, 29, 53–55, 57, 60, 75, 107

DM

DIffusion Model. 40, 43, 53, 54, 56, 124

DNN

Deep Neural Network. 29

E2E VarNet

End-to-End Variational Network. 56

ELBO

Evidence Lower Bound. 36, 38, 39

FID

Free Induction Decay. 14, 16, 18, 19, 21

FLAIR

Fluid Attenuated Inversion Recovery. 25, 76, 89

FR

Full-reference. 62

FT

Fourier Transform. 19

GAN

Generative Adversarial Networks. 36–38, 40, 42, 43, 53–55, 110–113, 118

GPU

Graphics Processing Unit. 29, 43

GRAPPA

Generalized Autocalibrating Partially Parallel Acquisitions. 11, 23, 45–47, 50, 54, 58, 76–80, 83, 88, 89, 92–108, 113, 115, 117, 118, 120, 122, 124

GRE

Gradient echo. 13, 25, 81–83, 95, 97, 99, 100, 102, 104, 108

HR

High-Resolution. 41, 42

INR

Implicit Neural Representation. 51

KL

Kullback-Leibler. 39

LBCS

Learning-Based Compressed Sensing. 56

LDM

Latent Diffusion Models. 43

LR

Low-Resolution. 41–43, 110

LSGM

Latent Score-based generative models. 40, 43

MAE

Mean Absolute Error. 63, 64

MI

Mutual Information. 125

ML

Machine learning. 10, 12, 29, 40, 41, 59

MLP

Multi Layer Perceptron. 29–31, 56, 59

MPRAGE

Magnetization Prepared Rapid Gradient Echo. 25, 81–83, 100, 101, 104, 122

MR

Magnetic Resonance. 10, 12, 15, 19, 26, 29, 41, 43, 45, 52, 54, 58, 59, 68, 122, 125

MRF

Magnetic Resonance Fingerprinting. 14, 26, 57, 58

MRI

Magnetic Resonance Imaging. 6, 10–17, 19, 23, 24, 26–29, 35, 40–44, 51–56, 58–62, 64, 67, 69, 75, 76, 111–113, 117–119, 122–125

MS

Multiple Sclerosis. 27

MSE

Mean Square Error. 62–66, 78–81, 86, 88, 104–107, 118, 124, 125

NLLS

Non Linear Least Squares. 57

NLP

natural language processing. 29

NMAE

Normalized Mean Absolute Error. 11, 61, 62, 64, 76, 98, 107, 115, 118

NMR

Nuclear Magnetic Resonance. 13, 15

NMSE

Normalized Mean Square Error. 63

NR

No-reference. 62

NRMSE

Normalized Root Mean Square Error. 11, 61–63, 76, 83, 86, 91, 94, 98, 107, 115, 118

PD

Proton Density. 59

PET

Positron Emission Tomography. 10

PI

Parallel Imaging. 22, 45, 46, 51, 52, 122, 123

PnP

Plug and Play. 53, 111

ppm

part-per-million. 16

PSNR

Peak Signal-to-Noise Ratio. 11, 61–63, 65, 66, 69, 76, 98, 107, 115, 118, 119, 123

qMRI

Quantitative Magnetic Resonance Imaging. 12, 26, 29, 41, 57, 59, 60

QSM

Quantitative Susceptibility Mapping. 11, 27, 35, 100, 104

RAKI

Robust artificial-neural-networks for k-space interpolation. 47, 77–80, 105, 107, 113, 115, 117, 118, 120

RDN

Residual Dense Networks. 42

ReLU

Rectified Linear Unit. 47

RF

Radiofrequence. 16, 17

RL

Reinforcement Learning. 56

RM

Residual Map. 63, 66–68, 107

RNN

Recurrent Neural Network. 52

RR

Reduced-reference. 62

rRAKI

residual RAKI. 11, 68, 76–78, 80–83, 86, 88–100, 104–106, 108, 122

SENSE

Sensitivity Encoding. 23, 45, 58, 124

SGD

Stochastic Gradient Descent. 32

SGM

Score-based generative models. 40, 43

SISR

single-image super-resolution. 42, 43

SNR

Signal-to-Noise Ratio. 21, 22, 45, 58, 123, 124

SPIRIT

Iterative Self-consistent Parallel Imaging Reconstruction. 23, 50–52

SR

Super-Resolution. 12, 41–43, 66, 110, 113, 125

SSIM

Structural Similarity Index. 11, 56, 61, 62, 64–66, 68, 69, 72, 76, 98, 107, 115, 116, 118, 119, 123–125

STE

Straight-Through Estimator. 55

SWI

Susceptibility Weighted Imaging. 27

T1

T1 relaxation time. 25

T1W

T1-Weighted. 12, 25, 43, 53, 76, 89, 122

T2W

T2-Weighted. 12, 25, 53, 58, 76, 80, 81, 109, 122

TE

Echo Time. 25, 27

TR

Repetition Time. 25, 27

TV

Total Variation. 50

TVP

Training–Validation Partition. 11, 47, 76, 89, 108, 123

VAE

Variational Autoencoders. 35, 36, 38–40, 42, 43, 54, 111

WGAN

Wasserstein Generative Adversarial Networks. 42, 112

WGAN-GP

Wasserstein Generative Adversarial Networks with Gradient Penalty. 42, 43

ZS-SSL

zero-shot self-supervised learning. 51, 53, 54