

Advanced Methods for Radial Data Sampling in Magnetic Resonance Imaging

Dissertation

zur Erlangung des mathematisch-naturwissenschaftlichen Doktorgrades

“Doctor rerum naturalium”

der Georg-August-Universität Göttingen

vorgelegt von

Kai Tobias Block

aus Mainz

Göttingen 2008

angefertigt in der

Biomedizinischen NMR Forschungs GmbH
am Max-Planck-Institut für biophysikalische Chemie

unter Betreuung durch das

Institut für Numerische und Angewandte Mathematik
der Georg-August-Universität Göttingen

D7

Referent: Prof. Dr. R. Schaback

Korreferent: Prof. Dr. J. Frahm

Tag der mündlichen Prüfung: 16.09.2008

Abstract

In magnetic resonance imaging (MRI), the received signal corresponds to the Fourier transform of the object's spin density, which can be sampled selectively by switching time varying gradient fields. Most conventional MRI techniques perform a line-by-line sampling of the Fourier transform, although in general arbitrary sampling schemes can be employed. In this thesis, the use of radial trajectories is discussed, which sample the Fourier transform along coinciding spokes. The imaging properties of this alternative acquisition technique are investigated in detail: while it offers several promising advantages over the conventional Cartesian approach, these salient properties are accompanied by a number of serious complications that so far hamper a more frequent use of radial techniques in routine applications.

One of the major problems consists in the higher complexity of the image reconstruction. In particular, existing approaches require 57% more Fourier lines relative to the Cartesian technique in order to sufficiently sample the data space. To overcome this limitation and the corresponding prolongation of the measurement time, this work describes the development of a novel method that allows for reasonable image reconstructions from a clearly reduced number of acquisitions. The reconstruction approach is based on the formulation as an inverse problem, which is solved with an iterative non-linear optimization technique. To compensate for the omitted acquisition steps, prior knowledge is incorporated using a set of penalty functions that prevent the selection of implausible solutions. The total variation of the estimate is employed as the main penalty term, which effectively suppresses radial undersampling artifacts for objects that are piecewise-constant. In addition, the approach exploits localized sensitivities of multi-coil arrays in a parallel imaging manner and copes with experimental phase variations to ensure the applicability in real MRI scenarios. The coil sensitivities are estimated from the same data in a preceding step, which avoids the recording of additional reference information. Global smoothness of the coil profiles is assumed and incorporated using a quadratic penalization of the pixel differences.

The effectiveness of the reconstruction method is demonstrated in studies of water phantoms and the human brain in vivo. The results reveal that the technique yields reasonable images from only a few k-space lines, which is not possible with current

Cartesian techniques. As a first real-world application, the performance of the new method is evaluated for the case of a single-shot MRI sequence based on stimulated echoes, which was adapted to the radial acquisition scheme. Because in single-shot MRI the number of obtainable k-space lines is limited, the combination with the proposed image reconstruction method results in a much higher spatial resolution than for the corresponding Cartesian case. Experimentally, a resolution enhancement up to a factor of four could be demonstrated for studies of the human brain. In a further step, the reconstruction method is extended to handle complex acquisition scenarios where data inconsistencies arise from spin relaxation or saturation effects. By employing a non-linear model of the received MRI signal, the extended approach is capable of merging such data in a reasonable way and enables exploiting respective acquisition techniques to reduce the overall measurement time. In the specific case of a radial fast spin-echo acquisition, the extended approach directly yields quantitative proton-density and relaxivity information, which offers motion-robust T_2 mapping from only a single radial data set.

Finally, it is shown that a minimization of the total variation can be exploited not only to fill k-space gaps in-between the measured lines but also to extrapolate the data beyond the measured part of k-space. As demonstrated for Cartesian sampling, this idea leads to a suppression of Gibbs ringing artifacts without noticeable loss of resolution. These results confirm that constraining the total variation is an effective concept to compensate for incomplete MRI acquisitions.

Zusammenfassung

Bei der Magnetresonanztomographie (MRT) wird eine räumliche Auflösung der induzierten Protonenresonanz durch magnetische Gradientenfelder erreicht, die die Resonanzfrequenz der angeregten Protonen mit ihrer räumlichen Position verknüpfen. Das Messsignal entspricht daher der Fourier-Transformierten der Protonendichte, die durch zeitabhängige Schaltung der Gradientenfelder selektiv abgetastet werden kann. Üblicherweise wird eine zeilenweise Abtastung der Fourier-Transformierten verwendet, jedoch können grundsätzlich auch andere Schemen eingesetzt werden. Diese Arbeit behandelt die bisher wenig verbreitete radiale Abtastung, bei der die Fourier-Transformierte entlang überlappender Speichen abgetastet wird. Im ersten Teil der Arbeit werden die Eigenschaften dieser Technik im Hinblick auf praktische Anwendungen untersucht. Es ergibt sich, dass die geänderte Abtastgeometrie teilweise deutlich vorteilhaftere Abbildungseigenschaften bietet - allerdings treten im Gegenzug verschiedene Komplikationen auf, die einen breiten Einsatz der radialen Abtastung behindern.

Ein Kernproblem besteht in der höheren Komplexität der Bildrekonstruktion. In den letzten Jahren haben sich zwar Standardtechniken zur Rekonstruktion aus radialen Daten etabliert, jedoch benötigen diese Methoden im Vergleich zur zeilenweisen Abtastung 57% mehr Linien im Fourier-Raum. Dies führt zu einer verlängerten Messzeit, die im Hinblick auf klinische Anwendungen einen klaren Nachteil darstellt. Daher wird im zweiten Teil der Arbeit eine neuartige Rekonstruktionstechnik vorgestellt, die eine Bildberechnung aus einer deutlich reduzierten Datenmenge ermöglicht. Die Methode basiert mathematisch auf der Formulierung der Bildrekonstruktion als ein Inverses Problem, das mit einer iterativen nicht-linearen Optimierungstechnik gelöst wird. Um fehlende Messdaten auszugleichen, wird A Priori Wissen in den Rekonstruktionsprozess einbezogen. Dies erfolgt in Form mehrerer Strafterme, die eine Berechnung unplausibler Lösungen verhindern. Als Hauptstrafe wird die Totale Variation (TV) der Bildintensität verwendet, die unter Annahme eines stückweise konstanten Objektes zu einer effektiven Unterdrückung von Unterabtastungsartefakten führt. Hierbei wird das mögliche Auftreten von Phasenvariationen innerhalb des Objektes berücksichtigt, um einen zuverlässigen Einsatz mit experimentellen MRT Daten zu gewährleisten. Bei Datensätzen, die mit mehreren Empfangsspulen aufgenommen

werden, erfolgt zudem eine Ausnutzung der lokalisierten Sensitivitätsprofile im Sinne der parallelen MRT Bildgebung. Hierfür werden die Spulenprofile direkt aus den Messdaten bestimmt, sodass keine zusätzliche Aufnahme von Referenzdaten erforderlich ist. Dabei wird angenommen, dass die gesuchten Sensitivitätsprofile glatte Funktionen sind, was mathematisch durch eine quadratische Bestrafung der Intensitätsdifferenzen benachbarter Pixel ausgedrückt wird.

Die Effektivität der Rekonstruktionsmethode wird in Simulationen sowie mit experimentellen Messdaten eines Wasserphantoms und des menschlichen Gehirns *in vivo* evaluiert. Die Ergebnisse zeigen, dass die Methode eine Bildberechnung aus nur wenigen Fourier-Linien ermöglicht, was mit einer Zeilenabtastung in vergleichbarer Weise bisher nicht erreichbar ist. Als erste konkrete Anwendung wird die ultraschnelle Bildgebung mit stimulierten Echos (STEAM) vorgestellt. Diese Aufnahmetechnik wird in der Arbeit auf die radiale Abtastung angepasst und optimiert. Da die Anzahl der messbaren Fourier-Linien bei der Technik physikalisch beschränkt ist, bietet die Kombination mit der vorgestellten Rekonstruktionsmethode eine deutliche Verbesserung der erreichbaren räumlichen Auflösung. Bei *in vivo* Untersuchungen am menschlichen Gehirn konnte experimentell eine Auflösungsverbesserung um einen Faktor von bis zu vier gezeigt werden. Im anschließenden Teil der Arbeit wird die Rekonstruktionsmethode für radiale Multi-Echo Aufnahmetechniken erweitert, bei denen erhebliche Dateninkonsistenzen durch ortsabhängige Relaxations- oder Sättigungseffekte der Magnetisierung auftreten. Durch Verwendung eines nicht-linearen Modells für das empfangene MRT Signal können diese Inkonsistenzen bei der Rekonstruktion berücksichtigt werden, sodass bisherige Bildartefakte vermieden und die Einsatzmöglichkeiten radialer Multi-Echo Sequenzen verbreitert werden. Im speziellen Fall der radialen schnellen Spin-Echo Aufnahmetechnik (FSE) liefert die Methode eine direkte Quantifizierung der lokalen Protonendichte und Relaxivität, womit sich ein effizientes Verfahren für die Bestimmung der transversalen Relaxationzeit mit niedriger Bewegungsempfindlichkeit ergibt.

Im letzten Teil der Arbeit wird gezeigt, dass eine Minimierung der Totalen Variation nicht nur dafür ausgenutzt werden kann, fehlende Speichen im Fourier-Raum auszugleichen, sondern auch dafür, die gemessenen Daten im Fourier-Raum zu extrapolieren. Es wird exemplarisch für die zeilenweise Abtastung vorgeführt, dass diese Idee zu einer effektiven Reduktion von Gibbs Artefakten führt, wobei keine sichtbare Verschlechterung der räumlichen Bildauflösung eintritt. Somit erweist sich die Beschränkung der Totalen Variation als leistungsfähiges Konzept um fehlende Messdaten bei der MRT Bildrekonstruktion zu kompensieren.

Contents

1	Introduction	1
2	Basics of Magnetic Resonance Imaging	4
2.1	The Pulsed NMR Experiment	4
2.1.1	Quantum Mechanical Description	5
2.1.2	Relaxation Effects	7
2.1.3	Conclusions for MRI	8
2.2	Signal Localization	8
2.2.1	Slice Selection	9
2.2.2	Spatial Encoding and k-Space Formalism	10
2.2.3	k-Space Sampling and Image Reconstruction	11
2.2.4	Sampling Requirements	13
2.3	Contrast Mechanisms	16
2.3.1	PD Contrast	16
2.3.2	T_2 Contrast	18
2.3.3	T_1 Contrast	19
3	Methodological Issues	20
3.1	MRI System	20
3.2	Simulation Principles	23
3.3	Software Framework	25
4	Radial Sampling of k-Space	27
4.1	Basic Sequence Design	27
4.2	Conventional Image Reconstruction from Radial Data	29
4.2.1	Projection Reconstruction	29
4.2.2	Regridding	31
4.2.3	Comparison of the Methods	35

4.3	Advantages and Disadvantages of Radial Sampling	37
4.3.1	Point-Spread-Function	37
4.3.2	Sample Distribution	40
4.3.3	Readout Oversampling	41
4.3.4	Motion Robustness	42
4.3.5	Gradient Deviation Sensitivity	43
4.3.6	Off-Resonance Sensitivity	44
4.3.7	Consistency Criteria	47
4.4	Variants of the Sampling Scheme	48
4.4.1	Multi-Echo Acquisitions	48
4.4.2	Partial Acquisitions and Ramp Sampling	49
4.4.3	3D Acquisitions	50
4.4.4	Propeller Trajectories	50
4.5	Summary	51
5	Iterative Reconstruction from Incomplete Radial Data	53
5.1	Reconstruction from Undersampled Data	53
5.2	Formulation as Inverse Problem	55
5.3	Incorporation of A Priori Knowledge	57
5.4	Experimental Complications	61
5.5	Proof-of-Principle Application	65
5.6	Experimental Results	68
5.7	Limitations and Extensions	71
5.8	Relation to Compressed Sensing	75
5.9	Summary	77
6	Rapid Imaging Using Stimulated Echoes	79
6.1	Single-Shot STEAM MRI	79
6.2	Combination with Radial Sampling	81
6.3	Experimental Results	85
6.4	Unresolved Issues	89
6.5	Segmented Data Acquisition	91
6.6	Summary	94

7	Efficient T_2 Mapping Based on Radial Fast Spin-Echo MRI	95
7.1	Radial Fast Spin-Echo Imaging	95
7.2	Extended Iterative Reconstruction	96
7.3	Experimental Results	101
7.4	Limitations and Extensions	107
7.5	Summary	109
8	Compensation of Truncation Artifacts for Cartesian Sampling	111
8.1	Cartesian Sampling and the Gibbs-Ringing Effect	111
8.2	TV-Constrained Data Extrapolation	114
8.3	Experimental Results	117
8.4	Limitations	119
8.5	Summary	124
9	Summary and Conclusions	125
A	Equivalence of Regridding and Filtered Backprojection	130
	Bibliography	135
	Abbreviations	146
	Curriculum Vitae	148
	List of Publications	149
	Acknowledgments	151

Chapter 1

Introduction

Magnetic resonance imaging (MRI) is a fascinating technique to image the human body non-invasively. It emerged from several fundamental inventions made in the 1970s and underwent a tremendous development during the last two decades. In fact, today's images have an astonishing quality, which is sometimes even comparable to that of a photograph. Because MRI offers not only high spatial resolution but also an excellent soft-tissue contrast, it is nowadays recognized as the leading modality for diagnostic imaging of numerous common diseases, in particular cancer and stroke.

As a clear advantage over other imaging techniques like computed tomography (CT) or positron emission tomography (PET), MRI uses neither ionizing radiation nor radioactive materials. Therefore, the examination procedure does not pose a hazard to the patient's health, which is especially valuable for patients receiving a long-term monitoring after therapy. Moreover, because the absence of any known side-effects allows for harmless experiments with healthy volunteers, MRI has found major applications in neuroscience and clinical research. The breakthrough of MRI was further supported by its powerful multi-contrast ability. Instead of mapping only a single physical property, as for example in x-ray imaging, MRI allows to observe a wide range of different parameters – ranging from proton density, diffusion, flow, and temperature up to complex quantities like tissue perfusion or brain activation.

On the other hand, the outstanding properties of the MRI technique are countered by a number of unresolved limitations. First, because a strong magnetic field is used, MRI is not suited for patients with a cardiac pacemaker as well as most patients with metallic implants. Second, the examinations are very expensive due to significant investment and operational costs of the MRI system. Third, the data acquisition is time-consuming, resulting in lengthy examinations relative to other imaging modalities. The latter causes substantial problems if the patient is unable to remain motionless during the measurements and, in particular, when imaging body sections with unavoidable

movements from blood circulation or breathing. For this reason, MRI studies of the abdomen and thorax are challenging. Similar arguments apply to measurements with low signal amplitude that require averaging of repeated acquisitions.

Consequently, most methodological developments of the past years aimed at an acceleration of the MRI acquisition speed. Considerable scan time reduction was achieved with the recent development of parallel imaging techniques, which exploit multiple localized receive coils to gain complementary information. Another promising concept consists in the use of alternative sampling schemes, which is referred to as *non-Cartesian* MRI. This idea is attracting growing interest, although a non-Cartesian sampling technique has already been proposed 35 years ago by Lauterbur in his seminal paper on MRI [1]. However, at that time the technique was soon replaced by the *Fourier encoding* or *Cartesian sampling* [2] because of a lower sensitivity to the strong inaccuracies of the early MRI systems.

In this thesis, the sampling method used by Lauterbur is revisited, which is today known as *radial sampling* or *radial encoding*. Although radial sampling did not find a wide application in the past, it offers a number of interesting advantages over the Cartesian technique, which result from the specific geometry of the sampling pattern. In the first part of the work, these properties are analyzed in detail along with the inherent disadvantages of the technique. One major complication arises from the non-trivial problem to reconstruct images from radially sampled data. In particular, conventional approaches require an inefficiently high amount of measured data, rendering the technique less attractive for potential applications. This aspect is addressed in the second part, which develops a novel approach based on an inverse problem formulation that compensates for missing data by incorporation of prior knowledge. It is shown that this concept can be exploited to obtain reasonable images in a significantly reduced measurement time, which is not achievable with Cartesian sampling at all. In the final part, further applications and extensions of the reconstruction approach are presented, which demonstrate that the underlying concept is also utilizable in more complex imaging scenarios.

During the development phase of these methods, numerous simulations were conducted to analyze their performance and to identify potential sources of artifacts. However, in order to rely not only on simulated data, respective acquisition techniques were also implemented on a standard MRI system. This allowed for an evaluation of the proposed approaches with real data from phantom and volunteer studies. Although the experimental results were always in close agreement with predictions from the simulations, it proved that the methods are indeed applicable in realistic imaging situations. In summary, the work shows that radial sampling has a high potential for various applications when using modern MRI hardware together with dedicated reconstruction methods.

Organization of Thesis

The thesis is organized as follows. Chapter 2 gives a brief introduction to the principles of MRI. Chapter 3 discusses common methodological aspects of the work presented in Chapters 4 - 8. Chapter 4 introduces the radial sampling scheme and analyses main properties, including consequences for the image reconstruction. Chapter 5 presents a novel approach for the reconstruction from incomplete radial data, which is the primary contribution of the thesis as it enables to utilize radial sampling in a much wider range of applications. In Chapter 6, this is demonstrated for the single-shot STEAM MRI technique, yielding a significant resolution enhancement due to the combination with radial sampling. Chapter 7 presents an extension of the approach from Chapter 5, which incorporates the acquisition time to handle data from multi-echo sequences and offers an implicit quantification of the local relaxation rate. Chapter 8 shows that the problem of incomplete radial data is closely related to the truncation effects arising in Cartesian sampling, and a compensation method based on the concepts of Chapter 5 is presented. Finally, Chapter 9 summarizes the main achievements of the thesis and gives an outlook for future work.

Chapter 2

Basics of Magnetic Resonance Imaging

This chapter gives a brief overview of the principles of magnetic resonance imaging from a signal processing perspective. The discussion is confined to main aspects that are relevant to the following chapters of the thesis. For more comprehensive information, the reader is referred to textbooks by Haacke et al. [3], Liang et al. [4], and Bernstein et al. [5].

2.1 The Pulsed NMR Experiment

The basic principle underlying magnetic resonance imaging is a measurement of signals induced by the nuclear magnetic resonance (NMR) effect. The NMR phenomenon was initially discovered by Purcell and Bloch in 1946 using a technique which is today known as continuous-wave NMR [6, 7]. In 1950, Hahn demonstrated that the NMR effect can also be observed with a modified experimental method based on finite radio frequency pulses [8]. The concept of this pulsed NMR experiment is exploited in the MRI technique to obtain a signal from the sample.

Figure 2.1 shows a schematic illustration of the experimental apparatus. The NMR effect can only be observed for nuclei with nonzero spin quantum number such as ^1H , ^{13}C , ^{19}F , or ^{31}P . These particles have a magnetic moment which is related to their angular momentum by

$$\boldsymbol{\mu} = \gamma \boldsymbol{J} , \quad (2.1)$$

where the gyromagnetic ratio γ is a nucleus-dependent property. The dominant nucleus in MRI applications is the proton in hydrogen, and most MRI techniques focus only on resonance effects of protons. To measure nuclear resonance, the sample is placed in a

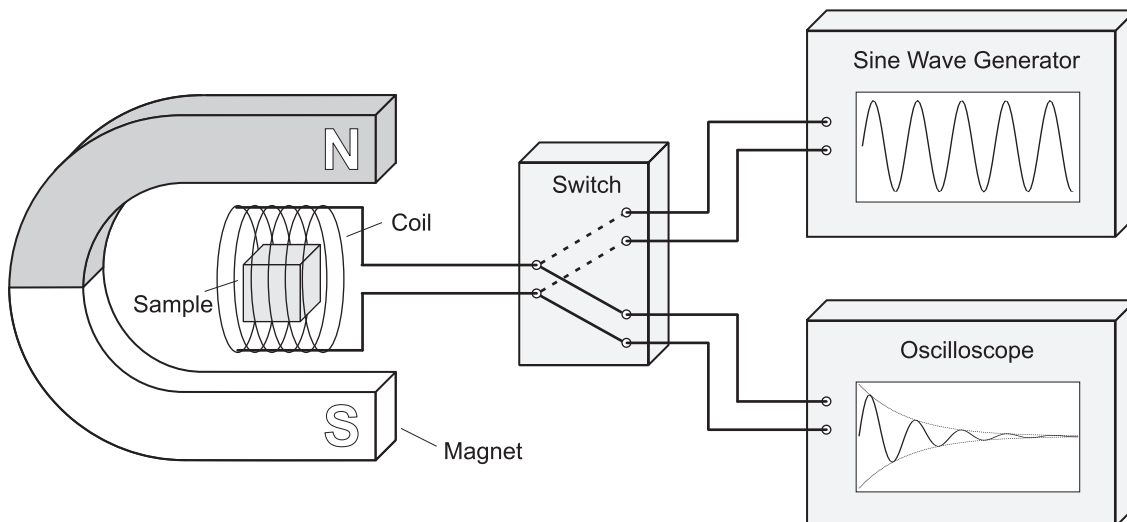


Figure 2.1: Schematic illustration of the apparatus for the pulsed NMR experiment.

static magnetic field which is aligned in z -direction $\mathbf{B}_0 = B_0 \mathbf{e}_z$. For a short duration, an alternating electromagnetic field is applied perpendicular to the static field. The electromagnetic field is generated with an inductive coil surrounding the sample, which is driven by an oscillator with an angular frequency matching the *Larmor frequency*

$$\omega_0 = \gamma B_0 . \quad (2.2)$$

For protons exposed to a magnetic field strength of $B_0 = 2.89 \text{ T}$, which was used in the present work, it corresponds to an angular frequency of $\omega_0 = 7.74 \cdot 10^8 \text{ rad/s}$ or, respectively, $f_0 = \omega_0/2\pi = 123.2 \text{ MHz}$. Because this frequency belongs to the spectrum used for radio transmission, the temporary generation of the electromagnetic field is called a radio frequency (RF) pulse.

Immediately after application of the RF pulse, a small voltage can be detected at the terminals of the induction coil. The received signal oscillates with the same Larmor frequency ω_0 , and its envelope decays exponentially in the order of milliseconds. It is induced by the magnetic moment of resonating protons as a consequence of the perturbation or *excitation* with the RF pulse.

2.1.1 Quantum Mechanical Description

The description of the NMR phenomenon requires a quantum mechanical analysis as the effect originates from nuclear spin interaction with the magnetic field. Because the derivation of the NMR theory is described extensively in the literature [9, 10], only the main results are discussed here.

The interaction with the static magnetic field produces an energy $-\boldsymbol{\mu} \cdot \mathbf{B}$, which leads

to a discrete number of eigenvalues of the Hamiltonian. For a proton with spin quantum number $j = \frac{1}{2}$, these energy levels are

$$E_{\uparrow} = -\gamma \frac{\hbar}{2} B_0 \qquad E_{\downarrow} = \gamma \frac{\hbar}{2} B_0 . \quad (2.3)$$

The difference between both states corresponds to the energy of an oscillation at Larmor frequency

$$\Delta E = \gamma \hbar B_0 = \hbar \omega_0 \quad (2.4)$$

and explains why the resonance phenomenon occurs only at this discrete frequency.

Analysis of the expectation value of the magnetic moment $\langle \boldsymbol{\mu} \rangle$ reveals that the z-component remains fixed in the static field, while the transversal component rotates in the xy-plane at angular frequency ω_0 . Hence, the expectation vector acts like a precessing gyroscope. If an additional time-varying field with frequency ω_0 is created in the xy-plane

$$\mathbf{B} = \mathbf{B}_0 + \mathbf{B}_1(t) = \mathbf{B}_0 + \sin(\omega_0 t) \mathbf{e}_x + \cos(\omega_0 t) \mathbf{e}_y , \quad (2.5)$$

then the expectation vector $\langle \boldsymbol{\mu} \rangle$ is tilted with respect to the z-axis. In the NMR experiment, this additional field is generated by the RF pulse. Derivation of the expectation vector with respect to time yields

$$\frac{d\langle \boldsymbol{\mu} \rangle}{dt} = \langle \boldsymbol{\mu} \rangle \times \gamma \mathbf{B} , \quad (2.6)$$

which obeys the classical equation of motion of a gyroscope and holds true for the static and time-dependent magnetic field.

Equation (2.6) shows that the expectation vector for a single proton can take an arbitrary orientation – regardless of the spin quantization. However, because the expectation value is of a statistical nature, observing the value would require a high number of measurements. On the other hand, in practice a high number of protons is excited at the same time. Therefore, it is convenient to introduce the *bulk magnetization* $\mathbf{M} = \sum_i \boldsymbol{\mu}_i$, which sums over all magnetic moments within a macroscopic volume of the sample. Assuming that the protons do not interact, Eq. (2.6) is also valid for the bulk magnetization. Since many “independent” protons are observed simultaneously in the experiment, the measured bulk magnetization corresponds to the expectation value of the magnetic moment of a single proton. This justifies to describe the NMR experiment based on Eq. (2.6).

Finally, although the protons tend to reach the lower energy state E_{\uparrow} when exposing the sample to the static field, in practice both states E_{\uparrow} and E_{\downarrow} are occupied to some extent

due to energy absorption from thermal contact. According to statistical physics, the population follows a Boltzmann distribution. In thermal equilibrium, the probability of finding a proton in either state is given by the Boltzmann factors

$$p(E_{\uparrow}) = \frac{e^{-E_{\uparrow}/kT}}{e^{-E_{\uparrow}/kT} + e^{-E_{\downarrow}/kT}} \quad p(E_{\downarrow}) = \frac{e^{-E_{\downarrow}/kT}}{e^{-E_{\uparrow}/kT} + e^{-E_{\downarrow}/kT}}, \quad (2.7)$$

where k is the Boltzmann constant and T is the temperature. The Boltzmann factors can be used to derive the population difference for a macroscopic volume with proton density ρ . It yields that the bulk magnetization points along the positive z -direction in thermal equilibrium $\mathbf{M} = M_0 \mathbf{e}_z$, while the magnitude is given by

$$M_0 = \rho \frac{\gamma^2 \hbar^2}{4kT} B_0. \quad (2.8)$$

2.1.2 Relaxation Effects

The signal detected in the pulsed NMR experiment declines rapidly after the RF excitation. Early NMR experiments showed that there are actually two different relaxation mechanisms affecting the transverse bulk magnetization [11]. The first mechanism is called *spin-lattice* or *longitudinal relaxation* and describes an exponential recovery of the longitudinal magnetization M_z after a preceding excitation. The effect is attributed to energy exchange between the protons and their environment, returning the excited system to its thermal equilibrium state.

The second mechanism, the *spin-spin* or *transverse relaxation*, corresponds to a dephasing of the moments inside a macroscopic volume. Because the bulk magnetization averages over all moments in the volume, its amplitude decays from destructive interference. The effect originates from frequency fluctuations caused by proton interactions.

To account for these two mechanisms, Bloch extended Eq. (2.6) by respective relaxation terms, yielding the *Bloch equation*

$$\frac{d\mathbf{M}}{dt} = \mathbf{M} \times \gamma \mathbf{B} + \begin{pmatrix} -M_x / T_2 \\ -M_y / T_2 \\ (M_0 - M_z) / T_1 \end{pmatrix}, \quad (2.9)$$

which allows to describe the evolution of the magnetization observed in NMR experiments [11]. In contrast to Eq. (2.6), it is written for the bulk magnetization as the relaxation terms are entirely based on empirical findings, without considering the physical mechanisms on a detailed level. The recovery rate of the longitudinal magnetization is given by the T_1 relaxation time, while the transversal relaxation is characterized by the T_2 relaxation time. Both time constants are properties of the individual sample

material. The T_2 relaxation time is always shorter than the T_1 relaxation time, and, therefore, the detected signal decays with the T_2 relaxation time. However, when performing fast repetitive excitations with incomplete recovery of the longitudinal magnetization, the signal amplitude becomes dependent on the T_1 relaxation time. Therefore, both relaxation times can be estimated from NMR experiments.

2.1.3 Conclusions for MRI

The analysis of spectroscopic NMR experiments usually requires detailed knowledge of nuclear spin physics and spin interactions. In contrast, most NMR-based imaging techniques can be properly explained with only the following four statements that summarize the results from the previous sections.

- (i) The bulk magnetization of the protons aligns in the positive z -direction when inserting the sample to the magnetic field.
- (ii) It can be tipped to the xy -plane using a RF pulse at Larmor frequency ω_0 , generated by a coil perpendicular to the static field. The flip angle is determined by the amplitude and duration of the RF pulse.
- (iii) The tipped magnetization acts like a gyroscope and precesses at Larmor frequency ω_0 in the xy -plane. It induces an alternating voltage in the coil proportional to its transversal component.
- (iv) The transversal component decays exponentially with the T_2 relaxation time, while the longitudinal component relaxes with the T_1 relaxation time.

2.2 Signal Localization

The basic NMR experiment enables to receive a signal from the sample, but it is impossible to spatially assign the resonating protons. However, spatial information can be obtained with the use of additional magnetic gradient fields. These gradient fields can be switched separately from each other in the x -, y - and z -direction

$$\mathbf{B}_G(\mathbf{x}, t) = (\mathbf{G}(t) \cdot \mathbf{x}) \mathbf{e}_z = (G_x(t) \cdot x + G_y(t) \cdot y + G_z(t) \cdot z) \mathbf{e}_z, \quad (2.10)$$

where $\mathbf{x} = (x, y, z)$ is a position with respect to the isocenter of the magnet, and G_x denotes the slope of the gradient field in x -direction. It should be noted that the orientation of the gradient fields is always parallel to the static field. Therefore, the gradients only alter the field strength experienced at a position \mathbf{x} . However, as the Larmor frequency ω_0 is proportional to the field strength, it becomes dependent on

the location when a gradient is switched. In other words, the gradient fields allow to link the Larmor frequency to the spatial location inside the magnet. This key idea was introduced in 1973 by Lauterbur [1], and two complementary concepts emerged to exploit the gradient fields for signal localization.

2.2.1 Slice Selection

Due to the quantization of the energy levels, protons can only be excited if the frequency of the RF pulse coincides with the Larmor frequency ω_0 . For a given RF pulse with frequency ω , it requires that the protons experience a field strength of $B = \omega/\gamma$. Therefore, it is possible to spatially limit the excitation by switching a gradient field during the RF pulse. If, for example, the z-gradient G_z is switched, then the resonance condition is solely fulfilled at the position

$$z = (\omega/\gamma - B_0)/G_z. \quad (2.11)$$

Consequently, only a single slice of the sample will be excited, reducing the localization problem from three to two dimensions. This concept is called *slice excitation* or *slice selection*.

In practice, however, finite generation of a single frequency is infeasible. In order to get a suitable waveform with limited support, it is necessary to multiply the desired carrier frequency ω with a compact window function. Of course, multiplication with a compact window is accompanied by excitation of a respective slice profile, centered around the position given by Eq. (2.11). Because the sinc-function has a rect-shaped frequency spectrum, common waveforms are based on a truncated version with additional filtering.

For sinc-based windows, the width of the slice profile is reciprocal to the width of the sinc function. Thus, the excitation of a thin slice requires a broad sinc function. However, because the extent of the RF pulse window is limited, more power of the sinc function will be truncated if the sinc is broadened. This causes increased excitation of side lobes, which degrade the quality of the slice profile. Therefore, the applicability of the slice selection technique is limited when very thin slices are needed. On the other hand, it is important to note that a slice selective excitation allows for interleaved acquisitions. Because the excitation is spatially limited, data from different slices can be acquired during the recovery phase of a previously excited slice, while for a non-selective excitation it is necessary to pause the measurements until the longitudinal magnetization is fully recovered. This offers a significant reduction of the overall measurement time, in particular when high flip angles are used.

2.2.2 Spatial Encoding and k-Space Formalism

The second concept for signal localization, which is referred to as *spatial encoding*, is based on the generation of gradient fields after the RF excitation pulse, i.e. during the precession of the magnetization. Because the protons experience a change of the field strength, it causes a shift of the precession frequency $\Delta\omega$ according to the spatial position inside the magnet.

If no gradients are switched, the rotation of the transverse magnetization can be described in the complex plane with

$$M_{\perp}(\mathbf{x}, t) = M_0(\mathbf{x}) \cdot e^{i\omega_0 t} , \quad (2.12)$$

where relaxation effects are ignored for simplicity. Here, the magnitude M_0 depends on the position \mathbf{x} because most samples have a spatially varying proton density $\rho(\mathbf{x})$. To account for frequency changes caused by gradient fields, an additional phase term is needed that depends on the time and location

$$M_{\perp}(\mathbf{x}, t) = M_0(\mathbf{x}) \cdot e^{i(\omega_0 t + \varphi(\mathbf{x}, t))} = M_0(\mathbf{x}) \cdot e^{i\omega_0 t} \cdot e^{i\varphi(\mathbf{x}, t)} . \quad (2.13)$$

The phase φ can be derived from the time course of the switched gradients

$$\begin{aligned} \varphi(\mathbf{x}, t) &= \int_0^t \Delta\omega(\mathbf{x}, \tau) d\tau = \int_0^t \gamma \Delta B(\mathbf{x}, \tau) d\tau = \int_0^t \gamma \mathbf{G}(\tau) \cdot \mathbf{x} d\tau \\ &= \gamma \int_0^t \mathbf{G}(\tau) d\tau \cdot \mathbf{x} = 2\pi \cdot \mathbf{k}(t) \cdot \mathbf{x} , \end{aligned} \quad (2.14)$$

where $\mathbf{k}(t)$ is the *k-space position* that is introduced with $\varphi = \gamma/2\pi$ for convenience

$$\mathbf{k}(t) = \varphi \int_0^t \mathbf{G}(\tau) d\tau . \quad (2.15)$$

Combination of Eq. (2.13) and (2.14) yields

$$M_{\perp}(\mathbf{x}, t) = M_0(\mathbf{x}) \cdot e^{i\omega_0 t} \cdot e^{2\pi i \mathbf{k}(t) \cdot \mathbf{x}} . \quad (2.16)$$

All excited protons contribute to the voltage induced in the coil $U(t)$, which is, therefore,

proportional to the total transverse magnetization

$$\begin{aligned} U(t) &= c \cdot \int M_{\perp}(\mathbf{x}, t) d\mathbf{x} = c \cdot \int M_0(\mathbf{x}) \cdot e^{i\omega_0 t} \cdot e^{2\pi i \mathbf{k}(t) \cdot \mathbf{x}} d\mathbf{x} \\ &= \hat{c} \cdot e^{i\omega_0 t} \cdot \int \rho(\mathbf{x}) \cdot e^{2\pi i \mathbf{k}(t) \cdot \mathbf{x}} d\mathbf{x} . \end{aligned} \quad (2.17)$$

Here, the proton density $\rho(\mathbf{x})$ was separated from the magnetization magnitude M_0 (see Eq. (2.8)), while the other components were assumed to be spatially invariant and included into the constant \hat{c} . Because the Larmor frequency ω_0 is known, it can be demodulated from the voltage, and the received signal is given by

$$S(t) = e^{-i\omega_0 t} \cdot U(t) = \hat{c} \cdot \int \rho(\mathbf{x}) \cdot e^{2\pi i \mathbf{k}(t) \cdot \mathbf{x}} d\mathbf{x} . \quad (2.18)$$

This equation reveals that the MRI signal corresponds to the Fourier transform of the proton density $\rho(\mathbf{x})$. In other words, the spatial information of the object is encoded as Fourier transform, which can be selectively sampled using time varying gradient fields. Therefore, MRI is in essence Fourier imaging, and most MR image characteristics can be deduced from the mathematical properties of the Fourier transformation. The coordinate in the Fourier space, i.e. the information where the Fourier transform is sampled at time t , is given by the k-space position $\mathbf{k}(t)$ as defined in Eq. (2.15). For a non-selective excitation, this Fourier space has three dimensions. However, when the spatial encoding is combined with a slice selective excitation, the Fourier transform is reduced to a two dimensional function. Unless stated otherwise, the latter is assumed for simplicity in the rest of the thesis.

In practice, the signal is additionally affected by relaxation effects as described in Section 2.1.2 and, thus, accurate description of the received signal requires an extension of Eq. (2.18) by respective relaxation terms. For most objects, the T_1 and T_2 relaxation times are spatially dependent. Therefore, the relaxation terms have to be included into the integral, and the received signal actually corresponds to the Fourier transform of the modulated spin-density.

2.2.3 k-Space Sampling and Image Reconstruction

It can be seen from Eq. (2.14) that the position in Fourier space at time t depends on all gradient fields switched after the RF excitation. Hence, for a given time-course of gradient fields, Eq. (2.15) can be used to calculate a *k-space trajectory*, i.e. the path along which the object's Fourier transform (2.18) is sampled during the experiment. To ensure that a later reconstruction of the object is possible, it is required to switch the gradient fields in such a way that sufficient Fourier data is collected along the trajectory.




Figure 2.2: (Top) Timing diagram of a generic FLASH sequence and (bottom) the respective sampling trajectory (bold line corresponds to one cycle of the sequence).

In the conventional case, the k-space is sampled with a line-by-line scheme, which is often referred to as *Cartesian sampling scheme*. For instance, in the FLASH method, one of the widely used MRI techniques, a single row of k-space is sampled after one RF excitation, while for each repetition a different row is acquired [12]. Hence, a total of, for example, 64 repetitions is required for an image with a size of 64×64 pixels. Figure 2.2 shows the k-space trajectory of the FLASH sequence together with the sequence timing diagram, which indicates the required time-course of gradient fields. The sequence starts with a slice-selective RF excitation in the z-direction, which includes a *rewinder gradient* to compensate for the undesired phase evolution caused by the *slice-selection gradient* during the second half of the excitation pulse. Because the excitation is spatially limited in the z-direction, the received signal corresponds to the two-dimensional Fourier transform of the slice in the xy-plane. Directly after the RF excitation, the phase of all proton spins is equal, which relates to the center position

in k-space. The magnetization is then dephased with *prephasing gradients* in the x- and y-direction, causing a shift to a peripheral position in k-space. Starting from this position, a single row of the k-space is sampled by switching the x-gradient with opposed amplitude, which yields a constant movement in the positive x-direction of k-space. During the flat-top time of this *readout gradient*, the signal is recorded at a fixed sampling rate. As a last step of the sequence, the remaining magnetization is strongly dephased with *spoiler gradients* to vanish any residual signal. For each repetition, the amplitude of the prephasing gradient in y-direction is varied so that different rows with respect to the y-direction are sampled. It should be noted that it is also possible to use a different assignment of gradient directions, for example slice-selection in the y-direction and readout gradient in the z-direction. In fact, any linear combination of gradients can be used, which allows to obtain arbitrary slice orientations.

When the acquisition of all rows is finished, the object can be reconstructed from the samples by inverting the Fourier encoding. A major advantage of the Cartesian sampling scheme is that all sample points lie directly on an equidistant grid. Therefore, it is possible to perform a simple reconstruction with the use of an inverse fast Fourier transformation (FFT). This is achieved by arranging the samples on a *raw data* matrix with respect to their k-space position and calculating a FFT of the matrix in each dimension. In the two dimensional case, the FFT has to be performed for each row of the matrix and then for each column, which yields an image of the object with identical extent, i.e. same number of entries as the raw data matrix. Because the received signal is complex (see Eq. (2.18)), the image matrix has complex entries, and usually the magnitude values are presented to the viewer.

2.2.4 Sampling Requirements

Because the spatial object information is encoded with the Fourier transformation, several requirements for the trajectory design result from the mathematical properties of the Fourier transformation. These mathematical properties are also fundamental for the analysis of image artifacts and can be summarized as follows.

(i) The Fourier transformation is a linear operation

$$\mathcal{F}\{a \cdot g(x) + b \cdot h(x)\} = a \cdot \mathcal{F}\{g(x)\} + b \cdot \mathcal{F}\{h(x)\}$$

where $a, b \in \mathbb{C}$, $x \in \mathbb{R}^n$, and $g, h : \mathbb{R}^n \rightarrow \mathbb{C}$.

(ii) Multiplication of two functions in the image space causes a convolution of their transforms in the Fourier space

$$\mathcal{F}\{g(x) \cdot h(x)\} = \mathcal{F}\{g(x)\} * \mathcal{F}\{h(x)\} ,$$

and vice versa

$$\mathcal{F}\{g(x) * h(x)\} = \mathcal{F}\{g(x)\} \cdot \mathcal{F}\{h(x)\} .$$

(iii) Translation of a function in the image space causes a linear phase modulation in the Fourier space and vice versa

$$\begin{aligned} \mathcal{F}\{g(x + c)\}(\omega) &= \mathcal{F}\{g(x)\}(\omega) \cdot e^{ic\omega} \\ \mathcal{F}\{g(x) \cdot e^{id x}\}(\omega) &= \mathcal{F}\{g(x)\}(\omega - d) , \end{aligned}$$

where $c, d, \omega \in \mathbb{R}^n$.

(iv) Scaling of a function with respect to the variable yields a reciprocally scaled Fourier transform

$$\mathcal{F}\{g(u \cdot x)\}(\omega) = \frac{1}{u} \cdot \mathcal{F}\{g(x)\} \left(\frac{\omega}{u} \right) ,$$

where $u \in \mathbb{R}$.

To understand the imaging properties of MRI techniques, it is always very helpful to decompose the complete data processing into these building blocks. During the data acquisition, the continuous Fourier transform of the object is sampled discretely at a certain sampling rate. The sampling can be seen as multiplication of the Fourier transform with a comb- or shah-function. According to property (ii), in image space this corresponds to a convolution with the Fourier transform of the comb-function, which is a comb-function with reciprocal interval width. Hence, the discrete sampling leads to occurrence of periodic object copies, where the distance between the copies is reciprocal to the sample distance in k-space. If the sample distance is chosen too large, then neighboring copies overlap in the image space, which makes it impossible to recover the object properly.

Therefore, it is convenient to introduce the field of view (FOV), which describes the extent of the object to be imaged, and to select the sample distance Δk with respect to a given FOV value

$$\Delta k = \frac{1}{\text{FOV}} . \quad (2.19)$$

One way to adjust the distance between the measured samples is to change the temporal sampling rate of the receiver. Hence, the higher the sampling rate is, the smaller is the sample distance in k-space. Further, it can be seen from Eq. (2.15) that the speed of the movement in k-space depends on the amplitude of the gradient fields. Thus, for a fixed sampling rate, the sample distance can also be influenced by changing the gradient strength. In other words, the sample distance can be reduced either by increasing the sampling rate or by lowering the gradient strength. While the latter prolongs the total duration of the sampling, it offers an increased signal-to-noise ratio (SNR). This is because a reduced spectrum of Larmor frequencies is used for the spatial encoding

when lowering the gradient strength, while the noise floor can be assumed to be equally distributed in the frequency space. For this reason, it is useful to introduce the *readout bandwidth* or *receiver bandwidth*

$$\text{BW} = \frac{1}{\Delta t}, \quad (2.20)$$

where Δt denotes the sampling interval or *dwelltime* of the receiver. Often, the value is normalized by the number of acquired samples n , yielding the *bandwidth per pixel*

$$\text{BW/pixel} = \frac{1}{\Delta t \cdot n} = \frac{1}{T_s}, \quad (2.21)$$

where T_s is the total duration of the sampling period. Basically, a high value of the bandwidth corresponds to a short acquisition with low SNR, while a low value indicates a long readout time with high SNR.

In practice, usually the FOV, the number of samples, and the bandwidth per pixel are given by the operator, while the other parameters are deduced. The number of samples n is also referred to as *base resolution* because it is directly related to the number of image pixels. The spatial resolution of the reconstructed image is then simply given by

$$\Delta x = \frac{\text{FOV}}{n}. \quad (2.22)$$

Of course, the resolution has to be calculated separately for each image direction if the number of samples or the FOV size is not equal for all directions. The required dwelltime of the receiver follows from

$$\Delta t = \frac{1}{\text{BW/pixel} \cdot n}, \quad (2.23)$$

and the amplitude of the readout gradient is given by

$$G_{\text{read}} = \frac{1}{\gamma \cdot \text{FOV} \cdot \Delta t}. \quad (2.24)$$

Finally, the prephasing gradient in the x-direction has to be selected in such a way that the data recording starts at the position

$$k_{\text{max}} = -\Delta k \cdot \frac{n}{2}. \quad (2.25)$$

Further, the prephasing gradient in the y-direction has to be switched such that the out-most position is equal to k_{max} and that a distance of Δk is ensured between neighboring rows.




Figure 2.3: Different image contrasts obtained for a study of the human brain. (Left) Proton density weighted contrast, (middle) T_2 weighted contrast, and (right) T_1 weighted contrast.

2.3 Contrast Mechanisms

As discussed in Section 2.1.2 on the relaxation mechanisms, three quantities have main influence on the signal amplitude: the proton density (PD), the T_1 relaxation time, and the T_2 relaxation time. In 1971, several years before the development of MR imaging techniques, Damadian discovered that the T_1 and T_2 values of tumor tissue differ significantly from those of normal tissue [13]. In fact, many diseases cause a local change of at least one of the three quantities, and common diagnostic procedures involve a comparison of values from suspicious tissue with surrounding regions. Therefore, different measurement sequences have been developed to create a specific weighting of the image contrast, which allows for a local estimation of the aforementioned quantities. Examples of such contrast weightings can be seen in Figure 2.3. In the following, the basic concepts for different contrast weightings are discussed. However, it should be noted that a large number of other and more advanced methods exists [5].

2.3.1 PD Contrast

Because the magnitude of the magnetization is proportional to the local proton density, regions with only few protons induce a low signal, while a strong signal is received from regions with a high proton density. Thus, every MR image is subject to a basic weighting with the proton density, which corresponds mainly to the water content of the tissue. However, to obtain a pure PD contrast, the data has to be acquired in a way that any additional signal modulation from relaxation or saturation mechanisms is avoided. As illustrated in Figure 2.4, this is achieved by minimizing the duration between the RF excitation and the data readout, while maintaining a rather long repetition delay to allow for a homogeneous recovery of the longitudinal magnetization. The readout delay




PD Weighting**T2 Weighting****T1 Weighting**

Figure 2.4: Sequence timing diagrams for (top) PD weighted, (middle) T_2 weighted, and (bottom) T_1 weighted image contrast (time axis is not to scale).

is usually specified by the *echo time* TE, which is formally defined as time distance between the middle of the excitation pulse and the sampling of k-space center. The delay between two repetitions of the sequence is described using the *repetition time* TR, which is formally defined as the time distance between subsequent RF excitation pulses. Hence, to obtain a PD weighting, the echo time should be set to the minimum value achievable, and a large TR value should be selected to ensure that maximum magnetization is available for each RF excitation.

In practice, however, it is often infeasible or highly inefficient to wait for a complete return of the longitudinal magnetization, and a shortened repetition time is used. In this case, the image contrast gets contaminated by a slight weighting with the T_1 relaxation time, which particularly affects regions with a slow recovery of the longitudinal magnetization like, for example, the cerebrospinal fluid (CSF). Moreover, because the RF excitation and the data readout itself need some time, the minimum echo time is limited, and the images exhibit also a minor weighting with the T_2 relaxation time. Thus, the proton density can only be approximately determined from an image with PD contrast, which, however, is sufficient for many diagnostic applications.

2.3.2 T_2 Contrast

A contrast weighting with the T_2 relaxation time can be obtained by increasing the echo time so that an explicit delay is introduced between the excitation and the signal detection. In this time period, the signal decays due to spin-spin relaxation, where the strength of the signal loss depends on the local T_2 relaxation time. Hence, the signal declines particularly in regions with a small T_2 value, while more signal remains in regions with a large T_2 value. This creates a modulation of the proton density with the T_2 relaxation time.

However, when employing such strategy in practice to estimate the relaxation times, the observed values are often much smaller than the T_2 values expected in theory. This effect is caused by local deviations of the effective field strength, which originate from unavoidable inhomogeneities of the static magnetic field as well as from susceptibility differences that arise, for instance, at air-tissue-interfaces. Thus, some of the protons precess at a slightly different frequency, which manifests as an additional phase evolution of these protons relative to the assumed resonance frequency and is referred to as *off-resonance* effect. If the phase differences become significant within some volume or voxel, local signal cancellation occurs from the destructive interference of the magnetic moments. Therefore, the *effective* T_2^* relaxation time has been introduced, which describes the experimentally observed signal decay and accounts for both, the dephasing from spin-spin interactions as well as the dephasing caused by off-resonance effects. Although both mechanisms originate from a loss of the phase coherence, a major difference is that the spin-spin interactions evoke irreversible phase fluctuations, which are only determinable in a statistical sense, while the off-resonances lead to a systematic linear dephasing. The latter can be revoked with the use of the *spin refocusing* technique or *spin echo* technique.

In the spin echo technique, a RF pulse with a flip angle of 180° is applied at the time $TE/2$ after the RF excitation pulse, as shown in Figure 2.4. The 180° refocusing pulse causes a spontaneous inversion of the rotating direction of the magnetic moments. Therefore, the undesired phase dispersion, which develops during the first $TE/2$ period, is rewound, and at the time $TE/2$ after the refocusing pulse all magnetic moments are in-phase again. Hence, at the time distance TE after the excitation pulse the signal amplitude depends only on the proton density and the true T_2 relaxation time, while dephasings from the off-resonance effects are canceled. If the sampling of k-space center is synchronized with this time point, then the image contrast reflects the true T_2 values. Noteworthy, the inversion of the rotating direction corresponds to a reflection of the k-space position, i.e. a change of the coordinate sign. It is, therefore, necessary to switch the readout gradient in the same direction as the prephasing gradient, which differs from the gradient switching shown in Figure 2.2.

The T_2 weighted image contrast is very sensitive to malignant changes of the tissue and often used for cancer detection. However, several important applications exist also for the T_2^* contrast. In particular, the functional MRI technique (fMRI) is based on this contrast as brain activation leads to a local change of the hemoglobin concentration, which causes a slight spin dephasing and becomes visible in T_2^* weighted images.

2.3.3 T_1 Contrast

Because the spin-spin relaxation is always faster than the spin-lattice relaxation, the MRI signal decays with the T_2 or T_2^* relaxation time, which yields no information about the T_1 relaxation time. However, when performing fast repetitive excitations with incomplete recovery of the longitudinal magnetization, then the signal amplitude becomes dependent on the T_1 relaxation. If the T_1 relaxation time in a region is long, only a small amount of the longitudinal magnetization is recovered at the end of the sequence and available for a next excitation. Consequently, this region creates a weak signal in the following run of the sequence, which is referred to as *spin saturation* effect. In contrast, a region with a small T_1 value induces a strong signal, because here a high amount of longitudinal magnetization is employed in the subsequent excitations. Thus, a modulation of the image contrast with the T_1 relaxation time is obtained by using a small TE and a small TR value, as illustrated in Figure 2.4. Further, the T_1 weighting can be amplified by incrementing the flip angle of the excitation pulse, which prolongs the time needed for a complete recovery of the longitudinal magnetization and leads to a higher degree of saturation. Since full longitudinal magnetization is available on the first run of the sequence, the initial signal is still independent from the T_1 value and only PD weighted. Therefore, it is necessary to drive the magnetization into a steady-state before recording data, which is done with several preparation excitations.

The T_1 contrast has an inverse character in comparison to the T_2 contrast, because a slow T_1 relaxation gives a weak signal, while a fast relaxation gives a strong signal. For instance, the CSF, which has a long T_1 and a long T_2 relaxation time, appears dark in the T_1 weighted contrast and bright in the T_2 weighted contrast (see Figure 2.3). The T_1 contrast mechanism allows to obtain images in a relatively short time because it is not required to pause the measurement for a recovery of the magnetization. Therefore, it is often utilized for anatomical measurements and provides a good contrast between the gray and the white matter of the brain. Further, the T_1 contrast is frequently used in combination with Gadolinium-based contrast agents, which cause a concentration-dependent shortening of the T_1 relaxation time. As a result, tissue with an abnormal uptake of contrast agent lights up in the T_1 contrast, which is helpful to identify malignant tumors and other suspicious lesions.

Chapter 3

Methodological Issues

In this chapter, general methodological aspects of the thesis are discussed to avoid redundancy in following chapters. It covers technical details about the MRI hardware used for the experiments, the concepts underlying the simulations performed, and it describes the software framework developed to conduct simulations and image reconstructions.

3.1 MRI System

All experiments were conducted with a commercially available whole-body MRI system, which is shown in Figure 3.1 (MAGNETOM Trio, A Tim System; Siemens AG, Erlangen, Germany). The system's core component is a superconducting magnet, operating at a field strength of $B_0 = 2.89$ T. The bore (with casing) has a length of 142 cm and a diameter of 60 cm, allowing for FOV sizes up to 50 cm in each direction. It is equipped with a gradient system that provides a maximum gradient field strength of 38 mT/m per axis, switchable on a $10 \mu\text{s}$ raster with a maximum slew rate of 170 mT/(m · ms). Separate coils were used for the RF excitation and signal detection. The excitation was done with a large body coil, which is mounted to the casing of the bore, and a head coil array was used to receive the resonance signal (see Figure 3.2).

The coil array is composed of 12 individual coil elements, which are arranged as a ring and grouped into 4 clusters of 3 elements. Each receiver element has a specific sensitivity profile and yields a differently modulated view of the object

$$\rho_i(\mathbf{x}) = C_i(\mathbf{x}) \cdot \rho(\mathbf{x}) , \quad (3.1)$$

where $C_i(\mathbf{x})$ denotes the complex-valued sensitivity profile of the i th coil. The signals from the three coils of each cluster are combined into a set of *mode signals* before



Figure 3.1: Whole-body MRI system.



Figure 3.2: Head coil array.

digitization. This is done by an integrated hardware device called *Mode Matrix*, which generates a primary mode P , a secondary mode S , and a tertiary mode T from the coil signals L , M , and R

$$P = \frac{L - R}{2} - i \cdot \frac{M}{\sqrt{2}} \quad S = \frac{L + R}{\sqrt{2}} \quad T = \frac{L - R}{2} + i \cdot \frac{M}{\sqrt{2}}. \quad (3.2)$$

The mode signals are digitized at a fixed sampling rate of 10 MHz using separate 24bit analog-to-digital converters. A subsequent digital processing includes demodulation of the resonance frequency, low-pass filtering, and downsampling to the desired dwelltime. Because a phase sensitive quadrature demodulation is employed, it yields a real and imaginary component, corresponding to the complex MRI signal described by Eq. (2.18). Any further signal processing steps were disabled in the experiments, and all measured data was exported from the MRI system for offline processing.

Figure 3.3 shows images reconstructed from the different mode signals of the head coil array. Due to the specific design of the coil elements, a combined image (with nearly homogenous sensitivity) can be obtained by calculating the sum-of-squares of the individual mode images

$$I(\mathbf{x}) = \sqrt{\sum_i I_i(\mathbf{x}) \overline{I_i(\mathbf{x})}}. \quad (3.3)$$

Because most intensity is concentrated in the primary modes, it is often sufficient to use only the primary signals for the image reconstruction and to discard the other signals.




Figure 3.3: Reconstructions from the individual mode signals of the head coil array and (sum) the sum-of-squares combination, obtained for a water phantom with a slice-selective FLASH sequence (P = primary mode, S = secondary mode, T = tertiary mode, number denotes the corresponding cluster of receive elements).

In this case, only 4 signals instead of the initial 12 signals have to be processed, which reduces the computational requirement of the reconstruction procedure.

All measurement sequences were implemented in the C/C++ language with the use of the IDEA sequence development environment (Siemens AG, Erlangen, Germany). The sequence modules are executed on a real-time unit running the VxWorks operating system (Wind River, Alameda, Canada) and take control of the signal processors used for RF pulse generation, data recording, and play-out of the gradient waveforms.

3.2 Simulation Principles

In practice, the received MRI signal can be significantly affected by various experimental inaccuracies. Among these effects are field inhomogeneities, object motion, inaccurate sequence timing, noise, and eddy currents. Because it is usually non-trivial to identify error sources from the observed image artifacts and, in particular, because different effects often translate into similar artifacts, it turned out to be highly efficient to perform simulations of the MRI signal generation and to mimic potential inaccuracies. This allows to systematically investigate the response behavior of acquisition and reconstruction techniques to a wide set of error sources. Further, the use of simulations is very helpful for understanding the imaging properties of different sampling schemes, which are usually non-intuitive due to the involvement of the Fourier transformation and rather complex for non-Cartesian patterns.

Several strategies can be employed to perform simulations of MRI experiments. The most general concept is based on the use of the Bloch equation (2.9), which describes the temporal evolution of bulk magnetization in time-varying magnetic fields. It allows to analyze all relevant mechanisms on a detailed physical level, including the RF excitation process and macroscopic interference effects. In a realistic imaging scenario, all variables of the Bloch equation are locally varying, which applies especially to the magnetic field strength when gradients are switched. Therefore, it is necessary to compute an ensemble of independent magnetization vectors at a sufficiently dense raster in the image space. The MRI signal is then obtained by summing over the transversal components of all vectors for each time point. However, such simulation procedure is computationally intensive, because a high number of vectors is required to observe realistic interference effects (for example, signal cancelation from dephasing).

For this reason, a different simulation concept was used in the present work. It is based on the higher-level finding from Eq. (2.18) that the MRI signal corresponds to the continuous Fourier transform of the object. Thus, if the analytical Fourier transform of a *numerical phantom* is known, a simulated MRI signal is obtained by calculating the k-space position for each time point and evaluating the Fourier transform at these positions. The analytical Fourier transform of a rectangle is given by

$$f_{\text{rect}}(k_x, k_y) = 2\pi \cdot s_x \cdot s_y \cdot \text{sinc}(s_x \cdot k_x) \cdot \text{sinc}(s_y \cdot k_y), \quad (3.4)$$

where s_x, s_y denote the size of the rectangle in the x- and y-direction. Further, the

Fourier transform of an ellipse is given by

$$f_{\text{circ}}(k_x, k_y) = \frac{a \cdot J_1 \left(b \cdot \sqrt{(a/b \cdot k_x)^2 + k_y^2} \right)}{\sqrt{(a/b \cdot k_x)^2 + k_y^2}}, \quad (3.5)$$

where $J_1(x)$ is the first-order Bessel function of the first kind, and a, b denote the short and long axes of the ellipse. Due to the linearity of the Fourier transformation, an analytical Fourier transform of any composition of rectangles and ellipses can be derived by superposition of these functions, where a displacement of a single element can be realized with a linear phase modulation (see Section 2.2.4). Therefore, arbitrary numerical phantoms can be constructed for the simulations. This concept offers a fast computation of the signal, but it is in turn less powerful. For instance, it does not allow to model the effects of RF pulses and, thus, can not be used to study excitation or refocusing profiles of different pulse waveforms. Further, it yields a global signal in the sense that it is impossible to directly analyze the effect of a local change of, for example, the field strength, because the signal formation itself is not covered by the simulation procedure. However, this limitation can be in part eluded by making use of the linearity of the Fourier transformation. For example, it is known that a deviation of the field strength causes a phase drift of the signal induced from the respective area. Therefore, the effect of a field deviation can be investigated by adding a temporal phase drift to the Fourier transform of, for example, a single circle and combining it with the Fourier transforms of the other object elements. This yields the MRI signal corresponding to a magnetic field with a discrete jump at the location of the circle. Although deviations of the field strength are rather smooth in practice, such analysis serves for understanding the resulting image artifacts on a qualitative level. In fact, many of the effects observed in an experimental setting, such as relaxation, saturation, or non-rigid object motion, can be analyzed in a similar way with dedicated numerical phantoms. Hence, the simulation concept offers a high flexibility and proved to be very valuable during the development and implementation of the techniques presented in the work.

Nevertheless, a clear limitation arises if it is necessary to include coil sensitivity profiles into the simulation, as for investigating parallel imaging techniques. Because a varying coil sensitivity causes a modulation in the image space, it corresponds to a convolution of the numerical phantom with the sensitivity profile's Fourier transform in k-space. Deriving an analytical solution for this equation is difficult for realistic sensitivity profiles. Therefore, a more convenient strategy for this situation is to conduct a discretized simulation, which can be computed by taking a discrete source image with high spatial resolution, multiplying it with the sensitivity profile, performing a discrete fast Fourier transformation, and evaluating the transform at the desired k-space locations.

Of course, when a non-Cartesian k-space sampling pattern is used, it is necessary to interpolate the Fourier transform from the grid to the desired k-space positions, which introduces certain inexactness to the simulation. Although the effect can be reduced by increasing the resolution of the source image so that the grid points get denser, the approach has to be used with care to avoid an interference with the artifacts under investigation. Therefore, this concept was only employed if it was indispensable to account for localized coil sensitivities.

3.3 Software Framework

To perform the aforementioned simulations and to reconstruct images with the proposed techniques, a comprehensive software framework, entitled MRISim, has been developed. The design of MRISim was inspired by the data flow on a modern MRI system, and its architecture is highly object oriented. User-defined measurement sequences, numerical phantoms, and image reconstruction techniques can be implemented as encapsulated modules, where core functionality for data and image processing is provided by the framework. As a specific feature, all data is handled in a global container system, which makes it possible to look at every intermediate result and, thus, is very helpful for debugging the individual modules. Moreover, MRISim provides an extensive graphical user interface, including various visualization and analysis tools for time-series and image data.

To conduct a simulation, the user has to select a sequence module, a phantom module, and a reconstruction module. MRISim then requests the sequence module to prepare an event table, which describes the sequence timing and may comprise RF events, data sampling, and gradient switchings. Here, MRISim follows the timing specification of the MRI system used ($10\ \mu\text{s}$ gradient raster, $\geq 100\ \text{ns}$ sampling interval) to account for synchronization difficulties observed in practice. The sequence module is also responsible for configuring the individual event objects (for example, selection of a gradient waveform and gradient amplitude). In a second step, MRISim calculates a k-space trajectory from the event table by integrating over all gradient events. If selected by the user, this includes a modeling of timing errors or imperfect gradient responses. In the next step, the k-space trajectory is translated into a MRI signal. This is achieved by estimating the exact sampling time points from the readout events and querying a signal sample for each of these time points from the phantom module, where the corresponding k-space positions and gradient time courses are passed to the module. Thus, the signal generation is completely encapsulated in the phantom modules, and a module may employ the concepts of the foregoing section or may use a totally different simulation technique. Because the phantom module is additionally notified about all

RF pulse events, it is possible to consider relaxation mechanisms as well as other time dependent effects. Further, the phantom module may generate a single signal or multiple channels to simulate the use of coil arrays. The signal samples are then appended to individual arrays for each readout event, which are stored in the data container. If selected, noise is added to the signals, as well as a phase drift to mimic an inaccurately chosen resonance frequency. As a last step, MRISim instructs the reconstruction module to calculate an image from the sample arrays in the data container. The results can then be inspected with MRISim's image viewer, which provides analysis functions like windowing, zooming, image subtraction, and profile calculation.

In order to employ MRISim for the reconstruction of real measurement data, it is first necessary to run a normal simulation as described above. In a subsequent step, the simulated signal in the container is replaced with data from the MRI system, and the reconstruction procedure is simply restarted. Because in this concept the same source code is used for processing simulated and real MRI data, it can be easily ensured that artifacts observed in experimental images are not due to implementation failures of the reconstruction algorithms. Further, the option to restart the image reconstruction can be utilized to reconstruct a single data set with a number of different settings (or techniques) and to compare the resulting image quality in a convenient way.

MRISim has been written in the C/C++ language with use of the QT4 class library (Trolltech ASA, Oslo, Norway). Time critical sections were parallelized with the OpenMP interface provided by the Intel C/C++ compiler (Intel Corporation, Santa Clara, USA). The GNU Scientific Library (GSL)¹ was used for numerical calculations, and Fourier transformations were done with the FFTW3 library². Further, the Blitz++ library³ was employed for data handling. The optimization routines were in part derived from source code of the restoreInpaint project⁴. All calculations were done on PowerEdge 2900 systems (Dell Inc., Round Rock, USA), equipped either with two Intel Xeon 5060 dual core processors (running at 3.2 GHz) and 4 Gb of memory, or with two Intel Xeon E5345 quad core processors (running at 2.33 GHz) and 8 Gb of memory.

¹ <http://www.gnu.org/software/gsl>

² <http://www.fftw.org>

³ <http://www.oonumerics.org/blitz>

⁴ <http://restoreinpaint.sourceforge.net>

Chapter 4

Radial Sampling of k-Space

This chapter provides a comprehensive discussion of the radial sampling scheme. It presents basic ideas for the sequence design and introduces the two methods that are conventionally used for the image reconstruction. Further, the main imaging characteristics are analyzed and appraised in view of potential applications, where particular emphasis is put on discussing the implicit sensitivity to certain image artifacts. Finally, several promising variants of the radial sampling scheme are reviewed.

4.1 Basic Sequence Design

In the radial acquisition scheme, first proposed by Lauterbur in 1973 [1], the k-space signal is sampled along spokes instead of parallel rows as illustrated in Figure 4.1. Gradient switchings to generate a corresponding trajectory can be derived from the time course used to sample the central row in the Cartesian scheme. As described in Section 2.2.3, it is achieved with a prephasing gradient in the negative x-direction and a readout gradient in the opposite direction. Because the central row is equal to a spoke with an angle of zero degrees, this gradient sequence directly yields the first spoke of the trajectory. Other spokes can then be acquired with a combination of respective gradients in the x- and y-direction, which produces a rotation of the spoke around the k-space center. Usually, the same gradient waveforms are used for all spokes, while the gradient amplitudes in the x- and y-direction are varied according to

$$G_x = G_0 \cdot \cos(\phi) \qquad G_y = G_0 \cdot \sin(\phi) , \qquad (4.1)$$

where ϕ is the desired angle of the spoke, and G_0 denotes the amplitude required for sampling the central k-space row. This strategy ensures a fixed sequence timing and allows for a straightforward implementation on the MRI system. Thus, many existing Cartesian sequences can be adapted to radial sampling with only mild effort.




Figure 4.1: Illustration of (left) the Cartesian sampling scheme and (right) the radial scheme.

As in the Cartesian case, the distance between two samples along a spoke is conveniently selected from a given FOV size

$$\Delta k = \frac{1}{\text{FOV}}, \quad (4.2)$$

while the number of samples per spoke is obtained from a given base resolution n . In the radial case, these two settings do not directly translate into a spatial resolution, because the latter depends additionally on the number of spokes n_s , which can be selected freely as a further sequence parameter. Since the relation between the number of spokes and the resolution is more complex, it will be discussed in more detail after introducing the reconstruction methods. Nevertheless, in the literature it is usually suggested to select the number of spokes according to

$$n_s = \frac{\pi}{2} \cdot n, \quad (4.3)$$

which ensures that the distance between samples on neighboring spokes (indicated by Δq in Figure 4.1) is less or equal to Δk , as provable with a simple geometrical derivation [5,3]. Thus, with this choice an isotropic spatial resolution of $\Delta x = \text{FOV}/n$ is warranted.

The spokes can be acquired by increasing the angle ϕ either from 0° to 180° , or from 0° to 360° . Both schemes result in an identical spoke distance of $\Delta\phi = 180^\circ/n_s$, but differ in the orientation of the spokes. In the 180° case, all neighboring spokes are equally aligned, whereas for 360° the orientation of neighboring spokes is opposed, which results in a modified response to experimental inaccuracies. The latter can also be obtained for the 180° scheme by inverting the orientation of every second spoke,

which, however, yields a different temporal ordering of the spokes. As an alternative to serially incrementing the angle, an arbitrary temporal ordering may also be used, which offers a more robust recording of the k-space information in certain situations as discussed later.

4.2 Conventional Image Reconstruction from Radial Data

Although changing the sampling pattern from parallel rows to spokes appears trivial at first glance, it has major consequences for the image reconstruction procedure. Basically, every MRI reconstruction technique comprises a discrete inverse Fourier transformation of the measured values S_j in order to extract the spatial information encoded in the MRI signal

$$I(\mathbf{x}) = \sum_j S_j \cdot e^{-2\pi i \mathbf{x} \cdot \mathbf{k}_j} . \quad (4.4)$$

For the Cartesian scheme, all k-space samples (and image pixels) lie on a regular grid and, thus, the above equation can be efficiently evaluated with the use of the FFT algorithm. This reduces the problem from N^2 to $N \cdot \log N$ and permits a fast calculation even for large image sizes.

However, if the k-space is instead sampled along spokes, the sampling positions are neither on a grid nor equidistant, which precludes a straightforward use of the FFT. Unfortunately, direct evaluation of Eq. (4.4) is very time consuming and practically infeasible for routine applications, even with modern high-performance computer systems. As a second complication, the sample density of the radial trajectory is much higher in the center of k-space than in the outer areas, which leads to severe image blurring. This is because the DFT in Eq. (4.4) does not account for a locally varying distance between the sample points. Thus, the signal from the dense samples accumulates and produces an overestimation of the low spatial frequencies. For this reason, it is necessary to cope with both, the non-gridded sampling positions as well as the varying sample density, and two different approaches are conventionally used for the reconstruction from radial sampling.

4.2.1 Projection Reconstruction

The *projection reconstruction* (PR) method is based on the projection-slice or central-slice theorem [1, 4, 3], and for many years it was the only practicable technique for reconstructing radial MRI data. Applied to the MRI situation, the projection-slice

theorem states that the inverse Fourier transform of a measured spoke is equal to the projection profile of the object, where the projection angle (beam direction) is perpendicular to that of the spoke [14]

$$\begin{aligned}\mathcal{F}^{-1}\{S_\phi\}(h) &= P_\phi(h) = \int \int \rho(x, y) \cdot \delta(h - x \cos \phi - y \sin \phi) \cdot dx dy \\ &\equiv \mathcal{R}\{\rho\}(\phi, h).\end{aligned}\tag{4.5}$$

Here, h denotes a position along the projection profile. Thus, by performing inverse FFTs of the individual spokes, a set of projection profiles can be obtained that corresponds to the discrete two-dimensional Radon transform of the object $\mathcal{R}\{\rho\}$.

This finding reveals that the reconstruction problem arising in transmission tomography (for example, CT or C-arm x-ray imaging) is closely related, although these techniques exploit a totally different signal generation mechanism. As a slight difference, however, the projection profiles from the MRI data are usually complex-valued. Nevertheless, in both cases images can be computed from the profiles P_ϕ using the backprojection procedure, which is the adjoint of the 2D Radon transformation

$$\text{BP}(x, y) = \sum_{\phi} P_\phi(x \cos \phi + y \sin \phi).\tag{4.6}$$

Simply speaking, the backprojection procedure “smears” the individual profiles back onto the object plane (in the opposed direction to the projection), and the values from different backprojected profiles accumulate and form an image. Because in practice the profiles and the image matrix are discrete, an interpolation is required for evaluating the profiles at the projected positions of the image pixels. Different methods can be used for this step, where a linear approach is the simplest option to interpolate between the samples of the profiles.

When performing the image computation with the plain backprojection procedure, the reconstructions still exhibit a severe blurring and do not recover the object properly. The blurring arises because the Radon transformation is not a unitary operation and, thus, it is not inverted by applying only the adjoint operator, which was already discussed by Radon in a paper from 1917 [15]. However, an inversion of the transformation can be achieved by filtering the profiles with the M- or ramp filter $|\mathbf{k}|/n_s$ in front of the backprojection step. For this reason, the combination of these two steps is called the filtered backprojection (FBP) reconstruction method.

In the MRI case, the raw data is already given in the frequency space, and the filtering step can be incorporated simply by multiplying the samples with the filter window before performing the FFTs. Because the signal amplitude in k-space declines away from the center while the noise level stays rather constant, the multiplication with the

ramp-filter can lead to an amplification of image noise. Therefore, the ramp-filter is often combined with a Hanning filter (or comparable) to suppress noise patterns, which is in turn accompanied by a loss of spatial resolution. Further, because the profiles are complex-valued, it is necessary to perform the FBP separately for the real and imaginary part. In early MRI applications, the profiles were sometimes recast to real values by taking the magnitude, which prevents a dual processing and compensates for certain phase artifacts [16, 17]. However, when using modern coil arrays with a local phase modulation, this approach leads to inconsistent reconstructions. In fact, for a spatially dependent phase offset, a removal of the phase variations can only be accomplished after the signal has been resolved into its spatial components.

4.2.2 Regridding

The second reconstruction method is called *regridding* or just *gridding*, and it has been adapted to medical imaging in 1985 by O’Sullivan [18]. Here, the measured spokes are interpolated onto a Cartesian grid, and afterwards a normal FFT based reconstruction of the gridded data is performed. This strategy seems not very sophisticated, but complications arise from the fact that the interpolation step has to be conducted in the frequency space. Basically, any two-dimensional interpolation can be seen as some kind of convolution procedure and, consequently, it leads to certain modulation effects in the image space. Therefore, the interpolation technique has to be selected in such a way that the modulation effects are tolerable, which hampers the use of a bilinear interpolation, for instance.

In the regridding approach, the interpolation is achieved by convolving the measured samples with a radial interpolation kernel and evaluating the convolved or “smeared” samples at the desired k-space positions. The interpolated signal can be written as a weighted sum over all sample points S_j

$$\hat{S}(\mathbf{k}) = \sum_j S_j \cdot K(|\mathbf{k} - \mathbf{k}_j|), \quad (4.7)$$

where $|\mathbf{k} - \mathbf{k}_j|$ denotes the distance between the j th sample and the interpolated position, and $K(d)$ is the window function of the interpolation kernel.

Because the multiplication of a compact signal with a rect function does not distort the signal (provided that the support of both functions is equal), the optimal kernel for interpolations in the frequency space is given by a properly scaled sinc function. However, due to the unlimited support, it is infeasible to compute a true convolution with the sinc function in practice. Hence, a compact kernel has to be used that approximates the properties of the sinc function, and Jackson et al. [19] investigated different

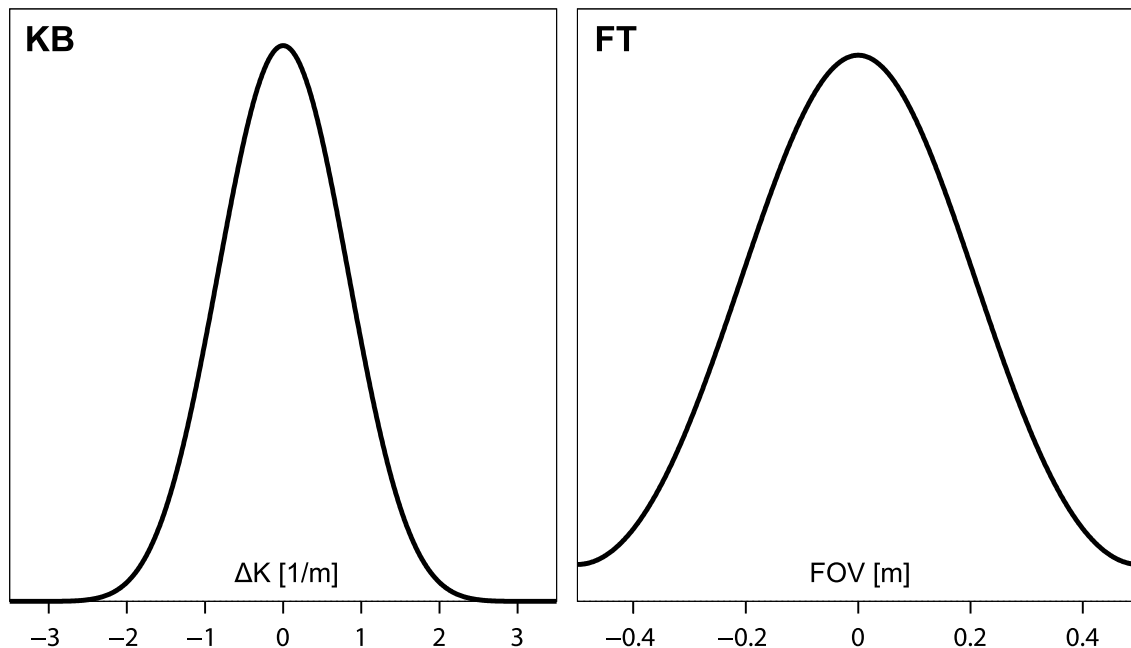


Figure 4.2: (Left) Plot of the Kaiser-Bessel window for $L = 6$ and $\beta = 13.8551$, where the FOV has been normed to 1 m. (Right) Fourier transform of the KB window (obtained using a DFT of the discretized function).

kernel functions in view of the achievable interpolation quality. Based on their results, usually a radial Kaiser-Bessel kernel is employed, which provides a good image quality at a reasonable window size. This ratio is an important property, because the larger a kernel is, the more sample points contribute to the interpolated signal and the higher is the computational cost of the operation.

The Kaiser-Bessel window [20] is given by

$$K_{\text{KB}}(d) = \begin{cases} \frac{1}{L} I_0(\beta \sqrt{1 - (2d/L)^2}) & |d| \leq \frac{L}{2} \\ 0 & |d| > \frac{L}{2} \end{cases}, \quad (4.8)$$

where L is the desired kernel width, $I_0(d)$ denotes the zero-order modified Bessel function of first kind, and β is a shape factor, which should be selected according to an equation reported by Beatty et al. [21]. A plot of the function is shown in Figure 4.2.

To interpolate the spokes onto the grid, Eq. (4.7) is evaluated at $n \times n$ grid points with a distance of Δk . Afterwards, a two-dimensional inverse FFT of the grid can be performed as in the Cartesian case. Due to the convolution with a finite interpolation kernel, the obtained image exhibits an undesired modulation with the Fourier transform of the kernel, which is called *roll-off* effect. A compensation of the modulation can be achieved by dividing the image by the kernel's Fourier transform, which is consequently referred to as *roll-off correction*. The Fourier transform of the Kaiser-Bessel window is




Figure 4.3: Illustration of the oversampling procedure. (Left) Regridding reconstruction without oversampling, showing clearly visible aliasing effects. (Right) Oversampled reconstruction with artificially increased FOV, which can be cropped in the end. Due to the FOV enlargement, the artifacts are shifted away from the object.

commonly approximated by

$$\mathcal{F}\{K_{\text{KB}}\}(x) \approx M_{\text{KB}}(x) = \frac{\sin \sqrt{(\pi Lx)^2 - \beta^2}}{\sqrt{(\pi Lx)^2 - \beta^2}}. \quad (4.9)$$

Thus, the rows and columns of the image have to be divided by the above equation to remove the roll-off effect.

As a second effect, the use of a finite interpolation kernel leads to side lobes outside of the image plane. Because Eq. (4.7) is evaluated discretely on a grid, which can be written as multiplication with a comb-function, a periodic function is obtained in the image space, and the side lobes are aliased back into the image plane from neighboring object copies. Although the initial amplitude of the side lobes is typically small for proper interpolation kernels, their intensity is amplified by the roll-off correction, and the aliasing patterns usually become visible in the image. The artifacts occur especially near the image border as these areas initially have low intensity due to the roll-off effect, but they might also coincide with the object. For this reason, a two-times oversampling of the gridding matrix is frequently employed [19]. Hence, Eq. (4.7) is evaluated on a $2n \times 2n$ matrix with a sample distance of $\Delta k/2$. This doubles the distance between the copies and shifts the aliased side lobes away from the object, so that visible artifacts in the image plane are reduced. The extended FOV can be cropped at the end of the reconstruction procedure, yielding an image with again $n \times n$ pixels. The oversampling procedure is illustrated in Figure 4.3, and its effect can be clearly seen in Figure 4.4.

Further, it is necessary to compensate for the varying sample density of the radial

trajectory, as discussed in the introduction of the section. Because the exact position of the individual samples is lost after the interpolation onto the grid, this procedure has to be done as an initial step. It is achieved by weighting each of the acquired samples with a density compensation function (DCF), which represents the inverse density of the sampling positions in k -space. For an arbitrary trajectory, a DCF can be estimated by computing a Voronoi diagram or Dirichlet tessellation of the trajectory [22], where the size of the individual Voronoi cells serves as a measure of the local sample density [23]. However, for radial trajectories it is possible to derive an analytic DCF, which yields the Ram-Lak filter [24, 25]

$$D_{\text{RL}}(\mathbf{k}) = \begin{cases} |\mathbf{k}|/n_s & |k| \neq 0 \\ 1/(2n_s) & |k| = 0 \end{cases}. \quad (4.10)$$

In summary, the subsequent steps of the regridding approach are

1. Density compensation
2. Convolution and evaluation on a grid
3. Inverse 2D FFT
4. Roll-off correction
5. Image cropping

In operator notation, the complete processing of the continuous k -space signal $S(\mathbf{k})$ can be written as

$$\text{GRD} \cdot S(\mathbf{k}) = (\text{M}_{\text{KB}}^{-1} \cdot \mathcal{F}^{-1} \cdot \text{III} \cdot \text{K}_{\text{KB}} * \text{D}_{\text{RL}} \cdot \text{T}) \cdot S(\mathbf{k}), \quad (4.11)$$

where III is the comb-function (with spacing Δk or $\Delta k/2$), and $*$ denotes that a convolution with K_{KB} is performed. Further, T denotes the sampling operator of the trajectory

$$\text{T} = \sum_j \delta(|\mathbf{k} - \mathbf{k}_j|). \quad (4.12)$$

It is important to note that the processing includes two different discretization steps: the discrete sampling along the trajectory T and the gridding of the convolved samples III. Both operations lead to certain periodicity in the image space, but must not be confused. While the III operator leads to regular neighboring object copies that are specific to the regridding approach, the sampling operator T creates complicated periodic patterns, which are related to the geometry of the trajectory. Thus, the latter effect occurs also when employing a direct DFT reconstruction according to Eq. (4.4).

4.2.3 Comparison of the Methods

In recent years, the regridding technique is more frequently used than the filtered back-projection approach. This is mainly because it is more flexible and can be used to reconstruct data from other sampling schemes as well, including spiral and propeller trajectories [26, 27]. Recently, however, also the FBP approach is attracting attention again due to the development of the closely related highly constrained backprojection (HYPR) technique [28].

Nevertheless, when looking closely, both approaches are quite similar. In the regridding approach, all spokes are first interpolated onto a grid and then an inverse Fourier transformation is performed. Due to the linearity of the FT, the same result would be obtained by interpolating the spokes separately onto individual grids, performing separate inverse FTs, and summing over all transforms. Further, a rotation of the coordinates in the image space yields a rotation in the Fourier space, which is a general property of the Fourier transformation. Thus, instead of gridding the spokes at the respective angles, a similar result would be obtained by gridding all spokes at 0° (which requires no interpolation), performing inverse FTs, and then rotating the images with respect to the individual spoke angles. Next, because for a single spoke at 0° all other grid columns are zero, the inverse 2D FT yields the backprojection or smearing of the spoke's inverse 1D FT over the image. Finally, for a single spoke the multiplication with the Ram-Lak DCF (4.10) is basically identical to the filtering with the ramp-filter in the FBP approach. The slight difference at the k-space center originates from the fact that the Ram-Lak filter was derived for a discrete sampling, while the ramp-filter was obtained for inverting the continuous Radon transform.

This discussion shows that the regridding and FBP approaches are nearly equivalent, which is verified on a more formal level in Appendix A. The main difference between both reconstruction techniques consists in the interpolation procedure. In the regridding method, the data is interpolated in the frequency space, whereas in the FBP approach the interpolation is done in the image space when backprojecting the profiles at different angles. Although the image-space interpolation might appear less complex in view of the discussion from the foregoing section, rotating images at high precision, in fact, faces similar problems as in the frequency-space interpolation [29]. Thus, also in terms of the complexity both approaches are quite comparable. Figure 4.4 compares reconstructions of the Shepp-Logan phantom obtained with FBP, regridding (with and without oversampling), and a direct DFT that is free of any interpolation artifacts (Eq. (4.4) with the Ram-Lak DCF). While some minor artifacts are visible for regridding without oversampling, all other methods recover the phantom properly and offer a sufficient reconstruction accuracy for most clinical applications – provided that enough spokes are acquired. Due to the similarity of the regridding and FBP method, the




Figure 4.4: Comparison of the image quality for different reconstruction techniques (Shepp-Logan phantom, base resolution 256 pixels, 402 spokes). (FBP) Filtered backprojection, (1x regr) regridding without oversampling, (2x regr) regridding with 2x oversampling, and (DFT) straightforward discrete Fourier transformation. Images are shown with (nc) a normal contrast and brightness as well as with (hc) a high contrast and overdriven brightness to emphasize the differences in the background patterns.

following parts of the thesis focus only on the regridding method, but it should be kept in mind that the regridding step can be replaced by a projection-based reconstruction. Further, because the regridding approach is also employed for reconstructing the data from the simulations, it should be noted that the presented results are always additionally affected by respective interpolation artifacts. However, as demonstrated, these artifacts are minor when using oversampling and do not interfere with the effects under investigation in a qualitative way.

4.3 Advantages and Disadvantages of Radial Sampling

Due to the distinct geometry of the trajectory as well as the different temporal ordering of the k-space acquisition, radial sampling offers unique imaging properties. Some of the properties are advantageous over the Cartesian sampling pattern, and some properties translate into drawbacks. In the following, these salient properties are discussed in view of potential applications and experimental inaccuracies that are commonly encountered in practice.

4.3.1 Point-Spread-Function

To understand the characteristics of any imaging system, it is very helpful to examine the *point spread function* (PSF), which formally describes the impulse response of the system and allows to conclude how an object is imaged by the system. Because in the MRI case the Fourier transform is sampled along the trajectory, the PSF is closely related to the sampling trajectory. When ignoring relaxation and other effects for simplicity, the sampling corresponds to a projection of the FT onto the trajectory, which can be written as multiplication with a function that is one for the sampling positions and zero otherwise (comparable to Eq. (4.12)). Therefore, the PSF of a MRI sequence can be obtained by setting all sample values to one and processing the data with the usual image reconstruction procedure.

Figure 4.5 shows images of the PSF and corresponding reconstructions of the Shepp-Logan phantom, obtained for a radial trajectory with a base resolution of $n = 256$ pixels and different numbers of spokes n_s . According to Eq. (4.3), a full dataset comprises $n_s = 402$ spokes, and in this case the PSF has a distinct peak in the center, which is surrounded by minor circular oscillations that decline away from the center. The circular oscillations result from the finite coverage of the k-space, which is known as truncation effect and discussed in detail in Chapter 8. If the number of spokes n_s is




Figure 4.5: Point spread function (top) and corresponding reconstruction of the Shepp-Logan phantom (bottom), obtained for a radial trajectory with 402, 64, and only 24 spokes (base resolution 256 pixels).

reduced to a much lower value, for example to 64 spokes, the central part of the PSF remains unchanged. However, it can be seen that gaps arise at a certain radius, which produce a streak pattern in the outer areas. Because every “point” of the object is convolved with the PSF, the streaking patterns appear in the reconstructed image at a respective distance from the individual points, which is referred to as *streaking artifact* and clearly visible in the lower row of Figure 4.5. When reducing the number of spokes to only 24, the diameter of the streak-free disc is further reduced and, consequently, more pronounced streaking artifacts arise in the reconstructed image. Thus, the lower the number of spokes is, the more intense and coarse are the streaking artifacts, and the closer do artifacts appear to a source point.

A comprehensive mathematical analysis of the radial PSF has been presented by Lauzon and Rutt [30, 31]. However, in a simple manner the PSF can be analyzed using the linearity of the Fourier transform, similar to the discussion in Section 4.2.3. Here, each of the spokes can be written as a multiplication in k -space with the Ram-Lak filter, where the filter is truncated at the end of the spoke (see Figure 4.6). Hence, every spoke creates a backprojection of the filter’s Fourier transform in the PSF. Because the truncated filter can be constructed by adding a rectangle and a flipped triangle, which in turn can be composed by convolving two rectangles, the Fourier transform of the

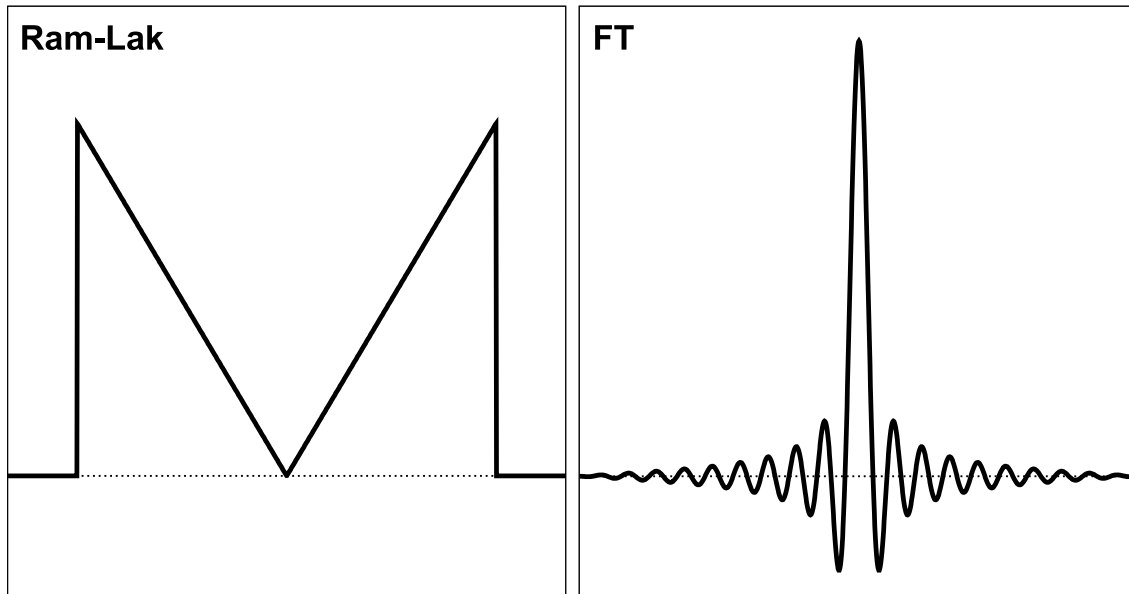


Figure 4.6: (Left) Truncated Ram-Lak filter and (right) its Fourier transform. The scaling of both functions has been selected arbitrarily.

Ram-Lak filter is basically given by $f(k) = 2 \cdot \text{sinc}(2k) - \text{sinc}^2(k)$, as shown in the right of Figure 4.6. This profile is backprojected over the image plane at different angles and, thus, the function maxima overlap in the center and form the central peak of the PSF. In the surrounding of the peak, the profiles create a densely covered disc with a radius that depends on the total number of spokes. Beyond this radius, the backprojections are not dense enough, and the maxima propagate as single streaks. Nevertheless, if the radius of the disc is equal to the FOV, it is prevented that the streakings may overlap with the reconstructed object. This condition is fulfilled if the number of spokes is selected according to Eq. (4.3).

However, in comparison to the Cartesian scheme, it means that 57% more k-space lines are required to obtain a fully sampled dataset, which is clearly unfavorable due to a prolonged measurement time. On the other hand, in practical imaging scenarios it is often tolerable to acquire only a reduced number of spokes. This is because the region of interest is frequently located inside the object so that streakings at the border are acceptable and, further, for mild reduction factors the streaking patterns are fine structured and rather benign. In fact, because the center of the PSF is basically unaffected by changes of the spoke count, radial sampling offers an interesting undersampling behavior. While in Cartesian sampling a reduction of the k-space lines leads to either a loss of resolution or the occurrence of aliasing effects that render the image useless, in radial sampling most object information remains visible even for a significant amount of streaking artifacts. For this reason, radial sampling is a very promising technique for applications where it is essential to acquire images in a short time and where the diagnostic information is not impaired by minor streaking patterns [32,33,34]. In other

words, the ability to select the number of spokes and, thereby, the degree of tolerable streaking artifacts independently from the desired in-plane resolution poses a clear advantage of the radial scheme, whereas it is less efficient than Cartesian sampling for conventional high-quality imaging.

4.3.2 Sample Distribution

Because all spokes pass the center of k-space, the number of sample points is much higher for the low spatial frequencies (in the central region) than for high frequencies. This is fundamentally different to the Cartesian case, where all frequencies are equally covered. Although the variable sample distribution poses certain complications for the image reconstruction (as discussed before), sampling the low frequencies more densely is actually a reasonable strategy for typical objects. In fact, many real-world images are characterized by an energy concentration around the center of the Fourier space, which applies not only to medical tomographic images but also to other natural images [35]. Therefore, it makes sense measuring the low spatial frequencies more “precisely”, while spending fewer samples on areas that contain less information. Moreover, it is possible to make explicit use of the oversampling of the k-space center. In particular, low resolution images can be calculated from only a subset of spokes, because within a certain distance from the k-space center the samples are sufficiently dense to calculate an artifact-free image even from a reduced number of spokes. This distance is called *Nyquist radius*, and it can be derived from Eq. (4.3) when inserting the reduced spoke count for n_s . Numerous applications of such subframe images are conceivable. For instance, a series of time-resolved subframes can be calculated to identify and to compensate for object motion that might have occurred during the acquisition of the full dataset [36,37].

As a further consequence of the radial geometry, all spokes carry an equal amount of low and high spatial frequencies, whereas for Cartesian sampling the important low spatial frequencies are contained in only few rows. This makes radial sampling an attractive option for applications where a continuous update of the object information is preferred, like in real-time MRI for tracking organs or devices during interventional procedures [38,39,34,40,41,42]. The use of a Cartesian scheme in such applications is often hindered by abrupt image changes or flickering when the central rows are updated, which is eliminated due to the more uniform image updates for radial sampling. On the other hand, because all spokes are equivalent, measurement errors in any of the spokes (for example, from blood inflow) directly manifest as visible artifacts. For Cartesian sampling, this effect is confined to only the central rows because the outer rows carry less important information. However, if one of these central rows is contaminated, the resulting artifacts are drastic since the incorrect information can not be countered by

other k-space lines, as it is in the radial case. Therefore, radial sampling is somewhat more robust and better suited for imaging dynamic information.

Another specific feature of the radial trajectory is the overlapping of all spokes in the k-space center, which again gives pros and cons. Because under ideal conditions the phase of the signal is zero in the k-space center, it can be used as a reference point to correct for static phase offsets of the individual spokes [43]. Such phase offsets can be caused by imperfect refocusing pulses or object motion and pose a serious complication when employing sequences with multiple RF-refocused acquisitions [44]. In the Cartesian case, the problem is usually addressed by recording additional navigator signals to estimate the initial phase offset [45]. This concept, however, delays the acquisition to some degree and can be omitted for radial sampling. Moreover, an advantage is that the phase offset is estimated in the middle of the acquisition time for each spoke, which can be exploited to further correct for an inaccurately assumed resonance frequency ω_0 . On the other hand, the fusion of data from different spokes leads to signal cancellation due to destructive interference if the phase of the encoded information is not coherent for all spokes. This effect is a general problem of intersecting trajectories and even more severe for patterns with many intersections like rosette or Lissajou trajectories [46, 47]. Although static phase offsets among the spokes can be easily removed as discussed, the procedure does not compensate for localized time-dependent phase evolutions, which are frequent in practice (see Section 2.3.2). Because the phase errors are varying within the object, the received samples reflect the averaged phase evolution and, thus, it is required to invert the spatial encoding before a compensation mechanism can be applied. However, because the destructive signal interference is encountered during the inversion procedure itself, the development of a robust compensation method is challenging. In terms of practical applications, it is, therefore, advisable to use radial sampling with acquisition parameters that minimize the formation of phase errors, for example, by employing short echo times or RF refocusing pulses.

4.3.3 Readout Oversampling

Because the k-space is sampled discretely in any MRI technique, the object reconstructions are always periodic. Therefore, aliasing effects arise if the measured sample points are too distant in k-space so that the neighboring copies overlap in the image space. For a fixed spatial resolution, this problem can be eliminated using a *readout oversampling*. Here, the sampling rate of the receiver is doubled while the gradients are unchanged, which increases the FOV and the sample count by a factor of 2 so that the spatial resolution is preserved. However, for the Cartesian scheme this strategy is restricted to only one spatial direction, and a reduction of the sample distance in the other direction requires additional acquisition steps (i.e. sampling of additional rows) that prolong the

examination time. In the radial case, this limitation does not exist, and readout oversampling can be employed in both directions. This is very interesting for examinations of the abdomen and for cardiac imaging [48], because in these cases the region of interest is very centrally located while aliasing effects from the extremities require that a large FOV is used. Hence, when imaging large objects at high spatial resolution, radial sampling offers a reduced acquisition time relative to Cartesian sampling.

4.3.4 Motion Robustness

The most prominent advantage of the radial sampling scheme consists in a lower sensitivity to object motion during the data acquisition [49, 50], and it is explained by two main reasons. First and foremost, the superior robustness is a consequence of the fact that conventional Cartesian sampling is rather vulnerable to motion. Due to the shift property of the Fourier transform, any motion in the image space translates into a phase modulation in the Fourier space. The exact form of the modulation depends on the type of motion, and in the literature it is usually distinguished between impulsive, periodic, and constant velocity motion [3, 17]. Although these motion types yield different phase relations, they all cause *ghosting artifacts* for Cartesian sampling [51], which exhibit as shifted copies of the object or, respectively, the moving structure. Ghosting artifacts appear exclusively in the direction perpendicular to the readout gradient (“phase encoding” direction), and depending on the motion character, either single focused ghostings or an ensemble of copies might appear, which is rather perceived as blurring. Because the spurious copies usually overlap with the main object representation, ghosting artifacts can lead to diagnostic misinterpretations [52], and numerous compensation techniques have been developed [5]. However, because most of the techniques rely on the validity of an assumed motion model, residual artifacts are often encountered in practice and pose a common problem in routine applications.

The signal inconsistencies emerge mainly among different k-space lines due to the rather long delay between the individual acquisitions, whereas motion within a single readout interval of several milliseconds is often negligible. For the Cartesian sampling geometry this results in a signal modulation along neighboring rows and explains the pronounced occurrence of the ghosting copies in the corresponding direction. In fact, the intrinsic sensitivity to ghosting artifacts is a general problem of the Cartesian scheme, which is not only perceivable for motion but also for other phase perturbations, for example, due to gradient deviations in sequences with multiple readouts [53]. This effect is completely prevented in the radial sampling geometry, and inconsistencies among different spokes translate into artifacts that are commonly appraised as less visually annoying [41, 34]. Here, the artifacts exhibit as blurrings or streakings which appear in a certain distance from the moving entity, so that the error energy is spread more homogeneously over the




Figure 4.7: Image artifacts from gradient deviations in radial sampling. (a) Ideal gradient response, (b) delayed gradient response as typically observed in experiments, and (c) exaggerated gradient delay. The gradient deviations lead to severe image modulations and blurring effects (arrows).

image plane and, thus, a higher visibility of the diagnostic information is retained [17]. Moreover, due to the oversampling of the k-space center, the spokes contain a certain degree of redundant information, which yields an averaging effect and balances for errors in single spokes. This property is a second reason for the improved motion robustness of radial sampling, which together with the general absence of ghosting artifacts promises a high potential for examinations of moving organs or uncooperative patients [54,36,39].

4.3.5 Gradient Deviation Sensitivity

In a typical MRI sequence, the magnetic gradient fields have to be switched within a time period of 10–100 ms, which involves rapid generation of currents up to 500 A. Such gradient switchings induce eddy currents [55,56], which arise in conducting structures inside the bore of the magnet and cause a time-varying perturbation of the magnetic field strength. Therefore, the time course of the actually experienced field strength deviates to some degree from the assumed time course, and respective inaccuracies have been discussed by several authors [26,57]. However, for the standard Cartesian scheme, the gradient deviations do not pose a major complication because identical gradient waveforms are generated in the readout direction for all repetitions and, thus, each acquired row is rather equally affected by the gradient deviations. When reconstructing the image with respect to the theoretical sampling positions, these deviations result only in an overall shift of the samples in the readout direction, which translates into a phase modulation in image space that disappears when the magnitude image is calculated. This is different in the radial sampling geometry because the readout direction is varying for each repetition. In this situation, a simple delay of the gradient response causes a non-uniform misalignment of the k-space center information along the individual spoke directions and, hence, inconsistent data is merged during the reconstruction procedure.

In the image space, this adds a varying phase modulation to the spatial information encoded in each spoke (i.e. the backprojection), which exhibits as signal cancelation and object blurring as demonstrated in Figure 4.7. Therefore, it is necessary to employ a compensation mechanism for the gradient deviations, which is also mandatory when using other non-Cartesian sampling schemes. Different techniques have been proposed for this task, which in part include a calibration measurement of the trajectory to assess the individual gradient deviations [58,59,60,61]. However, Speier et al. [62] showed that for the specific MRI system described in Chapter 3, deviations mainly occur as isotropic delays which can be approximated with

$$\Delta t = 0.42 \cdot t_{\text{os}} + 2.45 , \quad (4.13)$$

where t_{os} denotes the receiver dwelltime (including oversampling). This delay can be compensated for either by realigning the data in a postprocessing step or by adjusting the amplitude of the prephasing gradients. The latter concept was employed in all experiments presented in the work. It turned out that this procedure is sufficient to obtain high-quality images with radial sampling, which rendered the implementation of a more complex technique unnecessary.

4.3.6 Off-Resonance Sensitivity

As already discussed in Section 2.3.2 for the spin-echo technique, it is unavoidable in practice that a certain distribution of resonance frequencies exists, i.e. for some of the protons the frequency differs from the value that is expected under ideal conditions. The origins of the frequency deviations can be divided into three categories. First, it is technically challenging to build magnets with a high homogeneity (especially for high-field systems) and, therefore, the local field strength varies within a specified range. Second, magnetic field variations are induced at susceptibility boundaries of the object [63], in particular at air-tissue interfaces like the nasal cavities of the head. Third, depending on their (intra- and intermolecular) chemical environment, protons may experience microscopic field variations, which is called *chemical shift effect* and best known to occur in fat tissue. All these mechanisms are summarized as off-resonance effects and cause a local phase evolution in the affected regions of the object. To account for the off-resonances, the signal equation (2.18) has to be extended by a spatially- and time-dependent phase term

$$S(t) = \hat{c} \cdot \int e^{i\beta(\mathbf{x}) \cdot t} \cdot \rho(\mathbf{x}) \cdot e^{2\pi i \mathbf{k}(t) \cdot \mathbf{x}} d\mathbf{x} , \quad (4.14)$$




Figure 4.8: Off-Resonance artifacts obtained for radial sampling with 180° and 360° coverage in comparison to (ref) an on-resonant reference case. (Top) Magnified views of the PSF center, and (bottom) magnitude reconstructions of circles with varying deviations of the assumed resonance frequency.

where $\beta(\vec{x})$ is a function that describes the local deviation of the resonance frequency. Thus, in the presence of off-resonance effects, the received spatial information is phase modulated, and a proper reconstruction of the object information, i.e. the proton density $\rho(\mathbf{x})$, requires knowledge of the off-resonance map $\beta(\mathbf{x})$. If this condition is neglected and a conventional FT-based reconstruction is performed, the modulation translates into image artifacts whose appearance and strength depends on the trajectory shape and on the acquisition duration.

In the normal Cartesian scheme, the signal is recorded during a constant velocity movement in the readout direction, which is identical for all rows since the sequence repetitions differ only in the amplitude of the prephasing gradient. Therefore, also the signal modulation from the off-resonances is identical for all acquired rows, which simply emerges as an overall linear phase modulation along the readout direction. In image space, this corresponds to a translation in the readout direction, where the shift distance depends on the local offset of the resonance frequency. Consequently, object areas with strong off-resonance effects are shifted in the readout direction whereas areas with a match of the frequency remain unchanged, which results in a distortion of the reconstructed object [64]. However, because the readout time is short, the phase drifts encountered in practice are marginal and distortions are hardly noticeable. The situ-

ation changes when several k-space rows are acquired from a single RF excitation, as in the echo planar imaging technique [65]. Here, rather strong phase modulations can develop among the acquired rows, which produces significant geometrical distortions and causes a severe loss of spatial accuracy [66]. Nevertheless, for the basic Cartesian scanning techniques the off-resonance problem is not of major relevance.

For radial data acquisitions, the artifacts exhibit in a different and more perceivable way. Because in this sampling geometry the readout direction in k-space varies for all repetitions, the experienced phase evolutions cause a shift of the encoded spatial information with a different orientation for each spoke. This causes a blurring effect, which can be explained from the impact of the phase modulations on the PSF. As outlined in Section 4.3.1, each spoke creates a backprojection of a profile that is composed of sinc-functions, and the profile maxima overlap in the PSF center. In the presence of off-resonances, the linear phase modulation creates a shift of each backprojection profile and, therefore, the maxima of the individual backprojections do not coincide as a central peak anymore, which is demonstrated in Figure 4.8. Here, a difference exists between the radial sampling pattern with a 180° and a 360° coverage (see Section 4.1). Because for 180° the readout orientation and, thus, also the shift direction is uniformly aligned for all spokes, a U-shaped intensity concentration is obtained in the PSF, which causes unsymmetrical blurrings of the reconstructed object that spread over the entire image. In contrast, for the 360° coverage the opposing orientation of neighboring spokes yields a symmetrical ring-shaped widening of the PSF, which is advantageous because the blurring artifacts become more localized. As a slight drawback, it may lead to increased streaking artifacts, which, however, is secondary to the amelioration of the image blurring.

A correction of the off-resonance artifacts is rather complicated for two reasons. First, an unwinding of the phase modulations requires knowledge of the off-resonance map $\beta(\mathbf{x})$, which is a priori not given. Information about the local off-resonances can be gained using phase mapping techniques, which employ multiple measurements with different echo times to estimate the phase evolution [67, 68, 69]. However, such measurements are often time consuming, and the techniques face additional problems like phase wrappings. Second, because the modulation depends on both, acquisition time and spatial position, the calculation of a corrected image is very computationally intensive. To enable a practicable calculation, it is, therefore, necessary to introduce a segmentation of either the time or the frequency deviation [70, 71, 72]. Thus, the development of a robust compensation technique for routine applications is non-trivial, and a more rational strategy consists in the experimental reduction of off-resonance effects. This is possible with short echo times or RF refocusing techniques, allowing to eliminate most visible artifacts or, at least, to diminish their strength to a degree

that is tolerable. However, for applications that explicitly require long echo times, for example, to obtain a T_2^* contrast, radial sampling is less suited and certainly a second choice due to the more pronounced off-resonance sensitivity relative to the Cartesian approach.

4.3.7 Consistency Criteria

Based on the Fourier slice theorem, it is possible to assign a geometrical meaning to the information encoded in each individual spoke, i.e. it corresponds to a projection through the object. This is a unique property of the radial trajectory since for most other sampling schemes a full data set is required before information about the object shape can be extracted. Therefore, it is possible to formulate several consistency criteria for radial data sets, which can be exploited either to detect inaccuracies or to perform first-order artifact compensations [73, 74, 75].

For an ideal static object, the area under the projection profiles is invariant for different view angles. The validity of this statement can easily be seen for the x-ray analogon (refer to Section 4.2.1), because for x-ray transmissions the total amount of absorbed radiation apparently does not depend on the beam angle. The area corresponds to the zeroth moment (or sum) of the profile $P_\phi(h)$

$$Q_0(\phi) = \sum_h P_\phi(h) , \quad (4.15)$$

which, in theory, should give a constant value for all measured spokes. Thus, a significant deviation from the average Q_0 value indicates a corrupted or inconsistent k-space line, for instance due to blood inflow during the acquisition or spin saturation effects. For the latter type of effects, a first-order compensation can be achieved by weighting each spoke according to its zeroth-moment value, which balances intensity variations among all spokes. This allows to correct for slight artifacts that may arise in rapid measurements if the data recording is started before a full steady-state condition of the magnetization is reached (see Section 2.3.3). It is noteworthy, however, that the approach assumes homogeneous intensity changes throughout the object. Because saturation effects mainly depend on the T_1 relaxation rate, which is in general a locally varying quantity, it will be hardly effective or even fail if strong deviations are present within a local object section. Nevertheless, because the method is simple and showed not to worsen artifacts if the aforementioned assumption is tackled, it was employed in all radial reconstructions presented.

Furthermore, it can be shown that for a static object the first moment (or central mass)

of the projections is a sinusoidal function of the projection angle

$$Q_1(\phi) = \sum_h P_\phi(h) \cdot h = \sqrt{A^2 + B^2} \cdot \sin(\phi + \arctan(B/A)) , \quad (4.16)$$

where A and B are functions of the object [74]. Hence, deviations from the expected (object independent) sinusoidal frequency allow to identify spokes that suffer from object motion or other effects. It is then possible, for instance, to reject these spokes and to measure the affected k-space lines again before reconstructing a final image. Alternatively, the projections can be shifted in such a way that the first moment takes the expected value, yielding a first-order correction for translational motion. Similar criteria can also be derived for higher profile moments, which can be exploited to additionally account for rotational motion [74]. However, such correction approaches are expected to work only for rigid object motion and might be sensitive to concurrent effects like phase changes caused from a locally varying phase profile of the receive coil. In this regard, it should be reminded that for modern MRI systems the projection profiles are complex-valued and, therefore, require separate processing of the real and imaginary part. This might introduce complications if unexpected transitions between both components occur and, thus, respective techniques were not pursued in the present work. Nevertheless, the general possibility to assess the signal quality based on a set of consistency criteria might add a significant benefit in specific applications.

4.4 Variants of the Sampling Scheme

The majority of work presented in the thesis focuses on the radial sampling scheme in its basic form, where full spokes are acquired from individual slice-selective RF excitations. However, several interesting variants exist that offer complementary features while inheriting most of the salient properties discussed in the foregoing sections.

4.4.1 Multi-Echo Acquisitions

For applications that require a particularly fast and efficient acquisition of the image data, it is possible to measure multiple k-space lines after a single RF excitation pulse. This reduces the sequence overhead arising from the slice selection, the signal spoiling, the repetition delay for magnetization recovery, and other optional preparation steps. Such techniques are called *multi-echo* sequences and routinely used in combination with the Cartesian scheme [5]. For the radial sampling scheme, an adaptation of the multi-echo strategy has been demonstrated with purely gradient-based sequences [39], with sequences employing RF refocusing pulses [76, 77, 78], and with hybrid sequences

[79]. However, due to the pronounced off-resonance sensitivity of the radial scheme, the number of acquirable spokes per excitation is limited to a low value if no spin refocusing is applied. In particular, radial approaches that try to acquire the complete slice information from only a single excitation (similar to the EPI technique) yet failed to demonstrate convincing results [80, 81]. As a general side effect of the accelerated acquisition, the spokes sampled at different echo times exhibit a non-uniform contrast weighting, which might lead to image artifacts due to the merging of inconsistent data in the k-space center. On the other hand, when using a dedicated reconstruction approach, it is possible to exploit this situation to estimate the local relaxation rates, and a respective method is presented in Chapter 7.

4.4.2 Partial Acquisitions and Ramp Sampling

A different sequence variation aiming at higher acquisition speed consists in sampling only partial spokes, i.e. asymmetrically reduced spokes, which leads to a shorter pathway in k-space and allows for imaging at higher update rates [82, 83]. In addition, the reduction of the spoke length enables shorter echo times, as desirable, for example, for lung imaging with hyperpolarized gases [84]. To reduce the artifacts arising from the one-sided truncation, it is possible to exploit the Hermitian symmetry of k-space, which means that for real-valued objects the k-space information is symmetrical according to

$$S(\mathbf{k}) = \overline{S(-\mathbf{k})} , \quad (4.17)$$

where \overline{S} denotes the complex conjugate of the signal. Thus, under ideal conditions the information in both sides of a spoke is redundant, and the missing information of a partially sampled spoke can be compensated for using homodyne reconstruction techniques [85, 86, 83].

Taking the partial sampling idea a step further, it is also possible to sample only half spokes by moving solely into the positive k-space direction after the excitation. In this case, it is not feasible anymore to balance the missing information from the other half by exploiting k-space symmetries, so that separate measurements for both sides of a spoke have to be performed. However, because the data recording starts directly after the excitation (in the k-space center), the method offers ultra-short echo times (UTE) that are not achievable with any other sampling technique. This is well suited to observe fast moving objects and, in particular, it enables to image object components with very short (sub-millisecond) T_2 relaxation times if combined with specific excitation pulses [87, 88, 89]. The UTE imaging technique is attracting high interest recently, because it depicts tissue species that are invisible to other imaging sequences due to the extremely fast signal decay. Examples are collagen-rich tissues such as tendons, ligaments, and

menisci, as well as calcifications, myelin, periosteum, and cortical and trabecular bone [88, 90]. Therefore, it provides a complementary image contrast and promises to open new applications of MRI, in particular in orthopedics and musculoskeletal imaging. As a technical complication, however, the signal has to be recorded during the ramping phase of the readout gradient, whereas in the standard techniques the data is sampled only during the flat-top time. This causes an increased sensitivity to gradient deviations and eddy-current effects. Further, the opposing readout direction for both parts of a (combined) spoke leads to a pronounced off-resonance sensitivity [91]. Nevertheless, because it reveals information that is not accessible with Cartesian trajectories at all, the UTE technique is one of the most impressive applications of the radial sampling technique.

4.4.3 3D Acquisitions

The concept of radial sampling can also be employed for non-selective RF excitations, and two different strategies exist to handle the additional dimension of k-space. In the simpler form, a normal Cartesian phase encoding step is performed in the third direction, which yields a stack of discs in three-dimensional k-space [92, 93, 94]. An advantage of this scheme is that a usual FFT can be performed as an initial step for all individual samples along the stack, decoupling the reconstruction into separate two-dimensional problems. This circumvents a three-dimensional regridding procedure and, further, allows for the reconstruction of single slices instead of the complete data set at the same time. On the other hand, the approach does not benefit from the advantageous radial properties in the slice direction. Therefore, an interesting option is to acquire spokes that are rotated in all three dimensions, which yields a ball- or star-shaped trajectory. In this case, it is possible to fully exploit the undersampling abilities as well as the readout oversampling, and different ideas have been proposed for the selection and ordering of the spoke angles [95, 96]. For this reason, the technique offers a highly efficient coverage of large imaging volumes, and successful applications have been demonstrated in angiography and other contrast-enhanced studies where both, a high temporal and high spatial resolution is important [97, 98, 99].

4.4.4 Propeller Trajectories

The propeller technique is a hybrid approach between Cartesian and radial sampling, but it is still reasonable to list it as a variant of the radial scheme. In this technique, a set of parallel lines is sampled at a fixed angle, which is referred to as *blade* [27, 100]. Blades are acquired at different angles like in the radial scheme and, thus, the propeller trajectory combines complementary properties from Cartesian and radial sampling.

For instance, due to the Cartesian geometry of the individual blades, standard parallel imaging techniques can be applied to each blade in a straightforward manner [101], which is not the case for a pure radial trajectory. Moreover, it is possible to acquire a full blade from a single excitation using an EPI-alike readout, leading to a significant acceleration of the acquisition [102, 103]. However, the outstanding property of the propeller technique is that it allows for a highly efficient motion correction. Because low resolution images can be reconstructed from each individual blade, it is possible to perform a registration of the single images and to combine them afterwards, yielding a high-resolution motion-corrected image [104, 27]. Therefore, the technique arises as a very interesting option for situations where heavy object motion is unavoidable (for example, during examinations of uncooperative patients), and propeller sequences are now commercially available from most vendors of MRI systems. Noteworthy, however, the scanning time is remarkably longer than for normal Cartesian and radial sampling due to the stronger oversampling of the k-space [105].

4.5 Summary

In the radial sampling scheme, the k-space data is acquired along intersecting spokes instead of parallel rows. Corresponding measurement sequences can be derived rather easily from existing sequences with a Cartesian acquisition. However, due to the non-equidistant sampling positions a special reconstruction strategy is required, and the filtered backprojection method or the regridding technique are commonly used. Both approaches are closely related and differ only in the interpolation technique used, which is needed to map the acquired information from the spoke geometry onto a grid. Because the regridding technique offers a slightly higher flexibility, it is more frequently used in current practice.

Radial sampling offers several advantages over the Cartesian scheme that include a lower sensitivity to object motion and the ability to perform readout oversampling in all directions without additional acquisition steps, which eliminates any aliasing effects. Further, the k-space center is oversampled, which yields an interesting undersampling behavior. Although the reduction of the spokes leads to spurious streaking artifacts, a high degree of the object information remains visible even for stronger reduction factors, which is not the case for Cartesian sampling. Moreover, each sampled spoke captures an equal amount of low and high spatial frequency information, which offers more homogeneous image updates in dynamic MRI applications. Finally, it is possible to apply a number of consistency criteria for the detection and first-order correction of inaccuracies in the measured spokes. On the other hand, the number of k-space lines for a fully sampled data set is 57% higher than in the Cartesian scheme, which prolongs the

examination time. Radial sampling is also more sensitive to deviations of the gradient time courses, which, however, is a manageable problem with modern MRI systems. Further, the technique suffers from a pronounced sensitivity to off-resonance effects, which arises from the varying readout direction in k-space as well as the intersection of all spokes in the k-space center. Therefore, radial sampling is less suited to obtain images with T_2^* contrast, and practical applications should employ short echo times or RF refocusing pulses.

Chapter 5

Iterative Reconstruction from Incomplete Radial Data

In this chapter, the problem of the streaking artifacts is addressed, which appear when only a low number of spokes is acquired. A novel reconstruction technique is presented that compensates for the missing information by incorporating prior object knowledge. The feasibility of the method is demonstrated with experimental data from studies of phantoms and the human brain in vivo acquired using a radial spin-echo sequence. Finally, the relation to the recently proposed compressed-sensing concept is outlined.

5.1 Reconstruction from Undersampled Data

As discussed in Section 4.1, the radial acquisition scheme demands a higher number of sampled lines to fully cover all areas of k-space. In fact, $\pi/2 \cdot n$ spokes have to be acquired for an image with a base resolution of n pixels in order to ensure that the outmost samples of two neighboring spokes have a maximum distance of $\Delta k = 1/\text{FOV}$, which corresponds to the well-known Nyquist condition for conventional Fourier imaging. While this requirement enables to obtain high-quality images with conventional reconstruction methods like regridding, it prolongs the data acquisition by about 57% relative to that of a corresponding fully-sampled Cartesian data set. This factor is, of course, highly undesirable as it increases the total duration of the examination and hampers the use of radial techniques in clinical settings.

If the number of acquired spokes is reduced to a value far below the recommended value, the reconstructed image presents with two characteristic features: while most object information remains visible at good spatial resolution, the use of a regridding (or filtered back projection) approach results in streaking artifacts. This property has already been addressed in Section 4.3.1 on the PSF, and the artifacts are demonstrated




Figure 5.1: *Regridding reconstructions (Shepp-Logan phantom, 256×256 matrix) using simulated data from 402, 64, and 24 spokes (256 data samples). The lower right panel shows the Fourier transform of the image reconstructed from 24 spokes. It reveals unmeasured gaps in the k -space in-between spokes (arrows). The reconstructions from 64 and 24 spokes suffer from streaking artifacts caused by undersampling.*

again in Figure 5.1 for a reduction of the spokes from 402 to 64 and 24. Although the strength of the streaking artifacts increases with the extent of undersampling, it is remarkable how much information about the object can still be seen in an undersampled image from only 24 spokes.

The origin of the streaking artifacts may be best understood when considering the Fourier transform of an undersampled regridding image, which is shown in the bottom right of Figure 5.1. It can be seen that the resulting k -space pattern matches the acquired data at the spoke positions, but in-between the spokes the Fourier transform is zero (except for a small surrounding of the spokes resulting from the convolution with the interpolation kernel). Obviously, this solution with many gaps and jumps is not an accurate representation of the Fourier transform of the true object, which explains the failure of the regridding method for undersampled data sets. Therefore, the conventional reconstruction approaches are not appropriate for strongly undersampled acquisitions. To obtain improved reconstructions from such data, it is necessary to employ a dedicated technique that takes the undersampling into account, so that distracting streakings are removed while the visible object information is preserved.

5.2 Formulation as Inverse Problem

If the measured data from all spokes is stacked into a vector \mathbf{y} with each entry y_i corresponding to a single sample value, then the reconstruction procedure can be seen as estimating a stacked image vector \mathbf{x} with n^2 pixel intensities x_i from the given data vector \mathbf{y} . When only a limited number of spokes is acquired, the size of the data vector \mathbf{y} is usually smaller than the desired image vector \mathbf{x} . Because in this case the problem is underdetermined, one may address it in the opposite direction: suppose we have given an image \mathbf{x} and want to calculate the corresponding data vector \mathbf{y} . This can be achieved by a Fourier transformation of the image and an evaluation of the image's Fourier transform at the trajectory positions using a k-space interpolation (like in the regridding approach). Because these operations are linear, they can be combined into a single matrix \mathcal{A} , which is denoted as *system matrix*. Thus, the forward problem may be written as

$$\mathbf{y} = \mathcal{A}\mathbf{x} , \quad (5.1)$$

and, apparently, the reconstruction procedure corresponds to the inverse problem. However, instead of trying to directly invert this equation to obtain an image vector \mathbf{x} from a given data vector \mathbf{y} , it is more advantageous to iteratively estimate an image vector \mathbf{x} that fits best to the given data vector \mathbf{y} . This is because the problem is not only ill-posed but also very large, and the data vector might be contaminated by Gaussian noise.

How well the image estimate fits to the acquired data can be measured by calculating the L_2 norm of the residuum

$$\Phi(\mathbf{x}) = \frac{1}{2} \|\mathcal{A}\mathbf{x} - \mathbf{y}\|_2^2 , \quad (5.2)$$

which is referred to as *cost function* of the problem. Because an image should be found that best represents the measured data, it is necessary to look for a vector \mathbf{x} that minimizes the cost function (5.2)

$$\mathbf{x} = \underset{\mathbf{x}}{\operatorname{argmin}} \Phi(\mathbf{x}) . \quad (5.3)$$

Finding a solution to this equation requires a highly efficient optimization method due to the large size of the parameter space. A suitable approach for such problems is the conjugate gradient method. It has initially been presented by Hestenes and Stiefel in 1952 for the solution of linear systems and in the meantime successfully applied to MRI reconstruction problems [106]. The method has been extended to non-linear optimization by Fletcher and Reeves in 1964, and since then a number of optimized non-linear conjugate gradient approaches have been developed [107]. The optimization

procedure employed in the present work is based on a very recent variant with improved convergence properties [108], which proved capable of solving Eq. (5.3) in a more efficient way than the conventionally used versions.

The conjugate gradient method is an iterative two-step scheme, which is repeated until a satisfying solution has been found. First, a search direction is estimated in parameter space and, second, a line search into that direction is performed until the minimum of the cost function in this direction has been identified. The search direction is obtained by calculating the gradient of the cost function for the current estimate and by superposing it with the prior search direction, scaled by a factor that guarantees the conjugacy of successive search directions [109]. The gradient is a vector with n^2 entries, which correspond to the derivatives with respect to each entry (or pixel) of the image estimate. Obviously, an optimal image estimate is reached if the gradient of the cost function vanishes. For a cost function of type (5.2), it is well-known that the gradient is given by

$$\nabla\Phi(\mathbf{x}) = \mathcal{A}^\# \mathcal{A} \mathbf{x} - \mathcal{A}^\# \mathbf{y} = \mathcal{A}^\# (\mathcal{A} \mathbf{x} - \mathbf{y}) , \quad (5.4)$$

where $\mathcal{A}^\#$ denotes the adjoint matrix to \mathcal{A} , that is the transposed matrix with each entry replaced by its complex conjugate. As the matrix \mathcal{A} performs a Fourier transformation followed by an interpolation to the spokes, the matrix $\mathcal{A}^\#$ performs an interpolation from the spokes to a grid followed by an inverse Fourier transformation. Hence, applying the adjoint matrix to the measured data is similar to a regridding procedure without a density compensation. It is important to point out that $\mathcal{A}^\#$ is not the inverse matrix to \mathcal{A} , because the forward operation \mathcal{A} is in general non-invertible due to the projection of the Fourier transform to the spokes.

The right part of Eq. (5.4) gives insight into how the iterative reconstruction process works. At every step of the algorithm, the image estimate \mathbf{x} is mapped to the frequency domain by multiplication with \mathcal{A} . It is then compared how well the estimate fits to the measured data by calculating the difference to \mathbf{y} . If the estimate matches well, the residuum vector contains only small entries, otherwise it contains large entries. In the latter case, the algorithm needs to know how to modify the image estimate in order to improve the match of the samples in the frequency domain. This information is obtained by mapping the residuum back to the image space by a matrix multiplication with $\mathcal{A}^\#$.

The middle part of Eq. (5.4) allows for another view of the reconstruction process. Because applying the matrix \mathcal{A} to the estimate gives the corresponding samples in k-space, successive application of $\mathcal{A}^\#$, i.e. $\mathcal{A}^\# \mathcal{A} \mathbf{x}$, may then be understood as convolving the image estimate with the PSF of the trajectory. Due to the higher sample density in the k-space center, this yields a blurring of the image estimate \mathbf{x} for radial trajectories. Further, a multiplication of the adjoint matrix $\mathcal{A}^\#$ with the data vector \mathbf{y} gives




Figure 5.2: Illustration of two equally valid reconstructions of a rectangle from only two radial projections. (a) True solution and (b) solution suffering from streaking artifacts. Both solutions are identical at the measured positions (spokes) in k -space. Numbers indicate hypothetical pixel intensities.

an image comparable to that of a regridding solution (except for the missing density compensation). Thus, trying to match the blurred image estimate $\mathcal{A}^\# \mathcal{A} \mathbf{x}$ with the quasi regridding image $\mathcal{A}^\# \mathbf{y}$ is similar to a general deconvolution approach. This has previously been pointed out by Delaney et al. [110] for the case of iterative parallel-beam tomography reconstruction. Indeed, Eq. (5.2) has the same form as common approaches used in image restoration and image denoising. For this reason, an iterative solution of Eq. (5.2) implicitly yields a density compensation, where the accuracy of the deblurring depends on how well the system matrix \mathcal{A} models the true process underlying the generation of the data vector \mathbf{y} .

5.3 Incorporation of A Priori Knowledge

When reconstructing an undersampled radial data set by optimizing Eq. (5.2), the obtained image still exhibits streaking artifacts. This is not surprising as the procedure does not measure the accuracy of the estimate at any other position in k -space than at the acquired spokes. To illustrate the point, Figure 5.2 shows two reconstructions of a rectangle from just two spokes. While Figure 5.2a represents the true object, the image in Figure 5.2b is degraded by streaking artifacts. Equation (5.2) can not tell which solution is better because in k -space both solutions yield an identical pattern at the positions of the spokes, and differences occur only in-between the spokes. In other words, because the estimate’s FT is projected onto a limited set of spokes, which is an idempotent operation, it is impossible to recover the missing information with a “plain” deconvolution approach like Eq. (5.2).

In order to obtain a better estimate than the regridding solution, it is, therefore, necessary to extend the cost function (5.2) by some kind of complementary plausibility

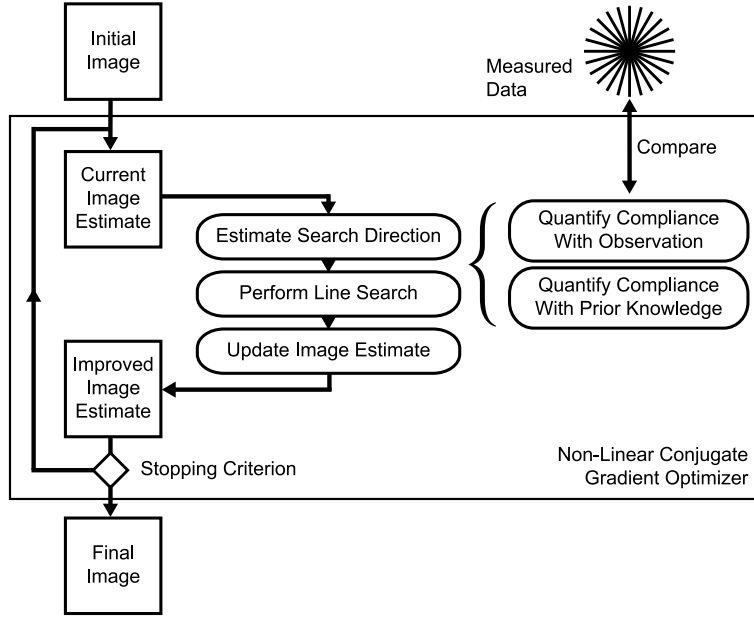


Figure 5.3: Schematic diagram of the proposed iterative reconstruction technique. The procedure has been formulated as an inverse problem. The solution employs a non-linear conjugate gradient method to obtain an image estimate that complies with the measured data as well as prior knowledge.

rating. This can be achieved by adding penalty functions to the cost function, which is known as *regularization*. Penalty functions assign a high value to solutions that are inconsistent with a priori knowledge about the true object and, thus, drive the algorithm to find a solution complying with both, the measured data and certain prior assumptions – as illustrated in the schematic flow chart shown in Figure 5.3. However, the challenge in selecting respective criteria is that they should not be too specific about the object and keep the problem optimizable. This requires the penalties to be convex functions, which can be handled with global optimization techniques. Accordingly, the regularized functional takes the form

$$\Phi(\mathbf{x}) = \frac{1}{2} \|\mathcal{A}\mathbf{x} - \mathbf{y}\|_2^2 + \sum_i \lambda_i \cdot R_i(\mathbf{x}), \quad (5.5)$$

where $R_i(\mathbf{x})$ are the penalty functions. The coefficients λ_i represent tuning factors that allow for shifting the preference from matching the measured data to satisfying the a priori knowledge. In fact, because the measured data is always contaminated by noise to some degree, the search for a perfect match to the measured data is usually not a good strategy as it drives the image estimate to render the experimental noise. To compensate for this effect, it is necessary to adjust the coefficients λ_i in accordance with the signal-to-noise ratio of the acquired data.

For radial acquisitions, there are several reasonable choices of how to restrict the solution space of the image estimation process. If knowledge about the size of the object is

available, it is possible to penalize image intensity outside of the assumed object. In particular, due to the rotational symmetry of radial sampling, all image intensity outside a circular FOV can be usually considered as artifactual. A corresponding penalty function can be formulated as

$$R_{\text{FOV}}(\mathbf{x}) = \sum_i \phi(x_i) , \quad (5.6)$$

where

$$\phi(x_i) = \begin{cases} |x_i|^2 & x_i \notin \text{cFOV} \\ 0 & x_i \in \text{cFOV} \end{cases} \quad (5.7)$$

and cFOV denotes the circular FOV.

Another penalty which turned out to be very effective in general image restoration is the restriction of the parameter space to positive values. It can be achieved by using the penalty function

$$R_{\text{pos}}(\mathbf{x}) = \sum_i \varphi(x_i) , \quad (5.8)$$

where

$$\varphi(x_i) = \begin{cases} x_i^2 & x_i < 0 \\ 0 & x_i \geq 0 \end{cases} . \quad (5.9)$$

Suppression of negative values prevents the algorithm from inserting negative fill values into the image, a tendency often performed by the unconstrained algorithm to better match the measured data. This, however, leads to inaccurate image estimates.

At first glance, the exclusion of negative values seems to ideally apply to the MRI situation where the measured physical quantity, that is the proton density modulated by some relaxation process, is a positive unit. Unfortunately, however, the use of coil arrays with complex-valued sensitivity profiles as well as the occurrence of phase variations within the object forbid a direct application of this criterion. In fact, in most imaging situations neither the real nor the imaginary part of the image can be restricted to positive values.

A third penalty, which has been successfully employed in image restoration, is the restriction of the total variation (TV), initially presented by Rudin et al. in 1992 for denoising applications [111]. The basic assumption of a TV penalty is that the true object consists of areas with constant (or only mildly varying) intensity, which applies quite well to a large class of medical tomographic images. Thus, if the true object is piecewise constant, then the best representation of all image estimates that match at the measured samples should be given by the one with the lowest derivatives at all pixel

positions, that is the one minimizing the total variation

$$R_{\text{TV}}(\mathbf{x}) = \sum_i |D_x x_i| + |D_y x_i|, \quad (5.10)$$

where D_x and D_y denote the derivatives in x and y direction, respectively. The first-order derivative at the pixel position $(u, v) \equiv x_{u+v \cdot n}$ can be calculated from the finite difference between neighboring pixels

$$\begin{aligned} D_x^{(1)}(u, v) &= (u, v) - (u - 1, v) \\ D_y^{(1)}(u, v) &= (u, v) - (u, v - 1). \end{aligned} \quad (5.11)$$

It is very important to note that the total variation in Eq. (5.10) depends on the modulus of the derivatives. This dependency, well-known as ℓ_1 norm, ensures edge preservation in the image and penalizes especially oscillations, which helps to suppress noise patterns as well as Gibbs ringing artifacts (refer to Chapter 8). Replacing the modulus by a square dependency leads to an image with global smoothness, because intensity changes between neighboring pixels become very strongly penalized.

The simple use of first-order derivatives for the TV constraint (5.10) tends to create images with a patchy-looking appearance, because in this case the TV value is minimized for regions with a constant intensity. These images are sometimes appraised as comic-alike and unnatural looking. In contrast, when employing second-order derivatives, the TV value gets small for image regions with constant intensity gradients

$$\begin{aligned} D_x^{(2)}(u, v) &= (u - 1, v) - 2 \cdot (u, v) + (u + 1, v) \\ D_y^{(2)}(u, v) &= (u, v - 1) - 2 \cdot (u, v) + (u, v + 1) \\ D_{xy}^{(2)}(u, v) &= (u, v) - (u - 1, v) - (u, v - 1) + (u - 1, v - 1). \end{aligned} \quad (5.12)$$

Therefore, it is often advantageous to use a combination of first-order and second-order derivatives

$$\begin{aligned} R_{\text{TV}2}(\mathbf{x}) &= \sum_i \sigma \cdot (|D_x^{(1)} x_i| + |D_y^{(1)} x_i|) \\ &\quad + (1 - \sigma) \cdot (|D_x^{(2)} x_i| + |D_y^{(2)} x_i| + |D_{xy}^{(2)} x_i|), \end{aligned} \quad (5.13)$$

where $\sigma \in [0, 1]$ is a weighting factor that allows for slightly tuning the smoothness of the image appearance. It was set to $\sigma = 0.77$ according to the work presented by Geman et al. [112].

The upper row of Figure 5.4 shows reconstructions of the numerical Shepp-Logan phantom from 24 spokes, obtained either by regridding or the proposed inverse formulation




Figure 5.4: Radial image reconstructions (Shepp-Logan phantom, 256×256 matrix) using simulated data from 24 spokes (256 samples). (Top left) Regridding and (top right) the proposed iterative technique with prior knowledge. (Bottom) Corresponding Fourier transforms. The iterative technique reconstructs the image of the object without streaking artifacts. Accordingly, its Fourier transform recovers the unmeasured gaps in k-space in-between spokes.

with penalties as presented in this section. A comparison of the images clearly demonstrates the superior performance of the iterative method in reducing streaking artifacts, which for the simulated data have been effectively removed. The lower row of Figure 5.4 depicts the corresponding Fourier transforms. It turns out that the incorporation of a priori information by appropriate penalty functions leads to a proper recovery of k-space information in-between the spokes.

5.4 Experimental Complications

When trying to employ the described iterative strategy for the reconstruction of real MRI data instead of simulated data, two additional difficulties arise. First, the observed object is usually complex-valued due to off-resonance effects and other phase perturbations. Although the desired result in the end is a real-valued image, it is impossible to remove the phase variations before resolving the spatial encoding. The removal of the phase is conventionally accomplished in a last step by calculating the magnitude im-

age, but this operation is non-linear and can not be integrated into the system matrix. Therefore, it remains necessary to deal with the complex-valued nature of the object during the reconstruction procedure. Second, modern MRI systems use an array of receive coils each having a varying intensity and phase profile, as described before. Again, a suitable combination of the signals from all individual coils can only be achieved after spatially resolving the object. As a consequence, any image reconstruction approach has to cope with the separate coil signals.

To this end, an iterative two-step reconstruction approach is proposed. The first step attempts to estimate sensitivity profiles for all coil elements, which are then utilized in the iterations of the second step to combine the individual coil channels and to remove the phase variations each time when mapping between the frequency domain and image space. Thus, the second step renders a combined and real-valued image. Noteworthy, because in this step a single image is found that complies with the observations from all individual coil elements at the same time, the method implicitly makes use of the locally varying coil sensitivities to gain additional object information in a similar manner to parallel imaging techniques like SMASH, GRAPPA, or SENSE [113,114,115,106]. However, while most of these techniques try to extract the information in a concrete form (for example, as synthesized k-space samples), in the present approach the exploitation of the coil sensitivities manifests as an improved conditioning of the optimization problem. Moreover, an attractive feature of the approach is that neither reference data nor shared data from prior measurements is needed to estimate the coil profiles, whereas in most parallel imaging methods additional acquisition steps are required for this task.

In the first step of the approach, the signals from all coil channels are handled separately and, thus, a set of individual images is calculated. Here, the real and imaginary parts are treated as independent parameters, leading to a complex-valued image estimate for every coil channel. It is well-known that MRI coil profiles are smooth functions that vary only slowly and do not have sharp edges. This knowledge is incorporated by using a quadratic regularization of the image derivatives, which leads to globally smooth images as discussed in the foregoing section

$$R_{\text{coil}}(\mathbf{x}) = \sum_i (D_x^{(1)} x_i)^2 + (D_y^{(1)} x_i)^2, \quad (5.14)$$

where $D_x^{(1)}$ and $D_y^{(1)}$ are the derivative operators of first order according to Eq. (5.11). After finishing the iterations for all coils, a sum-of-squares image is computed. A division of the single channel images by the sum-of-squares image yields the respective coil profiles. Noteworthy, the estimated coil profiles also include other phase variations inside the object as the real-valued sum-of-squares image has been taken as a reference. Because the penalty function $R_{\text{coil}}(\mathbf{x})$ depends quadratically on \mathbf{x} , the line search, which




Figure 5.5: Schematic diagram of the procedural implementations of (left) the system matrix A and (right) the adjoint system matrix $A^\#$ that are used to map from image domain to frequency domain and vice versa. For details see text.

is part of the conjugate gradient iteration, requires only one step and, thus, only a low number of function evaluations is needed to obtain a reasonable image. Therefore, the coil profile estimation step takes only moderate computational time.

For the second, combined reconstruction step, the raw data from all coil channels is stacked into one data vector \mathbf{y} . Further, the system matrix \mathcal{A} is extended by a multiplication with the corresponding coil profile before performing the Fourier transformation for every channel. Figure 5.5 shows a flow diagram of the operations that are executed by the system matrix \mathcal{A} and the adjoint matrix $\mathcal{A}^\#$ to map between frequency and image space. By combining the data from all coil channels into the data vector \mathbf{y} , the algorithm now renders one image estimate that matches to the observations from all coils. Moreover, removing the phase variations with the use of the estimated coil profiles allows to discard the imaginary part of the image estimates and enables to apply constraints on negative values, which otherwise would not be possible. Noteworthy, a combined coil reconstruction also ensures that the TV constraint remains applicable in a multi-coil setup. Otherwise, the intensity modulation from the coil profiles would conflict with the idea of piecewise-constant images.

Outline of the Complete Algorithm

To estimate the coil sensitivity profiles, individual complex-valued images \mathbf{x}_j are reconstructed for each channel j by iteratively solving the optimization problem

$$\mathbf{x}_j = \underset{\mathbf{x}_j}{\operatorname{argmin}} \left\{ \frac{1}{2} \|\mathcal{A}\mathbf{x}_j - \mathbf{y}_j\|_2^2 + \lambda_{\text{FOV}} \cdot R_{\text{FOV}}(\mathbf{x}_j) + \lambda_{\text{coil}} \cdot R_{\text{coil}}(\mathbf{x}_j) \right\}, \quad (5.15)$$

where \mathbf{y}_j denotes a vector containing all measured values from channel j . The penalty functions R_{FOV} and R_{coil} are used as defined in Eq. (5.6) and (5.14), where the weights of the penalty terms λ_{FOV} and λ_{coil} have to be adjusted according to the value range of the measured data (see next section). The system matrix \mathcal{A} comprises three operations

$$\mathcal{A} = \mathbf{G}_{\text{KB}} \cdot \mathcal{F} \cdot \mathbf{M}_{\text{KB}}^{-1}, \quad (5.16)$$

where $\mathbf{M}_{\text{KB}}^{-1}$ denotes the pre-compensation for the roll-off effect according to Eq. (4.9), \mathcal{F} is the two-dimensional fast Fourier transformation, and \mathbf{G}_{KB} denotes the interpolation to the sampling locations in k-space using a convolution with the Kaiser-Bessel window and subsequent evaluation along the trajectory. A sum-of-squares image is calculated from these images with

$$\mathbf{C}_{\text{sos}} = \sqrt{\sum_j \mathbf{x}_j \cdot \overline{\mathbf{x}_j}}, \quad (5.17)$$

which is then used to obtain the complex-valued coil sensitivities

$$\mathbf{C}_j = \mathbf{C}_{\text{sos}}^{-1} \cdot \mathbf{x}_j . \quad (5.18)$$

Here, \cdot denotes a component-wise multiplication of the two vectors. In the final stage, the real-valued image \mathbf{x} is calculated by solving the optimization problem

$$\begin{aligned} \mathbf{x} = \operatorname{argmin}_{\mathbf{x}} \{ & \sum_j \frac{1}{2} \|\mathcal{A}_j \mathbf{x} - \mathbf{y}_j\|_2^2 + \lambda_{\text{FOV}} \cdot R_{\text{FOV}}(\mathbf{x}) \\ & + \lambda_{\text{pos}} \cdot R_{\text{pos}}(\mathbf{x}) + \lambda_{\text{TV2}} \cdot R_{\text{TV2}}(\mathbf{x}) \} , \end{aligned} \quad (5.19)$$

where \mathcal{A}_j denotes that the system matrix in the final stage includes a multiplication with the j th coil profile according to

$$\mathcal{A}_j \mathbf{x} = \mathcal{A}(\mathbf{C}_j \cdot \mathbf{x}) . \quad (5.20)$$

The penalty functions R_{pos} and R_{TV2} are given by Eq. (5.8) and (5.13). Because the total-variation term R_{TV2} depends on the modulus of the estimate components, a non-linear optimization technique is required for solving Eq. (5.19). When utilizing a technique based on the conjugate-gradient method, as done in the proof-of-principle implementation, it is necessary to evaluate the cost functions (i.e. the functions inside the brackets of Eq. (5.15) and (5.19)) as well as their gradients. The gradients of the ℓ_2 -norm terms can be deduced from Eq. (5.4), and the gradients of the penalty terms can be obtained by deriving the penalty functions with respect to all components of the estimate vector.

5.5 Proof-of-Principle Application

Acquisition Parameters

To demonstrate the feasibility of the proposed technique for reconstructing undersampled data sets, a radial 2D spin-echo MRI sequence has been implemented on the MRI system described in Chapter 3. RF-refocused spin echoes rather than gradient echoes were chosen to avoid putative complications from the sensitivity to off-resonance effects, which are not related to the reconstruction process studied here. The MRI sequence was derived from a slice-selective spin-echo sequence of the manufacturer and modified to radial acquisitions with a readout oversampling factor of two. Figure 5.6 shows the corresponding sequence timing diagram. The RF refocusing pulse has been enclosed by crusher gradients to cancel spurious signals from magnetization that is reexcited if the actually generated flip angle differs from 180° . To receive the signal, the head coil array

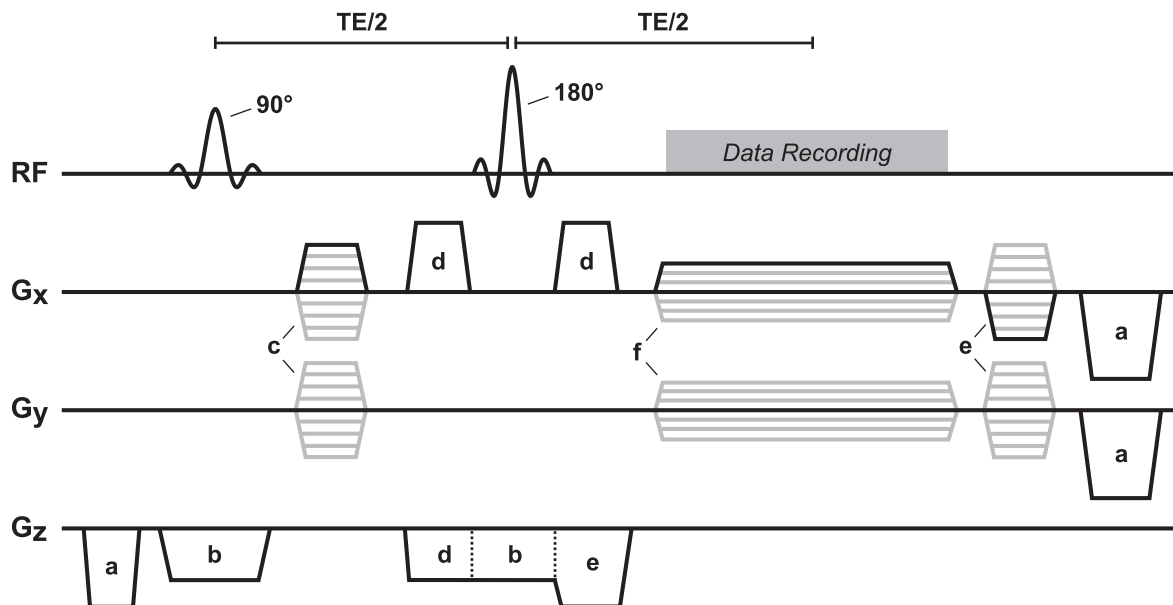


Figure 5.6: Timing diagram of the radial spin-echo sequence. Gray lines indicate that the gradient amplitudes vary for different cycles of the sequence. (a) Spoiler, (b) slice-selection, (c) prephasing, (d) crusher, (e) rewinder, and (f) readout gradients.

was used in triple mode, yielding all 12 channels with combinations of the coil elements (see Eq. (3.2)). Gradient timing errors were corrected using the technique presented by Speier et al. [62]. The measurements were conducted for water phantoms as well as the human brain in vivo to get experimental data from realistic imaging situations. Written informed consent was obtained from all subjects prior to the examination.

All images were acquired with a base resolution of 256 pixels, covering a 230 mm FOV (slice thickness 2 mm). The number of spokes varied from 8 to 96. The phantom images were acquired with a repetition time $TR = 4000$ ms and echo time $TE = 11$ ms (bandwidth 180 Hz/pixel), while the in vivo images of the human brain were acquired with $TR/TE = 2500/50$ ms (bandwidth 180 Hz/pixel) and $TR/TE = 3000/80$ ms (bandwidth 90 Hz/pixel) for T_2 contrast.

Image Reconstruction

The image reconstruction procedure was performed using the software package MRISim. To speed up the interpolation operations that are repeatedly carried out within the iterations, a look-up table was calculated in a preparation step which contains all coefficients needed to interpolate from spoke to grid data and vice versa. The coefficients were calculated using the Kaiser-Bessel window given by Eq. (4.8) with $L = 6$ and $\beta = 13.8551$. Further, a matrix containing the values for the roll-off correction was precalculated according to Eq. (4.9).

Prior to starting the iterations for a particular image, a preprocessing of the measured




Figure 5.7: Estimated MRI coil profiles (phantom, 256×256 matrix) using experimental data from 48 spokes (256 samples) and the proposed iterative technique. The profiles correspond to the four primary modes of the 12-channel receive-only head coil. The MRI sequence was a radial spin-echo sequence ($TR/TE = 4000/11$ ms, 230 mm FOV, bandwidth 180 Hz/pixel).

k-space samples was performed. The procedure included an unwrapping of phase offsets in the center of k-space (as described in Section 4.3.2) to compensate for interference artifacts that might arise, for example, from through-plane motion. In a subsequent step, the consistency criterion for the zeroth moment was employed to obtain a first-order correction for spoke intensity deviations (see Section 4.3.7). Such effects can occur when measuring in a transient phase of the magnetization, i.e. during the approach to steady-state conditions, and might result in smearing artifacts. Therefore, the data of each spoke was Fourier transformed, the zeroth moment was calculated, and the resulting spoke intensity was used in the iterations to weight the calculated spokes before matching them to the measured data.

In the initial coil estimation step, a moderate penalty for image intensity outside the circular FOV was used by setting $\lambda_{\text{FOV}} = 1$ in Eq. (5.5). Edges were strongly penalized in this step by quadratically constraining the first-order derivatives of the image intensity using Eq. (5.14) with $\lambda_{\text{coil}} = 10$. In the final reconstruction step, image intensity outside the circular FOV was strongly penalized by setting a high value for the corresponding coefficient λ_{FOV} or by even rejecting all intensity in this area. To incorporate the total variation constraint, the magnitude of the image derivatives was penalized with a weighting of $\sigma = 0.77$ for the first order and 0.23 for the second order as suggested by Geman et al. [112]. A value of $\lambda_{\text{TV}^2} = 0.0001$ for the total variation penalty turned out as a robust choice for the studies presented. Further, negative values were

penalized by setting $\lambda_{\text{pos}} = 5$, which stabilized the convergence process but in turn also slowed it down. Nevertheless, it should be noted that these weights were chosen with respect to the value range of the described spin-echo acquisition technique. Thus, the absolute values listed here can not be used in general but have to be adapted to the signal strength and the signal-to-noise ratio of the individual imaging application.

In the proof-of-principle implementation of the algorithm, a fixed number of iterations was used in both reconstruction steps. All images presented were rendered using 30 iterations for the coil estimation and 120 iterations for the final image reconstruction step. However, in many cases a reasonable image quality was already obtained with fewer iterations. Moreover, when initializing the optimizer with a properly scaled regridding image instead of a void estimate, the number of iterations can be reduced to a significantly smaller value, typical numbers being 10 iterations for the coil estimation and 20 iterations for the final reconstruction step. Thus, when launching the optimizer with an already reasonable initial guess, the algorithm can directly continue to improve the regridding solution, and fewer iterations are needed until a proper object reconstruction is reached.

5.6 Experimental Results

Phantom Studies

Figure 5.7 shows experimental coil profiles that were determined by the coil estimation step from a data set of 48 spokes obtained for a water phantom. The profiles are smooth inside the object and do not contain visible object features as expected. The algorithm is unable to determine the coil profile outside the object due to the absence of any signal in these areas, but this poses no major problem for the image reconstruction. The corresponding images of the phantom are summarized in Figure 5.8 together with reconstructions for 96 and only 24 spokes. For comparison, the upper row shows the results of the regridding approach with a sum-of-squares combination of the multiple coil images.

Obviously, the regridding reconstructions suffer from streaking artifacts that increase with decreasing number of spokes. In contrast, the iterative approach renders images without any visible or at least strongly reduced streaking artifacts while maintaining sharp edges. It can be seen that the proposed method is able to reconstruct a high-quality image of the object from only 48 spokes, and there is only a slight gain in image quality if the number of spokes is further increased. In the case of only 24 spokes, the algorithm again outperforms the regridding solution, but it fails in fully recovering the true object and residual streaking artifacts remain visible. However, it should be noted




Figure 5.8: Radial image reconstructions (phantom, 256×256 matrix) using experimental data from 96, 48, and 24 spokes (256 samples). (Top) Regridding and (bottom) the proposed iterative technique. Parameters as in Figure 5.7

that a reconstruction from 24 spokes corresponds to a data reduction factor of more than 16 compared to the 402 spokes recommended for a 256×256 image.

In vivo Studies

Similar to the phantom studies, Figure 5.9 compares reconstructions of radial images from the human brain in vivo using regridding and the proposed iterative method for 96, 48, and 24 spokes. Again, regridding suffers from pronounced streaking artifacts outside as well as inside of the brain, although the artifacts are less clearly visible than in the phantom images due to the occurrence of more complex structural details. Conversely, streaking artifacts are removed (48 spokes) or at least noticeably reduced (24 spokes) when using the iterative approach. Further, the application of the total variation penalty leads to a marked denoising of the images.

The improvement in image quality of the iterative reconstruction technique relative to regridding is even more visible in Figure 5.10 magnifying parts of the brain sections from Figure 5.9 by a factor of three. In order to demonstrate the limits, Figure 5.11 compares brain sections obtained by iterative reconstructions from 48, 32, 24, 16, 12, and 8 spokes. Of course, the data reduction from 48 to 8 spokes is accompanied by a loss of resolution. The effect is best appreciated for selected fine structures, whereas gross anatomical




Figure 5.9: Radial image reconstructions (human brain, 256×256 matrix) using experimental data from 96, 48, and 24 spokes (256 samples). (Top) Regridding and (bottom) the proposed iterative technique. Parameters as in Figure 5.7 except for $TR/TE = 2500/50$ ms.




Figure 5.10: Magnified views of the same data as shown in Figure 5.9




Figure 5.11: Radial image reconstructions (human brain, 256×256 matrix) using experimental data from 48, 32, 24, 16, 12, and 8 spokes (256 samples) using the proposed iterative technique. The MRI sequence was a radial spin-echo sequence with fat suppression ($TR/TE = 3000/80$ ms, 230 mm FOV, bandwidth 90 Hz/pixel).

features such as the ventricles are less affected. Simultaneously, total variation ensures a pronounced reduction of the noise, so that the overall image appearance even for extreme undersampling is of surprising quality.

5.7 Limitations and Extensions

Remaining Artifacts

The results shown in Figures 5.8 to 5.11 demonstrate that the proposed reconstruction technique for undersampled radial MRI yields images with clearly improved quality over the conventional regridding approach. Nevertheless, the algorithm is not able to fully remove the streaking artifacts in the heavily undersampled case of 24 spokes. This effect can be explained by closer inspection of the total-variation constraint, which plays a central role for the removal of such artifacts or, respectively, the recovery of the inter-spoke k-space information in the frequency domain.

The total-variation concept is based on the assumption that the true object is piecewise constant, which implies that only a limited number of edges and intensity jumps are




Figure 5.12: Radial image reconstructions (human brain, 256×256 matrix) using experimental data from 48 spokes (256 samples). (a) Regridding and (b-d) the proposed iterative technique with (b) a low weight, (c) an appropriate weight and (d) an overweight of the total variation constraint. While a proper choice of the total variation penalty yields an efficient denoising without compromising resolution, any overweighting causes a loss of object detail. Parameters as in Figure 5.9.

present in the image. For this reason, an object can be recovered from incomplete data if it can be represented by a limited or sparse number of edges. This condition is obviously fulfilled by the Shepp-Logan phantom shown in Figure 5.4. In this case, the major contribution to the total-variation value of the regridding solution comes from the streaks that overlap to form a texture-like pattern in the image. These undersampling artifacts can be removed by minimizing the total variation, so that the object can be perfectly recovered from only 24 spokes – also stated by Candes et al. [116]. However, the experimental data presented in Figures 5.8 to 5.11 fulfills the needed condition less optimal as the object itself contributes remarkably to the total variation of the image, or conversely, the total variation can not be dominantly ascribed to the undersampling artifacts. Hence, the more complex the true object is, the less accurate can missing information be recovered by restricting the image variation. In fact, it turned out that the algorithm works efficiently for images of the human brain, whereas the efficiency is significantly lower for cardiac images because for this class of images the assumption of a limited total variation is less appropriate.

Moreover, Figure 5.1 illustrates that the width of the streaks from the undersam-

pling increases with decreasing number of spokes. The total variation penalty given by Eq. (5.10) corresponds to the ℓ_1 norm of the image derivative, and, therefore, it is penalizing especially intensity oscillations while maintaining sharp edges. Accordingly, the removal of undersampling artifacts by a total variation constraint is most effective if the overlapping streaks create a strongly varying texture, as in the case of 64 spokes. In contrast, if the width of the streaks is wide and the object itself has a certain complexity, then after some iterations the total variation of the image becomes dominated by the true complexity of the object. Further attempts to minimize the total variation then lead to a removal of actual object features. In other words, there is a tradeoff between residual streaking artifacts and the preservation of object details that has to be considered when reconstructing a complex structured object from a low number of spokes. This is demonstrated in Figure 5.12 comparing a regridding solution to three iterative reconstructions with an increasing weight on the minimization of the total variation. Because the strongest weight led to a visible removal of object detail, it is recommended to choose the weight of the total variation constraint – given by the coefficient λ_{TV} – with respect to the imaging parameters and the object’s complexity. Regardless of this limitation for a complete removal of the artifacts, the proposed reconstruction technique provided visually improved image quality over regridding in all cases tested. Further, it is of course possible to integrate additional or possibly more advanced penalties to support the recovery of unmeasured information in k-space using prior object knowledge. These constraints might be based on recent multi-scale transformations like curvelets and contourlets [117,118], or could be motivated by a Bayesian formulation.

A more general problem that arises when reducing the number of spokes (or k-space lines) is a concomitant decrease of the signal-to-noise ratio (SNR), which also applies to partial Fourier imaging and parallel MRI. It turned out that low SNR poses a more severe limitation for the reconstruction of undersampled radial MRI data sets than the putative loss of resolution. Although the use of the total-variation constraint ensures a pronounced denoising while maintaining borders (compare Figure 5.12), it only allows to smooth noise textures but is, of course, incapable to recover object information that is not visible at all due to a low SNR. Moreover, the appearance of noise patterns reduces the general separability of the undersampling artifacts because in this case the “observed” object itself yields a high total-variation value. Therefore, the assumption that the fully-sampled object has a limited total variation is impaired, and the algorithm fails in identifying a reconstruction without undersampling artifacts by finding a solution with minimum total variation.

Computational Requirements

Without doubt, the proposed iterative reconstruction method is by far more computationally demanding than the conventional regridding or filtered backprojection technique. In fact, only a single evaluation of the cost function (5.5) already doubles the computational load required for regridding, and usually multiple evaluations are required during one iteration of the conjugate gradient algorithm. The duration of a single evaluation and the number of needed iterations depend on the degree of the undersampling and on the desired reconstruction quality. Thus, reconstructions from a very low number of spokes usually require a higher number of iterations than reconstructions with only mild undersampling factors. On the other hand, the number of operations needed to interpolate between the grid and spoke data is lower for high reduction factors and, thus, the interpolation can be performed in a faster time. Therefore, it is rather difficult to give general information on the achievable reconstruction times.

The images presented in this work were computed on a system equipped with two Intel Xeon 5060 3.2 GHz dual core processors (see Chapter 3.3). Using the proposed method with the aforementioned reconstruction parameters (256 pixels base resolution, 30 iterations for each coil estimation, 120 final iterations), the calculation of a radial image from 48 spokes took about 520 s on this system. However, running the reconstruction with only 10 iterations for the coil estimation and 20 iterations for the final reconstruction resulted in a suitable image quality already within about 120 s. Moreover, the use of only 4 instead of 12 channels further reduced the reconstruction time to 43 s. Finally, there is still considerable potential for optimizing the speed of the implementation, in particular by parallelizing many of the subroutines. Nevertheless, while the current reconstruction times are still too long for a routine clinical setting, steady progress in processor technology will not take long to render iterative reconstruction techniques more generally suitable for MRI.

Extensions

An attractive feature of the proposed method is that it can be easily adapted to meet different imaging scenarios by integrating more specific knowledge about the object with the use of additional penalty functions. Based on the Bayes theorem, basically every kind of a priori knowledge may be incorporated. A tough limitation, though, is that it is necessary to formulate this knowledge such that the problem remains optimizable, which implies at least convex penalty functions.

Furthermore, the system matrix can be extended to model the generation of the measured MRI signal in more detail and for specific acquisition modes. An example of such

an extension is presented in Chapter 7 for the case of a RF-refocused multi-echo acquisition. A related idea would be to include a modeling of off-resonance effects, which pose a significant problem for non-refocused radial acquisitions with long echo times. This could be done using a time-segmented approximation of the local phase evolution based on field maps, where the field map could be estimated within the same acquisition by slightly shifting the echo time of every other spoke. Thus, the reconstruction would render coil profiles, a field map, and the final combined image from a single data set.

Although penalizing the total variation is particularly well suited for radial trajectories due to the strongly varying patterns created by radial undersampling, the idea can of course be applied to other trajectories as well. Because the technique does not need a density compensation as required for regridding, it allows to reconstruct images from arbitrary trajectories without a prior estimation of the sample density using Voronoi diagrams or comparable methods. A second advantage is that the iterative approach reconstructs objects with absolute values that are independent of the amount of data measured. In contrast, for regridding the absolute values of the object usually depend on the total intensity inserted into the raw data matrix.

5.8 Relation to Compressed Sensing

Very recently, Donoho [119] and Candès et al. [116] developed a mathematical theory for the recovery of signals from highly incomplete data, which is by now commonly known as *compressed sensing* (CS). The CS concept is closely related to the proposed iterative reconstruction approach, and it can be seen as a theoretical foundation for the improved reconstruction quality over regridding. The theory essentially states that a signal (or image) \mathbf{x} can be exactly recovered with certain probability from a highly incomplete, linear measurement

$$\mathbf{y} = \mathcal{A}\mathbf{x} \tag{5.21}$$

if two conditions are fulfilled. (i) The true signal \mathbf{x} is compressible in the sense that it has a sparse representation in some transform basis \mathcal{M} . (ii) The measurement basis \mathcal{A} is *incoherent* to the sparsity basis \mathcal{M} .

For a highly incomplete measurement, the vector \mathbf{y} has fewer entries than the vector \mathbf{x} , and the problem (5.21) is clearly underdetermined. Thus, Eq. (5.21) is fulfilled by a hyperplane of vectors $\hat{\mathbf{x}}$, which all would likewise explain the observation \mathbf{y} . One of these vectors is the minimum energy or minimum ℓ_2 -norm solution

$$\hat{\mathbf{x}} = \underset{\hat{\mathbf{x}}}{\operatorname{argmin}} \|\mathcal{A}\hat{\mathbf{x}} - \mathbf{y}\|_2, \tag{5.22}$$

which corresponds to the regridding solution in the radial MRI problem. However, in most cases this solution does not properly recover the true signal \mathbf{x} . Therefore, the basic idea of compressed sensing is to exploit the compressibility of the true signal to identify a reasonable reconstruction out of all possible solutions. Obviously, because it is assumed in condition (i) that the true signal can be represented sparsely in the transform basis \mathcal{M} , also the reconstruction should have a sparse representation in this basis. Thus, a reasonable reconstruction can be found by searching for the solution with the highest sparsity after performing the transformation \mathcal{M} . The sparsity of a solution can be measured using the ℓ_0 norm, which is defined as the number of non-zero entries of the vector. Unfortunately, searching for a solution with minimum ℓ_0 value would require to evaluate values for all possible solutions, which is intractable in practice. However, Donoho [120] was able to show that for most large underdetermined linear systems also the minimum ℓ_1 -norm solution yields the sparsest solution, which can be found using a convex optimization approach. Therefore, a reasonable reconstruction can be obtained by finding the solution with minimum ℓ_1 -norm value, which can be written as a constrained optimization problem

$$\hat{\mathbf{x}} = \underset{\hat{\mathbf{x}}}{\operatorname{argmin}} \|\mathcal{M}\hat{\mathbf{x}}\|_1 \quad \text{s.t.} \quad \mathcal{A}\hat{\mathbf{x}} = \mathbf{y}, \quad (5.23)$$

where the sparsity transformation \mathcal{M} can be, for example, a wavelet transformation. Here, the side condition ensures that the solution matches to the measured values \mathbf{y} . However, in practice the measured samples are often contaminated by noise and, thus, it is better to search for a solution that matches to the observation in a least-squares sense rather than enforcing an exact match of the samples. This can be achieved by replacing the equality constraint with an inequality constraint

$$\hat{\mathbf{x}} = \underset{\hat{\mathbf{x}}}{\operatorname{argmin}} \|\mathcal{M}\hat{\mathbf{x}}\|_1 \quad \text{s.t.} \quad \|\mathcal{A}\hat{\mathbf{x}} - \mathbf{y}\|_2 < \epsilon, \quad (5.24)$$

where ϵ denotes the standard deviation of the noise.

As stated above, a sparse transform is a necessary condition for a reasonable reconstruction, but this criterion does not exclude that unreasonable solutions might exist, which might also have a sparse representation in the transform basis. This problem is prevented by the incoherence requirement (ii), which essentially means that the samples have to be acquired in such a way that the information content of a single sample can not be represented sparsely in the transform basis. In other words, the artifacts that arise in the minimum ℓ_2 -norm reconstruction for an incomplete measurement have to be of such a nature that they can not be represented with sparse coefficients in the transform basis, whereas the true signal is sparse. Therefore, unreasonable solutions are expected to be non-sparse, and a reasonable reconstruction can be found by solving

Eq. (5.23) or Eq. (5.24). This has been shown to work with overwhelming probability for a random sampling of the signal's Fourier transform. For such measurements, Candès et al. [116] derived the probability of finding the exact solution, so that they were able to quantify respective sampling requirements.

From this point, it is quite easy to see the relation of CS to the iterative reconstruction approach for radial MRI data. In the method proposed here, the Fourier transform of the object is sampled along spokes, which can be seen as an approximation of a random sampling scheme. Further, due to the penalty function Eq. (5.10), a solution is found that minimizes the total-variation value, corresponding to the minimum ℓ_1 -norm solution in the finite-difference basis. Thus, the sparsity transformation consists in the calculation of differences between neighboring pixel values, and the requirement (i) that the true object has to be sparse in the transform basis translates into the assumption of a piecewise-constant object. Moreover, it can be seen from Figure 5.1 that the minimum ℓ_2 -norm artifacts for radial undersampling are of such a type that they create a high TV value from the arising streaks, as required in condition (ii). In fact, it turned out that the reconstruction quality obtained for finite differences is better than for a wavelet transformation because of the high incoherence between the finite differences and the radial sampling basis. Hence, the proposed approach can be seen as a practical application of the CS theory and demonstrates the ability of obtaining improved reconstructions in a real-world scenario. Nevertheless, it is important to note that the equations derived in the seminal CS publications for estimating the reconstruction error are not directly applicable to the present situation, because here the sampling pattern is strictly non-random. Further, in the present approach the solution is found by recasting the constrained optimization problem Eq. (5.24) into an unconstrained optimization problem according to Eq. (5.5). Although a solution of the unconstrained problem is often also a solution to the constrained problem, this is not necessarily the case. In future work, it should, therefore, be investigated whether dedicated constrained optimization techniques like interior point methods [109] offer an improved effectiveness for finding a reasonable solution in the proposed reconstruction approach.

5.9 Summary

If only a low number of spokes is acquired, spurious undersampling artifacts arise when using conventional image reconstruction techniques. The artifacts present as streakings that overlap with the object and originate from the unmeasured gaps in k-space between the spokes. Therefore, conventional techniques like regridding fail to provide a proper reconstruction for incomplete data sets.

For incomplete data, the reconstruction problem is clearly underdetermined, and it is

reasonable to address it in the form of an inverse problem, which can be solved iteratively using the non-linear conjugate gradient method. Because the k-space sampling corresponds to a projection to the spokes, it is further necessary to introduce a priori knowledge about the object to compensate for the missing information. This can be achieved by adding suitable penalty terms to the cost function, which assign a high value to implausible solutions and drive the optimizer to find a solution complying with both, the measured data and the prior knowledge. In particular, constraining the total variation of the reconstructions leads to an effective reduction of streaking artifacts even for high undersampling factors.

In practice, however, the use of modern multi-coil arrays as well as experimental phase variations prevent a direct application of this strategy to MRI. To this end, a two-step procedure is proposed, which first estimates sensitivity profiles for all coils and then renders a final, combined image that complies with the observations from all coils. This strategy enables to employ a total-variation constraint for experimental MRI data, and it allows to exploit localized coil sensitivities without additional calibration steps.

The feasibility of the method has been demonstrated with experimental data from a radial spin-echo sequence for phantom and human brain studies. The results show that the proposed approach is able to obtain images from only a very limited number of spokes with markedly improved quality compared to conventional radial reconstructions. For a very low number of spokes, residual artifacts remain visible because the assumption of a piecewise-constant object is only approximately valid for real-world objects and, thus, a complete separation of the artifacts is not possible. A description of this effect can be derived from the compressed-sensing theory, which serves as a theoretical foundation for the proposed approach. Finally, the current computational speed of the proposed technique is already acceptable for scientific purposes, and foreseeable technical progress promises iterative approaches soon to become part of the MRI instrumentarium for more routine applications.

Chapter 6

Rapid Imaging Using Stimulated Echoes

In this chapter, it is first demonstrated that the adaptation of the single-shot STEAM MRI technique to the radial sampling scheme helps to overcome current limitations. Most importantly, when used in combination with the reconstruction approach proposed in Chapter 5, a remarkable enhancement of the image resolution can be achieved. Finally, the modification to a segmented acquisition is presented, which allows to trade-off between higher SNR and higher acquisition speed.

6.1 Single-Shot STEAM MRI

Many imaging applications require an ultra-fast data acquisition either because of strong object motion or because of the necessity to perform a large number of measurements in limited time. This applies, for instance, to cardiac imaging or to diffusion tensor imaging (DTI) where numerous diffusion directions have to be measured. In such situations, EPI sequences are often used [65], which acquire all k-space rows in a sequential manner after a single RF excitation. However, due to the long readout time in the EPI technique, significant phase evolutions arise in areas with strong off-resonance effects, which translate into remarkable object distortions (as discussed in Section 4.3.6). Such distortions pose a major problem of the EPI technique because the spatial correctness of the image can not be guaranteed, and it is unsafe to rely on the geometry of the reconstructed object. In particular, the use of distorted images in the planning of interventional procedures can obviously have fatal consequences. It is possible to reduce the strength of the distortions to some degree by combining the acquisition with a 180° refocusing pulse such that the sampling of the k-space center coincides with a spin echo (SE-EPI), but still the problem remains substantial.

Therefore, a highly interesting alternative to EPI is the single-shot STEAM (stimulated echo acquisition mode) or turboSTEAM technique [121], which is based on the generation of *stimulated echoes*. Basically, a stimulated echo can be generated from three RF pulses, where strong dephasing gradients have to be applied between the RF pulses [122]. However, it is also possible to replace the last 90° pulse by a low flip-angle or α pulse, which allows for the generation of a series of stimulated echoes by repeating the α pulse several times. At each repetition, a certain amount of magnetization is tipped into the transverse plane, which can be employed to sample a single line of k-space. This magnetization is taken from a pool of longitudinal magnetization prepared by the two preceding 90° RF pulses. Thus, it is possible to acquire all k-space lines at high speed using only two 90° pulses and a train of low-flip angle readout pulses.

A major advantage of the STEAM technique is that all echoes are RF refocused, making the sequence robust to artifacts from off-resonance effects. Moreover, after the initial preparation by the two 90° pulses, the magnetization is completely insensitive to off-resonances until a portion is tipped into the transverse plane by one of the α pulses. For this reason, all sampled k-space lines have a negligible T_2^* weighting - even the lines measured at the end of the echo train. Therefore, single-shot STEAM MRI can be used to acquire distortion-free images from brain areas with severe tissue-susceptibility differences, which are hardly accessible using EPI sequences. Further advantageous characteristics are a much lower RF power deposition in comparison to single-shot fast spin-echo sequences, in particular at higher magnetic field strengths, and the decay of the stimulated echo intensities with the T_1 relaxation time, which under in vivo conditions is usually much longer than the T_2 relaxation time.

Nevertheless, a major limitation of the STEAM technique is that half of the magnetization is unavoidably lost during the preparation of the longitudinal magnetization, i.e. after the application of the second 90° pulse. As a consequence, the resulting images have a clearly lower SNR than comparable EPI images. Moreover, for a typical T_1 relaxation time, receiver bandwidth, and low-flip angle RF pulses, the echo train length (ETL) is limited to about 50 usable stimulated echoes. Obviously, this value restricts the achievable spatial resolution for a conventional Cartesian sampling scheme, even when combining a rectangular FOV with partial Fourier or parallel imaging techniques. In practice, optimized applications to the human brain are limited to an isotropic resolution of 2 mm, as already shown for diffusion-weighted single-shot STEAM MRI [123, 124].




Figure 6.1: Schematic diagram for radial single-shot STEAM MRI. The acquisition part is repeated several times within one cycle of the sequence. Gray lines indicate that the gradient amplitudes vary for different repetitions. (a,h,m) Spoiler, (c,g,j) slice-selection, (d,k,l) prephasing, (n,o) readout, and (e,p,q) rewinder gradients. (b,f) 90° RF pulses, and (i) low flip-angle RF pulse.

6.2 Combination with Radial Sampling

It has been demonstrated in the foregoing chapters that radial trajectories offer salient undersampling properties, which can be exploited to obtain images from only few k-space lines relative to the Cartesian sampling requirements. In view of this finding, it makes sense to combine the single-shot STEAM technique with radial sampling in order to overcome the limited spatial resolution arising from the restricted number of stimulated echoes. Furthermore, the combination opens new applications for radial sampling. Radial trajectories suffer from an inherent sensitivity to off-resonance effects, causing serious complications in a non-refocused multi-echo scenario and precluding the application in an EPI manner. However, when employing the single-shot STEAM technique, this sensitivity is not an issue, and it allows to utilize radial sampling for rapid imaging applications. Thus, the combination of both methods adds complementary advantages: while stimulated echoes eliminate the off-resonance sensitivity of the radial trajectory, the undersampling abilities of radial sampling allow to overcome the resolution limitations of Cartesian single-shot STEAM MRI. In addition, the approach benefits from other advantages discussed in Section 4.3, in particular a low motion sensitivity, a central reference point, and a consistency criterion for the zeroth projection moment.

Sequence Design

Figure 6.1 shows the timing diagram of a generic single-shot STEAM MRI sequence with radial sampling. The sequence starts with a spoiler gradient (a) to ensure complete dephasing of any residual transverse magnetization components from preceding measurements. Slice-selective excitation is accomplished using a first 90° pulse (b) in the presence of gradient (c). Proper refocusing of this gradient involves gradient (e), the corresponding slice-selective gradients (g) and (j) for the second 90° pulse (f) and the low-flip angle read pulse (i), and the spoiler gradient (m), which is required for dephasing of the unwanted free induction decay elicited by pulse (i). During the first $TE/2$ interval, the dephasing of all transverse magnetizations, that is the stimulated echo condition [122], is accomplished by gradient (d) along the read direction. Any transverse magnetization during the first TM interval is effectively removed by a spoiler gradient (h) in the read direction.

The acquisition part of the sequence encompasses multiple repetitions of the final $TE/2$ interval with different orientations for the effective readout gradient (n,o). The starting point of a particular spoke in k-space is reached by the dephasing gradients (k) and (l) in read and phase encoding direction, respectively. After the data acquisition (ADC), the transverse magnetization components are returned to the center of k-space by rewinder gradients in read (p) and phase encoding (q) direction.

The waveforms of all varying gradients (gray lines in Figure 6.1) can be calculated once for the maximum moment needed and then played out with scaled amplitudes, which ensures a fixed sequence timing for all repetitions. The amplitude of the readout gradients for a spoke with angle ϕ is simply given by $A_{read} = A_{max} \cdot \sin \phi$ and $A_{phase} = A_{max} \cdot \cos \phi$ where A_{max} denotes the amplitude of the maximum moment gradient. To take full advantage of the initially prepared longitudinal magnetization and at the same time allow for a more homogeneous signal intensity for all stimulated echoes, variable flip angles were used for the α pulses as described previously [124]. Thus, the flip angle was increased for each repetition of the acquisition part to compensate for the successive consumption of longitudinal magnetization.

Ordering Schemes

Because all spokes of a radial trajectory are physically equivalent, the order at which specific spokes are acquired may follow different strategies. The simplest scheme would be to sample all spokes clockwise from zero to 180° as depicted in the top left of Figure 6.2. However, due to T_1 relaxation and deviations of the actually generated flip angles, this scheme will lead to a sharp edge in k-space between the first and last spoke.




Figure 6.2: (Top) Different schemes for the temporal ordering of acquired spokes in radial single-shot STEAM MRI (numbers indicate the time of sampling). (Bottom) Simulation of image artifacts caused by relaxation effects.

Simulations of a circular phantom demonstrate that this problem causes smearing artifacts in the reconstructed image as shown in the bottom left of Figure 6.2. It is, therefore, beneficial to arrange “early” and “late” spokes acquired along the stimulated echo train more uniformly in k-space. A scheme which totally avoids any intensity edges is shown in the top middle of Figure 6.2. Despite its smooth intensity transition in the angular direction, this scheme, however, leads to an unfavorable point-spread function and therefore causes object deformations in areas with a strong signal decay (bottom middle of Figure 6.2). In practice, the most preferential ordering emerged as an interleaved serial scheme depicted in the upper right of Figure 6.2. Here, the k-space is sampled clockwise for several times, but each time with a different angular offset. The interleaved scheme still leads to a number of edges in k-space, but with clearly reduced strength compared to a single clockwise sampling. Therefore, areas with a pronounced T_1 signal attenuation lead to only mild smearing artifacts without object deformations. In addition, this last scheme has the interesting property that it allows for the reconstruction of separate images from each interleave. For example, if 64 spokes are sampled in a total of 8 interleaves, it is possible to reconstruct 8 images with different degrees of mean T_1 attenuation along the train of stimulated echoes. Although such low-resolution images may be exploited to gain information about local

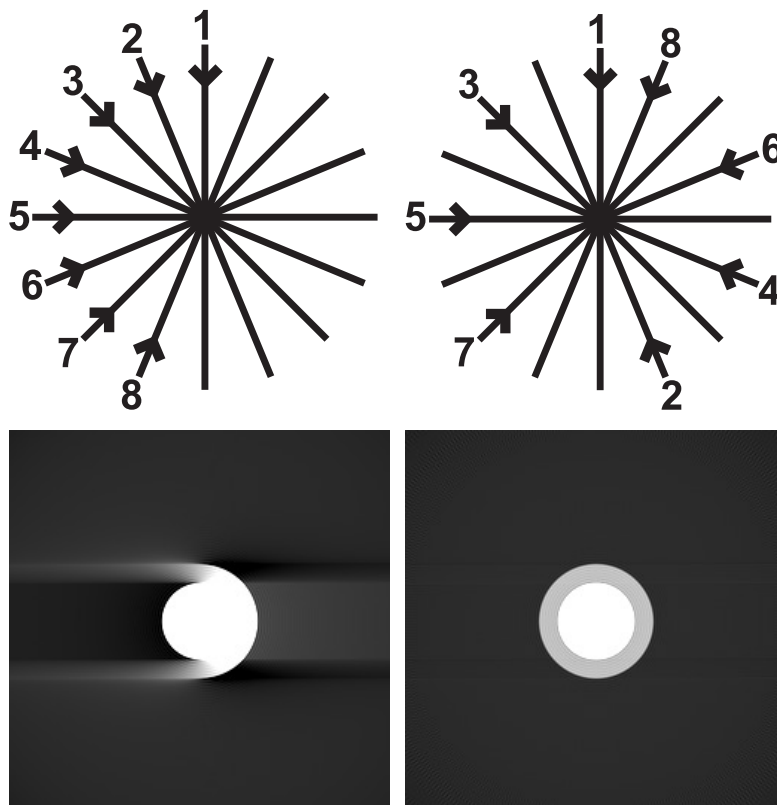


Figure 6.3: Radial sampling schemes with (top left) constant and (top right) alternate spoke directions reversing every second spoke (numbers indicate the time of sampling). (Bottom) Simulation of corresponding image artifacts caused by off-resonance effects.

T_1 relaxation effects, it remains to be seen in practice whether the reduced quality will be sufficient for the estimation of adequate T_1 maps.

A further variation of the ordering scheme can be achieved by reversing the sampling direction of every second spoke. Instead of monotonously scanning 180° of k-space as shown in the upper left of Figure 6.3, this modification yields angles ranging from zero to 360° (top right of Figure 6.3). The lower row shows corresponding image reconstructions from simulated data of an off-resonant circle. As already discussed in Section 4.3.6 on the general off-resonance behavior, unidirectional 180° coverage leads to unsymmetrical U-shaped artifacts that spread over the entire image (lower left of Figure 6.3). In contrast, much more tolerable and focused artifacts are obtained for reversed spokes with 360° coverage (lower right of Figure 6.3). Although the off-resonance problem is clearly mitigated in the STEAM technique, minor phase evolutions might still arise in practice, for example, from strongly off-resonant fat signals. Therefore, a sampling with alternating spoke directions is preferable for most applications.

Image Reconstruction

Due to the limited number of k-space lines obtainable along the echo train, streaking artifacts arise from the uncovered gaps in k-space when performing a conventional regridding procedure. In addition, the weighting of the raw data with the Ram-Lak filter leads to an amplification of image noise because of the rather low signal strength of the stimulated echoes. Therefore, it is essential to cope with the incompleteness of the data, for example, by employing the iterative reconstruction technique proposed in the Chapter 5. In this case, the TV constraint yields an effective removal of undersampling artifacts as well as an edge-preserving denoising of the image. For this reason, it is possible to suppress most of the spurious noise patterns that appear in regridding solutions, where the degree of denoising can be selected by varying the weight of the TV term (5.13) in the cost function Eq. (5.5).

Despite the use of variable flip angles for the readout pulses to compensate for the consumption of longitudinal magnetization, the amplitude of the stimulated echoes typically decreases along the echo train. This is due to imperfections of the generated RF pulses, especially for the higher flip-angles at the end of the train, as well as due to T_1 relaxation. While for Cartesian sampling the signal attenuation translates into blurring along the phase-encoding direction of the image, in the radial case smearing or streaking artifacts arise depending on the ordering scheme used. However, because the sum of the object's projection profile should be independent from the angle as explained in Section 4.3.7, it is possible to approximate the signal decay from the measured spokes and to utilize the information for a first-order correction. This is implicitly done when performing the preprocessing procedure of the iterative reconstruction approach as described in Section 5.5. Nevertheless, it is important to point out that the compensation mechanism assumes a homogeneous relaxation time throughout the object. As the T_1 relaxation time is a locally varying quantity, it will obviously fail if strong deviations are present within a particular section. However, the approach turned out to be quite effective in practice, and it has the advantage of being self calibrating.

6.3 Experimental Results

Acquisition Parameters

All experiments were conducted using the MRI system described in Chapter 3 with the receive-only head coil in triple mode, yielding all 12 channels with different combinations of the coil elements. Measurements were performed for water phantoms as well as the

human brain in vivo, where written informed consent was obtained in all cases prior to each examination.

Radial images were acquired with a base resolution ranging from 104 to 208 pixels and a FOV of 208 mm, which leads to a nominal in-plane resolution of 1 to 2 mm. The number of spokes varied from 32 to 80 and the receiver bandwidth from 120 to 480 Hz/pixel. Except for Figure 6.7, the section thickness was 4 mm and a preceding CHESS pulse was applied for fat suppression. The stimulated echo time TE and the duration of the acquisition interval TR were set to the minimum possible value in all cases. Typically, for 48 spokes, a 208 pixel base resolution, and a bandwidth of 160 Hz/pixel, the sequence resulted in TE = 9.6 ms, TR = 9.2 ms, and an overall measuring time of 464 ms. Noteworthy, in single-shot STEAM MRI, TR may be shorter than TE as the repetitive acquisition interval – which starts with a low-flip angle readout RF pulse – covers only the second half of the echo interval required for generation of a stimulated echo. The first half of the echo interval is the time between the leading two 90° RF pulses and used only once.

All radial acquisitions were performed using an ordering scheme with 8 interleaves and 360° coverage. A compensation mechanism was employed to avoid gradient timing errors and corresponding smearing artifacts due to a misalignment of the data in k-space [62]. Reconstructions of all radial images were performed offline using the software package MRISim.

For comparison, images were acquired with a Cartesian version of the single-shot STEAM MRI sequence [125], which employed variable flip angles for the readout RF pulses, 5/8 partial Fourier encoding, and image reconstruction with the use of a projection onto convex sets (POCS) algorithm [124]. POCS reconstructions are more appropriate than a TV minimization to handle incomplete data in the partial Fourier case. To allow for a fair comparison, the images covered a 160 × 256 mm FOV with an 50 × 128 acquisition matrix (receiver bandwidth 160 Hz/pixel) to yield a 2 mm in-plane resolution (best case) at 4 mm slice thickness. Spin-echo EPI images were acquired with a sequence supplied by the MRI vendor using comparable parameters where possible (256 mm FOV, 128 pixels base resolution, 5/8 partial Fourier encoding, receiver bandwidth 1345 Hz/pixel, 4 mm section thickness, effective echo time TE = 31 ms).

Results

Figure 6.4 compares the best image resolution obtainable with Cartesian single-shot STEAM MRI to the proposed radial version for cross-sectional images of a water phantom and human brain in vivo. While the Cartesian images have a 2.0 mm in-plane resolution, the radial images present with 1.0 mm resolution. This fourfold reduction




Figure 6.4: Image resolution achievable using single-shot STEAM MRI with (left) Cartesian sampling and (middle) radial encoding in comparison to (right) spin-echo EPI. (Top) Water phantom (radial encoding using 48 spokes, 160 Hz/pixel bandwidth) and (bottom) transverse section of human brain in vivo (64 spokes).

in voxel size is best appreciated in the images of the water phantom, where nearly all circles are resolved by the radial method. In the brain images, the improved resolution can also be appreciated, especially when focusing on the borders of the brain or the skull. Noteworthy, the Cartesian images suffer from pronounced blurring in the phase-encoding direction (anterior-posterior), which reflects the influence of decreasing stimulated echo intensities on the respective point-spread function. For comparison, the right column of Figure 6.4 shows corresponding spin-echo echo-planar images. Although presenting with a better SNR, both the phantom and brain image are affected by significant distortion artifacts. These problems obviously limit the usefulness of EPI whenever reliable anatomic accurateness is required.

The upper row of Figure 6.5 shows iterative image reconstructions from radial acquisitions with 48 spokes and different receiver bandwidths. Rather than affecting the SNR, the increased bandwidth leads to slightly decreased image resolution and image contrast. In fact, because the TV constraint efficiently removes noise during iterative reconstruction, a lower SNR translates into a lower resolution in the final image. The somewhat lower gray-white matter contrast is caused by a shorter duration of the acquisition part (shorter TR intervals) for higher bandwidths that reduce the effective T_1 weighting of the stimulated echoes. However, this is not necessarily a drawback because certain ap-




Figure 6.5: (Top) Iterative and (bottom) conventional regridding reconstructions (transverse sections of the human brain *in vivo*) for radial single-shot STEAM MRI (48 spokes) with a bandwidth of (left) 120 Hz/pixel, (middle) 320 Hz/pixel, and (right) 480 Hz/pixel (1.0 mm nominal in-plane resolution, 4 mm section thickness).

plications, including diffusion studies, do not require such image contrast. The bottom row of Figure 6.5 depicts reconstructions of the same data using a conventional regridding approach. Here, a higher bandwidth is accompanied by increased image noise. These images clearly demonstrate that a more advanced reconstruction method than regridding is required for this type of undersampled data acquisition.

Figure 6.6 shows radial images of the human brain with 32, 48, 64, and 80 spokes. Despite these differences, all iterative reconstructions recover the object with reasonable quality. The reconstruction from 32 spokes suffers from some residual undersampling artifacts, leading to a lower overall resolution. In contrast, the image from 80 spokes becomes mildly affected by artificial streaks in the horizontal direction that result from localized T_1 relaxation effects due to the very long echo train. In general, the images demonstrate that radial sampling offers considerable freedom for the trajectory design, allowing to tailor imaging protocols for specific examination needs.

The range of SNR and resolution achievable by radial single-shot STEAM MRI is summarized in Figure 6.7 for a transverse cross-section of the human brain with a nominal in-plane resolution of 1.0 to 2.0 mm and a section thickness of 2.0 to 4.0 mm. The reduced SNR for higher in-plane resolution is due to both, an increased bandwidth




Figure 6.6: Iterative reconstructions for radial single-shot STEAM MRI (transverse sections of the human brain in vivo) with 32, 48, 64, and 80 spokes (bandwidth 160 Hz/pixel, 1.0 mm nominal in-plane resolution, 4 mm section thickness).

needed to maintain the sequence timing and a larger degree of undersampling. Similar to the findings in Figure 6.5, increased noise in the data acquisition is not necessarily visible in the iteratively reconstructed images because of the TV constraint. A decrease of the section thickness is accompanied by a significantly higher noise level that lowers the reconstructed image resolution. This can be best seen for the highest in-plane resolution where proper windowing would reveal additional noise patterns in the background of the object for a 2 mm section thickness.

6.4 Unresolved Issues

As demonstrated, the use of radial trajectories allows to overcome the practical resolution limit for single-shot STEAM MRI with Cartesian sampling. While low SNR of the acquired data remains a persistent problem, the constrained variation of the image intensity, as part of the used iterative reconstruction method, yields images without visible noise artifacts. In fact, TV minimization techniques are among the most powerful methods currently available for general image denoising, especially due to the preservation of sharp edges. Nevertheless, the approach is only able to remove oscillating intensity fluctuations between neighboring pixels and, hence, only flattens noise




Figure 6.7: Iterative reconstructions for radial single-shot STEAM MRI (transverse sections of the human brain in vivo) with a nominal in-plane resolution of (top) 2.0 mm, (center) 1.4 mm, and (bottom) 1.0 mm and a section thickness of (left) 4 mm, (middle) 3 mm, and (right) 2 mm (48 spokes, bandwidth 160 Hz/pixel).

textures. While this certainly leads to a more easily “readable” representation of the object, it is of course impossible to recover object information that is lost due to a lack of signal strength. Further, if the noise level becomes too high, fine-structured information of the object will be removed as the TV constraint can not distinguish between actual object features and undesired noise patterns. A critical issue, therefore, is how much weight should be put on the TV minimization. If chosen too high, object details will be removed; if too low, the image remains noisy. For the reconstructions presented here, the weight was adapted manually to yield images with a comparable noise level as judged by visual inspection. A more routine use of the proposed method would probably need an automatic mechanism to find an appropriate weight. This,

however, requires a reliable noise estimator for the measured k-space data, which is a non-trivial problem. In particular, because the measured k-space data is incomplete, the noise estimator has to distinguish between actual noise patterns and image artifacts that arise from the undersampling. This precludes the use of conventional strategies for noise detection. An estimation of the noise strength might be achieved by identifying the support of the object and calculating the variance of the projection profiles outside the object support. Alternatively, wavelet transformations of the profiles could be performed to estimate the noise level from the high-band variance of the wavelet transforms. However, the development of a strategy that is robust enough for clinical imaging is outside the scope of the current proof-of-principle work.

Another open question arises from artifacts caused by spatially varying T_1 relaxation during long stimulated echo trains. In Cartesian sampling, the affected point-spread function leads to localized one-dimensional blurring along the phase-encoding dimension. For radial trajectories, the image resolution remains mostly unaffected, but instead smearing or streaking artifacts might appear inside the object (bottom right image of Figure 6.6). Unfortunately, such artifacts are sometimes hardly distinguishable from actual object features, and because the T_1 relaxation time is a locally varying quantity, it is impossible to fully compensate for the artifacts using a non-local approach. In practical terms, it might be best to avoid such problems by not using too many echoes in a train – a desirable feature anyhow as it reduces the measuring time by exploiting the radial undersampling properties. On the other hand, a time-segmented reconstruction approach might deal with the problem but would require a priori knowledge or estimation of a relaxation map to calculate snapshots of the object according to the T_1 weighting of each spoke. Possibly, this information could be extracted directly from the same data using individual reconstructions from each segment of an interleaved reordering scheme. However, the development of such an approach is expected to be highly challenging in practice.

6.5 Segmented Data Acquisition

A key problem of the single-shot STEAM MRI technique is the very low signal amplitude in comparison to other MRI techniques and, in particular, relative to the competing EPI technique. The low signal strength is caused by two different reasons. First, during preparation of the longitudinal magnetization with the two 90° pulses, half of the magnetization is lost. In fact, after exciting and dephasing the magnetization in the read direction ((b) and (d) in Figure 6.1), only those components are flipped to the longitudinal direction that are perpendicular to the tilting axis of the second 90° pulse (f), which corresponds to 50% of the magnetization. The remaining transverse magne-

tization components can no longer be utilized and have to be spoiled to avoid spurious signal contributions, which is accomplished by switching gradient (m). Although this loss of magnetization is highly undesirable, it is impossible to realize the stimulated-echo condition in a more efficient way, and as a matter of principle the STEAM technique delivers a relatively weak signal.

The second reason lies in the single-shot readout approach, i.e. the repeated generation of low flip-angle pulses for obtaining an echo train. At each repetition, only a small portion of the prepared longitudinal magnetization is employed to sample one line of k -space, which ensures a high phase coherence of the transverse magnetization for all lines. However, the use of “fresh” transverse magnetization for all k -space lines is rather inefficient in view of the total signal yield. After each line has been sampled, it is necessary to spoil the remaining transverse magnetization to vanish residual signals and, therefore, the prepared amount of longitudinal magnetization is exploited only up to a certain degree for acquiring k -space information. In contrast, in the EPI case all transverse magnetization is exploited for the data acquisition without losses from intermediate spoiling, which gives a higher signal yield (but at the expense of phase coherence). Therefore, several ideas have been proposed for gaining additional signal strength in order to make the rapid STEAM technique more competitive to EPI. For instance, in the GRASTE variant [126] multiple k -space lines (usually an odd number) are sampled during each cycle of the acquisition part, which decreases the echo train length and reduces the overall losses from the spoiling at the end of each acquisition cycle. However, because in this case only the k -space lines sampled in the middle of each cycle are RF refocused, the signal gain is accompanied by a stronger vulnerability to off-resonance effects as well as a non-uniform T_2^* weighting among the measured lines.

A different idea consists in using multiple preparation steps for the readout procedure, i.e. a block-by-block acquisition of the data instead of obtaining all k -space lines within a single run of the sequence. This strategy is called a *segmented* or *multi-shot* acquisition, and it offers a remarkable signal enhancement because the prepared magnetization has to be allotted to only a fractional number of echoes relative to the corresponding single-shot version. Hence, it allows to increase the flip angle for all stimulated echoes so that a stronger voltage is induced in the receive coils. However, because each additional preparation step is time-consuming and, further, a certain delay between the preparations is needed to ensure a proper relaxation to thermal equilibrium, the signal amplification comes at the cost of an overall longer measurement. Thus, by using a higher number of preparation steps the SNR of the images is increased, but in turn the examination time is prolonged, and the sequence loses its high-speed imaging advantage. Nevertheless, it has recently been demonstrated for cardiac imaging that a




Figure 6.8: Single-shot STEAM (1s) in comparison to multi-shot STEAM with (2s) two segments and (8s) eight segments (48 spokes, 1.0 mm nominal in-plane resolution, 2 mm section thickness, bandwidth 180 Hz/pixel). The images were reconstructed using (top) conventional regridding and (bottom) the TV-based iterative approach.

segmentation of the Cartesian STEAM method offers an essential improvement of the image quality and diagnostic information [127].

When using a radial trajectory as proposed here, it is possible to exploit several of the unique sampling properties for making the segmented approach more robust to experimental complications. For instance, the central intersection of all spokes can be used to remove phase offsets among the different preparation steps, which otherwise would need the recording of a navigator signal. Further, due to the undersampling ability it is possible to reconstruct individual low quality images from each segment, which enables to detect and to compensate for object movements that might occur between the measurements of the segments. Finally, because the center of k-space is oversampled in the radial case, successive acquisition of a segment yields a desirable averaging effect in areas that are already sufficiently covered, while missing information is added in the outer k-space areas. Therefore, radial data sampling is particularly well suited for the combination with segmentation, and the use of a segmented approach arises as straightforward option if the SNR of a single-shot acquisition is insufficient.

The effectiveness of this strategy is demonstrated in Figure 6.8, comparing the image quality obtained for a single-shot, two-shot, and eight-shot acquisition of the human

brain in vivo using otherwise identical sequence parameters (48 spokes, 2.0 mm section thickness, 1.0 mm in-plane resolution). The upper row shows regridding reconstructions, and it can be clearly seen that a higher number of segments leads to a significant reduction of the noise patterns (as well as a slight attenuation of the T_1 weighting). Noteworthy, the eight-shot image exhibits with pronounced streaking artifacts, whereas in the single-shot case the undersampling artifacts are hardly visible due to the high degree of noise. Iterative reconstructions of the data sets are shown in the lower row, and here it can be seen that the segmentation results in a better spatial resolution. This is because the assumption of a piecewise-constant object is much more appropriate for the eight-shot case with high SNR, so that the reconstruction algorithm is able to recover the object in a more reasonable way as discussed in Section 5.7.

6.6 Summary

Single-shot STEAM MRI is a rapid imaging technique that offers distortion-free reconstructions even in the presence of off-resonance effects, which poses a clear advantage over the commonly used EPI technique. Because the data acquisition is accomplished using a train of stimulated echoes, the signal strength is relatively weak, and the number of k-space lines obtainable within a sequence run is limited to a value of around 50. When using the Cartesian sampling scheme, this restriction translates into a limitation of the spatial resolution, so that it is impossible to achieve an in-plane resolution higher than 2 mm. In contrast, if the sequence is adapted to radial data sampling, the undersampling properties can be exploited to overcome this limitation because here the base resolution can be selected independently from the number of stimulated echoes. Moreover, the novel combination adds complementary advantages of both methods: on the one hand, the RF refocused stimulated echoes allow for the unrestricted utilization of radial k-space sampling for rapid imaging applications, which for non-refocused sequences would be hardly possible due to the pronounced off-resonance sensitivity. On the other hand, the undersampling capability of the radial sampling scheme makes optimal use of the limited magnetization available in single-shot STEAM MRI. Because the acquired data sets are strongly undersampled, it is essential to utilize the reconstruction technique proposed in the foregoing chapter, which copes with the incompleteness of the data and leads to an edge-preserving removal of the noise patterns that typically appear for rapid STEAM data. Finally, if the SNR of the measured data is very low so that the images exhibit a remarkable loss of resolution, a segmentation of the acquisition may be employed, which in turn prolongs the measurement to a certain degree.

Chapter 7

Efficient T_2 Mapping Based on Radial Fast Spin-Echo MRI

This chapter covers the image reconstruction from radial multi-echo acquisitions, where spokes are sampled at different echo times to accelerate the measurement. An extended iterative reconstruction approach is presented, which takes the time-dependency of the k-space lines into account by employing a signal model that combines data from different echo times. As a consequence, the method directly yields quantitative maps of the proton density and T_2 relaxivity.

7.1 Radial Fast Spin-Echo Imaging

Fast spin-echo (FSE) MRI is one of the most frequently used techniques in today's clinical practice [5]. In the FSE approach, a train of multiple spin-echoes is generated after each RF excitation. Therefore, it offers images with proton-density and T_2 contrast at significantly reduced measuring time compared to a spin-echo sequence with a single phase-encoding step per excitation. Although the k-space lines that are sampled during the train of spin echoes have different echo times and, thus, an increasing T_2 weighting, conventional Cartesian sampling strategies still allow for a straightforward image reconstruction. The underlying reason is the dominance of major image features by the low spatial frequencies located around the center of k-space, while the periphery of k-space defines the edges which are almost invariant to changes in contrast. Hence, in the Cartesian case, the contrast can be adjusted by reordering the acquisition scheme such that the central k-space lines are measured with the desired echo time. A noticeable disadvantage of the T_2 attenuation in k-space, however, is a certain degree of image blurring in the phase-encoding direction, similar to the situation in the rapid STEAM MRI approach.

As already mentioned in Section 4.4.1, the FSE concept can be also combined with radial sampling, which enables to acquire a set of spokes in a shorter measurement time and in turn makes the FSE technique more robust to object motion [17, 39, 54]. In this specific case, however, the use of a radial scheme has major implications for the image reconstruction. Because all spokes pass the center of k-space, each spoke carries an equal amount of low spatial frequency information which, for different echo times, exhibits pronounced differences in T_2 contrast. This situation is fundamentally different to the Cartesian case and poses severe complications when employing a conventional reconstruction method such as regridding or filtered backprojection: (i) The merging of spokes with different echo times may cause streaking artifacts around areas with pronounced T_2 relaxation, because the relative signal decay leads to jumps in the corresponding point-spread-function [77]. (ii) The contrast of the image always represents an average of the varying T_2 weightings. (iii) Differently ordered k-space acquisitions may no longer be used to distinguish between PD and T_2 relaxation.

On the other hand, due to oversampling of the k-space center, a single radial FSE data set implicitly contains information about the relaxation process along the echo train. For this reason, dedicated reconstruction methods have been developed that extract the embedded temporal information and enable to quantify the local T_2 relaxation within a significantly reduced measurement time relative to Cartesian-based approaches. Most of the existing techniques, such as k-space weighted image contrast (KWIC) [128, 129], attempt to calculate a series of time-resolved images by mixing the low spatial frequencies from spokes measured at a desired echo time with high frequency information from other spokes (at different echo times). Hence, each of the images corresponds to one echo time of the acquisition train, so that it is possible to determine the local relaxation rate using a pixel-by-pixel fitting procedure. A drawback, however, is that the mixing of unequal echo times tends to cause artifacts in the reconstructed images and, thus, it limits the accuracy of the T_2 estimates. In particular, the calculated values become dependent on the object size as the change in k-space is assumed to be located only in the low frequency area that is covered by spokes measured at an equal echo time [129].

7.2 Extended Iterative Reconstruction

To overcome the problems that arise when merging data with inconsistent acquisition times, a novel iterative reconstruction method is proposed, which can be seen as an extension of the approach presented in Chapter 5 for the reconstruction from highly undersampled acquisitions with multiple coils. In the former case, proper reconstruction is achieved by iteratively estimating an image that, on the one hand, is consistent with the data measured by all coils and, on the other hand, complies with prior knowl-

edge about the object. For multi-echo data from a FSE sequence, this strategy is not appropriate because it is impossible to find a single image that matches the different contrasts at the same time. Therefore, it is necessary to include the relaxation process into the modeling of the MRI signal used to compare the estimate to the actual observations. As the T_2 relaxation time is a locally varying quantity, this requires that the estimate consists of a spin-density and a relaxation component instead of just an intensity component. The objective of the present extended approach, therefore, is to find a spin-density map and a relaxivity map such that snapshots, calculated for each echo time from these maps, best match the spokes measured at the respective echo times. Hence, the approach completely circumvents the calculation of intermediate images, and it makes optimal use of all data acquired instead of relying on the assumption that the contrast changes are localized in the center of k-space. It should be noted, however, that the concept assumes a negligible signal relaxation during the readout interval of each echo to reduce the overall complexity of the problem, but this assumption is made in most MRI techniques.

Cost Function

In order to compute the maps, a cost function is needed that quantifies the accuracy or quality of the match to the measured data in a similar way to Eq. (5.2). Because the MRI noise can be assumed to be Gaussian, the cost function of the proposed method uses the ℓ_2 norm and has the form

$$\Phi(\boldsymbol{\rho}, \mathbf{r}) = \frac{1}{2} \sum_t \sum_c \|\mathbf{F}(\boldsymbol{\rho}, \mathbf{r}, t, c) - \mathbf{y}_{t,c}\|_2^2, \quad (7.1)$$

where $\boldsymbol{\rho}$ is a vector containing the values of the spin-density map, and \mathbf{r} is a vector containing values of the relaxivity map. For a base resolution of $n \times n$ pixels, both vectors have n^2 entries. Further, $\mathbf{y}_{t,c}$ is a vector containing the raw data from channel c of all spokes measured at echo time t , where c runs from 1 to the total number of channels and t runs over all echo times. Finally, \mathbf{F} is a vector function that calculates a snapshot from the given spin-density and relaxivity map at echo time t , and translates it to k-space using a Fourier transformation and subsequent evaluation of the spokes acquired at time t . Moreover, before Fourier transformation, this mapping includes a multiplication with the sensitivity profile of coil c , which is obtained from the same data in a prior step using the coil estimation procedure described in Section 5.4.

The function \mathbf{F} can be seen as the forward operation of the reconstruction problem and comprises a model of the received MRI signal. Noteworthy, this function is the equivalent to the multiplication with matrix \mathcal{A} in Eq. (5.2). In the multi-echo case, however, it is not possible to write the operation as a matrix multiplication because the

signal depends non-linearly on the components of the estimated maps. The j th entry, i.e. the j th sample of the synthesized data with k-space position \mathbf{k}_j at echo time t , is given by

$$F_j(\boldsymbol{\rho}, \mathbf{r}, t, c) = \sum_{\mathbf{x} \in \text{FOV}} \varrho(\mathbf{x}) \cdot e^{-R(\mathbf{x})t} \cdot C_c(\mathbf{x}) \cdot e^{2\pi i \mathbf{x} \cdot \mathbf{k}_j}, \quad (7.2)$$

where \mathbf{x} denotes a position in image space such that the sum runs over all (discrete) elements of the image matrix, and C_c is the complex sensitivity profile of the c th coil. Further, ϱ denotes a function which evaluates the spin-density vector $\boldsymbol{\rho}$ at image position \mathbf{x}

$$\varrho(\mathbf{x}) = \sum_i \rho_i \cdot \delta(\mathbf{x} - \mathbf{x}_i), \quad (7.3)$$

where ρ_i denotes the i th component of the spin-density vector with corresponding position \mathbf{x}_i in image space. Accordingly, the function R evaluates the relaxivity vector \mathbf{r}

$$R(\mathbf{x}) = \sum_i r_i \cdot \delta(\mathbf{x} - \mathbf{x}_i). \quad (7.4)$$

To reconstruct a given set of measured data \mathbf{y} , a pair of vectors $(\boldsymbol{\rho}, \mathbf{r})$ has to be found that minimizes the cost function Φ . This can be achieved by using a numerical optimization method that is suited for large-scale problems, like the non-linear conjugate-gradient method. To this end, the present implementation utilized the CG-Descent algorithm [108] as in Chapter 5, which can be employed in a straight-forward manner. Thus, it is only required to evaluate Φ and its gradient at given positions in the parameter space.

Evaluation of Cost Function

Evaluation of the cost function (7.1) at a given pair $(\boldsymbol{\rho}, \mathbf{r})$ can be done in a straight-forward manner. However, to calculate the value of Φ in a reasonable time, a more practical strategy is to perform a fast Fourier transformation (FFT) of the snapshots and to interpolate the transforms onto the desired spoke positions in k-space using a convolution with a radial Kaiser-Bessel kernel, similar to the regridding procedure discussed in Section 4.2.2. Because the kernel is finite, it is further necessary to pre-compensate for undesired intensity modulations by multiplying the snapshots with an approximation of the kernel's Fourier transform in front of the FFT.

The same strategy can be used for evaluating the gradient, which is a vector containing the derivative of the cost function with respect to each component of the spin-density and relaxivity vector. It is convenient to decompose the problem into a separate deriva-

tion of Φ with respect to components of $\boldsymbol{\rho}$ and \mathbf{r} , respectively,

$$\nabla\Phi = \begin{pmatrix} \nabla_{\boldsymbol{\rho}}\Phi \\ \nabla_{\mathbf{r}}\Phi \end{pmatrix}. \quad (7.5)$$

To simplify the notation, the calculation is shown only for a single time point and a single coil (indicated by ϕ instead of Φ)

$$\begin{aligned} \phi &= \frac{1}{2} \|\mathbf{F} - \mathbf{y}\|_2^2 = \frac{1}{2} \sum_j (F_j - y_j) \overline{(F_j - y_j)} \\ &= \frac{1}{2} \sum_j F_j \overline{F_j} + y_j \overline{y_j} - y_j \overline{F_j} - \overline{y_j} F_j. \end{aligned} \quad (7.6)$$

Derivation of Eq. (7.2) with respect to components of $\boldsymbol{\rho}$ gives

$$\frac{\partial}{\partial \rho_v} F_j = e^{-R(\mathbf{x}_v) \cdot t} \cdot C_c(\mathbf{x}_v) \cdot e^{2\pi i \mathbf{x}_v \cdot \mathbf{k}_j}, \quad (7.7)$$

where \mathbf{x}_v denotes the position of the v th component in image space. Using the chain rule, this yields

$$\begin{aligned} \frac{\partial}{\partial \rho_v} \phi &= \frac{1}{2} \sum_j F_j \frac{\partial}{\partial \rho_v} \overline{F_j} + \overline{F_j} \frac{\partial}{\partial \rho_v} F_j - \overline{y_j} \frac{\partial}{\partial \rho_v} F_j - y_j \frac{\partial}{\partial \rho_v} \overline{F_j} \\ &= \frac{1}{2} \sum_j (F_j - y_j) \frac{\partial}{\partial \rho_v} \overline{F_j} + (\overline{F_j} - \overline{y_j}) \frac{\partial}{\partial \rho_v} F_j \\ &= \frac{1}{2} \sum_j (F_j - y_j) \frac{\partial}{\partial \rho_v} \overline{F_j} + \frac{1}{2} \overline{\sum_j (F_j - y_j) \frac{\partial}{\partial \rho_v} F_j} \\ &= \Re \left\{ \sum_j (F_j - y_j) \frac{\partial}{\partial \rho_v} \overline{F_j} \right\} \\ &= \Re \left\{ \sum_j (F_j - y_j) e^{-R(\mathbf{x}_v) \cdot t} \cdot C_c(\mathbf{x}_v) \cdot e^{-2\pi i \mathbf{x}_v \cdot \mathbf{k}_j} \right\} \\ &= e^{-R(\mathbf{x}_v) \cdot t} \cdot \Re \left\{ C_c(\mathbf{x}_v) \cdot \sum_j (F_j - y_j) e^{-2\pi i \mathbf{x}_v \cdot \mathbf{k}_j} \right\}. \end{aligned} \quad (7.8)$$

Hence, the gradient with respect to $\boldsymbol{\rho}$ can be obtained by evaluating the cost function, calculating the residual, and performing an inverse Fourier transformation, which is followed by a multiplication with the complex conjugate of the coil profile and, finally, with the relaxation term. Derivation of Eq. (7.2) with respect to components of \mathbf{r} gives

$$\frac{\partial}{\partial r_v} F_j = -t \cdot \varrho(\mathbf{x}_v) \cdot e^{-R(\mathbf{x}_v) \cdot t} \cdot C_c(\mathbf{x}_v) \cdot e^{2\pi i \mathbf{x}_v \cdot \mathbf{k}_j} \quad (7.9)$$

and, in a similar way to (7.8), this yields

$$\frac{\partial}{\partial r_v} \phi = -t \cdot \varrho(\mathbf{x}_v) \cdot e^{-R(\mathbf{x}_v) \cdot t} \cdot \Re \left\{ \overline{C_c(\mathbf{x}_v)} \cdot \sum_j (F_j - y_j) e^{-2\pi i \mathbf{x}_v \cdot \mathbf{k}_j} \right\}. \quad (7.10)$$

Comparison with Eq. (7.8) shows that the gradient with respect to \mathbf{r} can be easily obtained by multiplying the gradient with respect to $\boldsymbol{\rho}$ with the components of the given $\boldsymbol{\rho}$ and the echo time t . Of course, Eq. (7.8) and Eq. (7.10) have to be summed over each channel and echo time occurring in the complete cost function (7.1).

Initialization and Scaling

In the single-echo reconstruction scenario, it is very efficient to initialize the optimizer with a properly scaled regridding solution. This choice significantly reduces the total number of iterations because the optimizer starts with a reasonable guess. In the multi-echo case, however, it is more difficult to obtain reasonable initial guesses for the spin-density and relaxivity map, and several options exist. For example, a curve fitting of the (strongly undersampled) regridding solutions from single echo times could be used to approximate the maps or, alternatively, an echo-sharing method like KWIC could be employed. While preliminary analyses confirmed that these strategies may lead to a certain acceleration of convergence, they also indicated complications if the initial guesses contain implausible values, for example, in relaxation maps which are obviously undefined in areas with a void signal intensity. It is therefore necessary to remove respective values from the initial guesses. The present work, however, simply used zero maps for the initialization, which require a higher number of iterations but ensure a straightforward convergence to a reasonable solution.

Another factor with essential impact on the convergence rate is the scaling of the time variable t . Although it makes sense to directly use physical units, a proper rescaling of the time variable significantly reduces the number of iterations. Equation (7.10) shows that the gradient with respect to the relaxivity depends linearly on t , while this is not the case with respect to the spin density. If the values of t for the different echoes are very small, then the cost function is much more sensitive to changes in $\boldsymbol{\rho}$, and the problem is said to be *poorly scaled* [109]. In contrast, large values of t lead to a dominant sensitivity to perturbations in \mathbf{r} . Because finding a reasonable solution requires a matching of both maps at the same time, the method is especially effective when the scaling of t is selected such that there is a balanced influence on the cost function. It turned out that a proper scaling, which depends on the range of the object's spin-density and relaxivity values, reduces the number of required iterations to about 80 for a typical data set. Of course, a rescaling of t is accompanied by a

corresponding scaling of the relaxivity values in \mathbf{r} , which can be corrected afterwards to allow for quantitative analyses.

Regularization and Snapshot Calculation

Because the proposed method employs a non-linear optimization technique, prior knowledge about the object can be incorporated with only low effort. This is accomplished, as detailed in Section 5.3, by complementing the cost function with suitable penalty terms that assign high values to implausible solutions and thereby lead to a better conditioning of the reconstruction problem. In this case, the cost function takes the form

$$\Phi(\boldsymbol{\rho}, \mathbf{r}) = \frac{1}{2} \sum_t \sum_c \|\mathbf{F}(\boldsymbol{\rho}, \mathbf{r}, t, c) - \mathbf{y}_{t,c}\|_2^2 + \sum_i \lambda_i P_i(\boldsymbol{\rho}, \mathbf{r}), \quad (7.11)$$

where P_i are the penalty functions, weighted by λ_i , which can act on both, the spin-density and relaxivity map. For example, because the spin density and relaxivity can take only positive values, negative entries in $\boldsymbol{\rho}$ and \mathbf{r} can be penalized with a quadratic term. Further, for a certain class of medical objects like the human brain, it is reasonable to assume that the images are approximately piecewise constant and characterized by a limited TV value. Thus, following the ideas of Chapter 5, it can be beneficial to introduce a TV constraint for the two maps. This allows for a suppression of remaining streaking artifacts if only a low number of total spokes is measured and results additionally in an edge-preserving denoising. However, to demonstrate the effectiveness of the main concept proposed in this chapter, that is the exploitation of a signal model to link samples measured at different echo times, no penalty functions were used in the reconstructions presented here.

Finally, after complete estimation of the spin-density map $\boldsymbol{\rho}$ and relaxivity map \mathbf{r} , snapshot images can be calculated for an arbitrary echo time with

$$I_t(\mathbf{x}) = \varrho(\mathbf{x}) \cdot e^{-R(\mathbf{x}) \cdot t}. \quad (7.12)$$

These images do not contain any additional information, but present the estimated temporal information in a more familiar view.

7.3 Experimental Results

Acquisition Parameters

All experiments were conducted using the MRI system described in Section 3.1, where the signal was detected with the receive-only head coil array. In the present case,

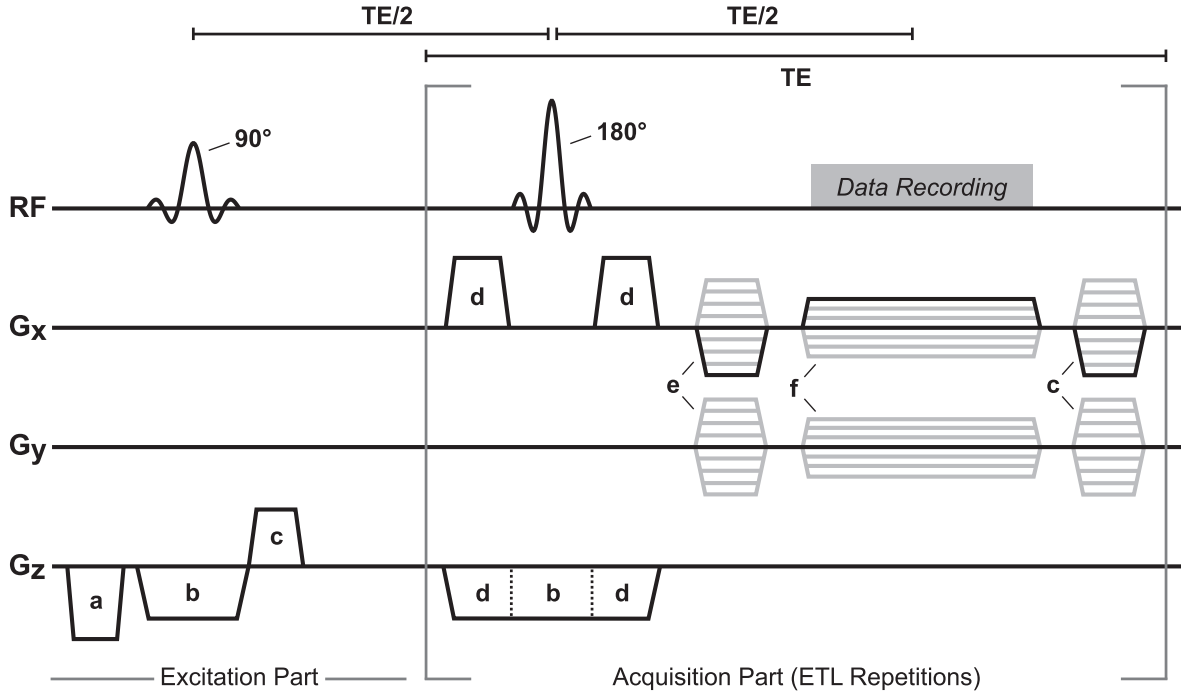


Figure 7.1: Timing diagram of the radial FSE sequence. The acquisition part is repeated ETL times within one cycle of the sequence. Gray lines indicate that respective gradient amplitudes vary for different repetitions. (a) Spoiler, (b) slice-selection, (c) rewinder, (d) crusher, (e) prephasing, and (f) readout gradients.

only the four channels corresponding to the primary mode signals P in Eq. (3.2) were recorded to reduce the high amount of data. The timing diagram of the sequence used for the acquisitions is shown in Figure 7.1. Measurements were performed for a water phantom doped with $MnCl_2$ as well as the human brain in vivo, where written informed consent was obtained in all cases prior to each examination.

Phantom data was acquired with a base resolution of 160 pixels covering a FOV of 120 mm (bandwidth 568 Hz/pixel), while human brain data was acquired with a base resolution of 224 pixels covering a FOV of 208 mm (bandwidth 438 Hz/pixel). A train of $ETL = 16$ spin echoes with an echo spacing of 10 ms was recorded after a slice-selective 90° excitation pulse (section thickness 3 mm). The spin echoes were refocused using a conventional 180° RF pulse, enclosed by crusher gradients to dephase spurious FID signals. A total number of 512 spokes per slice was acquired from 32 excitations, measured with a repetition time of $TR = 7000$ ms to avoid saturation effects of the CSF. The “angular bisection” view-ordering scheme was used as described for the KWIC method by Song and Dougherty [128], which ensures that spokes measured at consecutive echo times have a maximum angular distance. Noteworthy, this scheme is not required by the proposed method, but it was employed to permit reconstructions with the KWIC approach for comparison. Further, the sampling direction of every second repetition was altered in such a way as to generate opposing neighboring spokes, which yields more tolerable artifacts in the presence of off-resonances. Fat suppression was accomplished




Figure 7.2: *Regridding reconstructions from the complete FSE data sets, obtained for (left) a phantom (base resolution 160 pixels, FOV 120 mm, bandwidth 568 Hz/pixel bandwidth) and (right) a transverse section of the human brain in vivo (base resolution 224 pixels, FOV 208 mm, bandwidth 438 Hz/pixel). The data was acquired with a radial FSE sequence using 32 excitations and 16 echoes each.*

by a preceding CHES pulse, and the mentioned compensation mechanism was applied to avoid gradient timing errors [62].

Image Reconstruction

All data processing was done offline using the MRISim software package. In a first step, phase offsets were removed by aligning the phase of all spokes at the center of k-space. Coil sensitivity profiles were estimated from the data set using the procedure described in Section 5.5. In addition, a thresholding mask was obtained from the smoothed sum-of-squares image, so that areas with void signal intensity can be set to zero by applying the mask to all reconstructed images. For the interpolation in k-space from grid to spokes and vice versa, a Kaiser-Bessel window with $L = 6$, $\beta = 13.8551$, and twofold oversampling was used [21]. The optimizer for estimating the spin-density and relaxivity map was run for a fixed number of 100 iterations. The scaling of the time variable was chosen heuristically such that $t = 300 \cdot n$ for the phantom study and $t = 150 \cdot n$ for human brain data, where n is the echo number.

For comparison, regridding reconstructions of the spokes measured at each echo time were calculated using the same interpolation kernel. Here, the estimated coil sensitivity profiles were used to combine the different channels instead of taking a sum-of-squares. Further, time-resolved reconstructions employing the KWIC method were calculated. In the original KWIC approach, only 8 instead of 16 echoes per excitation were acquired. Therefore, two variants were implemented: either high frequency information from all spokes was used to fill the outer k-space area (kwic 16), or information from only 8 neighboring echo times was shared (kwic 8). To allow for a fair comparison, the




Figure 7.3: Spin-density maps (top) and relaxivity maps (bottom) estimated from the phantom data set using the proposed iterative method (*iter*), KWIC combining all 16 echoes (*kwic 16*), and KWIC combining only 8 neighboring echoes (*kwic 8*). Other parameters as in Figure 7.2. PD = proton density, R2 = T_2 relaxivity.

same interpolation kernel was used, and coil profiles were employed for the channel combination. Finally, spin-density and relaxivity maps were estimated from the KWIC images by a pixelwise curve fitting using the Levenberg-Marquardt algorithm.

Results

First of all, Figure 7.2 shows plain regridding reconstructions of the two datasets that were employed to demonstrate the effectiveness of the proposed method. The images exhibit an average contrast because all measured spokes (with varying echo times) were included into the regridding procedure. Consequently, these images do not allow to distinguish a high spin density with fast relaxation from a low spin density with slow relaxation. The object on the left is a phantom containing five water-filled tubes with different concentrations of MnCl_2 , which lead to different T_2 relaxation times of the compartments. The second data set, shown on the right, is a transverse section of the human brain in vivo.

Figure 7.3 compares spin-density and relaxivity maps for the phantom data set, which were estimated using the proposed method, the KWIC method sharing all echoes, and the KWIC method sharing 8 neighboring echoes. It can be seen that the sharing of




Figure 7.4: Spin-density maps (top) and relaxivity maps (bottom) estimated from the human brain data set using the proposed method (*iter*), KWIC combining all 16 echoes (*kwic 16*), and KWIC combining only 8 neighboring echoes (*kwic 8*). Other parameters as in Figure 7.2.

k-space data in the KWIC reconstructions leads to ring-like artifacts inside the tubes with fast T_2 relaxation, in line with the findings of Altbach et al [129]. The artifacts are more pronounced in the KWIC variant sharing all echoes, while the variant sharing only 8 echoes suffers from streaking artifacts due to incomplete coverage of the outer k-space. Such artifacts do not appear in the iteratively estimated maps. Here, the spin-density of the tube with the shortest relaxation time is slightly underestimated, which is probably caused by a higher amount of noise due to fast signal decay. Further, because the relaxivity is undefined in areas with a void spin density, the relaxivity maps are affected by spurious values outside of the tubes in all cases. It should be noted that this effect is limited to a narrow surrounding of the object due to the application of a thresholding mask.

Figure 7.4 shows corresponding reconstructions for a transverse section of the human brain in vivo. Again, the KWIC reconstruction using 8 echoes suffers from streaking artifacts, while the accuracy of the maps involving all echoes is limited. In the latter case, both maps are blurred and the spin-density map is contaminated by sharp structures with clear T_2 weighting. This results from padding the high frequencies with data from late echoes and poses a general problem when sharing data with varying contrast. The iteratively calculated maps present without these artifacts.




Figure 7.5: Snapshot reconstructions of the human brain (same data as Figure 7.4) at the time of the first (*ec1*), 6th (*ec6*), and last echo (*ec16*) using the proposed method (*iter*), direct regridding of the single echoes (*regr*), and KWIC combining 8 neighboring echoes (*kwic 8*).

Finally, Figure 7.5 compares snapshots of the first, 6th, and last echo reconstructed using the proposed method with Eq. (7.12), direct regridding, and KWIC with sharing of 8 echoes. The contrast of the regridding images can be taken as ground truth because in this case only spokes with identical echo time were used, i.e. 32 spokes per echo time point. It can be seen that the snapshots calculated with the proposed method have an identical image contrast without being affected by any streaking artifacts.

7.4 Limitations and Extensions

As demonstrated, existing reconstruction methods that simply merge data from different echo times always require a trade-off between the accuracy of the image contrast and an undersampling of the outer k-space areas, which results in streaking artifacts. The main advantage of the proposed method is that the use of a signal model allows for a proper combination of spokes measured at different times and, thus, exploits all sampled data without having to assume that contrast changes are limited to some central part of k-space.

Computational Load

A disadvantage of the present implementation, however, is a significantly higher computational requirement than for conventional non-iterative methods, as already stated for the single-echo approach in Chapter 5. Because in the multi-echo case an individual Fourier transformation with subsequent regridding is required for each echo time and receiver channel, a single evaluation of the cost function for the examples analyzed here involves 64 FFT and regridding steps. Evaluating the gradient requires even twice the number of operations, and one full iteration of the algorithm often needs several evaluations of the cost function and gradient. Thus, the optimization procedure for finding a solution creates a massive computational load, posing an obstacle even to modern high-speed processors.

However, many operations of the algorithm can be performed in parallel. In the current proof-of-principle implementation, the OpenMP interface was used to parallelize the calculations for different echo times. Hence, the evaluation of the cost function and gradient is executed on different cores at the same time, which decreases the reconstruction time remarkably. Using the system equipped with two Intel Xeon E5345 quad core processors, the 100 iterations took about one minute per slice (excluding the calculation of a look-up table). Despite foreseeable progress in multi-core processor technology which promises significant acceleration, a use of the method in near future is likely to be limited to applications where longer reconstruction times are tolerable. However, because in a typical radiological examination images are rarely viewed by the radiologist immediately after the scan, this limitation might be secondary in practice.

Accuracy

From a theoretical point of view, the proposed method should make optimal use of all data measured and, thus, deliver a high accuracy. Because the solution is found

in a least-squares sense, this should also hold true for data contaminated by Gaussian noise. Simulations performed to test the implementation showed that the estimate indeed converges to the correct solution. In addition, the experimental results were consistent with results obtained from the other approaches. In practice, however, there are a number of factors that might affect the achievable experimental accuracy.

First, the procedure used to determine the coil sensitivities is simple and might introduce a bias due to inappropriate characterization of the profiles. In particular, the procedure fails in areas with no or very low signal intensity, so that routine applications will probably require a more sophisticated procedure. Second, the Fourier transform of the object, as encoded by the MRI signal, is non-compact. Therefore, any finite sampling is incomplete, which makes it impossible to invert the spatial encoding exactly. Consequently, truncation artifacts arise when employing a discrete Fourier transformation (DFT), which present as a convolution of the object with a sinc function (see Chapter 8 for a detailed discussion of this effect). Because DFTs are used to compare the snapshots with the measured data and, further, because the truncation artifacts are different for each echo time, this effect might interfere with the estimation of a solution that is fully consistent with all measured data. In particular, ringing patterns around high-intensity spots might lead to a bias of surrounding pixels that possibly diverts the decay estimated in these areas.

Finally, if the relaxation process is so fast that the signal decay is insufficiently captured along the echo train, inaccurate spin-density and relaxivity values might be estimated. For example, if a signal intensity above noise level is received only at the first echo time, the algorithm will probably assume a too low spin-density and a too low relaxivity, which would likewise describe the observed signal intensities in a least-squares sense. However, this is a general problem of any T_2 estimation technique and can only be overcome by a finer temporal sampling. Also, inaccuracies that might occur when the actual relaxation process differs from a pure mono-exponential decay are not limited to the present method.

Extensions

Although focused on the reconstruction of FSE data, the proposed method can be used for multi-echo data from other sequences as well. Depending on the contrast mechanism of the individual sequence, it might be necessary to adapt the signal model (7.2). Further, for non-refocused multi-echo sequences the data can be significantly affected by off-resonance effects due to the pronounced sensitivity of radial trajectories. In this case, it might be possible to map also the off-resonances by replacing the relaxivity with a complex-valued parameter and adjusting the gradient of the cost function. How-

ever, due to the extended parameter space it is expected that this strategy will be only successful if suitable constraints for the estimates are incorporated.

Moreover, the reconstruction concept is not only applicable to data with different contrast due to spin relaxation or saturation, but it can be adapted to completely different imaging situations. In this regard, the work presented in this chapter should be seen as a demonstration of the feasibility to extend the inverse reconstruction scheme to more complex imaging problems that require a non-linear processing. A condition for any application, however, is that a simple analytical signal model, comparable to Eq. (7.2), can be formulated. Further, it is required that the derivative of the signal model with respect to all components of the parameter space can be calculated, and that the model allows for a relatively fast evaluation of the cost function and its gradient.

7.5 Summary

In the radial FSE technique, multiple spokes are acquired from a single RF excitation using a train of spin-echoes. It offers a remarkable acceleration of the measurement time, but is accompanied by a non-uniform T_2 weighting of the sampled data. When employing a simple regridding procedure, the images, therefore, exhibit an averaged contrast. On the other hand, due to the oversampling of k-space center, a radial FSE data set contains inherent temporal information, which can be extracted using special reconstruction techniques. Most existing methods approximate a series of time-resolved images by specifically mixing data with different echo times, which, however, tends to create ring-shaped artifacts in areas with pronounced relaxation.

Therefore, an improved method is proposed, which is based on the iterative reconstruction concept presented in Chapter 5. Instead of merging inconsistent data, the proposed method employs a non-linear model of the received MRI signal, which enables to combine the spokes measured at different echo times in a reasonable way. As a consequence, the approach directly yields a spin-density and relaxivity map from the k-space data without calculating intermediate images. Because a numerical optimization procedure is utilized for finding a solution, it makes optimal use of all data sampled and allows for an efficient T_2 quantification from a single radial data set. In comparison with Cartesian quantification techniques, such data can be acquired in a shorter time and with less motion sensitivity, so that the proposed method arises as an interesting option for rapid quantification applications. Experimental results for a phantom and human brain in vivo have been presented and demonstrate that the method yields spin-density and relaxivity maps that are neither affected by the typical artifacts from TE mixing, nor by streaking artifacts from the incomplete k-space coverage at the individual echo

times. As a drawback, the method is computationally intensive and presently limited to applications where a delayed reconstruction is acceptable.

Chapter 8

Compensation of Truncation Artifacts for Cartesian Sampling

In this final part, the rather general problem of truncation artifacts is addressed, which arise in all MRI techniques due to the finite sampling of the k-space information. It is demonstrated that a minimization of the total variation can be exploited to extrapolate the measured data in k-space, and a novel method is presented which utilizes this idea to reduce the truncation artifacts for Cartesian sampling schemes. In the presence of substantial noise, the approach can be extended to additionally yield an edge-preserving denoising.

8.1 Cartesian Sampling and the Gibbs-Ringing Effect

It was explained in Section 2.2.2 that spatial information can be encoded into the resonance signal using magnetic field gradients, which link the Larmor frequency of the excited spins to their spatial location. Because all excited spins of the object contribute to the voltage induced in the receive coil, the MRI signal corresponds to the continuous Fourier transform of the object's proton density

$$S(\mathbf{k}) = \int \rho(\mathbf{x}) e^{2\pi i \mathbf{k} \cdot \mathbf{x}} d\mathbf{x} , \quad (8.1)$$

where additional signal modulations from spin relaxation, off-resonance effects, and other mechanisms are neglected.

It is well known that functions with a compact support have a Fourier transform with non-limited support, conforming with the reciprocity property of the Fourier transformation. For instance, the Fourier transform of a finite rectangle is composed of sinc

functions in each dimension, which have an infinite support. Apparently, all measurement objects fall into this category and, thus, the encoded spatial information comprises an infinite number of k -space values. However, because only a single value can be measured at a time, it is impossible to fully sample such Fourier transform by traveling the k -space with gradient fields. For this reason, there are two general experimental restrictions for MRI acquisitions. First, the continuous Fourier transform is sampled at discrete positions, which can be seen as a multiplication with a comb-function in the frequency space. In image space, this corresponds to a convolution with a reciprocally spaced comb-function, and for the simple case of a Cartesian sampling scheme it leads to periodic object copies with a spacing inverse to the sample distance in k -space. Second, the Fourier transform can be sampled only within a finite region around the k -space center with all other information missing.

In the conventional reconstruction scenario, a discrete Fourier transformation of the finitely sampled data is performed to obtain an image. This strategy implicitly assumes that the Fourier transform is zero everywhere outside the sampled region, because the DFT does not distinguish whether frequencies are unmeasured or zero-valued. It is clear that the assumption is not appropriate for finite objects, although the corresponding reconstruction totally complies with all data measured. In fact, any solution that coincides at the sampling positions is a “valid” reconstruction because the finite sampling opens degrees of freedom from the null space of the projection operation. Setting this null space to zero is a simple and convenient solution. Unfortunately, however, the procedure corresponds to a multiplication of the true object’s Fourier transform with a rect-function along the sampled k -space lines. When using a Cartesian sampling pattern, this results in a convolution of the true object with a sinc function in both image directions, which is well known as truncation artifact or Gibbs ringing effect and mainly presents as an oscillating overshoot of the image intensity near discontinuities [4, 130]. Although the problem may be reduced by increasing the measured k -space, many practical applications still rely on acquisitions with a relatively low matrix resolution in at least one image dimension and, therefore, suffer from respective artifacts. In the case of a radial sampling pattern, the resulting image artifacts may be best understood when considering the FBP reconstruction model. Here, Gibbs ringing effects arise in each of the projection profiles, which translate into circular patterns when performing the backprojection operation, as best visible in the top left of Figure 4.8. Often, the artifacts exhibit also as a foggy surrounding of the object caused by backprojected undershoots from the object border, which are flipped to positive values in the magnitude view. Thus, truncation artifacts occur also for radial sampling, although the problem is of somewhat lower relevance because in the radial case it is possible to acquire a large number of samples for each spoke without additional measurement time. Therefore, the




Figure 8.1: (Top) One-dimensional profile of a rectangle reconstructed by Fourier transformation from 96 Fourier samples (solid line) in comparison to the true function (dotted). While the true function is piecewise constant, the Fourier reconstruction exhibits severe ringing artifacts due to truncation of the Fourier coefficients, which causes an increased total variation (TV) value. (Bottom) Magnified view.

following discussion is confined to Cartesian k-space sampling, but it can be adapted to the radial situation as well.

So far, various methods have been developed to ameliorate image disturbances due to finite k-space sampling [131, 130, 132, 133]. However, in the majority of MRI applications and, in particular, for most commercially available MRI systems, only a simple data filtering is routinely employed. In this case, visual reduction of the ringing artifacts is achieved by a smearing of the intensity oscillations, which leads to an undesired loss of image resolution. Alternative methods attempt to extrapolate the measured data and thereby avoid a sharp cut-off in k-space [134, 135, 136, 137]. A key difference to the filtering approach is that the actually measured data is not changed but supplemented with synthetic data – a reasonable strategy as the measured data is not incorrect but only incomplete. It can be achieved by exploiting a priori knowledge about the true object and, consequently, all extrapolation techniques rely on certain assumptions, where the existing methods follow different strategies. In this regard, the next sections demonstrate that also the assumption of a piecewise-constant object can be utilized to successfully extrapolate data in k-space and concomitantly reduce the ringing artifacts without compromising image resolution.

8.2 TV-Constrained Data Extrapolation

Figure 8.2 compares the one-dimensional profile of a rectangle reconstructed by Fourier transformation from only 96 Fourier samples to that of the original function. It clearly illustrates severe ringing artifacts, although the true function is piecewise constant and free of any oscillations. Apparently, such oscillations can be quantified using the total variation, which sums the modulus of jumps between all neighboring pixels of an image $I(x, y)$

$$\text{TV}(I) = \sum_{y=0}^N \sum_{x=0}^N |I(x, y) - I(x-1, y)| + |I(x, y) - I(x, y-1)|, \quad (8.2)$$

as already defined in Eq. (5.10). Thus, in a similar manner to the streaking artifacts that arise for radial undersampling, the truncation artifacts lead to an increased TV value relative to that of the true object, so that the TV value may also be taken as a measure of the artifact strength for finite k-space sampling. Therefore, the proposed idea is to exploit this measure for finding a solution that represents the true object in a more reasonable way than the solution obtained by setting all unmeasured frequencies to zero. In more detail, the idea is to add a set of synthetic frequencies \mathbf{v} to the measured data \mathbf{y} , which is specifically chosen such that the TV value of the image reconstructed from the combination of the measured and synthetic data is minimized

$$\mathbf{v} = \arg \min_{\mathbf{v}} \text{TV}(\mathcal{F}\{\mathbf{v} \oplus \mathbf{y}\}), \quad (8.3)$$

where \mathcal{F} denotes the discrete Fourier transformation. Interestingly, by searching for the set of synthetic frequencies \mathbf{v} , the unmeasured k-space data is recovered to a certain degree if the assumption of a piecewise-constant object is appropriate.

Estimation of the synthetic data can be achieved by minimizing Eq. (8.3) with a non-linear numerical optimization technique. The present proof-of-principle implementation is based on the CG-Descent algorithm [108] as it proved to rather efficiently solve large-scale problems. The algorithm can be employed in a straight-forward manner, requiring only the evaluation of a cost function and its gradient for given estimate vectors \mathbf{v} . For the problem defined in Eq. (8.3), the cost function simply has the form

$$\Phi(\mathbf{v}) = \text{TV}(\mathcal{F}\{\mathbf{v} \oplus \mathbf{y}\}). \quad (8.4)$$

The gradient of the cost function corresponds to the derivative of this function with respect to all components of the estimate vector \mathbf{v} . Because the discrete Fourier transformation is an unitary operation, it can be evaluated conveniently by calculating the gradient of the TV term in the image domain (i.e. estimating a vector that describes

how the TV value changes for modifications of the individual pixels), followed by an inverse Fourier transformation to the frequency domain.

Extended TV Formulation

Calculation of the TV value according to Eq. (8.2) uses only the first-order derivative of the image with respect to its x and y direction. This value is minimized if an image consists of areas with constant signal intensity, so that the extrapolation procedure yields a solution primarily with constant areas. While desirable for truly flat objects like numerical phantoms, it tends to create images with a slightly blocky or patchy appearance for real-world objects as already mentioned in Chapter 5. Therefore, it is advisable to additionally include second-order derivatives into the TV term, which then allows for intensity gradients in the images and yields more naturally looking solutions

$$\begin{aligned} \text{TV}_2(I) = \sum_{y=0}^N \sum_{x=0}^N \sigma \cdot (& |I(x, y) - I(x - 1, y)| + |I(x, y) - I(x, y - 1)|) \\ & + (1 - \sigma) \cdot (|I(x - 1, y) - 2 \cdot I(x, y) + I(x + 1, y)| \\ & + |I(x, y - 1) - 2 \cdot I(x, y) + I(x, y + 1)| \\ & + |I(x, y) - I(x - 1, y) - I(x, y - 1) + I(x - 1, y - 1)|) . \end{aligned} \quad (8.5)$$

Here, $\sigma \in [0, 1]$ is a weighting factor which can be used to tune the images between a slightly more blocky-looking and a slightly smoother appearance. For the reconstructions presented, it was set to $\sigma = 0.77$ based on the considerations by Geman et al. [112].

Edge-Preserving Denoising

In practice, experimental MRI data can significantly be contaminated by Gaussian noise. While the aforementioned approach is still able to reduce visible truncation artifacts under these circumstances, it does not reduce image noise because the measured k-space data remains unchanged. On the other hand, an additional denoising may be achieved by loosing the fixed bound on the measured data, that is by introducing a data fitting term. In this case, the algorithm not only adds synthetic frequencies to obtain a TV minimization but is also allowed to find a solution that slightly diverges from the measured data, which yields an effective edge-preserving denoising. Therefore, the estimate vector \mathbf{v} has to be extended such that it contains both, synthesized frequencies as well as frequencies from the measured part of k-space, which is indicated by writing \mathbf{v}_d instead.

In the denoising case, the cost function takes the form

$$\Phi(\mathbf{v}_d) = \lambda \cdot \|\mathbf{v}_d \ominus \mathbf{y}\|_2^2 + \text{TV}(\mathcal{F}\{\mathbf{v}_d\}) , \quad (8.6)$$

where \ominus denotes an operation that calculates the residual between the measured values and the corresponding entries of the estimate, which are now contained in the vector \mathbf{v}_d . Further, λ is a weighting factor that allows to select the desired denoising strength. While a low weight permits considerable divergences from the measured values and, thus, leads to an effective removal of noise, it can also cause a loss of object detail if selected too low. Therefore, the weight has to be adjusted with respect to the signal-to-noise ratio of the measurement sequence, where a reasonable strategy is to estimate a fixed value once for each protocol by computing a set of test images with different λ values and selecting the value yielding the desired degree of denoising.

Phase Variations

Although the basic physical quantity measured by MRI, i.e. the spin-density modulated by relaxation or saturation effects, should be real-valued and non-negative in theory, inherent experimental phase variations usually cause the observed object to be complex-valued. Moreover, nearly all modern MRI systems use multiple receive coils with complex-valued sensitivity profiles, yielding differently modulated views of the object. This issue has been already mentioned in Section 5.4, and as a consequence spatially varying transitions between the real and imaginary component as well as local intensity changes occur, which conflict with the assumption of a piecewise-constant quantity. In fact, the undesired phase variations pose a significant complication for the application of any TV-based processing technique in a practical MRI scenario, although the problem is often neglected. Therefore, it is important to point out that TV-based methods which are used with real MRI data require some mechanism to cope with both, the phase variations and the multi-coil situation.

In this proof-of-principle study, phase variations were removed in a preprocessing step by performing a Fourier transformation of the data from each coil and calculating the sum-of-squares of all channels in the image domain. Subsequently, an inverse Fourier transformation of the sum-of-squares data was performed to obtain a combined data set with real and non-negative values in the image domain, which enables a calculation of the TV value using only the real part of the image. While this simple technique turned out to be sufficient for demonstrating a removal of truncation artifacts by TV-constrained data extrapolation, routine applications will probably require a more sophisticated procedure, in particular when employing undersampling techniques such

as parallel imaging and when using complex coil configurations with more localized sensitivities of the individual receiver elements.

8.3 Experimental Results

Simulations were performed with the Shepp-Logan phantom by evaluating its non-compact analytical Fourier transform at the sampling positions along the trajectory, here yielding a matrix of 96×96 values. The experiments were conducted using the MRI hardware described in Section 3.1, where only the four primary mode signals of the head coil array were recorded. Measurements were performed for a water phantom and human brain in vivo, where written informed consent was obtained from all subjects prior to each examination. For demonstration purposes, the acquisitions were confined to a simple slice-selective spin-echo sequence with a 200×200 mm² FOV, covered by a 96×96 acquisition matrix. Different sequence settings were used to obtain data sets with a low and a high level of noise, where the latter was achieved by reducing the flip angle and slice thickness while increasing the receiver bandwidth.

All images were reconstructed on a 288×288 matrix corresponding to an extrapolation factor of 3. The proposed algorithm was run for a fixed number of 3000 iterations, which takes about 2–3 min on the dual quad-core system used. In cases where an additional data fitting term was used, the weighting factor λ was adjusted manually to yield a reasonable solution as judged by visual inspection. Zero-padded solutions with and without filtering were calculated for comparison. Here, a simple Lanczos sigma filter, i.e. multiplication with a sinc function, was applied, where the window width was selected such that the sinc function's first null coincides with the border of the measured k-space. Although other filters might perform better, it serves to demonstrate the general problem related to data filtering. Finally, all images were magnified and cropped to improve the visibility of the artifacts.

Results

Figure 8.2 shows different reconstructions of the Shepp-Logan phantom (left column) together with the respective Fourier transforms (right column). It is clearly visible that the zero-padded solution (zero) suffers from severe ringing artifacts around all edges of the phantom. The extent of the measured k-space can be seen in its Fourier transform. Most ringing artifacts disappear after filtering (filter), however, at the expense of a significant loss of image resolution. In contrast, the image reconstructed with the proposed method (TV) is neither affected by ringing artifacts nor by blurring, and it presents with considerably sharper edges relative to the zero-padded solution. Its




Figure 8.2: (Left) Images of the numerical Shepp-Logan phantom (96×96 samples, 288×288 reconstruction matrix) and (right) corresponding k -space representations reconstructed using zero-padding (zero), filtered zero-padding (filter), and the proposed extrapolation method (TV). For comparison, a data set with a fully sampled 288×288 matrix is shown (full). Arrow = truncation artifact.

Fourier transform reveals that the measured data has been properly extrapolated into the uncovered areas of k-space. For comparison, a full reconstruction from 288 x 288 samples is shown in the bottom row (full).

Figure 8.3 demonstrates the application of the method to experimental data obtained for a phantom (left column) and a human brain in vivo (right column) in comparison to zero-padded (zero) and filtered zero-padded solutions (filter). Again, the ringing artifacts obtained for zero padding (indicated by arrows) are significantly reduced when using TV-constrained data extrapolation with only first-order (TV) or additionally second-order derivatives (TV2). The blocky appearance of the TV reconstruction becomes much more smoother for the TV2 approach, although both solutions (TV and TV2) look somewhat more blocky than the zero-padding solution.

Figure 8.4 shows reconstructions of the Shepp-Logan phantom from noisy data using zero-padding (zero), the extrapolation approach (TV), and its combination with denoising (TVdns). While the basic extrapolation approach leads to a reduction of truncation artifacts also for noisy data, it does not reduce the noise patterns. However, when extending the TV penalty to the measured data, the method effectively flattens noise patterns in addition to the suppression of ringing artifacts.

Finally, corresponding reconstructions from experimental data with a high degree of noise are shown in Figure 8.5. Here, a combination of first-and second-order derivatives was used for the TV calculation. As in the simulations, the proposed method leads to a reduction of truncation artifacts (TV2), while the extension to data fitting yields an additional edge-preserving denoising (TV2dns).

8.4 Limitations

Accuracy

Both, simulations and experiments demonstrate that the proposed reconstruction technique allows to effectively reduce truncation artifacts due to finite sampling in MRI. This finding reveals that the concept of TV minimization can be used not only to interpolate between sample points for filling unmeasured k-space gaps (as done in Chapter 5), but also to extrapolate the data to the outside of the measured area. In both cases, the TV penalty yields a certain recovery of missing k-space information and, thus, the extrapolation approach is closely related to the methodology discussed in Chapter 5. As a difference, however, the unmeasured k-space positions are given directly on a Cartesian grid in the extrapolation case, so that the cost function can be evaluated without the need of a convolution operation.




Figure 8.3: Spin-echo images (96×96 samples, 288×288 reconstruction matrix) of (left) a phantom ($TR/TE = 4000/8$ ms, BW 243 Hz/pixel, FA 70° , 3 mm slice) and (right) a human brain in vivo ($TR/TE = 4000/25$ ms, BW 180 Hz/pixel, FA 90° , 2 mm slice) using zero-padding (zero), filtered zero-padding (filter), the proposed extrapolation method with first-order (TV), and additionally second-order derivatives (TV2). Arrows = truncation artifacts.




Figure 8.4: Images of the numerical Shepp-Logan phantom reconstructed from noisy data (96×96 samples, 288×288 reconstruction matrix) using zero-padding (zero), the proposed extrapolation method (TV), and the proposed method combined with denoising (TVdns). Arrow = truncation artifact.




Figure 8.5: Spin-echo images of (top) a phantom ($TR/TE = 4000/100$ ms, BW 789 Hz/pixel, FA 50° , 1 mm slice) and (bottom) a human brain in vivo ($TR/TE = 4000/15$ ms, BW 401 Hz/pixel, FA 70° , 1 mm slice) reconstructed from noisy data (96×96 samples, 288×288 reconstruction matrix) using zero-padding (zero), the proposed extrapolation method (TV2), and the proposed method combined with denoising (TV2dns). Arrows = truncation artifacts.

Usually, the extrapolated images exhibit a somewhat more blocky appearance compared to zero-padding. This effect reflects a rather general property of all TV-based image processing techniques and is often appraised as their main disadvantage. However, it should be noted that the smoothness observed for zero-padding originates to a significant degree from the inherent convolution with the sinc-function. As a consequence of this operation, a sharp edge of the object is mapped as a rather smooth pattern, which might appear more familiar to the viewer than a blocky image but strictly represents an image artifact. Hence, the extrapolation technique may even lead to a slight gain of resolution due to a sharpening of the point-spread function, following from the reciprocity property of the Fourier transformation. This effect can be best seen in Figure 8.5 when comparing the borders of the dark brain vessels obtained for zero padding (zero) with the proposed method (TV2).

Residual image artifacts are explained by multiple reasons. First, the method is based on the assumption that the true object is piecewise constant, which is only approximately valid for real-world objects. In the presence of additional experimental effects like flow artifacts, the assumption might be even less appropriate. Consequently, the extrapolation performance depends on the object's conformance with the assumption that it is piecewise-constant. Moreover, if the true object contains strongly varying patterns, the algorithm may erroneously soften such patterns by supplementing respective high frequencies. On the other hand, in the majority of cases the assumption of a piecewise-constant object seems to be more appropriate than that of all conventional DFT reconstructions, namely a Fourier transform of the object that is zero outside the sampled k-space area.

Second, the proposed method synthesizes only a finite number of additional frequencies, whereas an infinite number of k-space samples would be required to completely eliminate all truncation effects. In practice, however, it turned out that there is no perceivable benefit of extrapolating by a factor of higher than three. The underlying reason is that the method yields an implicit filtering of the extrapolated data: assuming that the extrapolation procedure would recover the unmeasured k-space samples exactly, then a new truncation effect would arise at the extended border and again lead to ringing artifacts in image space (though with a higher oscillation frequency). Because this would increment the TV value, the method automatically lowers outer frequencies during the extrapolation procedure to prevent the upcoming of new ringing artifacts. Hence, the extrapolated values diverge categorically from the true frequencies which, in this case, is a rather desirable feature as the prime target is to reduce visually annoying ringing artifacts rather than to gain super-resolution.

Third, if a completely artifact-free reconstruction of the object would be available, then respective frequency samples could be calculated with a discrete Fourier transformation

of the given image. Interestingly, these samples would diverge from the experimentally measured frequencies, because image pixels are discrete and, thus, the Fourier transform of the image is periodic such that outer frequencies from neighboring copies (of the true object's non-compact Fourier transform) overlap. This is different to the experimental situation where the object is continuous and the outer frequencies are missing instead of overlapping. Consequently, an artifact-free discrete reconstruction can only be obtained if the samples used for the reconstruction specifically diverge from the measured frequencies. A complete artifact removal, therefore, requires to alter the measured frequencies instead of keeping them unchanged. Unfortunately, the information how the samples have to be adjusted is not available, so that in practice a data fitting term might be the best solution when a complete removal of ringing artifacts is needed. However, this might cause a loss of object detail as described before.

Implementation Issues

The modulus function in the TV formula (8.2) has a fundamental role for the success of TV-based image processing, as discussed in detail in Section 5.8. Because the subtraction of neighboring pixels – performed before taking the modulus – can be seen as applying a difference operator to the estimate, TV minimization yields a solution with minimum ℓ_1 -norm in the difference basis. Due to the specific character of the modulus function, this solution tends to be sparse in the difference basis: it has few large jumps and most differences between neighboring pixels are near zero, which directly translates into a piecewise-constant image (and explains the edge preserving character of TV-based denoising). If the modulus would be replaced by a square function, then the optimizer would try to find a minimum ℓ_2 -norm solution with minimal jumps between all neighboring pixels. This corresponds to a globally smooth image, which is usually undesired due to a loss of sharp edges. While it is rather simple to obtain a minimum ℓ_2 -norm solution as its cost function is strictly convex, finding a minimum ℓ_1 -norm solution is much more challenging and many optimization algorithms fail if directly applied to the TV problem. One major reason is that the derivative of the modulus function is just ± 1 , which does not help to guess a reasonable step size toward the function's minimum. It turned out that the CG-Descent algorithm is capable to handle the problem as it comprises a powerful line-search procedure. However, it is fair to say that the method is probably not optimal for finding the solution. In particular, the convergence tends to be somewhat sensitive to the scaling of the data. In order to ensure convergence, it was, therefore, necessary to introduce a scaling factor that limits the modification strength for each iteration and to run the algorithm in turn for a high number of iterations (for example, 3000 iterations as arbitrarily chosen here). Nevertheless, this issue should not be seen as a drawback of the proposed extrapolation

approach itself, but rather as a technical aspect of the optimization method utilized in this proof-of-principle study. Employing a dedicated algorithm for TV minimization should render a scaling factor unnecessary and significantly improve the convergence rate. Although such enterprise promises reconstructions in a fraction of the current processing time, it is outside the scope of the present study.

8.5 Summary

The detected signal in MRI corresponds to the continuous Fourier transform of the object, which is always a function with infinite support. Because in any MRI experiment only a finite number of samples can be acquired, spurious image artifacts arise from the k-space truncation at the border of the covered area, which typically exhibit as oscillating overshoots near discontinuities of the object. These oscillations create a high TV value relative to that of the true object and, thus, the TV value can be used as a measure of the artifact strength if it can be assumed that the true object is piecewise constant to some degree.

Therefore, it is proposed to supplement the measured data with a set of synthesized k-space values which are chosen such that the TV value of the image reconstructed from the combined data is minimized. Estimation of the synthetic data is achieved simply by iteratively minimizing the TV value of the combined image with a non-linear optimization technique, where the measured values remain fixed. In contrast to commonly used filtering approaches, this procedure reduces truncation artifacts without a noticeable loss of the spatial resolution and rather leads to a mild resolution enhancement due to a sharpening of the point-spread function. The effectiveness of the method has been proved for phantom and in vivo human brain data, demonstrating the feasibility of TV-based data extrapolation in k-space and complementing the discussion from Chapter 5 on the use of TV minimization for MRI reconstructions.

If the measured data is seriously contaminated by noise, an extended approach offers edge-preserving denoising by slightly altering the measured data in addition to supplementing synthetic data, which is possible by introducing a data fitting term. Both variants can be implemented as a pure post-processing procedure and do not require any modification of the MRI acquisition technique. While the current implementation suffers from a relatively high computational load, the use of a dedicated TV optimization algorithm promises a processing speed suitable for routine applications.

Chapter 9

Summary and Conclusions

In this thesis, radial data encoding in MRI has been investigated as an alternative strategy to the conventional Cartesian technique and with respect to various practical applications. Several novel methods are proposed that promise a wider applicability of radial acquisition techniques in clinical imaging scenarios.

As an initial step, a comprehensive analysis of the radial sampling scheme was conducted, which revealed that radial trajectories offer certain advantages such as unique undersampling abilities and a lower sensitivity to motion. Other properties like the higher sampling requirement and the stronger off-resonance sensitivity turned out to be disadvantageous and hamper a routine usage. Based on these findings, it is unlikely that radial sampling will replace Cartesian sampling as the de facto standard for generic MRI acquisitions, but it might become the method of choice for a set of specific applications that benefit from its salient imaging properties. Such applications comprise situations where strong object motion is present or where high temporal resolution is required. Typical examples are interventional MRI, cardiac imaging, and magnetic resonance angiography. This appraisal is supported by many recent developments in these fields that deliberately utilize radial acquisition schemes.

The key problem of the higher sampling requirement was addressed in the second part of this work, which outlines the development of a novel reconstruction approach that allows to overcome this burden. A solution is achieved by exploiting prior object knowledge to compensate for missing k-space information so that certain acquisition steps can be skipped. The method is based on an inverse formulation of the reconstruction problem, offering a convenient way to incorporate prior knowledge in the form of penalty functions that prevent a selection of implausible estimates. Because radial undersampling creates streaking artifacts while medical objects tend to be rather piecewise constant, the total variation of the reconstruction was chosen as the main penalty function. Using simulations as well as experimental data, the total variation turned out

to be a highly effective measure for separating the undersampling artifacts from the true object. Moreover, the approach was designed to exploit localized coil sensitivities and to cope with experimental phase variations, which is an essential requirement for the use with real MRI data. The solution is found with a non-linear numerical optimization procedure, where a non-linear variant of the iterative conjugate-gradient method was employed in the proof-of-principle implementation. The effectiveness of the technique was demonstrated in studies of the human brain in vivo, showing that the approach yields reasonable reconstructions from only a low number of acquisition steps with significant violation of the Nyquist criterion. In other words, by employing a dedicated reconstruction technique, the initially higher sampling demand for radial acquisitions can be recasted to a very low requirement, clearly exceeding the sampling efficiency of Cartesian approaches. This offers a remarkable scan-time reduction and renders radial sampling highly attractive for time critical applications. Meanwhile, several related studies have been published in the context of compressed sensing [138, 139, 140, 141], which support the results presented in this thesis and confirm the gain of image quality arising from the advanced reconstruction strategy.

In the final part of this thesis, three subsidiary projects were carried out to complement the introduction of the reconstruction concept. First, by adapting the single-shot STEAM MRI technique to the radial readout scheme, a novel approach for rapid MRI has been presented which offers distortion-free images with significantly enhanced spatial resolution relative to the existing Cartesian version. Because the number of utilizable echoes is limited in the rapid STEAM technique, the approach requires a data processing strategy to handle the incomplete k-space information and, thus, represents a first real-world application that relies on the proposed reconstruction technique. Second, for the case of radial FSE acquisitions, it could be demonstrated that the reconstruction approach is extendable to handle more complex imaging scenarios with data inconsistencies from relaxation or saturation effects. This is achieved by including a modeling of the received MRI signal, which enables a comparison of the estimates to k-space samples measured at different time points. Therefore, the extended approach is capable of merging inconsistent raw data in a reasonable way, allowing to utilize more time-efficient strategies for the data acquisition. In the specific case of a FSE acquisition, the extended approach inherently yields a spin-density and relaxivity map, which offers an efficient quantification of the local T_2 relaxation time from a single radial FSE dataset. Finally, it has been demonstrated that a minimization of the total variation can also be exploited to extrapolate the measured data in k-space, leading to an effective reduction of ringing artifacts that conventionally arise from finite k-space sampling. This result evidences that a constraint of the total variation is effective to recover missing k-space information at a certain distance from acquired sample points

and supports the choice of a respective penalty function in the proposed reconstruction method for undersampled radial acquisitions.

Future Work

Although all techniques presented were experimentally evaluated using in vivo studies, the current work should be seen as a basic research contribution rather than a proposal of methods that are usable “out of the box.” In fact, several details are still insufficiently solved and need further investigation.

First and foremost, the reconstruction time of the iterative approaches is still too long for a routine usage, especially in the case of 3D acquisitions. While it is presently possible to utilize the techniques for research purposes or applications that involve additional postprocessing, the establishment in clinical settings will certainly require a significant acceleration, and respective options should be investigated in a next step. Acceleration might be achieved based on mathematical strategies like preconditioning or, alternatively, technical approaches like massive parallel processing. Further, the use of different numerical optimization techniques should be evaluated, which might be able to find the solution with fewer iterations. In the present work, all calculations were conducted using the CG-Descent algorithm, but the issue described in Section 8.4 indicates that this is probably not the optimal method to solve the total variation problem. Therefore, Newton-type methods [142] as well as modern interior point methods [143] should be investigated as an alternative to the conjugate-gradient approach, which probably offer a more robust convergence behavior for the non-linear cost functions used.

Moreover, future work should also focus on identifying additional or improved penalty functions to achieve even higher data reduction factors. While a penalty on the total variation proved to be astonishingly effective for transverse sections of the human brain, it turned out to be rather ineffective for objects where the assumption of a piecewise-constant object is less appropriate. Unfortunately, many cardiac images fall into this category, although radial sampling itself is well suited for these examinations. For this reason, it would be highly desirable to have a set of penalty functions which represent different assumptions about the true object, so that the incorporated prior knowledge can be adapted to individual applications. Further, several recent publications indicate that the effectivity of the total variation penalty can be increased remarkably by replacing the modulus function with functions that approximate the ℓ_0 norm in a more reasonable way [144, 145, 146]. This idea should be considered in future work as well.

Another open issue consists in the proper selection of weights for the penalty functions. Because an overweighting of the total variation causes loss of object detail, it would

be desirable to have a mechanism that automatically estimates appropriate weighting coefficients from the k-space samples, which is a non-trivial problem. It will also be worthwhile to study if a dynamic variation of the penalty terms during the iterations might lead to improved reconstruction results or faster convergence rates. Finally, when all these working packages have been completed, it will be necessary to evaluate the methods in clinical settings in order to appraise their added value for diagnostic procedures. If successful, there is little doubt that the techniques will be adopted to the commercial market and, thus, find a wide-spread application in radiology.

Taken together, the present work allows to conclude that radial sampling has a high potential to become a routine MRI technique for applications where either strong motion is present or very limited time is available for the data readout. It also shows that advanced numerical processing techniques, in particular iterative reconstruction methods, offer clearly improved image quality due to the flexible incorporation of strategies that compensate for missing or inconsistent data. It is unquestionable that these approaches will attract strong interest in the near future of MRI.

Appendix

Appendix A

Equivalence of Regridding and Filtered Backprojection

In this supplementary section, the equivalence between the regridding and filtered-backprojection technique is demonstrated on a more formal level using operator notation. It complements the discussion from Section 4.2.3 to justify that the proposed reconstruction methods were confined to mapping operations based on the regridding technique.

Notation and Basic Operators

In the following, it is assumed that n_s spokes were measured at the angles $\phi_1, \dots, \phi_{n_s}$ with a base resolution of n pixels, where n is an odd number for simplicity. Thus, the reconstruction procedure consists in mapping a raw-data vector \mathbf{y} with $n_s \times n$ entries onto an image vector \mathbf{x} with $n \times n$ entries, which can be written with the abstract reconstruction operator REC

$$\mathbf{x} = \text{REC}\{\mathbf{y}\}, \quad (\text{A.1})$$

where

$$\text{REC} : \mathbb{C}^{n_s \times n} \mapsto \mathbb{C}^{n \times n}. \quad (\text{A.2})$$

To facilitate the readability, a specific vector notation is used throughout the section. In this notation, the raw-data vector takes the form

$$\mathbf{y} = \begin{pmatrix} S_{1,1} & \cdots & S_{n_s,1} \\ \vdots & \ddots & \vdots \\ S_{1,n} & \cdots & S_{n_s,n} \end{pmatrix}, \quad (\text{A.3})$$

where $S_{i,j} \in \mathbb{C}$ is the measured signal at the j th sampling position of the i th spoke. Hence, each column contains the samples of one spoke. The image vector has the form

$$\mathbf{x} = \begin{pmatrix} I_{1,1} & \cdots & I_{n,1} \\ \vdots & \ddots & \vdots \\ I_{1,n} & \cdots & I_{n,n} \end{pmatrix}, \quad (\text{A.4})$$

where $I_{i,j} \in \mathbb{C}$ denotes the image intensity at the pixel position (i, j) .

For convenience, several basic operators are defined as follows. First, SEL_i is an operator that takes a vector of the raw-data form (A.3) and returns the i th column

$$\text{SEL}_i : \mathbb{C}^{n_s \times n} \mapsto \mathbb{C}^n, \quad \begin{pmatrix} a_{1,1} & \cdots & a_{n_s,1} \\ \vdots & \ddots & \vdots \\ a_{1,n} & \cdots & a_{n_s,n} \end{pmatrix} \mapsto \begin{pmatrix} a_{i,1} \\ \vdots \\ a_{i,n} \end{pmatrix}, \quad (\text{A.5})$$

where $i \in [1, n_s]$. ALN is an operator that takes a vector with n entries and inserts it to the column $(n+1)/2$ of a $n \times n$ vector of form (A.4) where all other entries are zero

$$\text{ALN} : \mathbb{C}^n \mapsto \mathbb{C}^{n \times n}, \quad \begin{pmatrix} a_1 \\ \vdots \\ a_n \end{pmatrix} \mapsto \begin{pmatrix} 0 & \cdots & 0 & a_1 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & a_n & 0 & \cdots & 0 \end{pmatrix}. \quad (\text{A.6})$$

Because n is assumed to be odd, the values are aligned in the center of the vector. Further, BPR is an operator that takes a vector with n entries and inserts it to all columns of a $n \times n$ vector of the form (A.4)

$$\text{BPR} : \mathbb{C}^n \mapsto \mathbb{C}^{n \times n}, \quad \begin{pmatrix} a_1 \\ \vdots \\ a_n \end{pmatrix} \mapsto \begin{pmatrix} a_1 & \cdots & a_1 \\ \vdots & \ddots & \vdots \\ a_n & \cdots & a_n \end{pmatrix}. \quad (\text{A.7})$$

Finally, ROT_i is an abstract operator that takes an $n \times n$ image vector of the form (A.4) and rotates the image by the angle ϕ_i

$$\text{ROT}_i : \mathbb{C}^{n \times n} \mapsto \mathbb{C}^{n \times n}. \quad (\text{A.8})$$

It is assumed for the moment that the rotation by an arbitrary angle can be conducted in an ideal manner without introducing any error. Obviously, this assumption does not hold true in practice because the rotation of a discrete image requires an interpolation procedure that is accompanied by a certain interpolation error.

Description of Reconstruction Procedures

The two reconstruction techniques are now constructed based on these building blocks. The regridding procedure can be written as an operator

$$\text{REGR} : \mathbb{C}^{n_s \times n} \mapsto \mathbb{C}^{n \times n} = \text{M}^{-1} \cdot \mathcal{F}_x^{-1} \cdot \mathcal{F}_y^{-1} \cdot \text{G} \cdot \text{DCF} , \quad (\text{A.9})$$

where DCF corresponds to the density compensation, M^{-1} to the roll-off correction, \mathcal{F}_x^{-1} to the discrete Fourier transformation along the rows of a $n \times n$ image vector, and \mathcal{F}_y^{-1} to the Fourier transformation along the columns, respectively. Further, G corresponds to the gridding operation itself, i.e. the mapping of the samples onto a $n \times n$ image vector. For the radial sampling geometry, this operation can be written using the rotation operator

$$\text{G} : \mathbb{C}^{n_s \times n} \mapsto \mathbb{C}^{n \times n} = \sum_i^{n_s} \text{ROT}_i \cdot \text{ALN} \cdot \text{SEL}_i . \quad (\text{A.10})$$

Thus, the samples from each individual spoke are selected, aligned on a matrix, rotated according to the spoke angle, and summed over all spokes. Because the gridding operation is constructed from ideal rotations that are assumed to perform with negligible interpolation error, the roll-off correction can be omitted for the moment, which yields

$$\text{REGR} = \mathcal{F}_x^{-1} \cdot \mathcal{F}_y^{-1} \cdot \sum_i^{n_s} \text{ROT}_i \cdot \text{ALN} \cdot \text{SEL}_i \cdot \text{DCF} . \quad (\text{A.11})$$

The filtered-backprojection reconstruction can be written as

$$\text{FBP} : \mathbb{C}^{n_s \times n} \mapsto \mathbb{C}^{n \times n} = \sum_i^{n_s} \text{B}_i \cdot \mathcal{F}^{-1} \cdot \text{RL} \cdot \text{SEL}_i , \quad (\text{A.12})$$

where RL corresponds to the weighting of the spoke samples with the Ram-Lak filter, \mathcal{F}^{-1} to the inverse Fourier transformation along the n entries of a vector, and B_i to the backprojection of a profile at the angle ϕ_i . It can be written as

$$\text{B}_i : \mathbb{C}^n \mapsto \mathbb{C}^{n \times n} = \text{ROT}_i \cdot \text{BPR} , \quad (\text{A.13})$$

which then yields

$$\text{FBP} : \mathbb{C}^{n_s \times n} \mapsto \mathbb{C}^{n \times n} = \sum_i^{n_s} \text{ROT}_i \cdot \text{BPR} \cdot \mathcal{F}^{-1} \cdot \text{RL} \cdot \text{SEL}_i . \quad (\text{A.14})$$

Equivalence of the Approaches

The DCF operator in the regridding approach performs a weighting of the measured values according to the local sample density. Due to the symmetry of the radial trajectory, it corresponds to a component-wise multiplication with a vector of the form

$$\mathbf{D}_{\text{DCF}} = \begin{pmatrix} |k_1| & \cdots & |k_1| \\ \vdots & \ddots & \vdots \\ |k_n| & \cdots & |k_n| \end{pmatrix}, \quad (\text{A.15})$$

where $|k_i|$ is the distance of the i th sample from the k -space center. Thus, each column corresponds to the Ram-Lak filter used in the FBP approach, and the following relation holds true

$$\text{SEL}_i \cdot \text{DCF} = \text{RL} \cdot \text{SEL}_i, \quad (\text{A.16})$$

which yields

$$\text{REGR} = \mathcal{F}_x^{-1} \cdot \mathcal{F}_y^{-1} \cdot \sum_i^{n_s} \text{ROT}_i \cdot \text{ALN} \cdot \text{RL} \cdot \text{SEL}_i. \quad (\text{A.17})$$

Using the linearity of the Fourier transformation gives

$$\text{REGR} = \sum_i^{n_s} \mathcal{F}_x^{-1} \cdot \mathcal{F}_y^{-1} \cdot \text{ROT}_i \cdot \text{ALN} \cdot \text{RL} \cdot \text{SEL}_i. \quad (\text{A.18})$$

Further, because a rotation in image space leads to a rotation of the Fourier transform, which follows from the definition of the two-dimensional transformation, the rotation operator commutes with the Fourier operators

$$\text{REGR} = \sum_i^{n_s} \text{ROT}_i \cdot \mathcal{F}_x^{-1} \cdot \mathcal{F}_y^{-1} \cdot \text{ALN} \cdot \text{RL} \cdot \text{SEL}_i. \quad (\text{A.19})$$

Applying the operator chain $\text{ALN} \cdot \text{RL} \cdot \text{SEL}_i$ to the raw-data vector \mathbf{y} yields a $n \times n$ vector of the form

$$\begin{pmatrix} 0 & \cdots & 0 & a_1 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & a_n & 0 & \cdots & 0 \end{pmatrix} \quad (\text{A.20})$$

Obviously, if the operator \mathcal{F}_y^{-1} is subsequently applied, it will only affect the central column and, therefore,

$$\mathcal{F}_y^{-1} \cdot \text{ALN} = \text{ALN} \cdot \mathcal{F}_y^{-1}, \quad (\text{A.21})$$

which gives

$$\text{REGR} = \sum_i^{n_s} \text{ROT}_i \cdot \mathcal{F}_x^{-1} \cdot \text{ALN} \cdot \mathcal{F}_y^{-1} \cdot \text{RL} \cdot \text{SEL}_i. \quad (\text{A.22})$$

Finally, if the operator \mathcal{F}_x^{-1} is applied to a vector of the form (A.20), which conducts an inverse Fourier transformation along the rows of the vector, it yields

$$\begin{pmatrix} a_1 & \cdots & a_1 \\ \vdots & \ddots & \vdots \\ a_n & \cdots & a_n \end{pmatrix}, \quad (\text{A.23})$$

because all Fourier coefficients are zero except the offset coefficient (it is assumed that the DFT uses an ordering such that the low frequency components are aligned in the center, which is a common convention). Thus,

$$\mathcal{F}_x^{-1} \cdot \text{ALN} = \text{BPR}, \quad (\text{A.24})$$

which leads to

$$\text{REGR} = \sum_i^{n_s} \text{ROT}_i \cdot \text{BPR} \cdot \mathcal{F}^{-1} \cdot \text{RL} \cdot \text{SEL}_i = \text{FBP}. \quad (\text{A.25})$$

This shows that regridding and filtered backprojection are equivalent reconstruction techniques under the assumption that images can be rotated in an ideal way, which, however, is an invalid assumption. In fact, any method for image rotations introduces certain artifacts, and, therefore, the two reconstruction techniques create slightly different image vectors in practice, where the differences depend on how the rotation or, respectively, interpolation operators are actually implemented (for instance, using a convolution with the Kaiser-Bessel kernel). To analyze differences in the reconstruction accuracy on a quantitative level, it is necessary to precisely derive the interpolation error arising from the concrete interpolation method as well as the error propagation along the operator chain (in particular with respect to the Fourier transformations). It is also important to consider potential dependencies on the object geometry because many interpolation techniques exploit certain properties of the input data (for example, smoothness of the intensity, which is a rather inappropriate assumption for k-space data). Nevertheless, based on the underlying equivalence of both reconstruction techniques as derived in this section, it is fair to state that these effects are secondary to the key problems addressed in the thesis and, in particular, in view of the severe inaccuracies that are introduced by the experimental measurement and signal-detection procedure.

Bibliography

- [1] P.D. Lauterbur. Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature*, 242:190–191, 1973.
- [2] A. Kumar, D. Welte, and R.R. Ernst. NMR Fourier Zeugmatography. *Journal of Magnetic Resonance*, 18:69, 1975.
- [3] E.M. Haacke, R.W. Brown, M.R. Thompson, and R. Ventatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. Willey-Liss, 1999.
- [4] Z.P. Liang and P.C. Lauterbur. *Principles of Magnetic Resonance Imaging*. IEEE Press Series on Biomedical Engineering, 2000.
- [5] M.A. Bernstein, K.F. King, and X.J. Zhou. *Handbook of MRI Pulse Sequences*. Elsevier Academic Press, 2004.
- [6] E.M. Purcell, H.C. Torrey, and R.V. Pound. Resonance absorption by nuclear magnetic moments in a solid. *Physical Reviews*, 69:37–38, 1946.
- [7] F. Bloch, W.W. Hansen, and M. Packard. Nuclear induction. *Physical Reviews*, 69:127, 1946.
- [8] E.L. Hahn. Nuclear induction due to free larmor precession. *Physical Reviews*, 77:297–298, 1950.
- [9] C.P. Slichter. *Principles of Magnetic Resonance*. Springer, 1996.
- [10] H. Haken and H.C. Wolf. *Atom- und Quantenphysik*. Springer, 2003.
- [11] F. Bloch. Nuclear induction. *Physical Reviews*, 70:460–474, 1946.
- [12] A. Haase, J. Frahm, D. Matthaei, W. Hänicke, and K.D. Merboldt. FLASH imaging: Rapid NMR imaging using low flip angle pulses. *Journal of Magnetic Resonance*, 67:258–266, 1986.
- [13] R.V. Damadian. Tumor detection by nuclear magnetic resonance. *Science*, 171:1151–1153, 1971.
- [14] H.H. Barrett and K.J. Myers. *Foundations of Image Science*. Wiley, 2004.
- [15] J. Radon. Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Ber. Sächs. Akad. Wissenschaft. Leipzig Math. Phys. Kl.*, 69:262–267, 1917.

- [16] K.J. Jung and Z.H. Cho. Reduction of flow artifacts in NMR diffusion imaging using view-angle tilted line-integral projection reconstruction. *Magnetic Resonance in Medicine*, 4:349–360, 1991.
- [17] G.H. Glover and J.M. Pauly. Projection reconstruction techniques for reduction of motion effects in MRI. *Magnetic Resonance in Medicine*, 28:275–289, 1992.
- [18] J.D. O’Sullivan. A fast sinc function gridding algorithm for Fourier inversion in computer tomography. *IEEE Transactions on Medical Imaging*, 4(4):200–207, 1985.
- [19] J. Jackson, C.H. Meyer, D.G. Nishimura, and A. Macovski. Selection of a convolution function for Fourier inversion using gridding. *IEEE Transactions on Medical Imaging*, 10(3):473–478, 1991.
- [20] J.F. Kaiser. Nonrecursive digital filter design using the I_0 -SINH window function. In *Proc. 1974 IEEE Int. Symp. on Circuits and Syst.*, pages 20–23, 1974.
- [21] P.J. Beatty, D.G. Nishimura, and J.M. Pauly. Rapid gridding reconstruction with a minimal oversampling ratio. *IEEE Transactions on Medical Imaging*, 24:799–808, 2005.
- [22] F. Aurenhammer. Voronoi diagrams - a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345–405, 1991.
- [23] V. Rasche, R. Proksa, R. Sinkus, P. Börnert, and H. Eggers. Resampling of data between arbitrary grids using convolution interpolation. *IEEE Transactions on Medical Imaging*, 18(5):385–392, 1999.
- [24] G.N. Ramachandran and A.V. Lakshminarayanan. Three-dimensional reconstruction from radiographs and electron micrographs: application of convolutions instead of Fourier transforms. In *Proc Natl Acad Sci USA*, volume 68, pages 2236–2240, 1971.
- [25] P.M. Joseph. Sampling errors in projection reconstruction MRI. *Magnetic Resonance in Medicine*, 40:460–466, 1998.
- [26] K.T. Block and J. Frahm. Spiral imaging: A critical appraisal. *Journal of Magnetic Resonance Imaging*, 21:657–668, 2005.
- [27] J.G. Pipe. Motion correction with PROPELLER MRI: Application to head motion and free-breathing cardiac imaging. *Magnetic Resonance in Medicine*, 42:963–969, 1999.
- [28] C.A. Mistretta, O. Wieben, J. Velikina, W. Block, J. Perry, Y. Wu, K. Johnson, and Y. Wu. Highly constrained backprojection for time-resolved MRI. *Magnetic Resonance in Medicine*, 55:30–40, 2006.
- [29] C.B. Owen and F. Makedon. High quality alias free image rotation. In *Conference Record of the Thirtieth Asilomar Conference on Signals, Systems and Computers*, volume 1, pages 115–119, 1996.

- [30] M.L. Lauzon and B.K. Rutt. Effects of polar sampling in k-space. *Magnetic Resonance in Medicine*, 36:940–949, 1996.
- [31] M.L. Lauzon and B.K. Rutt. Polar sampling in k-space: Reconstruction effects. *Magnetic Resonance in Medicine*, 40:940–949, 1998.
- [32] D.C. Peters, F.H. Epstein, and E.R. McVeigh. Myocardial wall tagging with undersampled projection reconstruction. *Magnetic Resonance in Medicine*, 45:562–567, 2001.
- [33] R.B. Thompson and E.R. McVeigh. Flow-gated phase-contrast MRI using radial acquisitions. *Magnetic Resonance in Medicine*, 52:598–604, 2004.
- [34] D.C. Peters, R.J. Lederman, A.J. Dick, V.K. Raman, M.A. Guttman, J.A. Derbyshire, and E.R. McVeigh. Undersampled projection reconstruction for active catheter imaging with adaptable temporal resolution and catheter-only views. *Magnetic Resonance in Medicine*, 49:216–222, 2003.
- [35] A. Srivastava, A. Lee, E. Simoncelli, and S. Zhu. On advances in statistical modeling of natural images. *Journal of Mathematical Imaging and Vision*, 18:17–33, 2003.
- [36] T. Schäffter, V. Rasche, and I.C. Carlsen. Motion compensated projection reconstruction. *Magnetic Resonance in Medicine*, 41:954–963, 1999.
- [37] A.C. Larson, P. Kellman, A. Arai, G.A. Hirsch, E. McVeigh, D. Li, and O.P. Simonetti. Preliminary investigation of respiratory self-gating for free-breathing segmented cine MRI. *Magnetic Resonance in Medicine*, 53:159–168, 2005.
- [38] V. Rasche, R.W. De Boer, D. Holz, and R. Proksa. Continuous radial data acquisition for dynamic MRI. *Magnetic Resonance in Medicine*, 34:754–761, 1994.
- [39] V. Rasche, D. Holz, and R. Proksa. MR fluoroscopy using projection reconstruction multi-gradient-echo (prMGE) MRI. *Magnetic Resonance in Medicine*, 42:324–334, 1999.
- [40] A.C. Larson, R.D. White, G. Laub, E.R. McVeigh, D. Li, and O.P. Simonetti. Self-gated cardiac cine MRI. *Magnetic Resonance in Medicine*, 51:93–102, 2004.
- [41] D.C. Peters, M.A. Guttman, A.J. Dick, V.K. Raman, R.J. Lederman, and E.R. McVeigh. Reduced field of view and undersampled PR combined for interventional imaging of a fully dynamic field of view. *Magnetic Resonance in Medicine*, 51:761–767, 2004.
- [42] A. Shankaranarayanan, M. Wendt, A.J. Aschoff, J.S. Lewin, and J.L. Duerk. Radial keyhole sequences for low field projection reconstruction interventional MRI. *Journal of Magnetic Resonance Imaging*, 13:142–151, 2001.
- [43] M. Shankaranarayanan, A. amd Wendt, J.S. Lewin, and J.L. Duerk. Two-step navigatorless correction algorithm for radial k-space MRI acquisitions. *Magnetic Resonance in Medicine*, 45:277–288, 2001.

- [44] J.G. Pipe, V.G. Farthing, and K.P. Forbes. Multishot diffusion-weighted FSE using PROPELLER MRI. *Magnetic Resonance in Medicine*, 47:42–52, 2002.
- [45] S. Mori and P.C.M. van Zijl. A motion correction scheme by twin-echo navigation for diffusion-weighted magnetic resonance imaging with multiple RF echo acquisition. *Magnetic Resonance in Medicine*, 40:511–516, 1998.
- [46] D.C. Noll. Multishot rosette trajectories for spectrally selective MR imaging. *IEEE Transactions on Medical Imaging*, 16:372–377, 1997.
- [47] H. Feng, H. Gu, D. Silbersweig, E. Stern, and Y. Yang. Single-short MR imaging using trapezoidal-gradient-based lissajous trajectories. *IEEE Transactions on Medical Imaging*, 22:925–932, 2003.
- [48] A. Shankaranarayanan, O.P. Simonetti, G. Laub, J.S. Lewin, and J.L. Duerk. Segmented k-space and real-time cardiac cine MR imaging with radial trajectories. *Radiology*, 221:827–836, 2001.
- [49] M. Katoh, E. Spuentrup, A. Buecker, W.J. Manning, R.W. Guenther, and R.M. Botnar. MR coronary vessel wall imaging: Comparison between radial and spiral k-space sampling. *Journal of Magnetic Resonance Imaging*, 23:757–762, 2006.
- [50] T.P. Trouard, Y. Sabharwal, M.I. Altbach, and A.F. Gmitro. Analysis and comparison of motion-correction techniques in diffusion-weighted imaging. *Journal of Magnetic Resonance Imaging*, 6:925–935, 1996.
- [51] D.G. Kruger, G.S. Slavin, R. Muthupillai, R.C. Grimm, and S.J. Riederer. An orthogonal correlation algorithm for ghost reduction in MRI. *Magnetic Resonance in Medicine*, 38:678–686, 1997.
- [52] L. Arena, H.T. Morehouse, and J. Safir. MR imaging artifacts that simulate disease: how to recognize and eliminate them. *Radiographics*, 15:1373–1394, 1995.
- [53] S.B. Reeder, E. Atalar, A.Z. Faranesh, and E.R. McVeigh. Referenceless interleaved echo-planar imaging. *Magnetic Resonance in Medicine*, 41:87–94, 1999.
- [54] M.I. Altbach, E.K. Outwater, T.P. Trouard, E.A. Krupinski, R.J. Theilmann, A.T. Stopeck, M. Kono, and A.F. Gmitro. Radial fast spin-echo method for T2-weighted imaging and T2 mapping of the liver. *Journal of Magnetic Resonance Imaging*, 16:179–189, 2002.
- [55] C.B. Ahn and Z. H. Cho. Analysis of eddy currents in nuclear magnetic resonance imaging. *Magnetic Resonance in Medicine*, 17:149–163, 1991.
- [56] C. Boesch, R. Gruetter, and E. Martin. Optimization of corrections and quantitative characterization of magnet/gradient systems. *Magnetic Resonance in Medicine*, 20:268–284, 1991.
- [57] M.T. Alley, G.H. Glover, and N.J. Pelc. Gradient characterization using a Fourier-transform technique. *Magnetic Resonance in Medicine*, 39:581–587, 1998.
- [58] D.C. Peters, J.A. Derbyshire, and E.R. McVeigh. Centering the projection reconstruction trajectory: Reducing gradient delay errors. *Magnetic Resonance in Medicine*, 50:1–6, 2003.

- [59] B.M. Dale and Duerk J.L. The use of measured k-space trajectory for reconstruction of radial MRI data. In *Proc Intl Soc Mag Reson Med*, volume 10, page 2334, 2002.
- [60] J.H. Duyn, Y. Yang, J.A. Frank, and J.W. van der Veen. Simple correction method for k-space trajectory deviations in MRI. *Journal of Magnetic Resonance*, 132:150–153, 1998.
- [61] Y. Jung, Y. Jashnani, R. Kijowski, and W.F. Block. Consistent non-cartesian off-axis MRI quality: Calibrating and removing multiple sources of demodulation phase errors. *Magnetic Resonance in Medicine*, 57:206–212, 2007.
- [62] P. Speier and F. Trautwein. Robust radial imaging with predetermined isotropic gradient delay correction. In *Proc Intl Soc Mag Reson Med*, volume 14, page 2379, 2006.
- [63] L. Li. Magnetic susceptibility quantification for arbitrarily shaped objects in inhomogeneous fields. *Magnetic Resonance in Medicine*, 46:907–916, 2001.
- [64] J.R. Reichenbach, R. Venkatesan, D.A. Yablonskiy, M.R. Thompson, S. Lai, and E.M. Haacke. Theory and application of static field inhomogeneity effects in gradient-echo imaging. *Journal of Magnetic Resonance Imaging*, 7:266–279, 1997.
- [65] R.R. Edelman, P. Wielopolski, and F. Schmitt. Echo-planar MR imaging. *Radiology*, 192:600–612, 1994.
- [66] H. Zeng and R.T. Constable. Image distortion correction in EPI: Comparison of field mapping with point spread function mapping. *Magnetic Resonance in Medicine*, 48:137–146, 2002.
- [67] C. Windischberger, S. Robinson, A. Rauscher, M. Barth, and E. Moser. Robust field map generation using a triple-echo acquisition. *Journal of Magnetic Resonance Imaging*, 20:730–734, 2004.
- [68] V. Roopchansingh, R.W. Cox, A. Jesmanowicz, B.D. Ward, and J.S. Hyde. Single-shot magnetic field mapping embedded in echo-planar time-course imaging. *Magnetic Resonance in Medicine*, 50:839–843, 2003.
- [69] K.S. Nayak and D.G. Nishimura. Automatic field map generation and off-resonance correction for projection reconstruction imaging. *Magnetic Resonance in Medicine*, 43:151–154, 2000.
- [70] D.C. Noll, C.H. Meyer, J.M. Pauly, D.G. Nishimura, and A. Macovski. A homogeneity correction method for magnetic resonance imaging with time-varying gradients. *IEEE Transactions on Medical Imaging*, 10(4):629–637, 1991.
- [71] D.C. Noll, J.M. Pauly, C.H. Meyer, D.G. Nishimura, and A. Macovski. Deblurring for non-2D Fourier transform magnetic resonance imaging. *Magnetic Resonance in Medicine*, 25:319–333, 1992.
- [72] D.C. Noll. *Reconstruction techniques for magnetic resonance imaging*. PhD thesis, Stanford University, 1991.

- [73] G.H. Glover and D.C. Noll. Consistent projection reconstruction (CPR) techniques for MRI. *Magnetic Resonance in Medicine*, 29:345–351, 1993.
- [74] E.B. Welch, P.J. Rossman, J.P. Felmlee, and A. Manduca. Self-navigated motion correction using moments of spatial projections in radial MRI. *Magnetic Resonance in Medicine*, 52:337–345, 2004.
- [75] N. Gai and L. Axel. Correction of motion artifacts in linogram and projection reconstruction MRI using geometry and consistency constraints. *Med Phys*, 23:251–262, 1996.
- [76] V. Rasche, D. Holz, and W. Schepper. Radial turbo spin echo imaging. *Magnetic Resonance in Medicine*, 32:629–638, 1994.
- [77] R.J. Theilmann, A.F. Gmitro, M.I. Altbach, and Trouard T.P. View-ordering in radial fast spin-echo imaging. *Magnetic Resonance in Medicine*, 51:768–774, 1999.
- [78] M.H.J. Seifert, P.M. Jakob, V. Jellus, A. Haase, and C. Hillenbrand. High-resolution diffusion imaging using a radial turbo-spin-echo sequence: Implementation, eddy current compensation, and self-navigation. *Journal of Magnetic Resonance*, 144:243–254, 2000.
- [79] A.F. Gmitro, M. Kono, R.J. Theilmann, M.I. Altbach, Z. Li, and T.P. Trouard. Radial GRASE: Implementation and applications. *Magnetic Resonance in Medicine*, 53:1363–1371, 2005.
- [80] A.C. Silva, E.L. Barbier, I.J. Lowe, and A.P. Koretsky. Radial echo-planar imaging. *Journal of Magnetic Resonance*, 135:242–247, 1998.
- [81] G.E. Sarty. Single trajectory radial (STAR) imaging. *Magnetic Resonance in Medicine*, 51:445–451, 2004.
- [82] D.C. Peters, F.R. Korosec, T.M. Grist, W.F. Block, J.E. Holden, K.K. Vigen, and C.A. Mistretta. Undersampled projection reconstruction applied to MR angiography. *Magnetic Resonance in Medicine*, 43:91–101, 2000.
- [83] Y. Crmillieux, A. Briguet, and A. Deguin. Projection-reconstruction methods: Fast imaging sequences and data processing. *Magnetic Resonance in Medicine*, 32:23–32, 1994.
- [84] J.M. Wild, M.N.J. Paley, L. Kasuboski, A. Swift, S. FICHELE, N. Woodhouse, P.D. Griffiths, and E.J.R. van Beek. Dynamic radial projection MRI of inhaled hyperpolarized ^3He gas. *Magnetic Resonance in Medicine*, 49:991–997, 2003.
- [85] D.C. Noll, D.G. Nishimura, and A. Macovski. Homodyne detection in magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, 10:154–163, 1991.
- [86] W.F. Block, D.C. Peters, and K.K. Vigen. Homodyne reconstruction for projection reconstruction trajectories. In *Proc Intl Soc Mag Reson Med*, volume 7, page 659, 1999.

- [87] D.J. Tyler, M.D. Robson, R.M. Henkelman, I.R. Young, and G.M. Bydder. Magnetic resonance imaging with ultrashort TE (UTE) PULSE sequences: Technical considerations. *Journal of Magnetic Resonance Imaging*, 25:279–289, 2007.
- [88] P.E.Z. Larson, S.M. Conolly, J.M. Pauly, and D.G. Nishimura. Using adiabatic inversion pulses for long-T2 suppression in ultrashort echo time (UTE) imaging. *Magnetic Resonance in Medicine*, 58:952–961, 2007.
- [89] P.E.Z. Larson, P.T. Gurney, K. Nayak, G.E. Gold, J.M. Pauly, and D.G. Nishimura. Designing long-T2 suppression pulses for ultrashort echo time imaging. *Magnetic Resonance in Medicine*, 56:94–103, 2006.
- [90] Gatehouse P.D. and G.M. Bydder. Magnetic resonance imaging of short T2 components in tissue. *Clinical Radiology*, 58:1–19, 2003.
- [91] J. Rahmer, P. Boernert, J. Groen, and C. Bos. Three-dimensional radial ultrashort echo-time imaging with T2 adapted sampling. *Magnetic Resonance in Medicine*, 55:1075–1082, 2006.
- [92] A.C. Larson, O.P. Simonetti, and D. Li. Coronary MRA with 3D undersampled projection reconstruction TrueFISP. *Magnetic Resonance in Medicine*, 48:594–601, 2002.
- [93] J. Du, T.J. Carroll, H.J. Wagner, K. Vigen, S.B. Fain, W.F. Block, F.R. Korosec, T.M. Grist, and C.A. Mistretta. Time-resolved, undersampled projection reconstruction imaging for high-resolution CE-MRA of the distal runoff vessels. *Magnetic Resonance in Medicine*, 48:516–522, 2002.
- [94] K.K. Vigen, D.C. Peters, T.M. Grist, W.F. Block, and C.A. Mistretta. Undersampled projection-reconstruction imaging for time-resolved contrast-enhanced imaging. *Magnetic Resonance in Medicine*, 43:170–176, 2000.
- [95] A.V. Barger, W.F. Block, Y. Toropov, T.M. Grist, and C.A. Mistretta. Time-resolved contrast-enhanced imaging with isotropic resolution and broad coverage using an undersampled 3D projection trajectory. *Magnetic Resonance in Medicine*, 48:297–305, 2002.
- [96] P.E.Z. Larson, P.T. Gurney, and D.G. Nishimura. Anisotropic field-of-views in radial imaging. *IEEE Transactions on Medical Imaging*, 27:47–57, 2008.
- [97] A. Lu, E. Brodsky, T.M. Grist, and W.F. Block. Rapid fat-suppressed isotropic steady-state free precession imaging using true 3D multiple-half-echo projection reconstruction. *Magnetic Resonance in Medicine*, 53:692–699, 2005.
- [98] C. Stehning, P. Börner, K. Nehrke, H. Eggers, and O. Dössel. Fast isotropic volumetric coronary MR angiography using free-breathing 3D radial balanced FFE acquisition. *Magnetic Resonance in Medicine*, 52:197–203, 2004.
- [99] P.T. Gurney, B.A. Hargreaves, and D.G. Nishimura. Design and analysis of a practical 3D cones trajectory. *Magnetic Resonance in Medicine*, 55:575–582, 2006.

-
- [100] A.B. Cheryauka, J.N. Lee, A.A. Samsonov, M. Defrise, and G.T. Gullberg. MRI diffusion tensor reconstruction with PROPELLER data acquisition. *Magnetic Resonance Imaging*, 22:139–148, 2004.
- [101] T. Chuang, T. Huang, F. Lin, F. Wang, C. Juan, H. Chung, C. Chen, and K.K. Kwong. PROPELLER-EPI with parallel imaging using a circularly symmetric phased-array RF coil at 3.0 T: Application to high-resolution diffusion tensor imaging. *Magnetic Resonance in Medicine*, 56:1352–1358, 2006.
- [102] J.G. Pipe and N. Zwart. Turboprop: Improved PROPELLER imaging. *Magnetic Resonance in Medicine*, 55:380–385, 2006.
- [103] S. Skare, R.D. Newbould, D.B. Clayton, and R. Bammer. Propeller EPI in the other direction. *Magnetic Resonance in Medicine*, 55:1298–1307, 2006.
- [104] K.P.N. Forbes, J.G. Pipe, C.R. Bird, and J.E. Heiserman. PROPELLER MRI: Clinical testing of a novel technique for quantification and compensation of head motion. *Journal of Magnetic Resonance Imaging*, 14:215–222, 2001.
- [105] K. Arfanakis, A.A. Tamhane, J.G. Pipe, and M.A. Anastasio. k-Space undersampling in PROPELLER imaging. *Magnetic Resonance in Medicine*, 53:675–683, 2005.
- [106] K.P. Pruessmann, M. Weiger, P. Boernert, and P. Boesiger. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magnetic Resonance in Medicine*, 46:638–651, 2001.
- [107] W.W. Hager and H. Zhang. A survey of nonlinear conjugate gradient methods. *Pacific J Optimization*, 2:35–58, 2006.
- [108] W.W. Hager and H. Zhang. A new conjugate gradient method with guaranteed descent and an efficient line search. *SIAM J Optimization*, 16:170–192, 2005.
- [109] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer, 2006.
- [110] A.H. Delaney and Y. Bresler. A fast and accurate Fourier algorithm for iterative parallel-beam tomography. *IEEE Transactions on Medical Imaging*, 5:740–753, 1996.
- [111] L.I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
- [112] D. Geman and C. Yang. Nonlinear image recovery with half-quadratic regularization. *IEEE T Image Processing*, 4:932–946, 1995.
- [113] D.K. Sodickson and W.J. Manning. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 38:591–603, 1997.
- [114] M.A. Griswold, P.M. Jakob, R.M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47:1202–1210, 2002.

- [115] K.P. Pruessmann, M. Weiger, M.B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42:952–962, 1999.
- [116] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE T Information Theory*, 52:489–509, 2006.
- [117] E.J. Candès and F. Guo. New multiscale transforms, minimum total variation synthesis: applications to edge-preserving image reconstruction. *Signal Process.*, 82:1519–1543, 2002.
- [118] M.N. Do and M. Vetterli. The contourlet transform: an efficient directional multiresolution image representation. *IEEE T Image Processing*, 14:2091–2106, 2005.
- [119] D. Donoho. Compressed sensing. *IEEE T Information Theory*, 52:1289–1306, 2006.
- [120] D. Donoho. For most large underdetermined systems of linear equations, the minimal l_1 norm solution is also the sparsest solution. *Communications on Pure and Applied Mathematics*, 59:797–829, 2006.
- [121] J. Frahm, A. Haase, D. Matthaei, K.D. Merboldt, and W. Hänicke. Rapid NMR imaging using stimulated echoes. *Journal of Magnetic Resonance*, 65:130–135, 1985.
- [122] J. Frahm, K.D. Merboldt, W. Hänicke, and A. Haase. Stimulated echo imaging. *Journal of Magnetic Resonance Imaging*, 64:81–93, 1985.
- [123] U.G. Nolte, J. Finsterbusch, and J. Frahm. Rapid isotropic diffusion mapping without susceptibility artifacts. whole brain studies using diffusion-weighted single-shot STEAM MR imaging. *Magnetic Resonance in Medicine*, 44:731–736, 2000.
- [124] S. Rieseberg, K.D. Merboldt, M. Küntzel, and J. Frahm. Diffusion tensor imaging using partial Fourier STEAM MRI with projection onto convex subsets reconstruction. *Magnetic Resonance in Medicine*, 54:486–490, 2005.
- [125] A. Karaus. Weiterentwicklung der schnellen Magnetresonanz-Tomografie mit stimulierten Echos. Diploma thesis, Universität Tübingen, 2007.
- [126] J. Finsterbusch, J. Frahm, and M.A. Koch. Gradient and stimulated echo (GRASTE) imaging. *Magnetic Resonance in Medicine*, 55:455–459, 2006.
- [127] A. Karaus, K.D. Merboldt, J. Graessner, and J. Frahm. Black-blood imaging of the human heart using rapid stimulated echo acquisition mode (STEAM) MRI. *Journal of Magnetic Resonance Imaging*, 26:1666–1671, 2007.
- [128] H.K. Song and L. Dougherty. k-space weighted image contrast (KWIC) for contrast manipulation in projection reconstruction MRI. *Magnetic Resonance in Medicine*, 44:825–832, 2000.

- [129] M.I. Altbach, A. Bilgin, Z. Li, E.W. Clarkson, T.P. Trouard, and A.F. Gmitro. Processing of radial fast spin-echo data for obtaining T2 estimates from a single k-space data set. *Magnetic Resonance in Medicine*, 54:549–559, 2005.
- [130] Z.P. Liang, F.B. Boada, R.T. Constable, E.M. Haacke, P.C. Lauterbur, and M. Smith. Constrained reconstruction methods in MR imaging. *Rev Magn Reson Med*, 4:67–185, 1992.
- [131] M.R. Smith and S.T. Nichols. A comparison of models used as alternative magnetic resonance image reconstruction methods. *Magnetic Resonance Imaging*, 8:173–183, 1990.
- [132] Gelb A. Archibald R. A method to reduce the Gibbs ringing artifact in MRI scans while keeping tissue boundary integrity. *IEEE Transactions on Medical Imaging*, 21:305–319, 2002.
- [133] M.F. Callaghan, D.J. Larkman, and J.V. Hajnal. Pade methods for reconstruction and feature extraction in magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54:1490–1502, 2005.
- [134] R.T. Constable and R.M. Henkelman. Data extrapolation for truncation artifact removal. *Magnetic Resonance in Medicine*, 17:108–118, 1991.
- [135] S. Amartur and E.M. Haacke. Modified iterative model based on data extrapolation method to reduce Gibbs ringing. *Journal of Magnetic Resonance Imaging*, 1:307–317, 1991.
- [136] S. Amartur, Z.P. Liang, F. Boada, and E.M. Haacke. Phase-constrained data extrapolation method for reduction of truncation artifacts. *Journal of Magnetic Resonance Imaging*, 1:721–724, 1991.
- [137] J. Tsao. Extension of finite-support extrapolation using the generalized series model for MR spectroscopic imaging. *IEEE Transactions on Medical Imaging*, 20:1178–1183, 2001.
- [138] M. Lustig, D. Donoho, and J.M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58:1182–1195, 2007.
- [139] M. Doneva, H. Eggers, J. Rahmer, P. Börnert, and A. Mertins. Highly undersampled 3D golden ratio radial imaging with iterative reconstruction. In *Proc Intl Soc Mag Reson Med*, volume 16, page 336, 2008.
- [140] A. Fischer, F. Breuer, M. Blaimer, N. Seiberlich, and P.M. Jakob. Accelerated dynamic imaging by reconstructing sparse differences using compressed sensing. In *Proc Intl Soc Mag Reson Med*, volume 16, page 341, 2008.
- [141] A. Samsonov, Y. Jung, A.L. Alexander, W.F. Block, and A.S. Field. MRI compressed sensing via sparsifying images. In *Proc Intl Soc Mag Reson Med*, volume 16, page 342, 2008.

-
- [142] M. Uecker, T. Hohage, K.T. Block, and J. Frahm. Image reconstruction by regularized nonlinear inversion - joint estimation of coil sensitivities and image content. *Magnetic Resonance in Medicine*, 2008 (in press).
- [143] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM J Sci Comput*, 20:33–61, 1999.
- [144] J. Trzasko, A. Manduca, and E. Borisch. Highly undersampled magnetic resonance image reconstruction via homotopic ell-0-minimization. *Preprint*, 2008.
- [145] E. Candès, M. Wakin, and S. Boyd. Enhancing sparsity by reweighted ell-1 minimization. *Preprint*, 2007.
- [146] A. Fischer, F. Breuer, M. Blaimer, N. Seiberlich, and P.M. Jakob. Introduction of a nonconvex compressed sensing algorithm for MR imaging. In *Proc Intl Soc Mag Reson Med*, volume 16, page 1487, 2008.

Abbreviations

CHESS	Chemical shift selective
CS	Compressed sensing
CSF	Cerebrospinal fluid
CT	Computed tomography
DCF	Density compensation function
DFT	Discrete Fourier transformation
EPI	Echo planar imaging
ETL	Echo train length
FOV	Field of view
FBP	Filtered backprojection
FSE	Fast spin echo
FT	Fourier transformation / transform
FFT	Fast Fourier transformation
KWIC	K-Space weighted image contrast
MRI	Magnetic resonance imaging
NMR	Nuclear magnetic resonance
PD	Proton density
PSF	Point spread function
RF	Radio frequency
R2	Spin-spin relaxivity

SE	Spin echo
SNR	Signal-to-noise ratio
STEAM	Stimulated echo acquisition mode
T_1	Spin-lattice relaxation time
T_2	Spin-spin relaxation time
T_2^*	Effective spin-spin relaxation time
TE	Echo time
TR	Repetition time
TV	Total variation

Curriculum Vitae

Personal Data

Name: Kai Tobias Block
Date of Birth: July 25, 1979
Sex: Male
Nationality: German
Degree: Dipl.Phys.

Educational Qualifications

1998 Abitur, Halepaghen-Schule Buxtehude
1999 - 2004 Studies of Physics (Diploma), University of Göttingen
Since 2005 PhD School of Mathematical Sciences, University of Göttingen

Professional Experience

2002 - 2004 Student Assistant at Biomedizinische NMR Forschungs GmbH, Göttingen.
2004 - 2005 Scientist at MeVis Research GmbH - Center for Medical Image Computing, Bremen.
Since 2005 PhD Student at Biomedizinische NMR Forschungs GmbH, Göttingen.

Awards / Stipends

2007 ISMRM Young Investigator Award (I.I. Rabi Award)
2007 ISMRM Educational Stipend
2008 ISMRM Educational Stipend

List of Publications

Journal Publications

Kai Tobias Block and Jens Frahm, Spiral imaging: A critical appraisal, *Journal of Magnetic Resonance Imaging*, 21(6):657-668, 2005.

Kai Tobias Block, Martin Uecker and Jens Frahm, Undersampled Radial MRI with Multiple Coils. Iterative Image Reconstruction Using a Total Variation Constraint, *Magnetic Resonance in Medicine*, 57(6):1086-1098, 2007.

Kai Tobias Block and Jens Frahm, Radial Single-Shot STEAM MRI, *Magnetic Resonance in Medicine*, 59(4):686-691, 2008.

Martin Uecker, Thorsten Hohage, Kai Tobias Block and Jens Frahm, Image Reconstruction by Regularized Nonlinear Inversion - Joint Estimation of Coil Sensitivities and Image Content, *Magnetic Resonance in Medicine*, in press, 2008.

Kai Tobias Block, Martin Uecker and Jens Frahm, Suppression of MRI Truncation Artifacts Using Total Variation Constrained Data Extrapolation, *International Journal of Biomedical Imaging*, in press, 2008.

Kai Tobias Block, Martin Uecker and Jens Frahm, Iterative Reconstruction for Radial Fast Spin-Echo MRI, *IEEE Transactions on Medical Imaging*, submitted, 2008.

Conference Contributions

Kai Tobias Block, Martin Uecker and Jens Frahm, Iterative Image Reconstruction of Undersampled Radial MRI Data from Multiple Coils, *ISMRM Workshop on Non-Cartesian MRI Sedona 2007*, Workshop Syllabus:Poster 41.

Kai Tobias Block, Martin Uecker and Jens Frahm, Undersampled Radial MRI with Multiple Coils. Iterative Image Reconstruction Using a Total Variation Constraint, *Joint Annual Meeting ISMRM-ESMRMB Berlin 2007*, *In Proc. Intl. Soc. Mag. Reson. Med.* 15 (2007):145.

Kai Tobias Block, Alexander Karaus and Jens Frahm, Radial Single-Shot STEAM MRI, Joint Annual Meeting ISMRM-ESMRMB Berlin 2007, In Proc. Intl. Soc. Mag. Reson. Med. 15 (2007):1663.

Martin Uecker, Kai Tobias Block and Jens Frahm, Parallel Imaging as a Non-linear Inversion Problem - Improved Reconstructions, Joint Annual Meeting ISMRM-ESMRMB Berlin 2007, In Proc. Intl. Soc. Mag. Reson. Med. 15 (2007):1740.

Kai Tobias Block, Martin Uecker and Jens Frahm, Iterative Reconstruction for R2 Mapping Based on Radial Fast Spin-Echo MRI, Annual Meeting ISMRM Toronto 2008, In Proc. Intl. Soc. Mag. Reson. Med. 16 (2008):1432.

Kai Tobias Block, Martin Uecker and Jens Frahm, Suppression of MRI Truncation Artifacts Using Total Variation Constrained Data Extrapolation, Annual Meeting ISMRM Toronto 2008, In Proc. Intl. Soc. Mag. Reson. Med. 16 (2008):1477.

Martin Uecker, Kai Tobias Block and Jens Frahm, Non-linear Inversion with L1-Wavelet Regularization - Application to Autocalibrated Parallel Imaging, Annual Meeting ISMRM Toronto 2008, In Proc. Intl. Soc. Mag. Reson. Med. 16 (2008):1479.

Martin Uecker, Kai Tobias Block and Jens Frahm, An Automatic Stopping Criterion for Iterative MRI Reconstructions, Annual Meeting ISMRM Toronto 2008, In Proc. Intl. Soc. Mag. Reson. Med. 16 (2008):1499.

Other Publications

Kai Tobias Block, Spiralförmige Abtastung des k-Raumes bei der Magnetresonanztomographie, Diploma thesis, Drittes Physikalisches Institut der Georg-August-Universität zu Göttingen, 2004.

Kai Tobias Block, Image Reconstruction from Incomplete Radial Data in Magnetic Resonance Imaging, MPIbpc News: Doktorandenseminar, Max-Planck-Institut für biophysikalische Chemie, Issue 12, 2007.

Acknowledgments

First of all, I would like to express my sincere gratitude to Prof. Dr. Jens Frahm, head of the Biomedizinische NMR Forschungs GmbH, for the continuous support of my work, for the extensive assistance in the preparation of manuscripts and presentations, and for the establishment of an excellent research facility.

I am very grateful to Prof. Dr. Robert Schaback from the Institut für Numerische und Angewandte Mathematik of the Georg-August-Universität Göttingen for the supervision of my thesis and for the constant interest in the work.

I am deeply indebted to my colleague Martin Uecker for the uncountable fruitful discussions, for the invaluable suggestions, and for the deep understanding of the topic.

Further, I would like to thank all staff members of the Biomedizinische NMR Forschungs GmbH for creating a relaxed working atmosphere, especially during the numerous Foosball sessions.

Needless to say, my warmest thanks go to my family for the endless support and for the outstanding help they provided me during the past years, which substantially contributed to the success of this venture. Finally, I wish to thank all my friends for encouraging me and for helping me to keep my life in context.