From Department of Clinical Neuroscience
Karolinska Institutet, Stockholm, Sweden

Acquisition Strategies for Fat/Water Separated MRI

Henric Rydén

Karolinska Institutet

Cover illustration: Single-species approximated NSA in pseudo color, ranging from zero (blue) to max (red), evaluated for a range of dephasing time pairs. The CSE phase is used as a surrogate for time to avoid a field strength dependency. The overlaid trajectories show typical dephasing time pairs of the developed techniques: Magenta - Paper I, Yellow - Paper II–III, White - Paper IV. Stars mark the dephasing time pairs at k-space center acquisition.

Published by Karolinska Institutet.
Printed by Universitetsservice US-AB.
ISBN 978-91-7831-883-4.

Acquisition Strategies for Fat/Water Separated MRI

THESIS FOR DOCTORAL DEGREE (Ph.D)

By

Henric Rydén

Dissertation presented at Karolinska Institutet to be publicly examined in *Rehabsalen*, Karolinska Universitetssjukhuset, Eugeniavägen 39, Solna, Norrbacka S2:01, on Friday 2020-09-18 at 09:00. The examination will be conducted in English.

Principal Supervisor:
Assoc. Prof. Stefan Skare
Karolinska Institutet
Department of Clinical Neuroscience

Co-supervisors:
Dr. Johan Berglund
Karolinska Institutet
Department of Clinical Neuroscience

Dr. Enrico Avventi
Karolinska Institutet
Department of Clinical Neuroscience

Opponent:
Dr. Holger Eggers
Philips Research Laboratory

Examination Board:
Prof. Lars Erik Olsson
Lund University
Department of Medical Radiation Physics

Assoc. Prof. Anders Garpebring
Umeå University
Department of Radiation Sciences

Dr. Andreas Sigfridsson
Karolinska Institutet
Department of Molecular Medicine
and Surgery

Abstract

This thesis focuses on new ways to more efficiently acquire the signal for fat/water separated MRI, also known as Dixon methods.

In paper I, the concept of dual bandwidths was introduced to improve SNR efficiency by removing dead times in a spin echo PROPELLER sequence. By correcting for the displacement of fat, we were able to improve the motion correction. This required additional considerations during reconstruction in order to avoid noise amplification, which was solved with a noise-whitening Tikhonov regularization.

Paper II explores the combination of fat/water separation in k-space with partially acquired data, i.e. partial Fourier sampling. With reduced sampling coverage comes the ability of increased spatial resolution, which is often limited in fat/water imaging, particularly in gradient echo sequences. A modified POCS routine was also developed with real-valued estimates, exploiting Hermitian symmetry to improve the inverse problem conditioning in the fully sampled region.

A single-TR dual-bandwidth RARE (fast/turbo spin echo) sequence without dead times was developed in Paper III, which uses partial Fourier sampling with late and early echoes to improve the chemical shift encoding. The proposed sequence can acquire images with 0.5 mm in-plane resolution without dead times, with image quality exceeding current state-of-the-art techniques. An automated selection of gradient waveforms based on Cramér-Rao bounds was implemented on the scanner.

In Paper IV, the dual-bandwidth concept was generalized to continuous bandwidths. Instead of the conventional shift of a trapezoidal readout gradient, we describe a new method of encoding chemical shift by using asymmetric readout waveforms. Asymmetric readouts were implemented in a RARE sequence to completely remove dead times from multi-TR acquisitions, with typical scan time reductions of 25 %.

The developed methods enable fat/water imaging with reduced scan times and increased spatial resolution, which has previously limited their use.

List of Papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

I. **Rydén H**, Berglund J, Norbeck O, Avventi E, Skare S. T1 weighted fat/water separated PROPELLER acquired with dual bandwidths. *Magn. Reson. Med.* 2018;80:2501–2513.

II. Berglund J, **Rydén H**, Avventi E, Norbeck O, Sprenger T, Skare S. Fat/water separation in k-space with real-valued estimates and its combination with POCS. *Magn. Reson. Med.* 2020;83:653–661.

III. **Rydén H**, Berglund J, Norbeck O, Avventi E, Sprenger T, van Niekerk A, Skare S. RARE two-point Dixon with dual bandwidths. *Magn. Reson. Med.* 2020 doi: 10.1002/mrm.28293.

IV. **Rydén H**, Norbeck O, Avventi E, Skorpil, M, van Niekerk A, Skare S, Berglund J. Chemical Shift Encoding using Asymmetric Readout Waveforms. *In revision for Magn. Reson. Med.*

Contents

1 Introduction

The ability to safely image soft tissues with multiple contrasts makes MRI an essential diagnostic tool. The downside of MRI is that the signal is weak, and signal acquisition must typically be repeated multiple times in order to yield images of diagnostic quality. This results in a continuous struggle between acquisition time and image quality, with a massive amount of research on how to acquire the signal as fast and efficiently as possible.

Fat plays a special role in MRI in that it often appears with high intensity and can cause signal loss or displacement due to chemical shift effects. This can obscure pathology and many ways of handling the fat signal have been proposed. One such method performs fat/water separation after acquiring the signal. This way, the signal from fat is dealt with mathematically instead of spoiled before acquisition, and the diagnostic potential of the fat signal itself is preserved. A downside of fat/water imaging is that it puts additional requirements on the signal acquisition, often resulting in prolonged acquisition time or reduced spatial resolution. This thesis proposes and investigates new acquisition methods aimed to overcome the limitations of fat/water separated MRI.

As there are many well-written introductions to MR physics, reference [1] in particular, no such introduction is given here. The intended reader is expected to be familiar with MR physics and pulse sequences. Chapter 2 summarizes existing techniques and phenomena related to the projects. Chapter 3 describes the project aims, why we sought to investigate these new acquisition methods in particular. Chapter 4 goes into detail of the research work and shows some results and methods that did not make it into the constituent papers. The appendix covers how the problem of coil combination was handled. Finally, chapter 5 discusses the clinical use of fat/water imaging, novel methods, and future research ideas.

2 Background and Theory

The diagnostic potential of MRI has been investigated widely since its invention, with numerous applications to investigate both morphological and functional properties of the human body. Since any nuclei with a *net* nuclear spin experience the nuclear magnetic resonance phenomenon, a multitude of elements can be imaged with an MRI scanner. The overwhelming majority of clinical MRI examinations are tuned to the nuclear magnetic resonance frequency of ^{1}H, i.e. protons, due to its high abundance in the human body compared to other isotopes with a net nuclear spin.

2.1 Chemical Shift

The resonance frequency ω is dependent only on the gyromagnetic ratio γ of ^{1}H, and the *effective* magnetic field strength B:

$$\omega = \gamma B \tag{2.1}$$

The molecular structure of triglycerides shields the magnetic field B, resulting in a lower resonance frequency of ^{1}H. This effect is known as *chemical shift* and is the basis for MR spectroscopy, where the relative concentration of several resonance frequencies are mapped, revealing information about the chemical structure.

Due to the fast decay or low concentration of other molecules, triglycerides and water make up the majority of the MR signal when imaging human subjects. While phospholipids also contain fatty acids, they are typically found in tightly constrained molecular structures, such as in cell membranes or myelin, resulting in a very rapid T2 decay. Consequently, phospholipids do not contribute to the MR signal unless sequences with

very short echo times are used [2]. It is therefore common to refer to the signal from ^{1}H bound to triglycerides as the fat signal. Since most imaging experiments are tuned to the resonance frequency of ^{1}H bound to water, there are no chemical shift related effects on the water signal. Therefore, the chemical shift displacement artifact is sometimes referred to as fat displacement or fat shift.

2.1.1 Displacement Artifact

In order to spatially encode the MR signal, linearly varying magnetic fields, gradients G, are applied in the spatial dimensions. The basis for spatial encoding is to vary the magnetic field, and thereby the resonance frequency, over the field of view in order to derive local spin densities, forming an image. The accrued phase from the gradient fields encode spatial harmonics. Once all spatial frequencies are encoded for a given resolution and field of view, the inverse Fourier transform of all sampled frequencies yields an image of the object. However, the gradient field is superimposed with the chemical shift effect, resulting in the net phase θ of a voxel at position x containing water (W) or fat (F) of

$$\begin{aligned} \theta_W &= 2\pi\gamma(G_{\mathrm{x}}x)t \\ \theta_F &= 2\pi\gamma(G_{\mathrm{x}}x + \sigma B_0)t \end{aligned} \tag{2.2}$$

In Equation 2.2, it is assumed that the receiver is tuned to the frequency γB_0, i.e. the water is *on resonance*, and that there is no initial phase after excitation. σ is the shielding constant, i.e. the fractional change in the static field B_0 from chemical shift effects, approximately 3.5 ppm for the main methylene peak of triglycerides [3]. t is the *dephasing time*, i.e. the time from excitation in gradient echo imaging and the time to/from the spin echo in spin echo imaging. The dephasing time can consequently be negative in spin echo imaging if sampling occurs prior to the spin echo.

In Cartesian spin warp imaging, an entire line in the frequency encoding direction is acquired while the readout gradient is applied. As data is sampled during gradient playout, t will differ between samples in the frequency encoding direction. When acquiring the next line, the readout timing is kept identical, and all samples have a constant dephasing time along the phase encoding direction.

For fat, this results in an additional linear phase ramp in the frequency encoding direction, and a constant phase offset in the phase encoding directions. The constant phase is modeled by a scalar multiplication in the sampled domain, corresponding to a scalar multiplication also in the spatial domain. The Fourier transform of a phase ramp, however, is a shifted Dirac delta function where the shift is proportional to the slope, resulting in a spatial translation of the fat signal in the frequency encoding direction. Since the slope is proportional to the difference in dephasing times between two neighbouring samples, a faster sampling, i.e. higher receiver bandwidth, reduces the shift.

A schematic representation of the resulting *chemical shift displacement artifact* is shown in Figure 2.1a-b, while Figure 2.1c shows a T2 weighted sagittal slice of the spine with fat displacement in the superior direction. This is most evident near the interfaces between the bone marrow and intervertebral discs, with signal pile-up in the overlapping region, and signal void on the opposing side. Notice that there is no fat shift in the phase encoding direction.

2.1.2 Phase Cancellation Artifact

Considering a voxel containing both water (W) and fat (F), the signal at dephasing time t is

$$S(t) = W + Fe^{2\pi i\gamma\sigma B_0 t} = W + Fe^{i\omega_f t}, \tag{2.3}$$

where ω_f is the chemical shift frequency of fat in rad/s. At 1.5T, this is approximately 220 Hz, and 440 Hz at 3T. The signal equation 2.3 reveals that fat is periodically oscillating *in and out of phase* with the water signal. At dephasing times $t_{\text{IP}} = \left\{\frac{2\pi n}{\omega_f}, n \in \mathbb{Z}\right\}$, the signal equation reduces to $S_{\text{IP}} = W + F$. An interesting phenomenon occurs at dephasing times when fat and water are opposed phase, i.e. $t_{\text{OP}} = \left\{\frac{2\pi n+\pi}{\omega_f}, n \in \mathbb{Z}\right\}$, namely that the signal equation reduces to the difference between fat and water: $S_{\text{OP}} = W - F$.

In spin echo imaging where the readout center coincides with the spin echo, the phase between water and fat is zero, resulting in an in-phase image. Shifting the readout gradient from the spin echo, or acquiring a gradient echo image with opposed phase

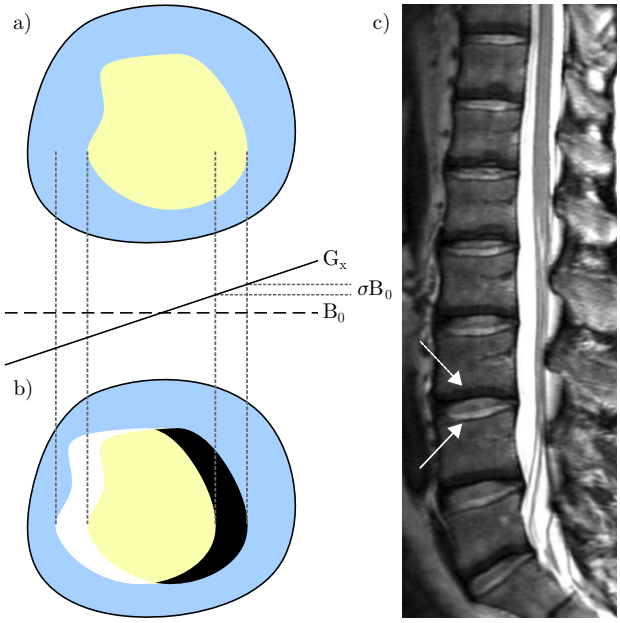

Figure 2.1: Chemical shift displacement. a) Object containing fat (yellow) and water (blue). b) The fat is shifted to the left as the chemical shift partially counteracts the gradient G_x. c) A sagittal slice of the lumbar spine, where the bone marrow is shifted in the superior direction, causing signal void (dashed arrow) and pile-up (solid arrow). The frequency encoding direction is superior-inferior.

echo times, results in opposed phase images [4]. Voxels containing a mixture of water and fat will interfere destructively, resulting in a signal loss, as seen in Figure 2.2.

2.2 Fat Suppression Techniques

The two artifacts described above are both sourced from chemical shift. In addition, the hyperintense signal from fat in most T1 and T2 weighted images can obscure pathology and reduce the dynamic range of the image. Numerous techniques have been proposed to overcome this problem by *suppressing* the fat signal.

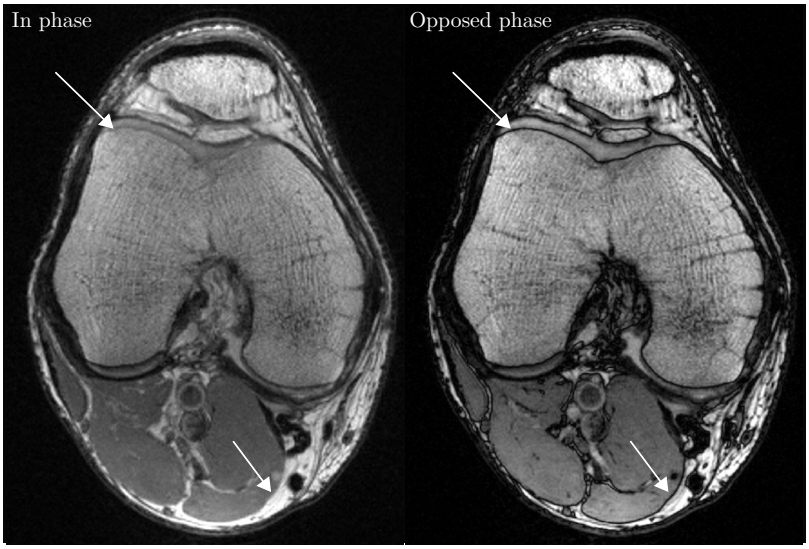

Figure 2.2: In and opposed phase images showing the phase cancellation artifact in voxels containing a mixture of water and fat.

2.2.1 CHESS

The CHESS technique involves a magnetization preparation to alter the longitudinal magnetization before the main sequence is played out. This preparation module consists of an RF pulse followed by gradient spoiling. Instead of a conventional RF pulse with a carrier frequency tuned to the water peak, the CHESS RF pulse is frequency offset according to the chemical shift. In combination with a narrow transmit bandwidth, only fat is excited and subsequent gradient spoiling will remove the fat signal. With CHESS RF pulses spectrally targeted on the fat peak, the technique is often called fat *saturation*, or "fat sat" in short.

The main sequence immediately follows the gradient spoiling, and at the time of excitation, an insignificant amount of longitudinal fat magnetization has recovered. Fat saturation is an attractive fat suppression alternative as the main sequence does not require any modifications to get rid of the fat signal. It is also short, only about 10/20 ms per CHESS playout (at 3/1.5T). However, there are a number of downsides associated with CHESS.

The tuning of the center frequency is solely determined by the center frequency of water. This is estimated in a short tuning stage at the start of the sequence, where some assumption must be made concerning the signal. One such assumption is that the majority of the signal originates from water. Once the nominal water resonance frequency is established, the chemical shift specifies the offset required for the CHESS pulse to saturate fat. Herein lies a weakness in that an erroneous measurement at the tuning stage will severely impact the fat saturation performance. For instance, in an anatomical region where fat is the dominant species, the *fat* peak is assumed to be water. The downfield chemical shift offset will then lead to saturation of frequencies at twice the chemical shift. This problem has the same origin as the fat/water "swap" artifact in Dixon methods: There is no certain way of telling the *source* of off-resonance, it can originate from field inhomogeneity, or chemical shift. Another downside of the CHESS technique is that it is centered around a single fat peak, and cannot account for the complex spectrum of fat with multiple resonances [5].

Contrary to modern Dixon methods, CHESS does not correct for spatial variations in the field heterogeneity other than the constant offset described above. Clinical scanners all try to achieve as homogeneous field as possible with active and passive shimming techniques. However, since the RF center frequency is spatially constant, fat saturation fails in regions where the field inhomogeneity is larger than the transmit bandwidth of the saturation pulse. This bandwidth is restricted by the chemical shift. CHESS fat suppression is therefore often insufficient when imaging large FOV or near implants [6]. An example of CHESS fat suppression with local failure is shown in Figure 2.3.

Challenging geometries such as the wrist, lungs, or nasal cavities, also have significant field inhomogeneities [7], [8]. Even more concerning is the risk of saturating the water signal, which happens when the field inhomogeneity is near the chemical shift. Since the signal is spoiled before acquisition, it is not possible to recover it during reconstruction. When imaging multiple chemical species, e.g. silicone, water, and fat, the problems above are amplified.

The scan time penalty of CHESS is related to the duration of the main sequence. For T2-weighted 2D RARE imaging, the scan time penalty is low as the main sequence is long ($\approx$ 200 ms) compared to the CHESS module ($\approx$ 10 ms). T1-weighted spin echo is a different scenario though, as the main sequence is much shorter, about 20 ms, resulting

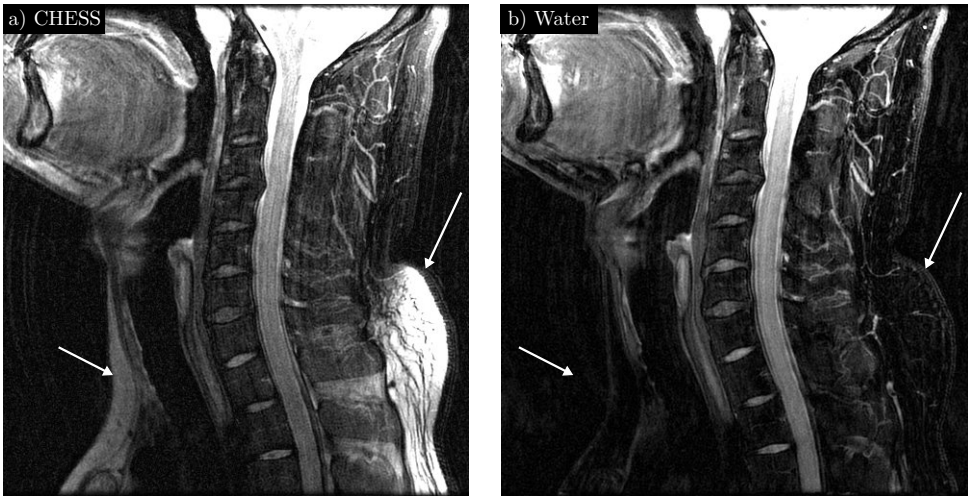

Figure 2.3: a) The CHESS technique fails in regions with too large field heterogeneity (arrows). b) A water image from a fat/water reconstruction with correct field inhomogeneity estimates can correct for this.

in a scan time increase of about 50 %. In 3D gradient echo imaging, the main sequence repetition time can be as short as a few milliseconds. To overcome a massive scan time penalty, the CHESS pulse is not played out every TR. This is reasonable as fat has insignificant T1 recovery between excitations, but this strategy has the downside of breaking steady state. Other ways of reducing the scan time penalty is to make the CHESS RF pulse shorter, but this comes at the cost of a degraded spectral selectivity, or to use water excitation pulses instead.

A disadvantage specific to T1-weighted brain imaging is that differentiation between gray and white brain matter is degraded by magnetization transfer effects from the additional RF pulses [9].

2.2.2 Water Excitation

While technically not a fat suppression technique, water excitation is closely related to CHESS in that both techniques depend on spectrally selective RF pulses. Instead of exciting fat prior to the main sequence, a spectrally selective RF pulse that is centered

around the water peak is used in the main sequence. Since no fat is excited, there is no need for spoiling. Water excitation is a viable option in some applications as the scan time penalty can be shorter compared to CHESS. However, the spectral profile requires a longer RF duration, which prolongs the RF echo spacing in RARE imaging.

One way of achieving spectral selection is through the use of binomial RF pulses [10]. Binomial pulses were originally used to avoid water signal in MR spectroscopy. The binomial RF pulse is composed of a set of sub-pulses, each with a temporal distance equivalent of a 180°dephasing between fat and water, e.g. 2.3 ms at 1.5 T. While the sub-pulses can have an arbitrary shape, the timing between them is determined by the chemical shift. Its simplest form uses only two hard sub-pulses, each with a flip angle of 45°. Increasing the number of sub-pulses, n, narrows the spectral profile at the the cost of prolonged RF duration.

The relative flip angle of a given sub pulse with index i (zero-based) is given by the binomial coefficient $\binom{n-1}{i}$. As long as this flip angle scheme is followed, targeted off-resonance spins will experience a net flip angle of 0°, regardless of the combined flip angle for on-resonance spins, making the binomial pulse sequence less susceptible to RF transmit field inhomogeneities, B_1^+. Binomial RF pulses are an extension of the "jump-and-return" strategy where two ±90°pulses are played out with a delay that determines the spectral response [11]. The phase of the sub-pulses can be alternated to excite an off-resonance species such as fat instead of water. Modifying the RF phase can also be used to reduce the inter-pulse delay [12].

It is possible to combine the spectral selectivity with spatial selectivity by replacing the hard sub-pulses with conventional band-limited RF pulses such as a Gaussian or sinc. In this case, sub-pulses are accommodated with slice selection gradients for spatial selection, while the binomial flip angle scheme and temporal spacing determine the spectral selection. These *spectral-spatial* pulses are suitable for spin echo EPI applications such as fMRI and diffusion imaging. Figure 2.4 illustrates a spectral-spatial RF pulse.

2.2.3 Inversion Recovery

A common method for suppressing fat signal without susceptibility to field inhomogeneities is by using STIR (Short Time Inversion Recovery [13]). Instead of an exci-

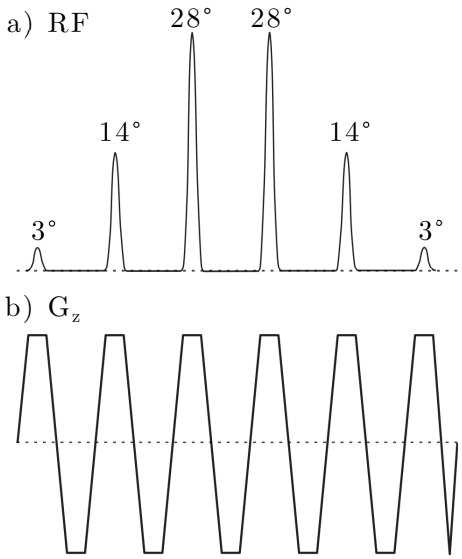

Figure 2.4: a) A spectral-spatial binomial RF pulse composed of six sub-pulses with an amplitude according to the binomial coefficients. The combined flip angle is 90°with a total duration of 12/6 ms at 1.5/3T. The slice-selection gradient is shown in b.

tation, STIR is based on an inversion RF pulse that is spectrally non-selective. By timing the inversion and excitation (inversion time, TI) in the main sequence with the longitudinal zero-crossing of fat, it becomes suppressed.

As the selectivity criterion in STIR is not chemical shift, it is not susceptible to field heterogeneity problems and saturates fat efficiently even at large field of views and near implants. However, there is a significant loss of signal from other tissues, about 40–50 % [14], since all available magnetization is inverted. Additionally, an inverse T1-weighting is imposed, regardless of the main sequence.

Another consequence is that, since STIR suppression is based on the T1 value, *any* tissue with T1 near that of fat will become suppressed. For instance, melanin, mucoid tissue and hemorrhages can be nulled by the inversion [7]. In some applications this might not be problem as fat has a significantly shorter longitudinal recovery rate compared to other tissues. However, tissue with gadolinium uptake can become suppressed due to T1 shortening, and STIR should therefore not be used in contrast-enhanced imaging.

While STIR is a powerful fat suppression technique that is insensitive to main field inhomogeneities, it is susceptible to inhomogeneities in the RF transmit field, B_1^+. With local deviations in B_1^+, the effective flip angle deviates from 180°, and residual transverse magnetization remains. That is, the inversion pulse can locally partially excite rather than invert the magnetization. Two methods have been suggested to overcome this problem:

The first method is called "Spectral Presaturation with Inversion Recovery", in short SPIR, where the inversion pulse is also spectrally selective, and a spoiler gradient is applied with the purpose of dephasing unintended transverse magnetization. SPIR can therefore be used with lower flip angles, typically around 100°. The second method, SPAIR, is similar to SPIR but uses adiabatic inversion pulses to further alleviate problems caused by B1 inhomogeneity.

2.3 Model Based Fat/Water Separation

2.3.1 The Two-point Dixon Method

Instead of suppressing the fat signal, one can acquire images with *chemical shift encoding* (CSE). The phase cancellation artifact described above is an example of this, where the encoding causes destructive interference in regions with both water and fat. Information about the chemical properties is encoded in each voxel, both in the magnitude and phase. In 1984, Dixon discovered that the water and fat signal could be separated through simple arithmetic operations on the in- and opposed phase echo [4]:

$$\begin{aligned} \frac{S_{\mathrm{IP}} + S_{\mathrm{OP}}}{2} &= W \\ \frac{S_{\mathrm{IP}} - S_{\mathrm{OP}}}{2} &= F \end{aligned} \tag{2.4}$$

The post-processing of two echoes results in a water and fat image, free from the phase cancellation artifact. However, the signal model does not account for phase accrual from field inhomogeneities ψ_q, signal decay, or initial phase after excitation ϕ_q. In addition, the fat spectrum contains multiple frequencies ω_p, each with a relative amplitude α_p,

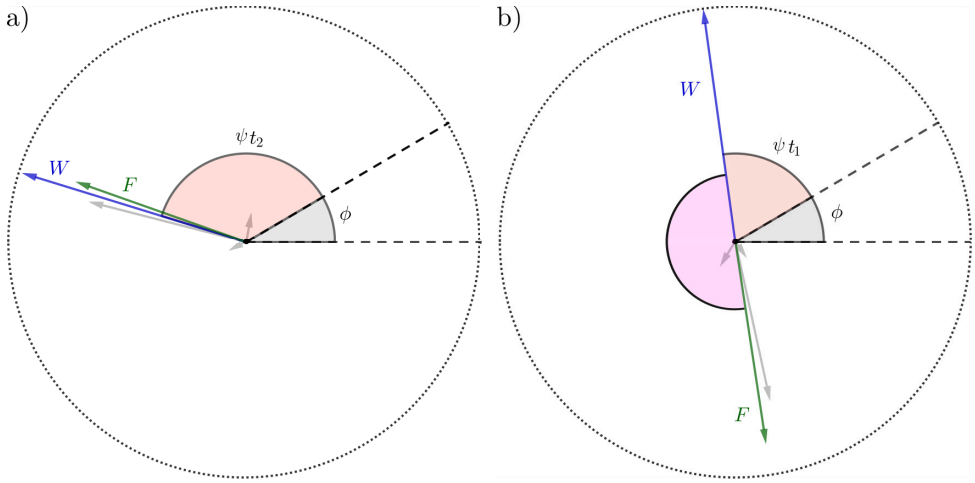

Figure 2.5: In (**a**) and opposed (**b**) phase sampling with the expanded signal model. The gray vectors represent components of the multi-peak fat spectrum, which sums to F. ψ is the off-resonance frequency caused by field inhomogeneities. The phase between W and F (magenta) is the chemical shift encoding. Note the amplitude modulation of F from the multi-peak model.

although the majority of the signal originates from the methylene peak at 3.5 ppm. Expanding Equation 2.3 for the above mentioned effects (excluding signal decay for simplicity) results in the modified signal Equation for a voxel q:

$$S_q(t) = e^{i(\phi_q + \psi_q t)} \left(W_q + \sum_p \alpha_p e^{i\omega_p t} F_q \right) \tag{2.5}$$

where $\sum \alpha_p = 1$. In the case of two echoes, $t_1 \in t_{\mathrm{OP}}$ and $t_2 \in t_{\mathrm{IP}}$, the expanded signal model above is illustrated in Figure 2.5, where it is clear that the arithmetic operations in Eq. 2.4 do not separate F and W.

The initial phase after excitation, ϕ_q, is time invariant and will cancel in the subtraction process, but plays an important role in the separation method, as described by Bydder et al. [15]. In the expanded model, fat and water signal cannot be separated using the original Dixon method without correction of the field inhomogeneity ψ_q. Unless corrected for, any unmodelled phase contribution causes signal leakage between fat and water.

Consider a voxel containing only fat, located in a region where the local field is stronger and the off-resonance exactly equals the chemical shift. This voxel is indistinguishable from the water as they have identical resonance frequencies. That is, there is no way of knowing the cause of a measured phase as it can originate from any combination of off-resonance sources.

Dixon originally proposed to take the absolute value of S_{IP} and S_{OP} to remove ψ prior to the arithmetic steps. As seen in in Equation 2.5, this merely results in a separation into a large and small component, with a remaining binary choice of assigning them to fat and water [16]. In other words, on a single voxel level, off-resonance water is indistinguishable from chemically shift encoded fat unless ψ is known.

2.3.2 Refined Dixon Methods

The introduction of the two-point Dixon method [4] initiated the research on model-based separation of fat and water in MRI. With the same spin echo sequence as Dixon proposed, Sepponen et al. suggested to sample 16 echoes by shifting the refocusing pulse, and Fourier transforming the signal along the echo dimension to yield a coarse spectrum[17]. The 3.5 ppm spectral image then reveals an image containing fat, while the on-resonance image contains the water signal. The echoes required to sufficiently resolve fat and water in the spectral domain results in long acquisition times, and the field inhomogeneity term is still unaccounted for.

By modifying the spin echo sequence and acquiring two echoes per excitation, in combination with an additional acquisition where the RF echo spacing changed but the readout position remain, Yeung and Kormos [18] were able to estimate ψ from the two in-phase echoes, and use the other two echoes for fat/water separation. Since ψ is estimated from the difference between the phase images, phase unwrapping is required to resolve field inhomogeneities larger than $\pm\frac{1}{2|t_{\mathrm{IP}}-t_{\mathrm{OP}}|}$. Given error-free phase unwrapping, the estimates are correct even in presence of larger field inhomogeneities. A major drawback of this method is that acquisition time is doubled without any contribution to the resulting image SNR.

An alternative to the original Dixon method, was suggested by Kim et al. [19]. The method uses three echoes, or points, where one S_{IP} and two S_{OP} are acquired. S_{OP}^{-} is

acquired before the spin echo, and S^{+}_{OP} after. Voxelwise division between S^{-}_{OP} and S^{+}_{OP} after phase unwrapping determines ψ, and the fat/water separated images are obtained from S_{IP} and the corrected S^{-}_{OP}. S^{+}_{OP} is only used to estimate ψ without contributing to the SNR of F and W.

One separation method was developed for acquisitions of one S_{IP} and two partially opposed signals, $S^{\pm}_{\mathrm{POP}}$ by Lodes et al. [20], where $S^{\pm}_{\mathrm{POP}}$ are symmetrically placed around the spin echo, acquired in multiple excitations. In this method, all echoes contribute to estimating F and W, providing a significant SNR improvement from the additional signals.

The three-point dephasing time constraints were further relieved by Xiang and An [21], enabling separation for acquisitions with an arbitrary initial dephasing time, t_0, and arbitrary but constant dephasing time difference, Δt. Reeder et al. generalized the method further to arbitrary echo times with a least-squares (LS) estimation [22], although LS estimation with any number of echoes was described earlier by An and Xiang [23].

Four-point methods allow further modelling of the signal, utilized by Glover [24] to correct for T2 decay. The signal decay is often negligible in spin echo imaging, since Δt is typically much smaller than T2 of water and fat. When the Dixon method is applied to gradient echo sequences, however, the underestimation of F due to $\mathrm{T2}^{*}$ decay may result in a residual lipid signal on the order of 10 % [25].

A common way of measuring the sampling efficiency is the *effective number of signal averages*, NSA, allowing comparison to the noise performance when multiple acquisitions without chemical shift encoding are averaged [24]. By definition, NSA $\in [0, N]$, where N is the number of echoes. Pineda et al. analysed the highest achievable NSA in the estimates for different dephasing times in the three-point case [26]. It was shown that sampling of $\mathbf{S}_{\mathrm{subopt}} = [S^{-}_{\mathrm{OP}}, S_{\mathrm{IP}}, S^{+}_{\mathrm{OP}}]$ is sub-optimal when ψ is estimated intrinsically. Instead, optimal NSA is achieved when the fat vectors are spread evenly in the unit circle, i.e. $\Delta t = 2\pi/3\omega$, about a $\pi/2$ phase shift, as shown in Figure 2.6.

Another important contribution to the field was made by Bydder et al. and Yu et al., where multiple resonances of the fat signal was modelled, resulting in superior fat suppression compared to single-peak or conventional fat saturation [3], [27]. In combination

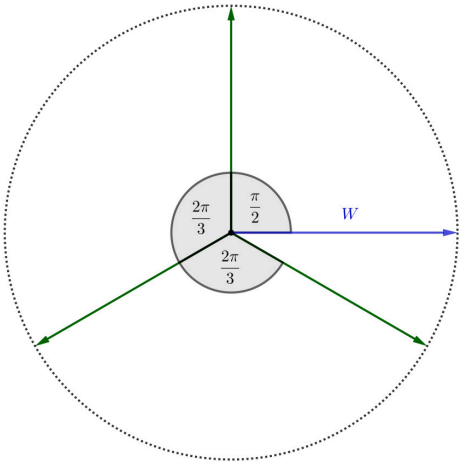

Figure 2.6: Optimal encoding in the three-point Dixon method with intrinsic field map estimation [26]. Green arrows show the fat dephasing of the three echoes.

with the k-space separation of species [28], [29], chemical shift displacement correction is performed for all spectral components.

Xiang et al. described a two-point separation method free from the opposed phase constraint. The two-point in-phase constraint was later removed by Eggers et al. and Berglund et al. [30], [31]. However, NSA is maximized with $\mathbf{S} = [S_{\mathrm{IP}}, S_{\mathrm{OP}}]$.

While multi-point ($N \geq 3$) methods offer an NSA improvement, the sequence must typically be prolonged in order to acquire the additional echoes. The two-point method measures two complex-valued signals, allowing estimation of the four real-valued unknowns (F, W, ψ, ϕ) in Equation 2.5. For a qualitative fat suppression, two echoes are sufficient and is the method of choice in this thesis as it imposes a minimum demand on the acquisition in terms of additional scan time.

2.4 Field Map Estimation

Besides the need to reliably reflect inhomogeneities, it is desirable if the field map can be established from the data itself. This is a difficult problem to solve, however, and

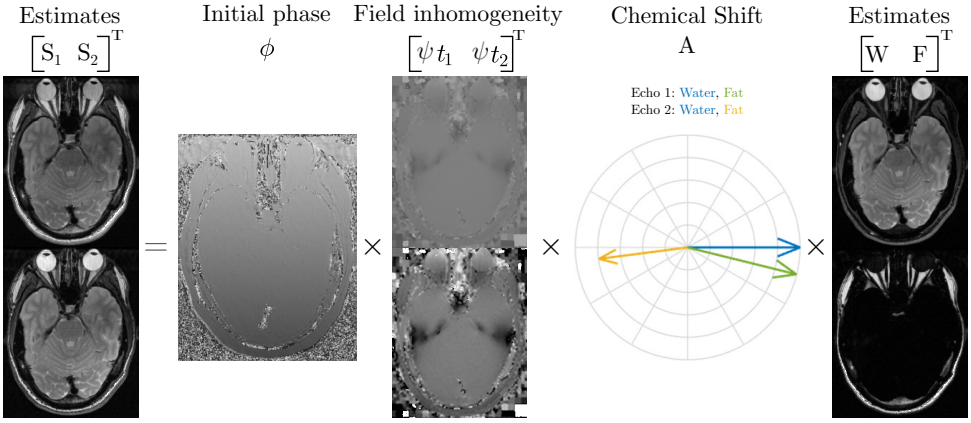

Figure 2.7: A pictoral representation of Equation 2.6.

was for a long time the focal point of model based fat/water separation research. With reasonable assumptions on the spatial smoothness of ψ, the problem is solvable. Herein lies both the weakness and strength of model based fat/water imaging: Any field inhomogeneity can in theory be corrected for in the reconstruction, but relies on an accurate field map estimation. As discussed in detail below, the ability to correctly estimate the field map is highly dependent on the dephasing times.

For a set of two echoes, the expanded signal Eq. 2.5 can be written on matrix form as

$$\underbrace{\begin{bmatrix} S_q(t_1) \\ S_q(t_2) \end{bmatrix} e^{-i\phi_q}}_{\mathbf{S}_q^{\phi}} = \underbrace{\begin{bmatrix} e^{i\psi_q t_1} & 0 \\ 0 & e^{i\psi_q t_2} \end{bmatrix}}_{\mathbf{B}(\psi_q)} \underbrace{\begin{bmatrix} 1 & \sum_p \alpha_p e^{i\omega_p t_1} \\ 1 & \sum_p \alpha_p e^{i\omega_p t_2} \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} W_q \\ F_q \end{bmatrix}}_{\mathbf{x}_q}. \tag{2.6}$$

A pictoral represention of Eq. 2.6 is shown in Figure 2.7. Since only two echoes are available, F and W are modelled as real-valued to avoid under-determination of Eq. 2.6. The initial phase ϕ can be estimated for a given ψ as described in Paper II, which is equivalent to that described by Bydder et al. [15]. Since Eq. 2.6 is nonlinear when ψ is unknown, ordinary least squares is not applicable. Various methods have been proposed to solve this problem, and all require some assumption to overcome the fat/water swap ambiguity described above.

The entire parameter vector $[\phi, \psi, W, F]$ can be solved with an iterative Gauss-Newton search on a voxel-by-voxel basis, as was suggested by Reeder et al. for three echoes and complex valued estimates [22]. This search is computationally expensive and dependent on the initial guess. The resulting field map is smoothed after converging, and since there are no safeguards against local convergence with this method, errors can propagate in the smoothing process. Yu et al. proposed to refine the method by estimating a low resolution field map followed by a selection of voxels where the field map estimates are in local agreement. Interpolation between these confident field map estimates serve as the initial guess in the region growing procedure, "snapping" voxels to the closest minimum at native scale [32]. This estimation, with an extension for flexible echo times, is known as the IDEAL method [33]. Although stated that the interpolation strategy *ensures* a correct field map estimate, no such guarantees can be made since it relies on field map varying smoothly between interpolation points, and that the confident points are correct.

Similar to Dixon's proposal to take the magnitude of the two signals before subtraction, one can derive analytical solutions for the two components, which results in a separation between fat and water on a voxel level but without identifying the two. In other words, two field map *candidates* can be derived for each voxel where one corresponds to a swapped solution between water and fat, and the other is correct. These candidates can be found even when the encoding differs from in- and opposed phase sampling [30], [31], [34] and for any number of echoes [35], [36].

An alternative method for obtaining the candidates is to calculate the ℓ^2 residual norm J between the measured and expected signal for each voxel q and field map candidate ψ:

$$J_q(\psi_q) = \left\| \mathbf{S}_q^{\phi} - \mathbf{B}(\psi_q)\mathbf{A}\,\mathrm{Re}(\mathbf{A}^{-1}\mathbf{B}(\psi_q)^H \mathbf{S}_q^{\phi}) \right\|_2 \tag{2.7}$$

where $\mathbf{S}_q^{\phi}$ is the measured signal vector $\mathbf{S}_q$ modelled by Equation 2.5, after demodulation with the initial phase after excitation ϕ.

Equation 2.7 deserves a close investigation as it is important to understand how the cost function is configured. Reading from right to left with the aid of Figure 2.7: For

a given ψ candidate, the signal vector is demodulated with the associated ϕ estimate. $\mathbf{B}^H$ is the conjugate transpose of $\mathbf{B}$, meaning that it acts to roll back the off-resonance phase accrued at dephasing times t_1 and t_2. After the signal vector has been corrected for both ψ and ϕ, the problem is linear, and applying the model matrix inverse $\mathbf{A}^{-1}$ to $\mathbf{B}(\psi_q)^H\mathbf{S}_q^\phi$ yields an intermediate estimate of F and W. Note that, in contrast to IDEAL, the search space here considers only ψ, made possible since the minimization problem for ϕ is a *separable* non-linear least squares problem [35], [37]. This separation simplifies the problem significantly, and linear least squares estimation can be used after field inhomogeneity correction:

$$\hat{\mathbf{x}} = \mathbf{A}^{-1}\mathbf{S}^{\phi_q,\psi_q} \tag{2.8}$$

where the ϕ and ψ superscripts of $\mathbf{S}$ indicates that the respective phase terms have been demodulated from the signal. The real-valued property of the estimates is enforced in the residual by the taking the real part in Eq. 2.7. The real-valued estimates are thereafter forward projected by applying the model matrix $\mathbf{A}$ and re-adding off-resonances in $\mathbf{B}$. The difference between the estimated and measured signal thereby form the residual vector. We then seek to minimize the ℓ^2 norm, i.e. find the least-squares fit to Eq. 2.6. It has been shown that the residual function J is periodic with a period of $\frac{1}{|t_1-t_2|}$ [38], so the residual can be calculated by sweeping over one period only. Although this restricts the resolved fieldmap range to $\pm\frac{1}{2|t_1-t_2|}$, it is sufficient for estimating $\mathbf{x}$. The two minima are exactly zero since Eq. 2.6 is exactly determined, so the ℓ^1 norm can also be used. Similar to the analytical solutions, the minima correspond to the true or fat/water swapped solution. Figure 2.8 shows Eq. 2.7 evaluated for two sets of dephasing times and varying fat fractions.

Since only two local minima in one period are considered, the search space for ψ in an image with Q voxels is reduced to 2^Q. This is still an insurmountable search space, and many alternatives exist for selecting the correct set from the candidates. One common method is to start from a solution in a seed point, and assume that the neighboring voxels have a similar field inhomogeneity. The field map can be resolved by continuing outwards from one or several seeds in a region growing manner, where the candidate closest to the resolved neighbor is chosen [23], [38], [39]. This is based on a spatial smoothness assumption of the underlying field. Region growing techniques share a considerable

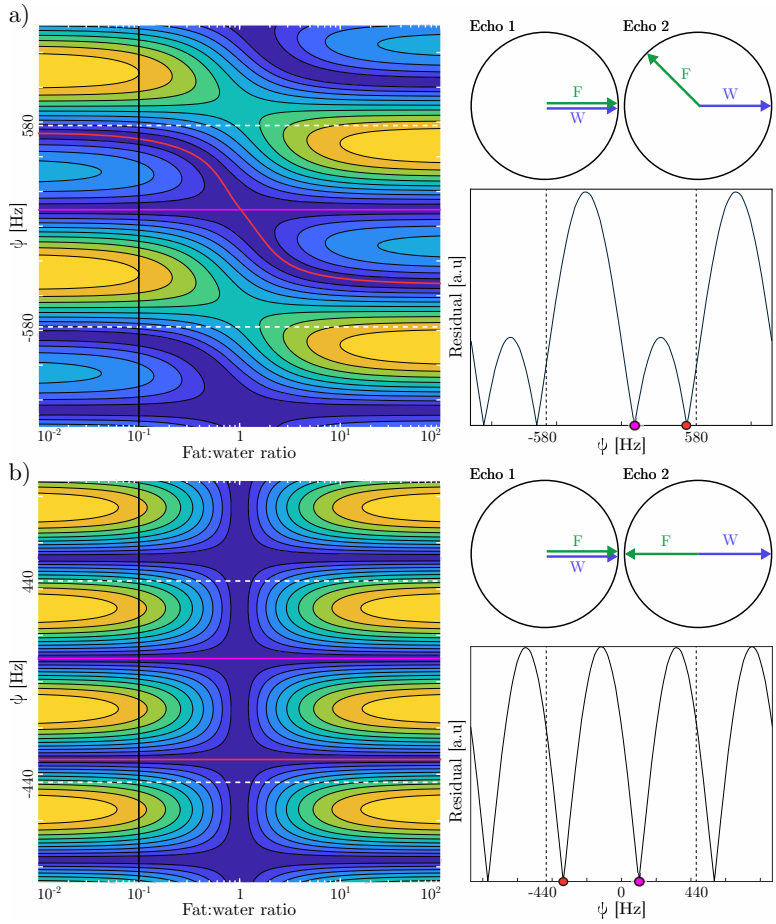

Figure 2.8: Contour plot of residuals calculated for a range of fat:water ratios and ψ candidates at 3T with a simulated off-resonance of $\psi_{\text{true}} = 100$ Hz (magenta solid horizontal line). The residuals are pseudo colored from blue to yellow (low to high). a) In-phase and a 135° echo. b) In-phase and a 180° echo. The residual is periodic over ψ, marked by the white dashed lines. For each candidate, there are two minima every period for a given fat:water ratio, as shown in the profile plot (black solid vertical line). ψ_{true} is shown in magenta, and the swapped minimum (red) is offset by the chemical shift (440 Hz) for a single-species voxel. When the sampling is exactly in and opposed phase, the minima are always 440 Hz apart.

weakness in that it requires an initial solution in the seeds. If the seed is wrong, that error propagates to other regions and causes the fat/water swap artifact.

The method used in this thesis was inspired by the formulation of the problem as a graph, where each voxel is a vertex, together with a source and a sink [40]. Each voxel has two edges associated with the current and proposed ψ candidates. The cost of cutting such a data edge is determined by the associated residual, for instance Eq. 2.7 There are also regularization edges between adjacent voxels, with an associated cost to penalize ψ differences in a local region.

The algorithm originally proposed is iterative. Initially, field map candidates are chosen with a starting guess of zero and the graph is solved with a graph-cut technique. Subsequent iterations propose ”jumps” to an adjacent minimum to avoid local convergence, i.e. the search space is not restricted to one residual period. After a certain number of iterations, constant offset are used as candidates, i.e. slightly off-resonance from local minimas, thereby ”fine-tuning” the estimated field map. A regularization parameter μ is introduced to impose a smoothness on the field map, which regulates the relative cost between regularization and data edges.

Even aliases of the true solution yield correct fat/water estimation, so the residual periodicity can be exploited to reduce the search space to 2^Q as previously mentioned. Berglund et al. proposed to restrict the search space and construct a single QPBO (quadratic pseudo-boolean optimization) graph, which can be solved to find an optimal solution without iterations [41]. For reasons described in detail in [42] and summarized in the theory section in Paper I, this technique can render unlabelled voxels. This is a problem particularly seen in low-signal regions, and can be adequately solved by downsampling the image, solve that graph, and extract the ψ estimate [43]. The field map in Figure 2.7 displays a blocky appearance outside of the head which originates from this multi-scale approach.

2.5 Pulse Sequences

A pulse sequence describe how the MR signal is acquired. It consists of a scheme of timed instructions played out on different boards. These boards include the three orthogonal

physical gradient boards G_X, G_Y, and G_Z, the RF board, and data sampling (ADC). Since all of the sequences described in this thesis are Cartesian (at least within a TR interval), the gradient boards are presented here as G_{read}, G_{phase}, and G_{slice} to indicate their primary purpose and avoid involving field-of-view dependencies. By this logical arrangement, the data sampling events can be overlaid on G_{read}. Although other details such as the ADC receive frequency and phase also need consideration when designing a pulse sequence, they are left out here for simplicity.

2.5.1 Multi-Echo Gradient Echo and the Receiver Bandwidth

The simplest pulse sequence, at least in terms of number of gradient events, is the gradient echo (GRE) sequence. A 2D multi-echo GRE sequence is shown in Figure 2.9. Two ways of acquiring the echoes are shown: Either the readout gradient polarity is reversed every other echo (Fig 2.9a), or an intermediate flyback gradient is sandwiched between readouts with the same polarity (Fig 2.9b). Fat starts dephasing due to chemical shift at the magnetic isocenter of the excitation pulse (blue dashed line). Since the dephasing of water and fat is not refocused in gradient echo sequences, the echo time and dephasing time are identical.

Before the readout, gradient prephasers are played out together with a phase encoding gradient to select the k-space line to be acquired. A readout gradient is then applied with an amplitude that governs the range of frequencies (or receive bandwidth, rBW) across the object. By choosing to acquire samples with a certain frequency, the resulting image will have the desired FOV, i.e. the sampling bandwidth is matched to the FOV and rBW. Through this relationship, rBW is linearly related to the readout gradient amplitude, and the two terms are sometimes used interchangeably. In other words, the receiver dwell time $\tau = 1/\text{rBW}$ of each k-space sample determines the FOV for a specific gradient amplitude. Since an increased dwell time directly relates to more data being averaged per sample (or equivalently, lower sample variance) the associated SNR relationship is

$$\text{SNR} = \frac{\mu}{\sigma} \propto \sqrt{\tau} = \sqrt{\frac{1}{\text{rBW}}} \tag{2.9}$$

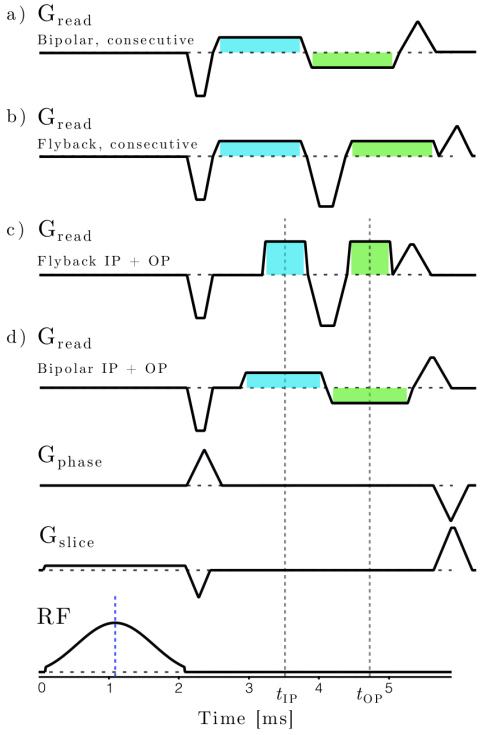

Figure 2.9: Multi-echo gradient echo. Four different sampling strategies are depicted in a-d. a-b) Consecutive acquisition where dead times are avoided. c-d) Specific echo times. Dead times are introduced before the first readout gradient (cyan).

Where μ is the mean signal value, and σ is the standard deviation of the signal. The expected value of the signal does not change by averaging. By the definition of SNR in Eq. 2.9, its relationship to $\sqrt{\tau}$ can be derived from the noise power reduction in n averaged uncorrelated measurements of a signal:

$$\mathrm{Var}\left(x_{\mathrm{avg}}\right) = \mathrm{Var}\left(\frac{1}{n}\sum_{i=1}^{n} x_i\right) = \frac{1}{n^2}\underbrace{\mathrm{Var}\left(\sum_{i=1}^{n} x_i\right)}_{=n\sigma^2} = \frac{\sigma^2}{n} \tag{2.10}$$

For instance, decreasing rBW by a factor of 4 means the gradient duration and τ is 4 times longer, and a two-fold increase in SNR. To cover the same k-space extent, or

equivalently, to keep the k-space step size Δk constant in order to maintain the FOV, the gradient amplitude is decreased by the same factor. This concept is key to the proposed techniques in this thesis where we aim to maximize the time spent sampling, and doing so as efficiently as possible.

Going back to the gradient echo sequence in Figure 2.9a, the gradient polarity changes between readouts and they are played out consecutively. The bipolar readouts cause the chemical shift displacement to appear in opposite directions between the two echoes, which can complicate the fat/water problem. Distortions due to field inhomogeneity is also further complicated as they alternate between the echoes. Eddy currents and gradient delays can cause phase contributions that in general must be resolved in fat/water applications [44]–[47]. This does not pose a problem when the field map is intrinsically determined in a double echo acquisition. The alternative of rewinding the k-space trajectory with the use of a *flyback* gradient has the downside of prolonged minimum dephasing time difference, and increase sequence dead time, as seen in Figure 2.9b.

Both scenarios in Figure 2.9a-b acquire the echoes as soon as possible, and the two echo times are consequently determined by the readout bandwidth and spatial resolution. In fat/water imaging, there is a strong relationship between the dephasing times and the noise performance as will be discussed later, hence the echo times should not be left to chance. Consider the case where we want to acquire the center of k-space with a *specific* CSE, i.e. specific dephasing times in Eq. 2.3. Instead of placing the readout gradients consecutively, they must be centered around chosen time points t_1 and t_2. In the common case of in- and opposed phase sampling, $t_2 - t_1$ is only 1.1 ms at 3T, meaning that a high receiver bandwidth is necessary to achieve the desired CSE. If a flyback gradient is used, it must also fit within $t_2 - t_1$, and even higher amplitudes are necessary. In practice, this restricts the resolution to such a degree that monopolar in- and opposed phase sampling is often not feasible at 3T unless very high performing gradient hardware is available. Another consequence of the echo time restriction is apparent in Figure 2.9c-d, namely that the short acquisition windows in relation to the sequence duration can result in an inefficient sequence. In general, dead times appear with fixed echo times, also evident in Figure 2.9c-d.

2.5.2 Multi-Echo Spin Echo and RARE

The spin echo sequence corrects for chemical shift dephasing by applying refocusing pulses. Acquisition can be sped up by applying multiple refocusing pulses to acquire several k-space lines per excitation, known generically as the RARE method (Rapid Acquisition with Relaxation Enhancement [48]), but also known as TSE or FSE in vendor implementations (Turbo/Fast Spin Echo). The number of lines per excitation defines the ETL (echo train length, or turbo factor). Since each echo is acquired with its own echo time, there is a T2 modulation along the phase encoding direction.

The time between two spin echoes is commonly referred to as the *echo spacing*, which is a confusing term in multi-echo sequences with CSE. A gradient echo is formed whenever the center of k-space is traversed, and a spin echo occurs midway between refocusing pulses. Then, the *RF echo spacing* is the duration between RF pulses, and the *gradient echo spacing* is the difference in dephasing times. The dephasing time in spin echo sequences must account for the RF refocusing effect, and is therefore defined as the negative/positive time to/from the nearest spin echo.

RARE sequences must comply with CPMG (Carr-Purcell-Meiboom-Gill) conditions. The Carr-Purcell conditions state that the timing between excitation and refocusing pulses must be half of that between two refocusing pulses [49]. Meiboom and Gill realized that by applying the refocusing pulses orthogonal to the excitation pulse, cumulative errors from non-180°refocusing pulses are cancelled [50]. Another condition commonly stated as part of the CPMG set is that the zero order magnetic moment between refocusing pulses should be constant. Although it is not mentioned in the original works, because they relate to MR spectroscopy where no imaging gradients are present, Hennig considered the between-RF gradient compensation as a method of fulfilling CPMG conditions [48], and this nomenclature has since caught on [51]. Any deviation from these conditions can cause stimulated echo pathways that interfere with the primary spin echo, as can be appreciated from extended phase graph analysis [52], [53]. Concomitant gradient fields and eddy currents can induce a net magnetic moment and pose a problem in RARE as that can break CPMG conditions.

For the purpose of chemical shift encoding, Dixon originally used a spin echo sequence where the two echoes are acquired in separate excitations, thereby doubling the scan time

[4]. Instead of acquiring another phase encode, the gradient echo was shifted so that is does not coincide with the spin echo. This way, the echoes are chemically shift encoded. The RARE sequence is often not acquired in a single-shot fashion, but rather from multiple excitations where the ETL is around 10-20 for T2 weighted imaging. Since the goal is to acquire multiple echoes, the total number of readouts must increase according to the desired echo count. One method of achieving this is to keep the chemical shift encoding consistent between refocusing pulses. For instance, the first excitation acquires ETL = 8 lines where the readout gradient is centered around the spin echo, while the second excitation acquire the same lines with an opposed-phase CSE. Such *sequential* sampling was used in Paper IV. Alternatively, CSE echoes can be *interleaved*, meaning that odd readouts are acquired opposed phase, while even readouts are in-phase echoes. This strategy was used in Paper I. Figure 2.10 shows both alternatives.

The high flip angles in RARE imaging can be restrictive due to patient heating, which can be alleviated with lower flip angles accompanied with strong crusher gradients to avoid stimulated echoes. As the time spent on crusher gradients, phase encoding, readout ramps, and the RF pulse itself is of comparable duration to the acquisition window, it is tempting to further speed up the acquisition by acquiring multiple gradient echoes between each refocusing pulse pair. This technique is known as GRASE (gradient and spin echo), where phase encoding blips are applied during gradient switching to cover multiple k-space lines [54]. The echoes are differently affected by T2′ decay, so a modulation is expected beyond that of T2 decay. An even number of gradient lobes can in theory be used by negating every other GRASE block to account for the k-space trajectory being mirrored, but associated gradient delays and eddy currents can cause CPMG violations. GRASE sequences are therefore most often acquired with an odd number of readouts composing the GRASE readout block.

In the case of fat/water separated GRASE sequences, the additional echoes can be used to acquire chemical shift encodings instead of phase encodings, which is the underlying method for the fTED technique (fast spin-echo triple-echo Dixon [55]) and its derivatives [56]–[58]. The fTED technique has three bipolar readout gradients, centered around the spin echo and timed so that the first and third echoes are acquired opposed phase. This requires a high performance gradient system, or too much time will be spent on gradient switching. More recently, Son et al. proposed an alternative to fTED where the shift

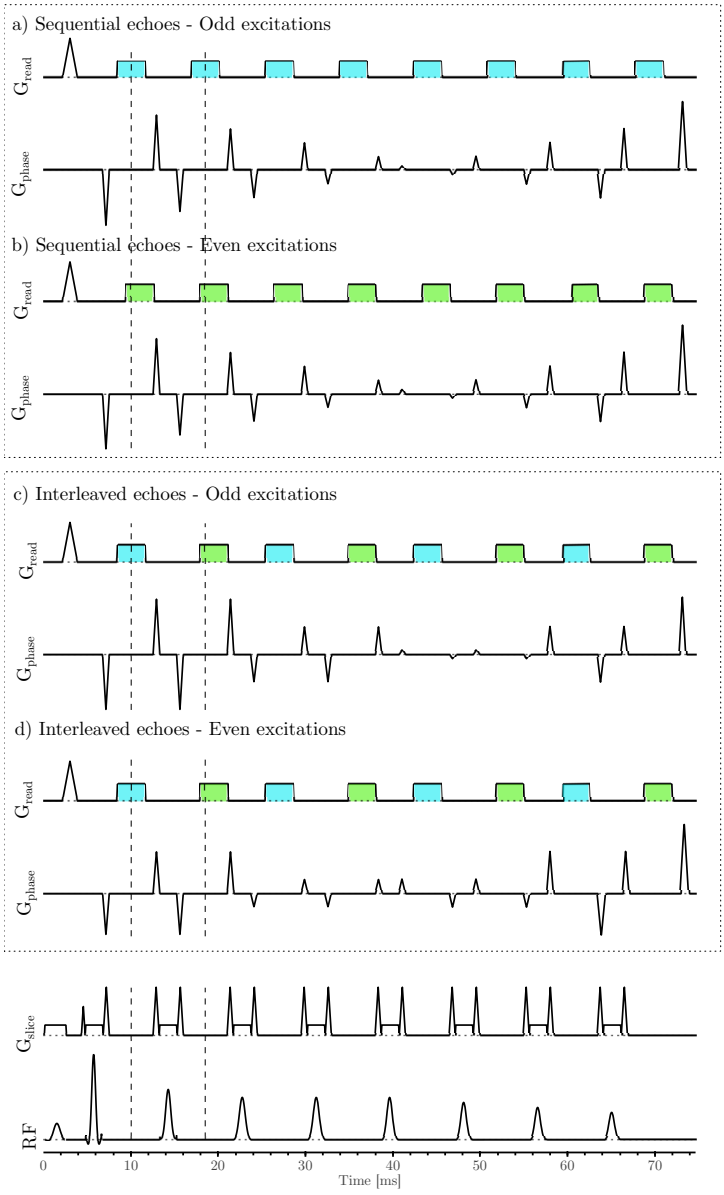

Figure 2.10: Pulse sequence diagram of RARE showing two methods acquiring chemical shift encoded signals: Sequential (a-b) where the phase encodes are the same between shots but the CSE varies. In interleaved mode (c-d), CSE is alternated between refocusing pulses and the phase encodes are varied between shots. Vertical lines mark the first two spin echoes, while the colored patches indicate sampling of the in-phase (cyan) and opposed phase (green) echo.

can be chosen flexibly, although the center echo is still acquired in-phase [57].

As described in detail in Paper III, the use of even number of readout gradients with equal rBW is not compatible with chemical shift encoded GRASE as the encoding will be ill-posed.

2.5.3 PROPELLER and Motion Correction

A consequence of the long acquisition times in MRI is the sensitivity to patient motion. Numerous sequences and reconstruction methods have been proposed to counter this issue, which can be categorized into prospective and retrospective methods.

Prospective methods alter the sequence in real time based on some feedback during acquisition. The feedback can be external, e.g. from a respiratory belt or a cardiac monitoring device such as an ECG. Internal feedback is an alternative, where for example, a short navigator is acquired between or even during sequence kernels.

External feedback often requires additional patient setup but no alterations to the sequence, as opposed to internal feedback where typically the sequence duration is prolonged. Since prospective methods apply the motion estimates in real time, expected magnetization from subsequent excitation (i.e. steady state) and inversion pulses are maintained. If rigid body motion occurs, the field of view can be rotated and translated to correct for motion. However, prospective methods rely on high quality feedback, and erroneous estimates will degrade the image quality.

Retrospective motion correction techniques are based on strategically acquiring data so that motion can be compensated for in the reconstruction process. One such strategy is the PROPELLER (Periodically Rotated Overlapping ParallEL Lines with Enhanced Reconstruction) sequence, where multiple rotated central strips, or *blades*, of k-space are acquired sequentially, as introduced by Pipe [59].

PROPELLER is based on RARE, where multiple lines are acquired every excitation. Blades are rectangular in k-space, with high k-space coverage in its unique frequency encoding direction, and a low coverage in its phase encoding direction. Since the center of k-space is redundantly sampled for every blade, retrospective rigid body motion

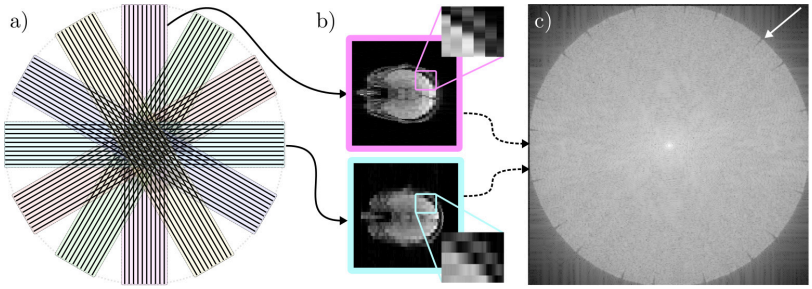

Figure 2.11: a) Illustration of the PROPELLER k-space coverage. For clarity, only eight blades are shown. b) Two blades in the image domain, with insets highlighting the difference in resolution. The magenta blade is acquired with high resolution in the left-right direction, while the orthogonal cyan blade has high resolution in the anterior-posterior direction. Dashed arrows represent motion correction, density compensation, and gridding. c) Density-compensated blade-combined k-space. Rotational motion correction has resulted in k-space wedges (arrow).

correction is possible by in-plane alignment between blades. A side effect of sampling the k-space center more often is that higher frequencies are less densely sampled, leading to an increased scan time factor up to $\pi/2$ to fulfill the Nyquist sampling criterion, depending on blade width [59]. However, this scan time increase does contribute to the SNR of the final image. The PROPELLER sequence is illustrated in Figure 2.11.

Once motion is estimated, blades are corrected by applying phase ramps and rotations in k-space, followed by resampling onto a Cartesian grid, i.e. *gridding*. Any misalignment of the k-space center results in a phase ramp in the image domain, which is why a a triangular high-pass phase filter is applied in the image domain to center the blades prior to gridding.

Since k-space is sampled non-uniformly, density correction is required to avoid overestimation of low spatial frequencies. The sample weights can be estimated iteratively [60], through Voronoi cells [61], or analytically for some sampling patterns [62]. Rotational motion correction can cause wedges in k-space that breaks the Nyquist criterion, as illustrated in Figure 2.11c.

Even without rigid body motion correction, the PROPELLER sequence is less sensitive to motion than Cartesian sampling due to the averaging effect from redundant sampling

of the k-space center. The PROPELLER sequence, and its variants, therefore also helps to reduce flow and abdominal motion artifacts [59], [63].

The problem of combining multi-shot PROPELLER, or any RARE-based sequence, with diffusion is the sensitivity to motion and spatially varying phase between shots. Pipe et al. extended the PROPELLER reconstruction to account for the phase sensitivity due to motion. In order to mitigate the CPMG violations imposed by the diffusion gradients, the phase of the refocusing pulse was cycled between the x and y axis, as described earlier by Shaka et al. [64], enabling diffusion weighted PROPELLER imaging [65].

Pipe and Zwart later combined GRASE with PROPELLER to acquire wider blades and shorten the minimum scan duration. The proposed technique also reduces the SAR scan time penalty [66]. Later work by Norbeck et al. has combined PROPELLER with simultaneous multi slice (SMS) imaging [67]. Three dimensional susceptibility weighted imaging [68] was combined with PROPELLER by Holdsworth et al. [69], using EPI readouts along the short-axis of the blade as described by Skare et al. [70].

2.5.4 Fat/Water Separated PROPELLER

The first combination of PROPELLER and the Dixon technique was demonstrated by Huo et al. [71], called Turboprop IDEAL. Instead of acquiring several phase encoding lines between refocusing pulses, as in the original Turboprop method, the GRASE readout was used to sample multiple CSE echoes. Essentially, Turboprop is a combination of PROPELLER and fTED with flexible CSE. In order to avoid long echo spacing and the associated T2 blurring effects, high readout bandwidths are required with significant dead time during gradient switching. To correct for gradient delays and eddy-current induced phase difference φ between echoes, all blades are acquired twice with opposite readout polarity. The center echo of the blade-pairs is subsequently phase unwrapped and matched to estimate φ. The second acquisition not only doubles the scan time, but also limits the retrospective motion correction capabilities since it is assumed that no motion occurs between the two acquisitions. In addition, a large RF echo spacing is unavoidable due to the GRASE design and the introduced shift.

He et al. combined compressed sensing [72] with the Turboprop IDEAL reconstruction, where system imperfections are incorporated in the model matrix [73]. Fat/water

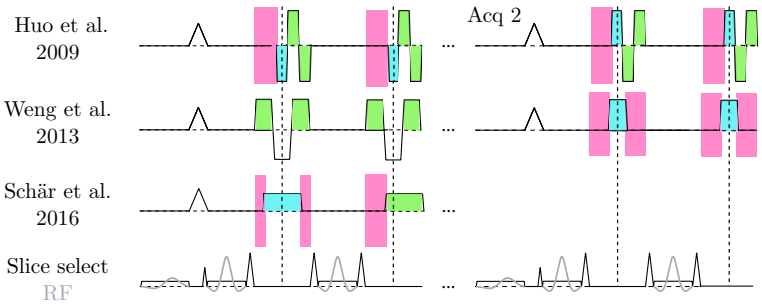

Figure 2.12: Fat/water separated PROPELLER sequences [71], [74], [75], showing the readout gradients and their common slice selection gradients and RF pulses. Cyan and green patches indicate sampling of in-phase and out-of-phase echoes, respectively. while magenta marks dead times. Dashed vertical lines mark spin echoes.

separation is performed on gridded data, after the compressed sensing reconstruction, thereby avoiding a second acquisition.

Weng et al. also proposed a modified Turboprop sequence, where one acquisition samples an in-phase echo, S_{IP}, while the other acquires two opposed phase echoes, S^{+}_{OP} and S^{-}_{OP}, symmetrically placed around the spin echo, between a flyback gradient [74]. All three echoes are then sampled with the same readout polarity, with φ estimated from S^{+}_{OP} and S^{-}_{OP}.

Blade-wise fat/water separation was recently described by Schär et al. [75], where two echoes are acquired in an interleaved manner. That is, the echo train is doubled and each line is acquired sequentially with two different CSE. A major benefit of the interleaved scheme is that both echoes are acquired within the echo train, therefore less susceptible to motion induced phase differences between echoes. φ does not need to be estimated as it is incorporated in the field map and echoes are acquired with the same readout polarity.

Figure 2.12 shows the above mentioned acquisition strategies.

As a result of changing the frequency encoding direction between blades, the chemical shift displacement varies between blades. In combination with the high fat signal inten-

sity in RARE sequences [76], this can have a significant impact on the image quality, with a fat signal radially smeared in the final image. PROPELLER sequences are often acquired with a high readout bandwidth to lessen the shift, at the cost of lower SNR. Removing the fat shift in the reconstruction process was done by Schär et al. and later in Paper I to improve both image quality and motion estimates [75].

2.6 Partial Fourier

Real-valued objects have by definition no imaginary component, and can be fully described by one side of k-space. This property of redundant information in k-space is powerful as it in theory enables imaging acceleration up to a factor of two. This can be achieved e.g. by acquiring only the top half of k-space.

In reality however, most objects do carry some phase information. A conventional spin echo image contains mostly smoothly varying phase information as the refocusing pulse corrects for off-resonance phase contribution. In such scenarios, a solution is to acquire a k-space region beyond the symmetry point, where the distance is chosen to cover the assumed low frequency phase distribution. For instance, the in-phase image in Figure 2.7 (page 17) with 320 phase encoding lines can sufficiently be reconstructed with 160+64 lines.

It is important to emphasize that no information is gained with partial Fourier. Although the scan time in the example above can be 30 % shorter, the SNR is approximately 16 % lower from the data ”loss” compared to the fully sampled case.

The conjugate symmetry is a property in all dimensions in k-space. This means that for 3D imaging, the acquisition times can be shortened by reduced phase encoding sampling in one phase encoding direction, but not both. Partial Fourier is commonly used to shorten gradient echo spacings or echo times in EPI.

There are several methods for synthesizing conjugate samples, where the homodyne and POCS technique, or variants thereof, are widely used.

2.6.1 Homodyne

Homodyne is a one-pass method to synthesize missing k-space samples [77]. Consider the sampling scenario in Figure 2.13a, where the symmetric part of k-space is marked with a rectangle.

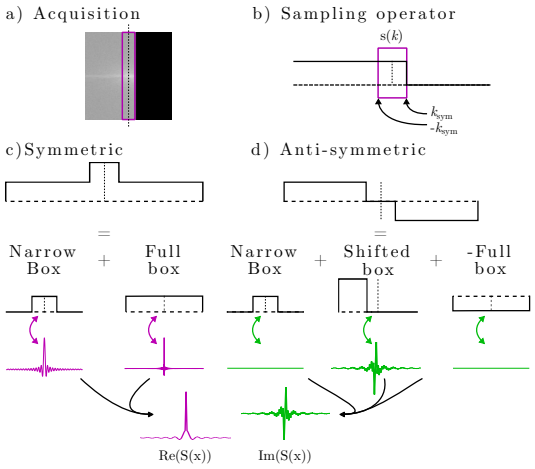

Figure 2.13: Partial Fourier acquisition. Magenta and green arrows indicate inverse Fourier transform followed by taking the real and imaginary part, respectively. a) The acquired data is the continuous Fourier representation multiplied with the sampling operator (b). The non-acquired region in (a) is black, and the symmetric center region is outlined in magenta. The sampling operator can be split into a symmetric and anti-symmetric component, which describe the real and imaginary transfer function. c) The symmetric component over-emphasizes the symmetric region. d) Since the anti-symmetric component is non-zero, an imaginary component leaks into the image.

A two-dimensional k-space is shown, but a single-dimensional explanation is given here for simplicity, where k and x are the k- and image space indices. Let k_{sym} be the sampling cut-off index as defined in Figure 2.13b. $K(k)$ is the Fourier representation of the object, $X(x)$. The sampling operator, $s(k)$, has a sharp transition between the acquired and zero-filled region, mathematically described by the step function:

$$s(k) = \begin{cases} 1, & k \leq k_{\text{sym}} \\ 0, & k > k_{\text{sym}} \end{cases} \tag{2.11}$$

The associated point spread function is given by the inverse Fourier transform of the sampling operator, $S(x) = \mathcal{F}^{-1}(s(k))$. $S(x)$ is complex since $s(k)$ itself is asymmetric, apparent from the real and imaginary representation of $S(x)$ in Figure 2.13c-d.

The sampling operator is split into a symmetric and an anti-symmetric part in Figure 2.13c-d, which contribute to the real and imaginary component of the zero-filled image, respectively. By such decomposition, one can see that the symmetric component is essentially the result of a "conjugate-reflect-add" operation. This method is composed of taking the conjugate transpose of all acquired data around the k-space origin, and adding it to the same. However, as seen in Figure 2.13c, the symmetric region is over-emphasized and the real image will be smeared as evident by the wide point spread function.

An improved point spread function for the symmetric part can be achieved by pre-weighting data with a step filter:

$$W_{\text{step}}(k) = \begin{cases} 0, & k > k_{\text{sym}} \\ 1, & -k_{\text{sym}} \leq k \leq k_{\text{sym}} \\ 2, & k < -k_{\text{sym}} \end{cases} \tag{2.12}$$

While this does help flatten improve the symmetric component, the sharp transitions in W_{step} cause Gibbs ringing in the image. For a more benign side-lobe behaviour of the filter in the image domain, a ramp filter has been proposed instead:

$$W_{\text{ramp}}(k) = \begin{cases} 0, & k > k_{\text{sym}} \\ 1 + \frac{k}{k_{\text{sym}}}, & -k_{\text{sym}} \leq k \leq k_{\text{sym}} \\ 2, & k < -k_{\text{sym}} \end{cases} \tag{2.13}$$

This ramp does alleviate much of the Gibbs ringing, at the cost of a non-uniform sym-

metric component. Other merging filters have been proposed to reduce ringing with a narrow point spread function [51].

A low resolution image X_{sym} is extracted from the symmetrically sampled region. The merging filter is applied to the acquired data, and the phase of X_{sym} is demodulated from the filtered image.

$$X_{\text{homodyne}} = \text{Re}\left(\underbrace{e^{-i\angle X_{\text{sym}}}}_{\text{low-res phase},\theta_L} \mathcal{F}^{-1}\left(\underbrace{sKW_{\text{ramp}}}_{\text{pre-weighted data}} \right) \right), \tag{2.14}$$

where $\angle$ is the argument (angle) operator. The removal of θ_L in Eq. 2.14 should, ideally, result in no imaginary residue and the real operator in Eq. 2.14 should not discard any signal. Note that homodyning effectively removes any phase, even low frequency components that are reliable from the original data. Therefore, it cannot be directly combined with fat/water imaging which relies on CSE in the signal phase. Instead, θ_L can be added to X_{homodyne}, although this is unnecessary for conventional magnitude display of source images.

2.6.2 Projection Onto Convex Sets (POCS)

In contrast to homodyning where the signal is rectified to the real channel, POCS is an iterative method aimed to synthesize data in agreement with the acquired sample set [78]. For each iteration, data consistency is enforced such that the phase of the estimated image is in line with the low-resolution symmetric region.

An overview of the POCS algorithm is shown in Figure 2.14. Initially, the low-resolution phase map $\theta_L = e^{-i\angle X_{\text{sym}}}$ is extracted, and the sampled k-space is zero-padded in the non-sampled region. The iterations consist of an inverse Fourier transform, followed by the magnitude operator to retrieve $|X_i|$, where i is the iteration index. The low-resolution phase map is applied to $|X_i|$ and forward transformed. To maintain agreement with the acquired data, only the non-sampled region of this intermediate matrix is extracted and assigned to K_i. Upon converging, which is determined by a threshold of change between iterations or after a fixed number of iterations, there might be discontinuities near k_{sym} which cause rippling in the image. These discontinuities can for instance appear with

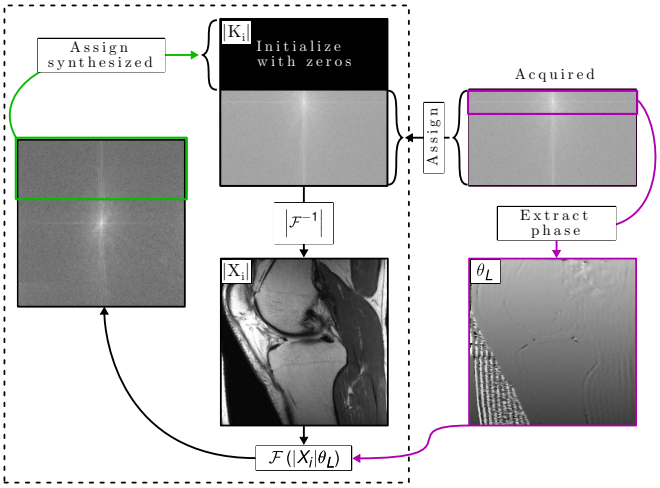

Figure 2.14: The POCS algorithm. Before iterating, the low-resolution phase map θ_L is extracted from the symmetrical region (magenta) of the acquired data, and a zero-padded matrix K_0 is created from all acquired data. The i:th iteration starts by transforming K_i to image space, followed by the magnitude operator, yielding X_i. The data-consistent phase θ_L is added to X_i prior to returning to k-space. In k-space, the non-sampled region (green) now contains synthesized data, which replaces the previous estimate (i.e. zeroes in the initial case). The iterative steps (dashed) are repeated, each time with a smoother transition between synthesized and acquired data, meaning that $|X_i|$ is less affected by Gibbs ringing as i increases.

RARE, where sectional k-space differences are present due to multiple echo times in the acquisition. Similar to homodyne, a transition filter such as W_{ramp} is proposed to reduce this error.

McGibney et al. compared several partial Fourier techniques, and concluded that if a true phase estimate is provided, i.e. the symmetric region is sufficiently large, POCS converges to the exact true image [79]. However, for CSE imaging such as an opposed phase echo, the true phase map often contains *all* spatial frequencies. To overcome this limitation, we proposed a new technique in paper II that allows POCS reconstruction from partial Fourier acquisitions in combination with model based fat/water separation. The technique was used in paper III to improve the sampling conditionality in a two-

point RARE sequence with dual bandwidths.

3 Project Aims

The primary aim of this thesis was to optimize the acquisition strategy for fat/water separated MRI. Modifications to the reconstruction routine, FWQPBO [43], were performed to suit the new sampling methods.

3.1 Study I - Fat/Water Separated Spin Echo PROPELLER

A T1 weighted spin echo sequence had been internally developed, where we noticed a motion estimation bias from the chemical shift displacement artifact. Schär et al. had recently published a PROPELLER sequence that circumvented this by separating water and fat blade-wise, followed by motion estimation on the water image. We hypothesized that the dead times from shifted readouts could be removed by acquiring the echoes with dual bandwidths, thereby increasing SNR. The use of dual bandwidths required regularization of the inverse problem, where the method by Lu et al. was improved to whiten noise [44].

3.2 Study II - Fat/Water Separation in k-Space with Real-valued Estimates and its Combination with POCS

A complete coverage of k-space along the readout dimension limits the resolution and gradient echo spacing in fat/water separated sequences, as the dephasing time difference

for well-conditioned sampling can be very short at high field strengths. Moreover, real-valued estimation in k-space would allow better conditioning of the inverse problem as the conjugate sample pairs in general have different CSE. We hypothesized that the underlying k-space symmetry could be enforced in the estimates rather than the signal in a modified POCS technique to allow reconstruction of fat/water estimates from partial Fourier acquisitions.

3.3 Study III - RARE Two-point Dixon with Dual Bandwidths

The sequence described in Paper I is only applicable to spin echo imaging due to CPMG violations. The vendor supplied fTED implementation available at our scanners is based on RARE. It is an attractive single-TR method as it does not prolong scan time, but is restrictive in terms of resolution. We noticed that this is mainly due to the requirement of sampling three echoes and a significant time is spent on gradient switching, which we hypothesized could be reduced by acquiring only two echoes. The dual bandwidth concept could then be used to avoid ill-posed sampling. Not only would this enable higher resolutions but also put less stress on the gradient hardware. This study highly benefited from Paper II by allowing an additional degree of freedom in the search for optimal gradient waveforms.

3.4 Study IV - Chemical Shift Encoding using Asymmetric Readout Waveforms

With Paper III successfully implemented, we sought to improve the multi-TR fat/water RARE sequence by removing the dead times. We hypothesized that a generalized concept of dual bandwidths, i.e. dynamic bandwidths could be used. This relied on the realization that sampling chemical shift encoding requires asymmetry around the spin echo, which we hypothesized could be achieved using asymmetric rather than shifted readout waveforms. This would enable significant reduction in scan time or increased SNR, and reduced T2 blurring.

4 Methodological Considerations

The pulse sequences in this thesis have all been developed in-house using the KS Foundation framework [80]. Due to technical restrictions, it was not possible to completely take control of the signal receive chain. Instead, the sequence proposed in Paper III and IV employed acquisition windows with the maximum sampling rate (2 μs) and as wide analog filtering as possible. Data was resampled data on a Cartesian grid during reconstruction. In paper I, this was not a concern as there was no ramp sampling involved and online filtering could be used.

An illustration of the different acquisition strategies proposed in this thesis is given in Figure 4.1. The background shows the single-species approximated NSA, derived in Paper III, and typical dephasing time pairs of each sequence is overlaid. These trajectories show the dephasing time pair of each sample, with a star indicating the k-space center acquisition, which are acquired in- and opposed phase. The dephasing angle is used as a surrogate for the dephasing time to avoid a field strength dependency. Single bandwidth sequences (not shown) are restricted to be parallel with the diagonal of this plot, with the offset determined by the shift and dead times cause a line shortening. Dual bandwidth sequences alter the slope, as seen by the magenta line which show the acquisition of Paper I. The dBW-RARE sequence in Paper III is bipolar, causing the slope of the yellow line in Figure 4.1 to be negative. Paper II describes fat/water estimation from partial Fourier data, which corresponds to a shift of the k-space center along the yellow line. Finally, the white line shows that with asymmetric readouts, the trajectory becomes curved.

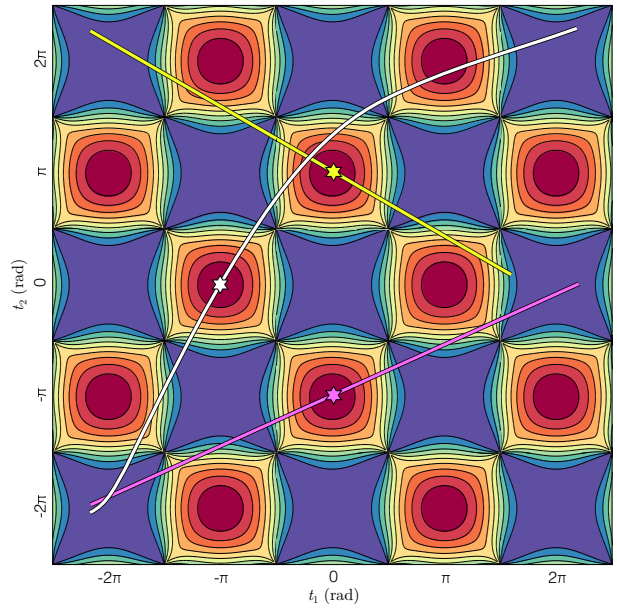

Figure 4.1: Single-species approximated NSA in pseudo color, ranging from zero (blue) to max (red). Dephasing time pairs of the proposed acquisition strategies are overlaid as lines with a star indicating the k-space center. The spin echo dual bandwidth sequence in Paper I is magenta colored. The dBW-RARE sequence (Paper III) is shown in yellow, with partial Fourier sampling shifting the k-space center from the center of the line (Paper II). Dual bandwidths change the slope of the line. Asymmetric readouts, described in Paper IV, are acquired with dynamic bandwidths, causing the corresponding white trajectory to become curved.

4.1 Spin Echo PROPELLER with Dual Bandwidths

The use of dual bandwidths allows the same k-space coverage with different sampling duration. This varying k-space traversal speed is flexible as it allows sampling to occur when single bandwidth sequences must wait due to timing reasons. Eggers successfully removed dead times from the in-phase echo in a shifted RARE sequence by stretching the gradient to cover the full available acquisition time, but dead times remained in the CSE echo in order to comply with CPMG conditions [81].

In Paper I, this concept was further investigated in a PROPELLER sequence, where

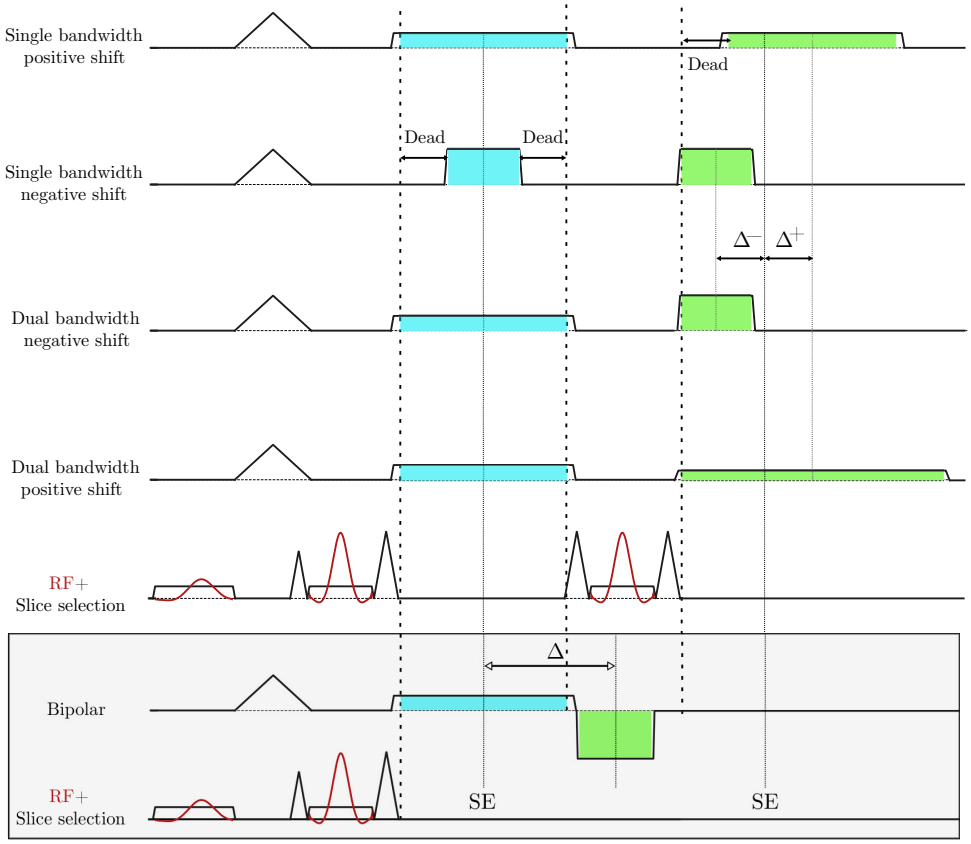

Figure 4.2: Acquisition strategies proposed for the fat/water spin echo PROPELLER sequence in Paper I. The bipolar acquisition mode shown in grey has only one refocusing pulse. Cyan and green patches indicate sampling of the first and second echo, respectively. Δ marks the dephasing times of the echoes, acquired before (minus) or after (plus) the second spin echo.

we completely removed dead times. Paper I investigates multiple ways of acquiring the second echo, either with refocusing pulses where the readout is acquired before ($-$) or after ($+$) the spin echo, or by adding a readout gradient with opposite polarity. These techniques are shown in Figure 4.2. A benefit of a flexible readout duration is that one can opt for a faster bipolar acquisition, or if sufficient time is available, an additional refocusing pulse can increase the sampling conditioning. In other words, combining spin echo with dual bandwidths can increase the SNR for a given protocol for a given

repetition time and number of slices.

Fat/water separation was performed blade-wise as suggested by Schär et al. [75], which allows chemical shift displacement correction by performing the separation in k-space before gridding. Motion estimation in image space is often driven by edge alignment, which introduces a blade-angle dependent bias as subcutaneous fat surrounds the head. In fact, this fat delineation itself is well suited to estimate motion from using fat navigators [82]–[84]. Using water images rather than conventional in-phase images, we investigated the blade-angle dependent bias in from chemical shift displacement and compared it to the theoretical displacement given the bandwidth and field of view. The results showed that a peak bias of 1.3 mm was present with rBW=±25 kHz, which was completely removed when using the water image for motion correction.

Dual bandwidths in combination with k-space separation where the specific CSE for each sample pair is modelled resulted in large noise artifacts, a consequence of ill-conditioned model matrices $\mathbf{A}$ for some sample pairs, which we solved using noise whitening regularization. An example of the effect on the water estimate is shown in Figure 4.3, with the proposed solution described in the section below.

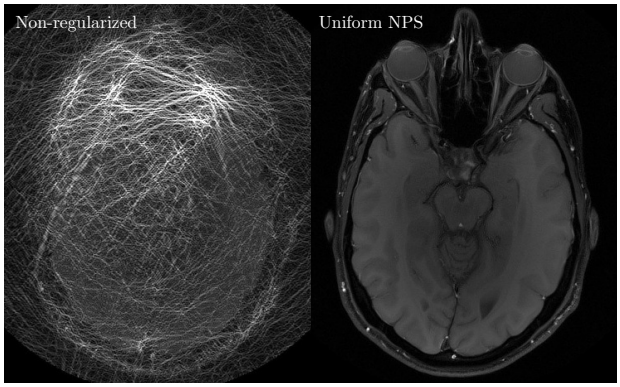

Figure 4.3: Effect of noise-whitening when performing k-space separation on a dataset where some sample pairs are ill-conditioned. Data was acquired with a PROPELLER sequence.

4.2 Regularization

4.2.1 Noise Propagation

In a linear inverse problem such as

$$\hat{\mathbf{x}} = \mathbf{E}\mathbf{S}, \tag{4.1}$$

the propagation of uncertainty from $\mathbf{S}$ to $\hat{\mathbf{x}}$ is given by

$$\mathbf{\Sigma}_{\hat{\mathbf{x}}} = \mathbf{E}\mathbf{\Sigma}_S(\mathbf{E})^H, \tag{4.2}$$

where $\mathbf{\Sigma_S}$ is the covariance matrix of $\mathbf{S}$ and $\mathbf{E}$ is the estimator (e.g. the model matrix inverse or regularized pseudo-inverse). Since the samples in $\mathbf{S}$ are independent and identically distributed, $\mathbf{\Sigma_S} = \sigma_s^2\mathbf{I}$, where σ_s^2 is the input variance. Noise can be amplified differently in the water and fat estimates, described for instance by the ratio between the water and input variance σ_W^2/σ_s^2. The noise amplification factor is defined as the average ratio over both estimates, which are the diagonal elements of $\mathbf{E}\mathbf{E}^H$ and can be expressed as the trace of this product:

$$\eta = \frac{N}{2}\sum_{m\in\{W,F\}}\frac{\sigma_m^2}{\sigma_s^2} = \frac{N}{2}\,\mathrm{tr}\left(\mathbf{E}\mathbf{E}^H\right) \tag{4.3}$$

From Eq. 4.3 we can evaluate the noise performance for any linear estimator $\mathbf{E}$. For instance, if the model matrix inverse is used, i.e. $\mathbf{E} = \mathbf{A}^{-1}$, and $N = 2$ echoes are acquired, η can be plotted with an encoding where the first echo is in-phase (as in Paper I), and the other echo varies. This is shown in Figure 4.4, solid line.

The ill-posed sampling case near the edges of Figure 4.4 will cause large noise propagation in the estimates. While sampling with ill-conditioned dephasing times at the k-space center should always be avoided, some dephasing time pairs are often poorly conditioned, in particular if the readout gradients are bipolar. When the fat/water separation is performed in k-space, these dephasing times are incorporated in the model matrix and local noise propagation can ruin the entire image, as evident in Figure 4.3. Lu et

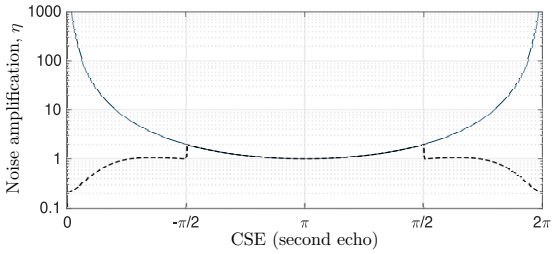

Figure 4.4: Noise amplification factor η evaluated for a two-point acquisition plotted against varying CSE where one echo is acquired in-phase without (solid) and with (dashed, according to [44]) regularization. In- and opposed phase sampling minimizes η. When both echoes are in-phase, η approaches infinity when no regularization is used. Note the logarithmic y-axis.

al. described this problem and proposed Tikhonov regularization of the model matrix pseudoinverse to dampen ill-conditioned frequencies [44].

$$\mathbf{A}^{\ddagger} = (\mathbf{A}^H\mathbf{A} + \kappa\mathbf{I})^{-1}\mathbf{A}^H \tag{4.4}$$

With κ chosen according to Eq. 4.5:

$$\kappa = \begin{cases} 0.3 & \eta > 2 \\ 0 & \text{else} \end{cases} \tag{4.5}$$

By evaluating η with the proposed regularization scheme, shown as the dashed line in Figure 4.4, we can see that the ill-conditioned sample pairs are suppressed. However, the noise is unevenly distributed, and it is clear from Figure 4.4 that these sample pairs are overly suppressed.

4.2.2 Choosing κ

As the encoding matrix $\mathbf{A}$ varies over k-space indices, so does the noise amplification factor. Figure 4.5a shows a typical bipolar acquisition at 3T, where the CSE is plotted for each sample pair. A plot of the noise amplification factor over the sampled spatial

frequencies make up the *noise power spectrum*, NPS. In Paper I, κ was chosen to achieve a uniform NPS, which was derived as follows.

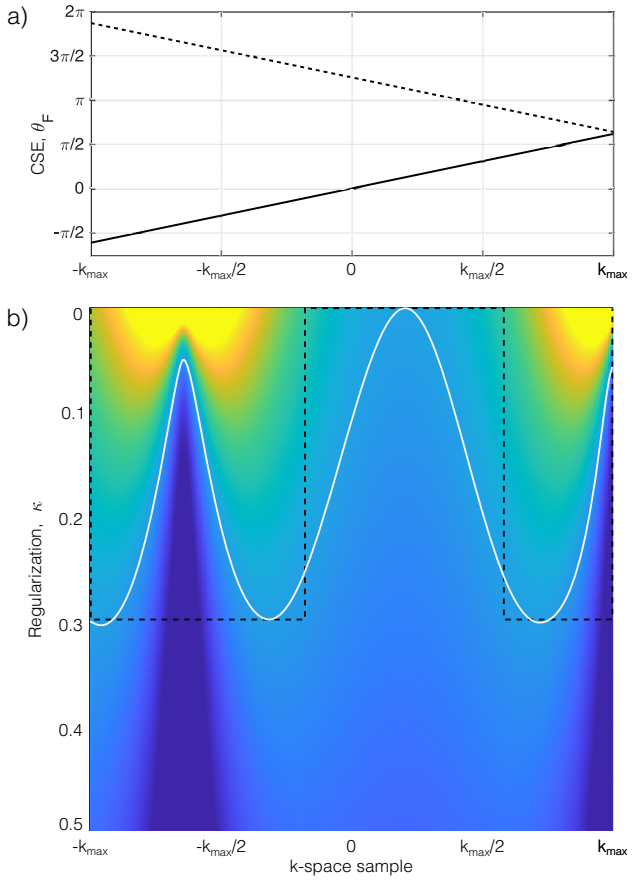

Figure 4.5: a) CSE for a bipolar acquisition with rBW=±90 kHz, 256 samples, dephasing times at k-space center are 0 and 1.45 ms. b) NPS for varying κ. The top row is the non-regularized case where the estimator is $\mathbf{A}^{-1}$. The derived κ follows a contour of the NPS (white line). Regularization according to Lu et al. [44] is shown for comparison (black dashed line).

The relationship between the singular values of $\mathbf{A}$ and $\mathbf{A}^{\ddagger}$ is

$$\sigma_i^{\ddagger} = \frac{\sigma_i}{\sigma_i^2 + \kappa}.$$

(4.6)

The singular value decomposition of $\mathbf{A}^{\ddagger} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H$, so

$$\mathbf{A}^{\ddagger}\mathbf{A}^{\ddagger H} = \mathbf{U}\mathbf{\Sigma}^H\mathbf{\Sigma}\mathbf{U}^H, \tag{4.7}$$

where $\mathbf{V}$ is unitary, i.e. $\mathbf{V}^H\mathbf{V} = \mathbf{I}$. The right hand side of Eq. 4.7 is the eigendecomposition of $\mathbf{A}^{\ddagger}\mathbf{A}^{\ddagger H}$, which proves that the eigenvalues of $\mathbf{A}^{\ddagger}\mathbf{A}^{\ddagger H}$ are the squared singular values of $\mathbf{A}^{\ddagger}$. This result is combined with Eq. 4.3 where we use the fact that the trace of a matrix is also the sum if its eigenvalues, which gives the following expression for water and fat estimation:

$$\eta = \left(\frac{\sigma_W}{\sigma_W^2 + \kappa}\right)^2 + \left(\frac{\sigma_F}{\sigma_F^2 + \kappa}\right)^2 \tag{4.8}$$

Rearranging Eq. 4.8 for κ with $a = \sigma_W^2$ and $b = \sigma_F^2$ (given by the singular value decomposition of $\mathbf{A}$) gives

$$\begin{aligned} 0 \;=\; & ab(ab - \tfrac{a+b}{\eta}) \\ & +2\kappa ab(a + b - 2/\eta) \\ & +\kappa^2(a^2 + 4ab + b^2 - \tfrac{a+b}{\eta}) \\ & +2\kappa^3(a + b) \\ & +\kappa^4. \end{aligned} \tag{4.9}$$

The quartic polynomial above can be solved for a desired η by finding the eigenvalues of the associated companion matrix

$$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ c_0 & c_1 & c_2 & c_3 \end{bmatrix} \tag{4.10}$$

where c_n is the n:th order coefficient in Eq. 4.9. The noise amplification factor, η, was chosen according to the best-conditioned sample pair, i.e. leaving one sample pair unregularized. The derived κ levels the NPS, as seen in Figure 4.5b. Even though Tikhonov matrices other than the identity matrix can be motivated, we found the proposed regularization scheme to yield high quality images even with poorly conditioned data.

4.2.3 Signal Loss

The downside of regularization is an introduced bias in the estimator. This can be described with transfer functions, where one metric is the loss in amplitude, or the *modulation transfer function*, MTF. In a linear model $\mathbf{S} = \mathbf{Ax}$, which is solved according to Eq. 4.1, the MTF is contained in the diagonal of the transformation $\mathbf{EA}$.

For instance, in the non-regularized case where $\mathbf{E} = \mathbf{A}^{-1}$, this transformation is the identity matrix, i.e. optimal MTF as long as $\mathbf{A}$ is invertible. With $\mathbf{E} = \mathbf{A}\ddagger$ with κ chosen to equalize NPS, $\mathbf{EA}$ is in general a non-diagonal matrix. The off-diagonal elements describe the signal leakage due to introduced estimator bias. Using the proposed NPS-flattening κ, a plot of the MTF of water and fat together with the signal leakage is shown in Figure 4.6.

4.3 Separation Domain

After seeing the significant problem with colored noise in Figure 4.3a, it may seem like a lot of trouble to retrieve an image quality similar to what is expected from solving the inverse problem in image space rather than in k-space. While it is more common to separate in image space, this is arguably related to the additional complexity of k-space separation and associated regularization, rather than being more robust. There

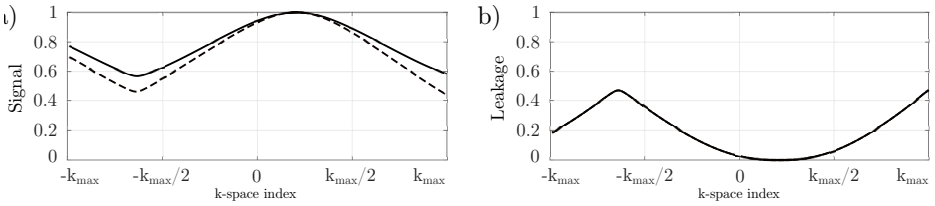

Figure 4.6: Transfer functions from the proposed regularization with a multi-peak fat model. The sampling is shown in Fig. 4.5. a) Modulation transfer function (MTF) for water (solid) and fat (dashed). b) Signal leakage, same for water and fat.

are several aspects to keep in mind in the decision of separation domain.

First, separation in image space does not offer the ability to model each sample point according to their known encoding. This includes not only dephasing times, but also sample weights, which varies significantly in the suggested dynamic bandwidth approach in Paper IV.

Second, there is no possibility for full correction of chemical shift displacement with image space separation, as the artifact is related to the dephasing time span in k-space (see section 2.1.1, page 4). While it is possible to retrospectively shift the fat in the image domain, this is not trivial in combination with multi-peak modeling and can cause resolution losses from the resampling procedure as, in general, the shift is not divisible by the voxel size.

Finally, one might be led to the conclusion that signal leakage and non-uniform MTF, discussed in section 4.2.3, is a problem which only appears when performing the separation in k-space. However, k-space separation and image-space separation are equivalent if the estimators are the same, since the Fourier transform is a linear operator acting on a linear system. Image-space separation could therefore, perhaps more accurately, be described as a constant estimator method. Noise-whitening follows from keeping $\mathbf{E}$ the same, given by Eq. 4.3 on page 45. The benefits of additional modelling should therefore not be considered a downside as one can always resort to the ”image space method” for any sample by choosing the estimator as such.

In conventional image space separation, the estimator is the model matrix (pseudo-)inverse with dephasing times at k-space center acquisition. Since the encoding of each

sample k is contained in the model matrix $\mathbf{A}_k$, it is interesting to investigate the bias from an image space estimator $\mathbf{A}_{\text{im}}^{-1}$ on the encoding presented in Figure 4.5. With image space separation $\hat{\mathbf{x}} = \mathbf{A}_{\text{im}}^{-1}\mathbf{S} = \mathbf{A}_{\text{im}}^{-1}\mathbf{A}_k\mathbf{x}$. Similar to what is described in section 4.2.3, the MTF can be extracted from $\mathbf{A}_{\text{im}}^{-1}\mathbf{A}_k$. The MTF calculation and leakage plots are shown in Figure 4.7. Although the MTF and NPS is perfect in the water estimate, there is significant loss in the fat signal. Additionally, the fat signal leaks considerably to the water image.

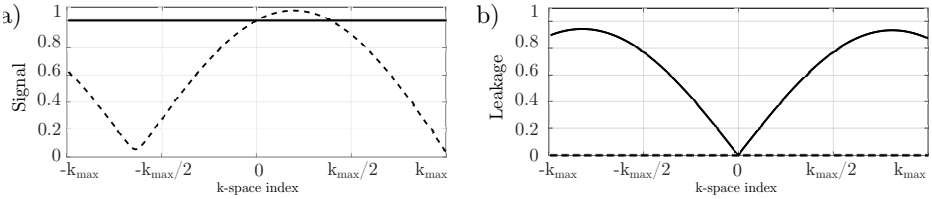

Figure 4.7: Transfer functions for an image space estimator with the sampling from Fig. 4.5. a) Modulation transfer function (MTF) for water (solid) and fat (dashed). b) Signal leakage. Solid line shows the fat leakage ending up in the water estimate.

For a more intuitive understanding of these results, an opposed-phase image that has been retrospectively filtered according the fat transfer functions in Figure 4.6 and 4.7 is shown in Figure 4.8. Note that signal leakage is not accounted for in these experiments.

To summarize, k-space separation is not is associated with more loss in image quality, but rather enables an opportunity to trade off noise behavior with resolution.

4.4 RARE with Dual Bandwidths

RARE sequences are restrictive in their pulse sequence designs due to CPMG conditions [49], [50]. Most challenging in the sequences developed in this work is the need for constant zero-order magnetic moment, or gradient area, between refocusing pulses. With any odd number of bipolar readouts, the gradient area is the same as in the conventional single-echo case, and no further modifications are necessary. Since we aimed to acquire only two echoes, which causes a net area of zero, the gradient pre- and rephasers were moved adjacent to the bipolar readout block. In terms of k-space trajectory, this means that it starts from the origin after every refocusing point, goes out, acquires two echoes

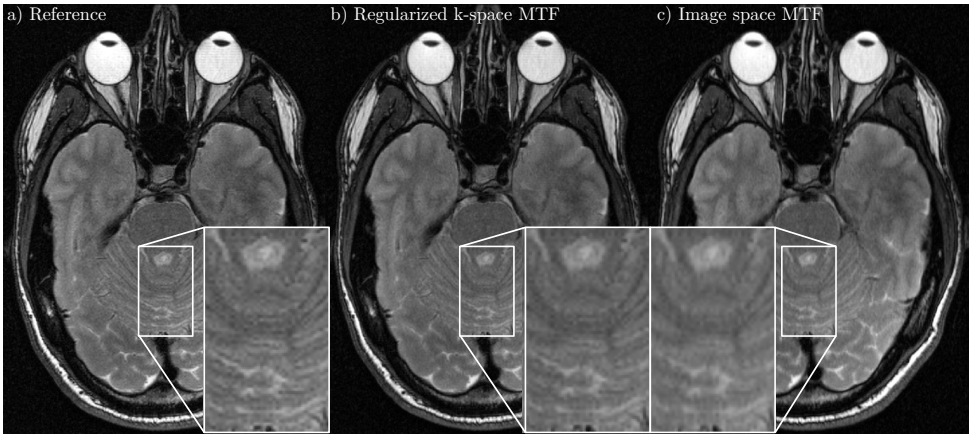

Figure 4.8: Smoothing effect from non-unity MTF. a) Reference image. b) Filtered according to the fat MTF in Figure 4.6, corresponding to a regularized k-space estimator. c) Filtered according to the fat MTF in Figure 4.7, corresponding to an image space estimator.

back-to-back, and returns to the center. While it may seem like this would incur dead times, the pre- and rephasers have time to be played out during crusher lobes and phase encoding gradients.

The fTED sequence has a constant NSA of 8/9 when acquired in- and opposed phase [55]. For an efficiency comparison between dBW-RARE and fTED, we used normalized NSA under the single-species approximation (Paper III) where the time spent sampling is taken into account as the efficiency metric. As three gradient lobes are played out with fTED, gradient ramping can occupy the entire available acquisition time, which restricts the spatial resolution, as seen in Figure 4.9. From this result, it is obvious that fTED is restricted by gradient hardware, and even state-of-the-art gradient systems struggles to achieve voxel sizes lower than 1 mm.

The use of dual bandwidths in the proposed sequence breaks the redundant sampling associated with a naive approach of single-bandwidth bipolar dual-echo RARE (Figure 4.10a). In other words, the use of dual bandwidths allows a flexibility when selecting echo shifts. The search space is so far one-dimensional as the bandwidth ratio determines the waveform for a given duration. It turned out that the use of dual bandwidths alone often

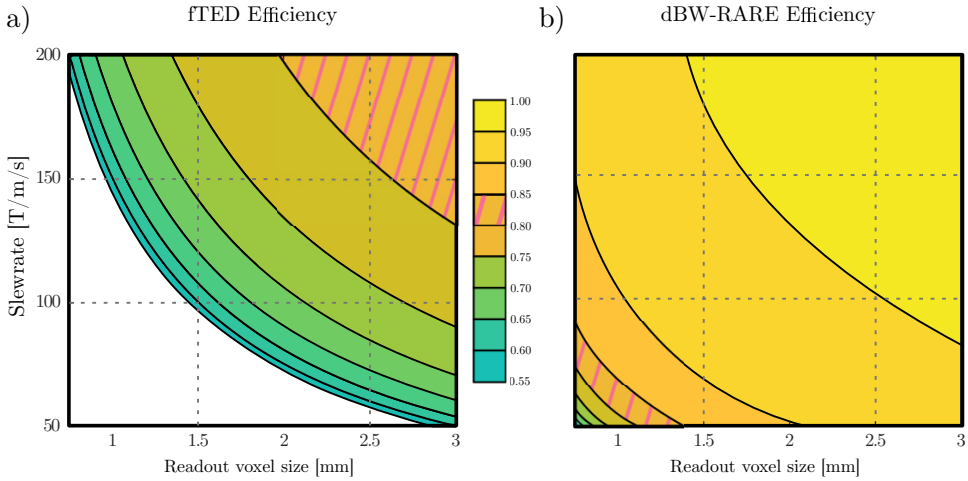

Figure 4.9: Efficiency comparison between fTED (a) and dBW-RARE (b). Partial Fourier sampling was used with dBW-RARE. The white region is non-feasible as ramps occupy all available acquisition time.

does not allow the desirable in- and opposed phase sampling as the readout durations were too long for spatial resolutions and RF echo spacings of clinical relevance. In order to shorten the gradient echo spacing, the use of partial Fourier sampling along the readout axis enables an additional search dimension. This higher flexibility in optimizing the gradient waveforms is evident from Figure 4.10c.

Consequently, the problem of separating fat/water images from partial Fourier data arises, which was investigated in Paper II. In addition to partial Fourier reconstruction, the use of real-valued two-point separation exploiting Hermitian symmetry had recently been described [85], which was further explored in Paper II.

4.5 Partial Fourier

As described in section 2.6, partial Fourier sampling is not well suited for CSE data, as the underlying real-valued assumption does not hold. The POCS algorithm has been shown to outperform homodyning [79], and while the latter had been adapted for fat/water imaging by Reeder et al. [86], we found no such modification to POCS in

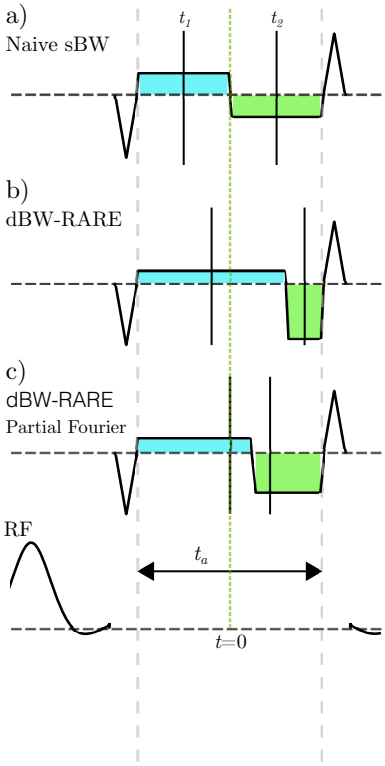

Figure 4.10: dBW-RARE pulse sequence diagram showing the readout board for three different readout couples in one RF echo spacing interval. a) Single bandwidth is badly conditioned, evident from the encoding redundancy given by $t_1 = -t_2$. b) Dual bandwidths break the symmetry and makes fat/water separation possible with only two echoes. c) Partial Fourier sampling improves the sampling by allowing the time of the gradient echo to vary from the center of the readout.

our literature search. In Paper II, we perform projections similar to the description in section 2.6.2, but with modifications to project the real-valued property on $\hat{\mathbf{x}}$ rather than $\mathbf{S}$.

The modified POCS algorithm is illustrated in Figure 4.11. It starts by extracting the symmetrically sampled k-space region, from which low-resolution field inhomogeneity and initial phase maps ψ and ϕ are estimated. The undersampled data is zero-filled to form the initial guess K_0, where the suffix indicates the iteration count. The i:th

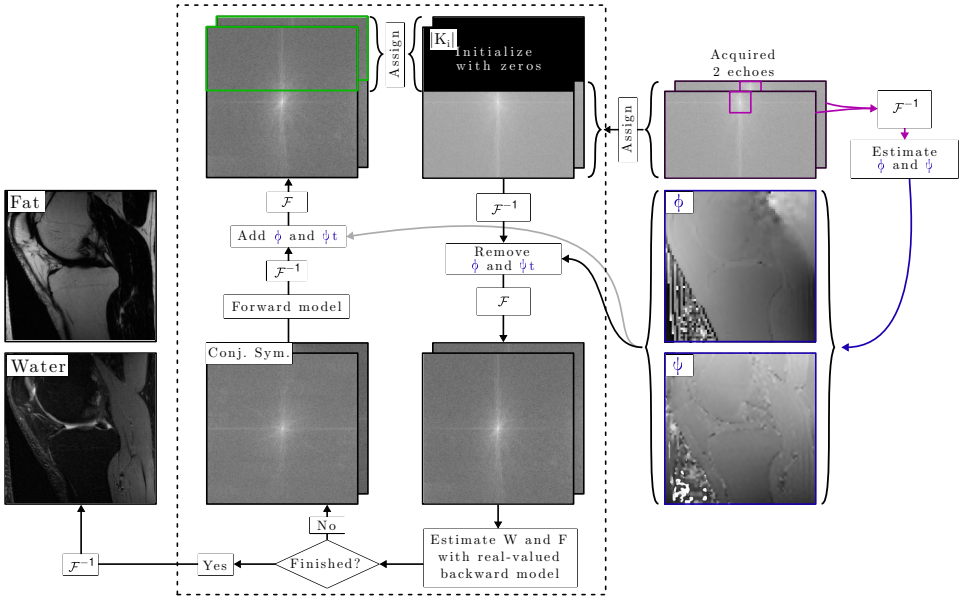

Figure 4.11: The modified POCS algorithm with real-valued fat/water estimates, starting from the rightmost column. The iteration loop is framed with a dashed line.

iteration starts by transforming K_i to image space, followed by demodulation of ϕ and ψ, yielding data were any remaining phase is sourced from chemical shift (bottom right in the dashed region). Water and fat are estimated as real-valued by applying the backward model to conjugate sample pairs (section 4.6) in k-space, forming a conjugate symmetric k-space of the estimates, W and F (bottom left in the dashed region). If the predetermined number of iterations or a convergence threshold has been reached, the process stops, yielding the estimated water and fat images. Otherwise, a signal is synthesized by applying the forward model, followed by re-adding the phase terms to both echoes in image space, according to the respective dephasing times. The synthesized data is Fourier transformed in order to extract the missing k-space samples, which is assigned to K_{i+1} with a transition filter.

Note that the estimated fat/water images from the initial iteration describes how real-valued estimation is performed in the fully sampled case, used e.g. in Paper IV. It is essentially the leftmost column in the dashed region together with the iterative scheme that constitutes the modified POCS method.

4.6 Real-valued Estimates

The two-point model in Eq. 2.6 assumes that fat and water share the time-invariant phase ϕ. This results in an exactly determined inverse problem of each sample. Real-valued signals possess a Hermitian symmetry in k-space, meaning that under the real-valued assumption on $\mathbf{x}$, only half of k-space need to be fat/water separated.

Under the common assumption that the spatial distributions of ϕ and ψ are smooth, the real-valued property of $\mathbf{x}$ give rise to an *over-determined* system after these phase contributions have been demodulated. The signal equation can then be re-written by stacking the model matrices $\mathbf{A}$ of sample index k with that of its conjugate sample index c, while also stacking the ϕ and ψ demodulated signals $\mathbf{S}' = \mathbf{S}^{\phi,\psi}$.

$$\begin{bmatrix} \mathbf{S}'_k \\ \mathbf{S}'^{*}_c \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{A}_k \\ \mathbf{A}^{*}_c \end{bmatrix}}_{\mathbf{A}_{\text{real}}} \mathbf{x}_k \tag{4.11}$$

Note the complex conjugate operator $*$ in Eq. 4.11. The over-determined inverse problem associated with Equation 4.11 can be solved either with the pseudoinverse of $\mathbf{A}_{\text{real}}$, or with a regularized pseudoinverse as estimator $\mathbf{E}$, as described in section 4.2.

Even in the simplest form with the estimator being the non-regularized pseudoinverse of the stacked model matrices, $\mathbf{A}^{+}_{\text{real}}$, the varying encoding between conjugate sample pairs tend to improve the sampling conditionality significantly. In Figure 4.12, the same sampling shown in Figure 4.5 was analyzed in terms of NPS without regularization using Eq. 4.3. As can be seen, the use of Hermitian samples alone prevents the strong noise amplification near singular sample pairs. Real-valued estimates in combination with regularization is suitable to achieve white noise without sacrificing image sharpness while keeping signal leakage small.

4.7 Asymmetric Readouts

The proposed acquisition strategies in Paper I and III alternate the bandwidth between the two readout gradients to achieve better encoding. This concept was generalized in Paper IV where chemical shift encoding is achieved by shifting the "center of gravity" of

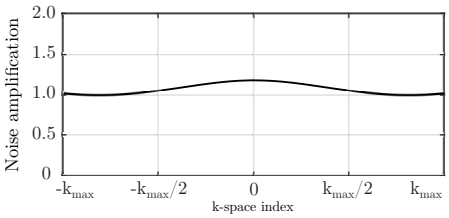

Figure 4.12: Noise power spectrum for the real-valued estimation of the sampling shown in Figure 4.5. Note the linear scale. While the NPS is rather flat, it is slightly larger in the center, which can be explained by the small CSE difference between these samples so that the conjugate sample pair does not contribute with complementary information. The singularities in Figure 4.5 are avoided.

the gradient. Alternatively, this can be viewed as a generalization of the dual-bandwidth approach where a dynamic rather than a static intra-echo bandwidth is used. The use of asymmetric readout waveforms comes with several benefits in pulse sequence design.

First, chemical shift encoding can be achieved without increasing the temporal footprint of the gradient waveform, since the CSE is achieved by skewing the waveform rather than shifting it. Consequently, no further sequence modifications are necessary in order to, for instance, change an in-phase echo to an opposed phase echo in an otherwise conventional spin echo sequence.

Second, asymmetric readout can be be used to ”fill the gap” in conventional gradient echo imaging. At least theoretically this will improve the SNR as more time is spent sampling the signal without prolonging the scan time. This way, dead times can be avoided while still achieving the desired echo time at the center of k-space. Further applications of this could include prolonged sampling duration in susceptibility weighted imaging, filling the rather long dead time required to achieve the desired $T2^*$ dephasing [68], although this has not been investigated. It should be mentioned that there are other ways of filling the dead time by sampling more lines with an EPI trajectory [69].

Finally, in sequences with strict timing restrictions such as RARE, the use of asymmetric instead of shifted trapezoidal readouts result in no dead times. This was investigated in Paper IV where we showed that this can shorten acquisition times or sample more efficiently to the benefit of SNR.

4.7.1 Resampling

With the benefits comes the shortcoming of non-uniformly sampled k-space along the readout dimension. While this is the case also between echoes for dual bandwidths, asymmetric readouts has additional non-uniformity *within* the echoes. Consider the readout waveform shown in Figure 4.13, which also displays the associated k-space trajectory over time. The trajectory is calculated from the cumulative sum of the waveform. The center of mass along the time dimension is shifted such that the center of k-space (vertical dashed line) is acquired with the desired CSE. Consequently, less time is spent on one side of the gradient echo compared to the other.

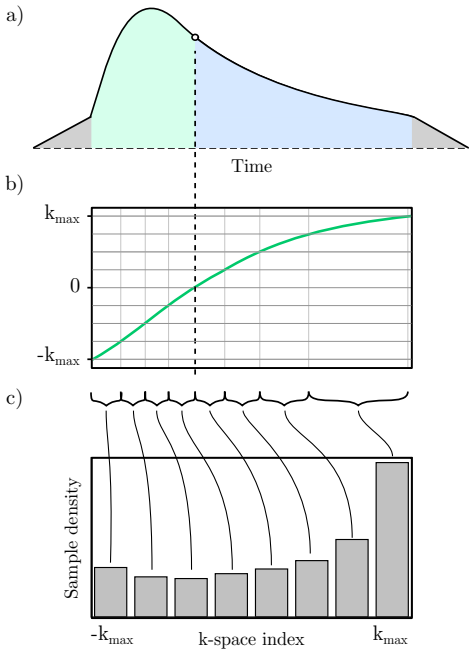

Figure 4.13: a) Asymmetric readout waveform. The vertical dashed line marks the time when the k-space center is acquired, which appear before half of the waveform is played out. The gray areas are used for prephasing and the colored areas each correspond to half of the sampled k-space area. b) k-Space trajectory, calculated from the cumulative sum of the waveform. c) Sample density for eight gridded points, calculated from the residence time (or effective dwell time) of the k-space trajectory.

Note that the dwell time of each *acquired* sample is kept constant, in this work at 2 μs as it was the shortest possible. This means that the acquisition bandwidth, i.e. the inverse of the ADC dwell time, does not match the FOV bandwidth, i.e. the frequency range of spins caused by the imaging gradients. Since the time spent within each to-be-gridded k-space region (i.e. the *effective* dwell time) varies, the local *sample density* should also be accounted for in the reconstruction.

A sample density plot can be derived from the trajectory (Fig. 4.13b) by sectioning the y-axis according to the desired number of samples after re-gridding, and measuring the duration of the trajectory within each segment, as shown in Figure 4.13c for eight samples after gridding. The sample density is also inversely related to the k-space velocity, governed by the gradient amplitude.

With a time-varying bandwidth, the imaged object is still band-limited, so the box function describes the desired truncation of the frequency range that achieves sharp field of view edges (i.e. no apodization). However, the box width changes. The Fourier transform of the box function is the sinc function, which has full support, meaning there is a non-zero contribution of every acquired sample to every gridded point. As the bandwidth changes dynamically, the sinc kernel must adapt accordingly following the inverse relationship between the width of the sinc and the box truncation window.

Interpolation with sinc kernels is generally not computationally feasible in 2D non-Cartesian imaging such as PROPELLER or spiral imaging, and kernels with compact support are often used instead. Truncation of sinc causes Gibbs ringing in the image domain and so other shapes with smaller side-lobes are preferred, such as the ubiquitous Kaiser-Bessel kernel which strikes a good balance between computational speed and ringing artifacts [87]–[89]. Using asymmetric readouts in (otherwise) Cartesian sampling as in Paper IV, the non-uniform sampling problem is one-dimensional, making it feasible to use sinc kernels with full support. More specifically, a sinc function is centered around each acquired sample index [90], meaning each sample will contribute equally to the resampled data. However, this will result in an uneven density after gridding and so *density compensation* is necessary. This can either be performed in a separate step after resampling, or the sinc kernels can be scaled according to the local sample density.

If no scaling is applied, neither the signal nor noise is levelled across k-space. Alterna-

tively, the gridded signal can be scaled according to the inverse of the effective dwell time, making each point in the gridded space receive equal signal contributions. However, neither of these two alternatives keep the *noise* level equal, which can be enforced by filtering the gridded signal with the inverse square root of the effective dwell time. The reason why this operation equalizes noise is described in section 2.5.1, page 22.

A resampling matrix describes the transformation from acquired to gridded space. It is formed by stacking the sinc kernels of each acquired sample, with examples shown in Figure 4.14 for the asymmetric waveform in Figure 4.13, where 2300 samples were acquired, and a target grid size of 256 was desired. For instance, the middle column (blue) describes how each sample contributes to the gridded k-space center, while the middle row (green) describes the contribution of the single sample acquired at mid-point of the waveform, as plotted in Figure 4.14b. The row-wise representation of a resampling matrix therefore describes the resampling in terms of how acquired data is *scattered* to the grid, while the column-wise representation describes how the gridded samples *gather* the signal. The above methods of density compensation can be incorporated as a filter in the resampling matrix.

Without any density compensation (Fig. 4.14a), each row contains the same energy which, in this example, leads to an overemphasis in high-frequency regions as more samples were acquired there. The plots below and to the right of each resampling matrix shows the sum along the rows and columns, respectively. They describe the relationship between filtering of the acquired data and the density after gridding. For instance, the equal contribution strategy in Figure 4.14a yields a resampled density equal to the acquired sample density, i.e. a finer version of that shown in Figure 4.13c.

In order to achieve uniform density after gridding, the acquired samples can be filtered according to the readout waveform itself, as it is inversely proportional to the sample density as stated above and evident in Figure 4.14c . This filter effectively reduces noise in regions with a long effective dwell time.

To achieve a uniform distribution of noise, i.e. a noise pre-whitening filter is described by the square root of the gradient waveform, as shown in Figure 4.14d. This has the effect of a non-uniform signal level. It should be noted that the choice of filter does not affect the ratio between the signal and noise profile, i.e. the spectral SNR is the same

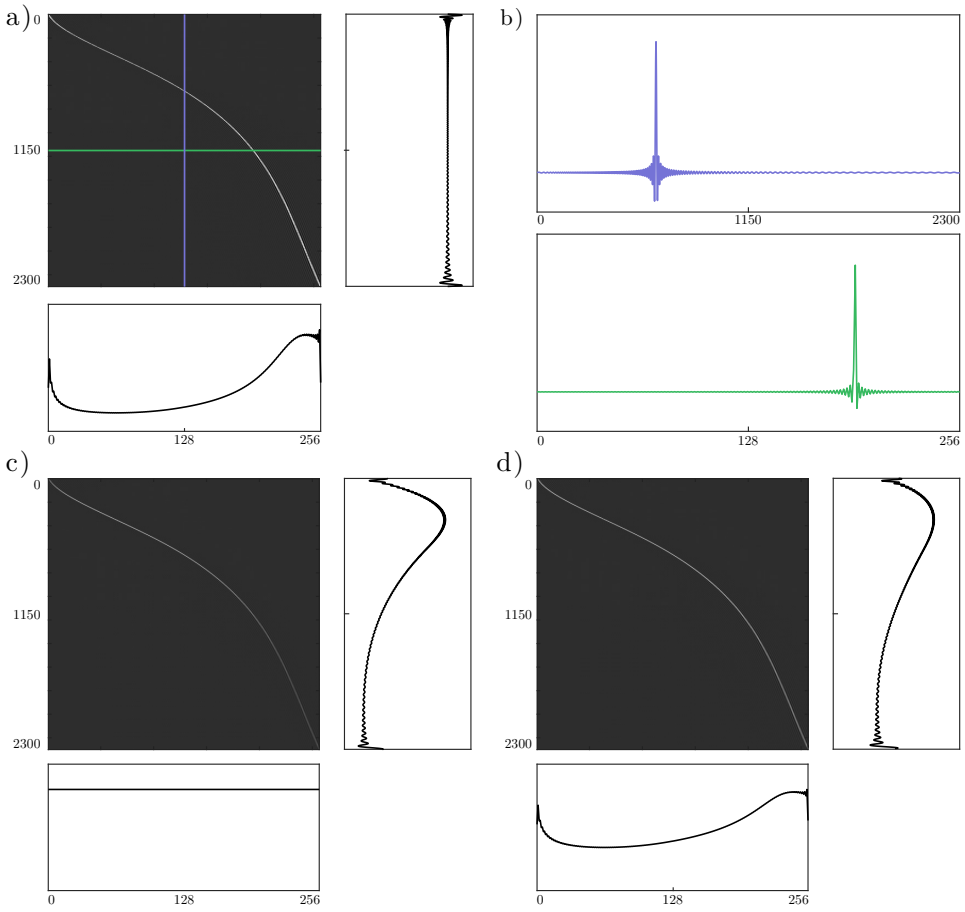

Figure 4.14: Resampling matrices of the waveform in Fig. 4.13 with gridded sample density plots shown below each matrix, and the associated pre-grid filter to the right. A profile along a single column and row of the resampling matrix is displayed as a blue and green line in a), and plotted in b). Three different methods of handling the non-uniform sampling are shown in a, c, and d. a) No filter. This results in large signal variation after gridding. b) Filtered by the inverse of the effective dwell time, $1/\tau$. This results in a uniform signal. c) Filtered by the inverse square root of $1/\sqrt{\tau}$. This results in a uniform noise level. Truncation of the sinc side lobes at the edges results in Gibbs ringing of the filters.

for all filters. However, methods for measuring SNR typically include some type of low-pass filtering, and imaged objects in general have a low-frequency dominant spectrum, leading to this measured "apparent" SNR almost entirely determined by the SNR of the DC signal.

In conclusion, the choice of filtering comes down to a preference of either white noise or a uniform MTF, similar to the trade-off discussed in section 4.2.

4.7.2 RARE Implementation

In a conventional multi-TR RARE sequence, dead times are necessary to achieve CSE. Paper IV describes a fat/water separated multi-TR RARE sequence without dead times, made possible by asymmetric readout waveforms. The efficiency gain by removing dead times can be either used to sample longer, i.e. keeping the RF echo spacing the same, or to shorten it and speed up the acquisition. A sequence illustration of the case of sampling longer is shown in Figure 4.15, together with the conventional shifted trapezoids for reference.

In Figure 4.15c, the alternative of shortening the scan time is shown. The reduction in RF echo spacing is $T/(T + 2\Delta)$, with Δ being the dephasing time at k-space center and T the readout duration. For the well-conditioned in- and opposed phase encoding acquired with a bandwidth of ± 35 kHz and a matrix size of 320, $T = 4.6$ ms at 3T, the dead times equal half of the readout duration and prolongs the scan time. Using asymmetric readout waveforms, the scan time can be shortened by approximately 25 % with maintained sampling duration.

4.7.3 Matching the In-Phase Waveform

Similar to how asymmetric readout result in a non-zero chemical shift encoding with a non-uniform sample density, it is interesting to investigate how the in-phase waveform can be designed to accommodate it. In particular, it is desirable to design the in-phase waveform in such a way that the *combined* sampling density is smoother, a method of "flattening the curve". Since we desire no chemical shift encoding in this echo, the waveform must be symmetric, and the readout area must be consistent with that of the

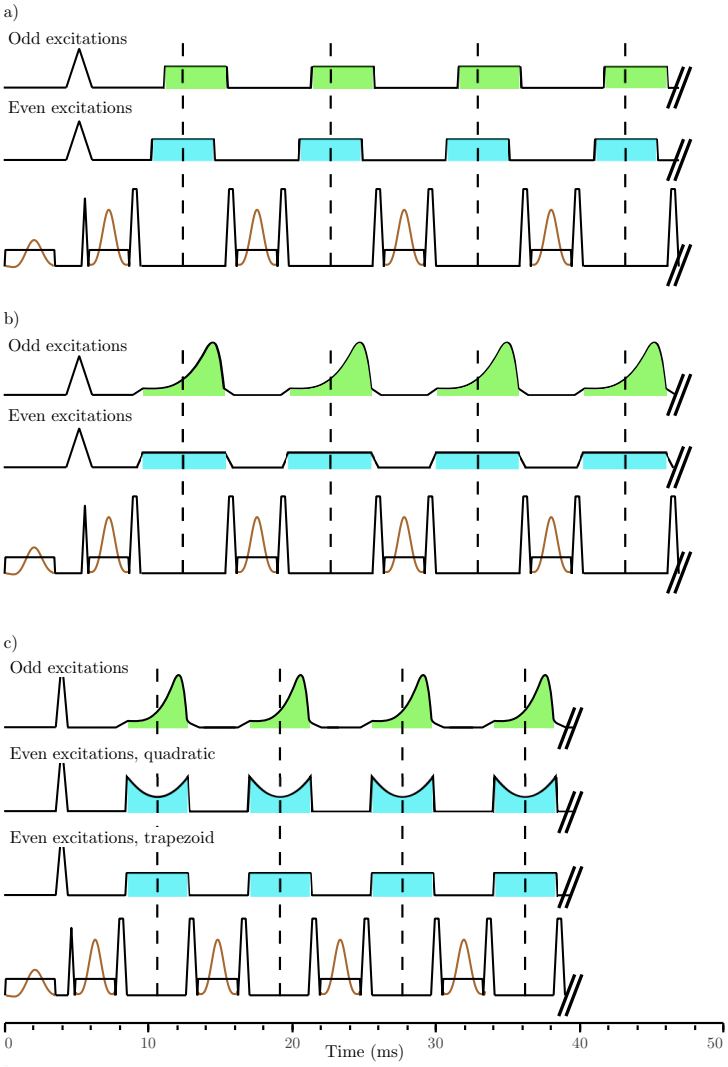

Figure 4.15: a) Conventional multi-TR fat/water separated RARE, where readouts are shifted to achieve CSE. b) By using asymmetric readouts, dead times are avoided and more time can be spent sampling. Alternatively, the sample duration can remain and instead, the scan time is shortened (c). Shaded regions indicate sampling, with cyan marking an in-phase echo and green a CSE echo.

asymmetric. This concept was tried in a simple form in Paper IV, using a quadratic waveform defined by the the offset q_0 and the quadratic coefficient q_2:

$$Q\left(t\right) = q_0 + q_2(t - t_c)^2, \tag{4.12}$$

where t_c is the center of the waveform, i.e. the time of k-space center acquisition. A certain sample density can then be enforced at the k-space center by choosing q_0 accordingly. q_2 then follows from the desired area. In paper IV, q_0 was chosen such that the combined effective dwell time at the k-space center was equal to that of two centered trapezoids. The waveform is shown in Figure 4.15c.

4.7.4 Dephasing Times and the SNR Transfer Function

All constituent papers in this thesis perform the fat/water separation in k-space, where the dephasing time of each sample is modelled. However, an ambiguity arises in determining the dephasing times as a consequence of the resampling procedure described above, where gridded samples are formed by contributions from every acquired sample. Since the sinc kernel is symmetric, its center peak is a good estimate of the effective dephasing time, as most of the resampled signal originates from there. Therefore, an obvious solution to this ambiguity is available in the resampling matrix, where the dephasing time of a gridded sample can be looked up according to the maximum value of its associated column vector, as shown in Figure 4.16 in the case of a quadratic waveform.

Given all dephasing times for the two echoes, it is possible to derive the *SNR transfer function*, which describes the efficiency of the entire acquisition and reconstruction. The SNR transfer function is the ratio between the MTF and NPS which incorporates the sample densities, dephasing times, real-valued estimates, and the regularization applied in the fat/water separation process. Note that the SNR transfer function does not account for the non-linear terms ϕ and ψ which are estimated from the low-frequency data. The inverse problem in k-space taking all of this into account is

$$\mathcal{F}\left(\hat{\mathbf{x}}\right)_k = \underbrace{\begin{bmatrix}(\mathbf{\Lambda A})_k \\ (\mathbf{\Lambda A})_c^*\end{bmatrix}^{\ddagger}}_{\mathbf{E}} \begin{bmatrix}\mathcal{F}(\mathbf{S}^{\phi,\psi})_k \\ \mathcal{F}(\mathbf{S}^{\phi,\psi})_c^*\end{bmatrix} \tag{4.13}$$

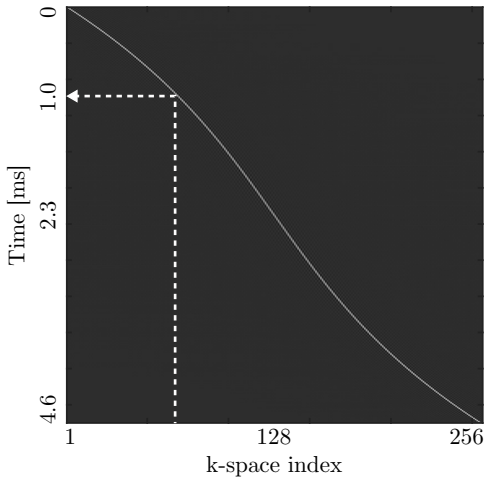

Figure 4.16: Resampling matrix of a quadratic in-phase waveform. The dashed arrow looks up the effective dephasing time at k-space index 64, which equals 1 ms. Note that the y-axis shows the elapsed time from start of acquisition, so the time to the nearest in-phase (i.e. excitation or spin echo) must be considered as well.

where $\mathbf{\Lambda}$ accounts for the local sample density after gridding (with τ being the effective dwell time),

$$\mathbf{\Lambda} = \begin{bmatrix} \sqrt{\tau_1} & 0 \\ 0 & \sqrt{\tau_2} \end{bmatrix}, \tag{4.14}$$

and the subscripts k and c are the sample index and its conjugate index as described in section 4.6. The ‡ superscript is the Tikhonov regularized pseudo inverse (section 4.2.2). Taking all of these factors into account yields the noise-whitening, density-compensated, real-valued regularized estimator $\mathbf{E}$, from which we can calculate the MTF and NPS as described in section 4.2.

An example of the SNR transfer function is seen in Figure 4.17, together with that of a shifted readout. The dotted line corresponds to the conventional shifted readout with dead times. Dead times are avoided by stretching the in-phase trapezoid according to Eggers (Ref. [85]), and using an asymmetric readout for the opposed phase echo, resulting in an increased SNR for almost all samples but those at high frequencies as

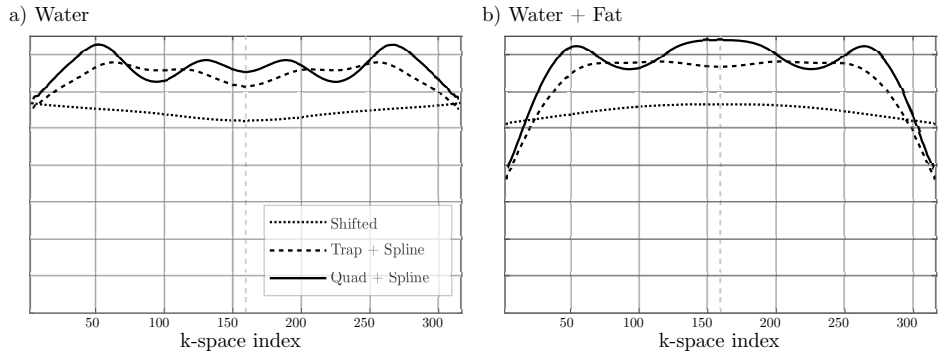

Figure 4.17: SNR Transfer function for water (a) and the combined water and fat image (b), evaluated for three acquisition strategies with 320 samples. The vertical dashed line marks the k-space center.

shown by the dashed line. The in-phase echo has been replaced with the quadratic waveform in the third acquisition strategy (solid line), which further increases SNR for most frequencies. Note the symmetric shape, a consequence of real-valued estimates.

5 Discussion

Dixon methods are unique in their approach of handling the fat signal in MRI. By aligning echoes to acquire specific chemical shift encodes, a model based reconstruction can separate the fat signal without any signal destruction otherwise present when using CHESS or STIR. The requirements on the signal are, however, more constrained, since multiple echoes are necessary. This increases scan time in multi-TR sequences. However, unlike CHESS and STIR, the additional scan time is used to acquire signal, which is the hard currency in any MR sequence design. At least in theory, the scan time penalty can be made void by applying other means of acceleration such as parallel imaging, meaning that the fat/water separation becomes ”free of charge”. If this was possible, there would essentially be no downside of acquiring fat/water separated images. In reality, however, this falls short as it is impossible to acquire data which can be reconstructed with both a uniform noise behavior and without any loss in signal or resolution.

The question is then: Are Dixon methods a competitive option compared to other fat suppression techniques? When it comes to inversion recovery based fat suppression in T1 imaging with contrast enhancement, model based fat/water separation is unquestionably superior as it avoids suppressing the administered contrast agent. Fat suppression using CHESS pulses or water excitation comes with a significant scan time penalty in T1-weighted imaging as the TR is short. Specifically in the case of brain imaging, magnetization transfer effects from the saturation pulses can degrade contrast between gray and white matter [9]. Although FLAIR sequences with fat saturation offer T1-weighted contrast without this large scan time penalty due to longer TR, there is an SNR loss incurred by the inversion pulse which affects all tissues.

The RARE sequence is a workhorse in clinical T2-weighted MRI, which comes with an inherently strong fat signal [76]. Due to the long TR and ETL, CHESS pulses can be

played out sparsely and still achieve fat suppression, and the rather small scan time penalty makes CHESS better suited here than for T1-weighted spin echo sequences. However, the CHESS technique is restricted to only suppressing a certain frequency band, so it fails in regions where the field inhomogeneities are large, as exemplified in Figure 2.3 on page 9.

One of the advantages of Dixon methods over fat suppression is that multiple image sets are generated. In contrast to fat suppression rendering only one ”water” image, Dixon methods generate water and fat images, from which synthetic in- and opposed phase images can be generated without chemical shift artifacts, in addition to the source echoes. These additional image sets carry diagnostic information. For instance, the phase cancellation artifact in opposed phase imaging can emphasize ankle fractures outlines due to mixed fat/water in regions with bone marrow edema [91], [92].

The fat images are also of clinical value, specifically so as a replacement of the T1-weighted image in detection of bone marrow metastases [93]–[95] and multiple myeloma [96]. Prior to the refinement of Dixon methods and their implementations becoming available on clinical systems, the signal ratio between in- and opposed phase gradient echo images were used to detect abnormal bone marrow [97]. Since bone marrow lesions are on-resonance and that they replace fat content [98], it makes sense to separate fat and water to remove ambiguous interpretation of these ratios. Other applications include using the amount of vertebral marrow fat as a proxy for evaluating bone density, a correlation previously reported using MR spectroscopy [99].

Not only can this shorten MR protocols, a driving force in the effort of increasing availability and effectiveness of scans, but all image stacks align with each other. The in-phase image, synthetic or acquired, spatially and temporally matches the water image, which cannot be achieved without fat/water separation. This simplifies diagnostic interpretation and reduces motion-related problems due to long scan sessions. For whole-body MR protocols, short scan times with resistance to field inhomogeneities are of particular importance, which has led to the inclusion of T2-weighted fat/water sequences [100].

Further processing of the estimates can give the relative proton density fat fraction, or PDFF. The use of PDFF as a biomarker has been extensively researched, and most vendors today offer multi-point Dixon sequences which render PDFF maps. Since the

fat fraction is a ratio between fat and the in-phase signal, any signal fluctuations from coil sensitivities are cancelled. Fat fraction measurements have been used to investigate obesity [101] and brown adipose tissue [102]. PDFF maps also provides a non-invasive way of classifying nonalcoholic fatty liver disease (NAFLD) [103], [104]. Fischer et al. used a two-point fat/water sequence to investigate the extent of fatty degeneration of the gastrocnemius muscle and found a strong correlation with the Goutallier classification [105].

The use of only two echoes can be questioned as the model is rather restricted, i.e. it is assumed that no T2′-decay occurs and that the initial phase of fat and water are equal, which might be false, at least when using certain RF excitation pulses. Conditioning of the inverse problem also depends on the fat:water ratio, introducing a bias in fat fraction measurements. While conventional IP/OP imaging is confounded in the presence of both fat and iron in the liver, multi-point fat/water separation can be extended to model T2* decay to overcome this problem [106], [107]. These models do come with additional demands on the acquisition in that more echoes are required in order to resolve all parameters. The price associated with this is generally either longer acquisition times or reduced spatial resolution.

A benefit of two-point Dixon methods is that the phase evolution does not have to be linear for correct fat/water separation. This is a side-effect of the two-point model as *any* phase difference between the two echoes will be incorporated in ψ. While any phase contribution from sources such as eddy currents or concomitant fields is demodulated, the ψ map describes a combination of these effects and the field inhomogeneity. The ψ map can therefore not be directly used for other off-resonance correction purposes such as distortion correction.

Despite the benefits of Dixon methods, CHESS is often preferred over fat/water separation in clinical MRI. Part of this is likely attributed to radiologists being familiar with CHESS, including its weaknesses. Dixon methods are often associated with the swap artifact, which was a problem in early vendor implementations, and not everyone is aware of its potential of correcting for chemical shift. When asked about why Dixon methods are not the primary method of dealing with fat, the response is often that it lacks ”image sharpness”. With the MTF analysis described in this thesis and its smoothing effect showed in Figure 4.8, this is understandable. It is not trivial to overcome this

problem, but real-valued estimation does appear to have significant impact.

Another problem associated with Dixon methods is that the technique involves many vital steps, all of which can render an image with artifacts particular to the choice of handling each of these steps. For instance, a region growing scheme for field map estimation can result in a complete swap between the water and fat signal. In our method of estimating field map, local field map errors can result in signal leakage of varying magnitude. The multiple ways of handling both acquisition and reconstruction makes it difficult, if not impossible, to draw general conclusions about the performance of fat/water imaging.

The dBW-RARE technique outperforms fTED in terms of achievable resolution and sampling efficiency, and it remains to be investigated how pushing the resolution frontier in single-TR imaging affects clinical utility. Both fTED and dBW-RARE are non-invasive sequence designs where the conventional single readout is replaced by a readout block to facilitate fat/water separation in the reconstruction. dBW-RARE puts less stress on the gradient system, making it a quieter sequence with smaller risk of peripheral nerve stimulation.

dBW-RARE uses partial Fourier sampling to improve its sampling efficiency. This makes it incompatible with partial Fourier sampling along any phase encoding direction. The symmetrically sampled region is used to estimate the field map and initial phase, and it must be sufficiently large to reflect the true distributions. A reasonable assumption would be that the performance of dBW-RARE is dependent on the resolution of these maps rather than its coverage as a fraction of the matrix size. For instance, if a 256×256 image is acquired and successfully reconstructed with a partial Fourier fraction of $\frac{3}{4}$, it would be expected that if the matrix size is increased to 512×512, the partial Fourier fraction could be lowered to $\frac{5}{8}$, i.e. the symmetric region is the same in k-space. During the development of dBW-RARE, we did not find this to be the case. Instead, ringing artifacts seemed to appear at partial Fourier fractions below 0.7, regardless of voxel size and FOV. Sampling with fractions above 0.7 resulted in almost complete removal of these artifacts, regardless of matrix size.

Given this heuristic limit in partial Fourier fraction, a resolution of 0.5 mm could be obtained along the readout dimension at 3T. If smaller voxels than that are required,

multi-TR is likely a better alternative, not only because the partial Fourier factor must be pushed too low, but also since SNR is a limiting factor in very high resolution imaging. More sampling is required to circumvent the low SNR, e.g. by signal averaging. However, signal averaging increases scan time. At that point, switching to multi-TR comes with little to no scan time cost and with the added benefit of avoiding partial Fourier sampling.

Multi-TR sequences do, however, come with dead times from shifted gradients and the scan time penalty is in fact more than doubled compared to single-TR sequences without signal averaging. With the new chemical shift encoding technique using asymmetric readouts, presented in Paper IV, these dead times can be avoided. This opens a new topic of designing gradient waveforms to acquire CSE signals as efficiently as possible, where the waveforms described in Paper IV serves as a proof of concept. With the only constraint being those imposed by the gradient systems and the symmetric/asymmetric property of the in/opposed phase CSE, further waveform designs could be improved by accounting for the entire fat/water separation model. For instance, instead of only designing waveforms by their dephasing time pairs at the k-space center, the *entire* SNR transfer function (Fig. 4.17, p. 66), which accounts for conjugate symmetry imposed by real-valued estimates, local sample densities, noise level differences between echoes, and more, can be incorporated in the optimization problem.

6 Appendix - Coil Combination

The performance of the fat/water separation in this thesis was highly affected by the coil combination algorithm. The eigenvector approach by Walsh et al. was used for all images produced in this thesis [108]. Below follows a rather verbose description of that method and a comparison with the commonly used but inferior sum-of-squares approach.

The measured MR signal from n coils $\mathbf{S} \in \mathbb{C}^n$ at a single voxel is

$$\mathbf{S} = \rho\mathbf{b} + \epsilon \tag{A.1}$$

where $\mathbf{b}$ is the coil sensitivity and ϵ is the noise of each coil element. To estimate the complex scalar $\hat{\rho}$ with minimum noise amplification, we are looking for the optimal weights $\mathbf{w} \in \mathbb{C}^n$:

$$\hat{\rho} = \mathbf{w}^H\mathbf{S} = \sum_{i=1}^{n} w_i^* s_i. \tag{A.2}$$

Sum of Squares Combination

Assuming no noise correlation between coil elements, ρ_{SOS} can be estimated, suboptimally, from the magnitude of the signal [109]:

$$\hat{\rho}_{\mathrm{SOS}} = \sqrt{\mathbf{S}^H\mathbf{S}}. \tag{A.3}$$

Expressing Eq. A.3 on the same form as Eq. A.2 allows inspection of the weights:

$$\hat{\rho}_{\mathrm{SOS}} = \mathbf{w}_{\mathrm{SOS}}^{H}\mathbf{S} = \sum_{i=1}^{n} \frac{s_i^* s_i}{\|\mathbf{S}\|_1} = \sum_{i=1}^{n} \underbrace{e^{-i\angle s_i} \frac{\|s_i\|_1}{\|\mathbf{S}\|_1}}_{w_i} s_i \tag{A.4}$$

From A.4 we see that the sum-of-squares weights are inversely proportional to the magnitude of the measured signal and the phase is removed. With this estimation, the weights are reliable in voxels with high SNR, but unreliable in noisy regions.

Correlation Matrices and SNR

The relationship between weights and the signal-to-noise ratio κ is described in [108] as the generalized Rayleigh quotient:

$$\kappa = \frac{\mathbf{w}^H \mathbf{R}_{\mathbf{s}} \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\epsilon} \mathbf{w}} \tag{A.5}$$

Where $\mathbf{R}_{\mathbf{s}}$ and $\mathbf{R}_{\epsilon}$ are correlation matrices of signal and noise. $\mathbf{R}_{\epsilon}$ can be estimated from a pre-scan or from a noisy region (assuming white noise). If no estimation is done, we assume $\mathbf{R}_{\epsilon}$ to be the identity matrix, i.e. no correlation of noise between coil elements and unit variance. We estimate $\mathbf{R}_{\mathbf{s}}$ from the sum of outer products between S and S^H in a local region.

Eigenvectors as Weights

According to [108], the optimal weights are the dominant eigenvector (corresponding to the largest eigenvalue, λ_1) of $\mathbf{R}_{\epsilon}^{-\mathbf{1}}\mathbf{R}_{\mathbf{s}}$. Recall the definition of the set of eigenvectors, i.e. the eigenspace:

$$\mathbf{E}_{\lambda}(\mathbf{R}_{\epsilon}^{-\mathbf{1}}\mathbf{R}_{\mathbf{s}}) = \left\{\mathbf{w}_{\lambda} : \mathbf{R}_{\epsilon}^{-\mathbf{1}}\mathbf{R}_{\mathbf{s}}\mathbf{w} = \lambda\mathbf{w}\right\} \tag{A.6}$$

Assuming no noise correlation, $\mathbf{R}_{\epsilon}^{-\mathbf{1}} = \mathbf{I}_n$, i.e. the eigenspace of $\mathbf{R}_{\epsilon}^{-\mathbf{1}}\mathbf{R}_{\mathbf{s}}$ is equal to the eigenspace of the signal correlation matrix. Naturally, any scaling of the weights will results in the same SNR. Setting $\|w\| = 1$ results in unit scaling of the combined signal, only the numerator of Eq. A.5 needs consideration since

$$\mathbf{w}^H \mathbf{R}_\epsilon \mathbf{w} = \mathbf{w}^H \mathbf{I}_n \mathbf{w} = 1. \tag{A.7}$$

$\mathbf{R_s w}$ is a linear transformation of $\mathbf{w}$. If $\mathbf{w} \in \mathbf{w}_\lambda$, this transformation is merely a constant scaling. $\mathbf{R_s}$ can be decomposed, or diagonalized into

$$\mathbf{R_s} = \mathbf{U}\mathbf{D}\mathbf{U}^H \tag{A.8}$$

Where $\mathbf{U}$ and $\mathbf{D}$ are eigenvectors and eigenvalues of $\mathbf{R}_s$, respectively. Consider some arbitrary weights $\mathbf{w}$ expressed as a linear combination of $\mathbf{U}$:

$$\mathbf{w} = \mathbf{U}\mathbf{x}. \tag{A.9}$$

Applying these weights to Eq. A.5 and using the fact that $\mathbf{U}$ is orthonormal due to $\mathbf{R_s}$ being Hermitian gives

$$\kappa = \mathbf{w}^H \mathbf{U}\mathbf{D}\mathbf{U}^H \mathbf{w} = \mathbf{x}^H \mathbf{U}^H \mathbf{U}\mathbf{D}\mathbf{U}^H \mathbf{U}\mathbf{x} = \mathbf{x}^H \mathbf{D}\mathbf{x} = \sum_{i=1}^{n} \lambda_i \|x_i\|_2 \tag{A.10}$$

With any eigenvector (corresponding to eigenvalue λ_k) as weights, the SNR is

$$\kappa_{\lambda_k} = \mathbf{x}_k^H \mathbf{U}\mathbf{D}\mathbf{U}^H \mathbf{x}_k = \lambda_k \mathbf{x}_k^H \mathbf{x}_k = \sum_{i=1}^{n} \lambda_k \|x_i\|_2 = \lambda_k \tag{A.11}$$

In the case of the dominant eigenvector as weights,

$$\lambda_1 \geq \lambda_i$$

$$\kappa_1 = \sum_{i=1}^{n} \lambda_1 \|x_i\|_2 \geq \sum_{i=1}^{n} \lambda_i \|x_i\|_2 \tag{A.12}$$

meaning that we are maximizing SNR by using the dominant eigenvector as weights.

Phase Ambiguity with Eigendecompisition

Since $\mathbf{R}$ is Hermitian, the decomposition $\mathbf{R} = \mathbf{U}\mathbf{D}\mathbf{U}^H$ is also a singular value decomposition. The eigenvectors are not unique, as an arbitrary phase can be added to the eigenspace:

$$\left(e^{i\phi}\mathbf{U}\right)\mathbf{D}\left(e^{i\phi}\mathbf{U}\right)^{H} = e^{i\phi}\mathbf{U}\mathbf{D}\mathbf{U}^{H}e^{-i\phi} = \mathbf{R} \tag{A.13}$$

This ambiguity results in an arbitrary phase shift between voxels, rendering a useless phase image. We can solve this by adding the constraint that the first element in $\mathbf{U}$ is real-valued. This constraint restricts the phase offset in weights to be equal to that of the first coil element.

7 Acknowledgements

Stefan Skare, my supervisor, for bringing me into your research group and giving me the opportunity to perform research on this topic that I really enjoy. The group you have created is outstanding and I am very grateful to be a part of it. With the risk of sounding too *categorical*, I would say it's the best.

My co-supervisor **Enrico Avventi** for helping me with all those compiling issues. Hopefully I now have the time to clean up my code.

Johan Berglund, my other co-supervisor. It is difficult to describe how much your support has meant to me. Not only are you an excellent scientist, but also the best supervisor I could ask for. It's always a pleasure working with you and spawning new ideas together.

Ola Norbeck, I've really enjoyed doing this PhD side-by-side with you and I already miss the days at KTH. I'm sorry for never remembering what tailored RF does.

Tim Sprenger and **Adam van Niekerk**, for being such helpful and happy colleagues.

Mikael Skorpil, I always enjoy talking to you and having you as my friend. Now we're both real doctors!

All my other colleagues at the neuroradiology department. Thank you for your patience and understanding of the problems that come with doing research in a clinical environment.

My love **Nathalie Thomé**, for being you. In the beginning of this journey, I said that it will be tough at times. Thanks to you, it never really was.

Joachim Nilsson, for being the best friend one can have.

Omar El-Mansoury, my childhood friend, and his fiancée **Madeleine Pontusson**, who are always in a good mood. I hope to someday learn that from you.

My friend **Christian Jamtheim Gustafsson**, who I wish I could see more often. We said it wasn't worth the trouble going through doctoral studies. We were wrong.

My father **Ingvar**, my sister **Isabelle**, and brother **Fredrik**. I miss you.

8 Bibliography

[1] R. B. Buxton, *Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques*. Cambridge University Press, Aug. 2009.

[2] C. J. Bergin, J. M. Pauly, and A. Macovski, "Lung parenchyma: Projection reconstruction MR imaging", en, *Radiology*, vol. 179, no. 3, pp. 777–781, Jun. 1991.

[3] H. Yu, A. Shimakawa, C. A. McKenzie, E. Brodsky, J. H. Brittain, and S. B. Reeder, "Multiecho water-fat separation and simultaneous r2* estimation with multifrequency fat spectrum modeling", en, *Magn. Reson. Med.*, vol. 60, no. 5, pp. 1122–1134, Nov. 2008.

[4] W. T. Dixon, "Simple proton spectroscopic imaging", en, *Radiology*, vol. 153, no. 1, pp. 189–194, Oct. 1984.

[5] T. A. Bley, O. Wieben, C. J. François, J. H. Brittain, and S. B. Reeder, "Fat and water magnetic resonance imaging", en, *J. Magn. Reson. Imaging*, vol. 31, no. 1, pp. 4–18, Jan. 2010.

[6] H. Guerini, P. Omoumi, F. Guichoux, V. Vuillemin, G. Morvan, M. Zins, F. Thevenin, and J. L. Drape, "Fat suppression with dixon techniques in musculoskeletal magnetic resonance imaging: A pictorial review", en, *Semin. Musculoskelet. Radiol.*, vol. 19, no. 4, pp. 335–347, Sep. 2015.

[7] E. M. Delfaut, J. Beltran, G. Johnson, J. Rousseau, X. Marchandise, and A. Cotten, "Fat suppression in MR imaging: Techniques and pitfalls", en, *Radiographics*, vol. 19, no. 2, pp. 373–382, Mar. 1999.

[8] T. Aoki, Y. Yamashita, H. Oki, H. Takahashi, Y. Hayashida, K. Saito, Y. Tanaka, and Y. Korogi, "Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL) of the wrist and finger at 3t: Compari-

son with chemical shift selective fat suppression images", en, *J. Magn. Reson. Imaging*, vol. 37, no. 3, pp. 733–738, Mar. 2013.

[9] W. T. Dixon, H. Engels, M. Castillo, and M. Sardashti, "Incidental magnetization transfer contrast in standard multislice imaging", en, *Magn. Reson. Imaging*, vol. 8, no. 4, pp. 417–422, 1990.

[10] P. J. Hore, "A new method for water suppression in the proton NMR spectra of aqueous soh tiom", *J. Magn. Reson.*, vol. 54, pp. 539–542, 1983.

[11] P. Plateau and M. Gueron, "Exchangeable proton NMR without base-line distorsion, using new strong-pulse sequences", *J. Am. Chem. Soc.*, vol. 104, no. 25, pp. 7310–7311, Dec. 1982.

[12] D. Thomasson, D. Purdy, and J. P. Finn, "Phase-modulated binomial RF pulses for fast spectrally-selective musculoskeletal imaging", en, *Magn. Reson. Med.*, vol. 35, no. 4, pp. 563–568, Apr. 1996.

[13] G. M. Bydder and I. R. Young, "MR imaging: Clinical use of the inversion recovery sequence", en, *J. Comput. Assist. Tomogr.*, vol. 9, no. 4, pp. 659–675, Jul. 1985.

[14] M. Dousset, R. Weissleder, R. E. Hendrick, D. D. Stark, C. J. Fretz, G. Elizondo, P. F. Hahn, S. Saini, and J. T. Ferrucci, "Short TI inversion-recovery imaging of the liver: Pulse-sequence optimization and comparison with spin-echo imaging", en, *Radiology*, vol. 171, no. 2, pp. 327–333, May 1989.

[15] M. Bydder, T. Yokoo, H. Yu, M. Carl, S. B. Reeder, and C. B. Sirlin, "Constraining the initial phase in water-fat separation", en, *Magn. Reson. Imaging*, vol. 29, no. 2, pp. 216–221, Feb. 2011.

[16] J. Ma, "Dixon techniques for water and fat imaging", en, *J. Magn. Reson. Imaging*, vol. 28, no. 3, pp. 543–558, Sep. 2008.

[17] R. E. Sepponen, J. T. Sipponen, and J. I. Tanttu, "A method for chemical shift imaging: Demonstration of bone marrow involvement with proton chemical shift imaging", en, *J. Comput. Assist. Tomogr.*, vol. 8, no. 4, pp. 585–587, Aug. 1984.

[18] H. N. Yeung and D. W. Kormos, "Separation of true fat and water images by correcting magnetic field inhomogeneity in situ", en, *Radiology*, vol. 159, no. 3, pp. 783–786, Jun. 1986.

[19] Y. S. Kim, C. W. Mun, and Z. H. Cho, "Chemical-shift imaging with large magnetic field inhomogeneity", en, *Magn. Reson. Med.*, vol. 4, no. 5, pp. 452–460, May 1987.

[20] C. C. Lodes, J. P. Felmlee, R. L. Ehman, C. M. Sehgal, J. F. Greenleaf, G. H. Glover, and J. E. Gray, "Proton MR chemical shift imaging using double and triple phase contrast acquisition methods", en, *J. Comput. Assist. Tomogr.*, vol. 13, no. 5, pp. 855–861, Sep. 1989.

[21] Q. S. Xiang and L. An, "Water-fat imaging with direct phase encoding", en, *J. Magn. Reson. Imaging*, vol. 7, no. 6, pp. 1002–1015, Nov. 1997.

[22] S. B. Reeder, Z. Wen, H. Yu, A. R. Pineda, G. E. Gold, M. Markl, and N. J. Pelc, "Multicoil dixon chemical species separation with an iterative least-squares estimation method", en, *Magn. Reson. Med.*, vol. 51, no. 1, pp. 35–45, Jan. 2004.

[23] L. An and Q. S. Xiang, "Chemical shift imaging with spectrum modeling", en, *Magn. Reson. Med.*, vol. 46, no. 1, pp. 126–130, Jul. 2001.

[24] G. H. Glover, "Multipoint dixon technique for water and fat proton and susceptibility imaging", en, *J. Magn. Reson. Imaging*, vol. 1, no. 5, pp. 521–530, Sep. 1991.

[25] J. Szumowski and D. B. Plewes, "Fat suppression in the time domain in fast MR imaging", en, *Magn. Reson. Med.*, vol. 8, no. 3, pp. 345–354, Nov. 1988.

[26] A. R. Pineda, S. B. Reeder, Z. Wen, and N. J. Pelc, "Cramér-Rao bounds for three-point decomposition of water and fat", en, *Magn. Reson. Med.*, vol. 54, no. 3, pp. 625–635, Sep. 2005.

[27] M. Bydder, T. Yokoo, G. Hamilton, M. S. Middleton, A. D. Chavez, J. B. Schwimmer, J. E. Lavine, and C. B. Sirlin, "Relaxation effects in the quantification of fat using gradient echo imaging", en, *Magn. Reson. Imaging*, vol. 26, no. 3, pp. 347–359, Apr. 2008.

[28] E. K. Brodsky, J. H. Holmes, H. Yu, and S. B. Reeder, "Generalized k-space decomposition with chemical shift correction for non-cartesian water-fat imaging", en, *Magn. Reson. Med.*, vol. 59, no. 5, pp. 1151–1164, May 2008.

[29] J. L. Honorato, V. Parot, C. Tejos, S. Uribe, and P. Irarrazaval, “Chemical species separation with simultaneous estimation of field map and t2* using a k-space formulation”, en, *Magn. Reson. Med.*, vol. 68, no. 2, pp. 400–408, Aug. 2012.

[30] H. Eggers, B. Brendel, A. Duijndam, and G. Herigault, “Dual-echo dixon imaging with flexible choice of echo times”, en, *Magn. Reson. Med.*, vol. 65, no. 1, pp. 96–107, Jan. 2011.

[31] J. Berglund, H. Ahlström, L. Johansson, and J. Kullberg, “Two-point dixon method with flexible echo times”, en, *Magn. Reson. Med.*, vol. 65, no. 4, pp. 994–1004, Apr. 2011.

[32] H. Yu, S. B. Reeder, A. Shimakawa, J. H. Brittain, and N. J. Pelc, “Field map estimation with a region growing scheme for iterative 3-point water-fat decomposition”, en, *Magn. Reson. Med.*, vol. 54, no. 4, pp. 1032–1039, Oct. 2005.

[33] S. B. Reeder, A. R. Pineda, Z. Wen, A. Shimakawa, H. Yu, J. H. Brittain, G. E. Gold, C. H. Beaulieu, and N. J. Pelc, “Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL): Application with fast spin-echo imaging”, en, *Magn. Reson. Med.*, vol. 54, no. 3, pp. 636–644, Sep. 2005.

[34] Q.-S. Xiang, “Two-point water-fat imaging with partially-opposed-phase (POP) acquisition: An asymmetric dixon method”, en, *Magn. Reson. Med.*, vol. 56, no. 3, pp. 572–584, Sep. 2006.

[35] D. Hernando, J. P. Haldar, B. P. Sutton, J. Ma, P. Kellman, and Z.-P. Liang, “Joint estimation of water/fat images and field inhomogeneity map”, en, *Magn. Reson. Med.*, vol. 59, no. 3, pp. 571–580, Mar. 2008.

[36] D. Wang, N. R. Zwart, Z. Li, M. Schär, and J. G. Pipe, “Analytical three-point dixon method: With applications for spiral water-fat imaging”, en, *Magn. Reson. Med.*, vol. 75, no. 2, pp. 627–638, Feb. 2016.

[37] G. Golub and V. Pereyra, “Separable nonlinear least squares: The variable projection method and its applications”, en, *Inverse Probl.*, vol. 19, no. 2, R1, Feb. 2003.

[38] W. Lu and B. A. Hargreaves, “Multiresolution field map estimation using golden section search for water-fat separation”, en, *Magn. Reson. Med.*, vol. 60, no. 1, pp. 236–244, Jul. 2008.

[39] J. Berglund, L. Johansson, H. Ahlström, and J. Kullberg, "Three-point dixon method enables whole-body water and fat imaging of obese subjects", en, *Magn. Reson. Med.*, vol. 63, no. 6, pp. 1659–1668, Jun. 2010.

[40] D. Hernando, P. Kellman, J. P. Haldar, and Z.-P. Liang, "Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm", en, *Magn. Reson. Med.*, vol. 63, no. 1, pp. 79–90, Jan. 2010.

[41] V. Kolmogorov and C. Rother, "Minimizing nonsubmodular functions with graph cuts - a review", en, *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 7, pp. 1274–1279, Jul. 2007.

[42] J. Berglund and J. Kullberg, "Three-dimensional water/fat separation and t2* estimation based on whole-image optimization–application in breathhold liver imaging at 1.5 T", en, *Magn. Reson. Med.*, vol. 67, no. 6, pp. 1684–1693, Jun. 2012.

[43] J. Berglund and M. Skorpil, "Multi-scale graph-cut algorithm for efficient water-fat separation", en, *Magn. Reson. Med.*, vol. 78, no. 3, pp. 941–949, Sep. 2017.

[44] W. Lu, H. Yu, A. Shimakawa, M. Alley, S. B. Reeder, and B. A. Hargreaves, "Water-fat separation with bipolar multiecho sequences", en, *Magn. Reson. Med.*, vol. 60, no. 1, pp. 198–209, Jul. 2008.

[45] H. Yu, A. Shimakawa, C. A. McKenzie, W. Lu, S. B. Reeder, R. S. Hinks, and J. H. Brittain, "Phase and amplitude correction for multi-echo water-fat separation with bipolar acquisitions", en, *J. Magn. Reson. Imaging*, vol. 31, no. 5, pp. 1264–1271, May 2010.

[46] P. Peterson and S. Månsson, "Fat quantification using multiecho sequences with bipolar gradients: Investigation of accuracy and noise performance", en, *Magn. Reson. Med.*, vol. 71, no. 1, pp. 219–229, Jan. 2014.

[47] H. Eggers and P. Börnert, "Chemical shift encoding-based water-fat separation methods", en, *J. Magn. Reson. Imaging*, vol. 40, no. 2, pp. 251–268, Aug. 2014.

[48] J. Hennig, A. Nauerth, and H. Friedburg, "RARE imaging: A fast imaging method for clinical MR", en, *Magn. Reson. Med.*, vol. 3, no. 6, pp. 823–833, Dec. 1986.

[49] H. Y. Carr and E. M. Purcell, “Effects of diffusion on free precession in nuclear magnetic resonance experiments”, *Phys. Rev.*, vol. 94, no. 3, pp. 630–638, May 1954.

[50] S. Meiboom and D. Gill, “Modified Spin-Echo method for measuring nuclear relaxation times”, *Rev. Sci. Instrum.*, vol. 29, p. 688, Aug. 1958.

[51] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI pulse sequences*. Elsevier Academic Press, 2004, pp. 550–551.

[52] J. Hennig, “Multiecho imaging sequences with low refocusing flip angles”, *J. Magn. Reson.*, vol. 78, no. 3, pp. 397–407, Jul. 1988.

[53] M. Weigel, “Extended phase graphs: Dephasing, RF pulses, and echoes - pure and simple”, en, *J. Magn. Reson. Imaging*, vol. 41, no. 2, pp. 266–295, Feb. 2015.

[54] K. Oshio and D. A. Feinberg, “GRASE (gradient- and spin-echo) imaging: A novel fast MRI technique”, en, *Magn. Reson. Med.*, vol. 20, no. 2, pp. 344–349, Aug. 1991.

[55] J. Ma, J. B. Son, Y. Zhou, H. Le-Petross, and H. Choi, “Fast spin-echo triple-echo dixon (fTED) technique for efficient t2-weighted water and fat imaging”, en, *Magn. Reson. Med.*, vol. 58, no. 1, pp. 103–109, Jul. 2007.

[56] Z. Li, A. F. Gmitro, A. Bilgin, and M. I. Altbach, “Fast decomposition of water and lipid using a GRASE technique with the IDEAL algorithm”, en, *Magn. Reson. Med.*, vol. 57, no. 6, pp. 1047–1057, Jun. 2007.

[57] J. B. Son, K.-P. Hwang, J. E. Madewell, E. Bayram, J. D. Hazle, R. N. Low, and J. Ma, “A flexible fast spin echo triple-echo dixon technique”, en, *Magn. Reson. Med.*, vol. 77, no. 3, pp. 1049–1057, Mar. 2017.

[58] X. Wang, C. Harrison, Y. K. Mariappan, K. Gopalakrishnan, A. Chhabra, R. E. Lenkinski, and A. J. Madhuranthakam, “MR neurography of brachial plexus at 3.0 T with robust fat and blood suppression”, en, *Radiology*, vol. 283, no. 2, pp. 538–546, May 2017.

[59] J. G. Pipe, “Motion correction with PROPELLER MRI: Application to head motion and free-breathing cardiac imaging”, en, *Magn. Reson. Med.*, vol. 42, no. 5, pp. 963–969, Nov. 1999.

[60] J. G. Pipe and P. Menon, “Sampling density compensation in MRI: Rationale and an iterative numerical solution”, en, *Magn. Reson. Med.*, vol. 41, no. 1, pp. 179–186, Jan. 1999.

[61] V. Rasche, R. Proksa, R. Sinkus, P. Börnert, and H. Eggers, “Resampling of data between arbitrary grids using convolution interpolation”, en, *IEEE Trans. Med. Imaging*, vol. 18, no. 5, pp. 385–392, May 1999.

[62] M. L. Lauzon and B. K. Rutt, “Effects of polar sampling in k-space”, en, *Magn. Reson. Med.*, vol. 36, no. 6, pp. 940–949, Dec. 1996.

[63] Y. Hirokawa, H. Isoda, Y. S. Maetani, S. Arizono, K. Shimada, and K. Togashi, “MRI artifact reduction and quality improvement in the upper abdomen with PROPELLER and prospective acquisition correction (PACE) technique”, en, *AJR Am. J. Roentgenol.*, vol. 191, no. 4, pp. 1154–1158, Oct. 2008.

[64] A. J. Shaka, S. P. Rucker, and A. Pines, “Iterative carr-purcell trains”, *J. Magn. Reson.*, vol. 77, no. 3, pp. 606–611, 1988.

[65] J. G. Pipe, V. G. Farthing, and K. P. Forbes, “Multishot diffusion-weighted FSE using PROPELLER MRI”, en, *Magn. Reson. Med.*, vol. 47, no. 1, pp. 42–52, Jan. 2002.

[66] J. G. Pipe and N. Zwart, “Turboprop: Improved PROPELLER imaging”, *Magn. Reson. Med.*, vol. 55, no. 2, pp. 380–385, Feb. 2006.

[67] O. Norbeck, E. Avventi, M. Engström, H. Rydén, and S. Skare, “Simultaneous multi-slice combined with PROPELLER”, en, *Magn. Reson. Med.*, Dec. 2017.

[68] E. M. Haacke, Y. Xu, Y.-C. N. Cheng, and J. R. Reichenbach, “Susceptibility weighted imaging (SWI)”, en, *Magn. Reson. Med.*, vol. 52, no. 3, pp. 612–618, Sep. 2004.

[69] S. J. Holdsworth, K. W. Yeom, M. E. Moseley, and S. Skare, “Fast susceptibility-weighted imaging with three-dimensional short-axis propeller (SAP)-echo-planar imaging”, *J. Magn. Reson. Imaging*, vol. 41, no. 5, pp. 1447–1453, May 2015.

[70] S. Skare, R. D. Newbould, D. B. Clayton, and R. Bammer, “Propeller EPI in the other direction”, en, *Magn. Reson. Med.*, vol. 55, no. 6, pp. 1298–1307, Jun. 2006.

[71] D. Huo, Z. Li, E. Aboussouan, J. P. Karis, and J. G. Pipe, "Turboprop IDEAL: A motion-resistant fat–water separation technique", *Magn. Reson. Med.*, vol. 61, no. 1, pp. 188–195, Jan. 2009.

[72] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging", en, *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, Dec. 2007.

[73] Q. He, D. Weng, X. Zhou, and C. Ni, "Regularized iterative reconstruction for undersampled BLADE and its applications in three-point dixon water–fat separation", *Magn. Reson. Med.*, vol. 65, no. 5, pp. 1314–1325, May 2011.

[74] D. Weng, Y. Pan, X. Zhong, and Y. Zhuo, "Water–fat separation with parallel imaging based on BLADE", *Magn. Reson. Imaging*, vol. 31, no. 5, pp. 656–663, 2013.

[75] M. Schär, H. Eggers, N. R. Zwart, Y. Chang, A. Bakhru, and J. G. Pipe, "Dixon water-fat separation in PROPELLER MRI acquired with two interleaved echoes", en, *Magn. Reson. Med.*, vol. 75, no. 2, pp. 718–728, Feb. 2016.

[76] R. M. Henkelman, P. A. Hardy, J. E. Bishop, C. S. Poon, and D. B. Plewes, "Why fat is bright in RARE and fast spin-echo imaging", en, *J. Magn. Reson. Imaging*, vol. 2, no. 5, pp. 533–540, Sep. 1992.

[77] D. C. Noll, D. G. Nishimura, and A. Macovski, "Homodyne detection in magnetic resonance imaging", en, *IEEE Trans. Med. Imaging*, vol. 10, no. 2, pp. 154–163, 1991.

[78] E. M. Haacke, E. D. Lindskogj, and W. Lin, "A fast, iterative, partial-fourier technique capable of local phase recovery", *J. Magn. Reson.*, vol. 92, no. 1, pp. 126–145, Mar. 1991.

[79] G. McGibney, M. R. Smith, S. T. Nichols, and A. Crawley, "Quantitative evaluation of several partial fourier reconstruction algorithms used in MRI", en, *Magn. Reson. Med.*, vol. 30, no. 1, pp. 51–59, Jul. 1993.

[80] S. Skare, E. Avventi, O. Norbeck, and H. Rydén, "An abstraction layer for simpler EPIC pulse programming on GE MR systems in a clinical environment", in *Proceedings of the 25th Annual Meeting of the ISMRM, Honolulu, Hawaii*, Honolulu, Hawaii, 2017, p. 3813.

[81] Eggers, Holger, "Variable bandwidth turbo Spin-Echo dixon imaging", in *Proceedings of the 23rd Annual Meeting of ISMRM, Milan, Italy*, 2014, p. 1657.

[82] S. Skare, E. Avventi, M. Mårtensson, O. Norbeck, M. Engström, M. Sandell, and C. Wang, "Propeller techniques for pediatric exams in the presence of large motion", in *Proc. Intl. Soc. Mag. Reson. Med. 23 (2015)*, 2015.

[83] M. Engström, M. Mårtensson, E. Avventi, O. Norbeck, and S. Skare, "Collapsed fat navigators for brain 3D rigid body motion", en, *Magn. Reson. Imaging*, vol. 33, no. 8, pp. 984–991, Oct. 2015.

[84] C. Glessgen, D. Gallichan, M. Moor, N. Hainc, and C. Federau, "Evaluation of 3D fat-navigator based retrospective motion correction in the clinical setting of patients with brain tumors", en, *Neuroradiology*, vol. 61, no. 5, pp. 557–563, May 2019.

[85] H. Eggers, "Fat shift correction in bipolar Multi-Echo dixon imaging using Water-Fat separation in k-space", in *Proceedings of the 26th Annual Meeting of ISMRM, Paris, France*, Paris, France, 2018, p. 4217.

[86] S. B. Reeder, B. A. Hargreaves, H. Yu, and J. H. Brittain, "Homodyne reconstruction and IDEAL water-fat decomposition", en, *Magn. Reson. Med.*, vol. 54, no. 3, pp. 586–593, Sep. 2005.

[87] A. Nuttall, "Some windows with very good sidelobe behavior", *IEEE Trans. Acoust.*, vol. 29, no. 1, pp. 84–91, Feb. 1981.

[88] J. D. O'Sullivan, "A fast sinc function gridding algorithm for fourier inversion in computer tomography", en, *IEEE Trans. Med. Imaging*, vol. 4, no. 4, pp. 200–207, 1985.

[89] J. G. Pipe and P. Menon, "Sampling density compensation in MRI: Rationale and an iterative numerical solution", en, *Magn. Reson. Med.*, vol. 41, no. 1, pp. 179–186, Jan. 1999.

[90] C. E. Shannon, "Communication in the presence of noise", *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.

[91] E. H. Park and K.-B. Lee, “Usefulness of black boundary artifact on opposed-phase imaging from turbo spin-echo two-point mdixon MRI for delineation of an arthroscopically confirmed small fracture of the lateral talar dome: A case report”, en, *Medicine*, vol. 96, no. 52, e9497, Dec. 2017.

[92] J. H. You, I. H. Kim, J. Hwang, H. S. Lee, and E. H. Park, “Fracture of ankle: MRI using opposed-phase imaging obtained from turbo spin echo modified dixon image shows improved sensitivity”, en, *Br. J. Radiol.*, vol. 91, no. 1088, p. 20 170 779, Jul. 2018.

[93] H. J. Yoo, S. H. Hong, D. H. Kim, J.-Y. Choi, H. D. Chae, B. M. Jeong, J. M. Ahn, and H. S. Kang, “Measurement of fat content in vertebral marrow using a modified dixon sequence to differentiate benign from malignant processes”, en, *J. Magn. Reson. Imaging*, vol. 45, no. 5, pp. 1534–1544, May 2017.

[94] Y. Maeder, V. Dunet, R. Richard, F. Becce, and P. Omoumi, “Bone marrow metastases: T2-weighted dixon Spin-Echo fat images can replace t1-weighted Spin-Echo images”, en, *Radiology*, vol. 286, no. 3, pp. 948–959, Mar. 2018.

[95] S. Hahn, Y. H. Lee, and J.-S. Suh, “Detection of vertebral metastases: A comparison between the modified dixon turbo spin echo T2 weighted MRI and conventional T1 weighted MRI: A preliminary study in a tertiary centre”, en, *Br. J. Radiol.*, vol. 91, no. 1085, p. 20 170 782, May 2018.

[96] T. J. P. Bray, S. Singh, A. Latifoltojar, K. Rajesparan, F. Rahman, P. Narayanan, S. Naaseri, A. Lopes, A. Bainbridge, S. Punwani, and M. A. Hall-Craggs, “Diagnostic utility of whole body dixon MRI in multiple myeloma: A multi-reader study”, en, *PLoS One*, vol. 12, no. 7, e0180562, Jul. 2017.

[97] M. Seiderer, A. Staebler, and H. Wagner, “MRI of bone marrow: Opposed-phase gradient-echo sequences with long repetition time”, en, *Eur. Radiol.*, vol. 9, no. 4, pp. 652–661, 1999.

[98] D. M. Plecha, “Imaging of bone marrow disease in the spine”, en, *Semin. Musculoskelet. Radiol.*, vol. 4, no. 3, pp. 321–327, 2000.

[99] J. F. Griffith, D. K. W. Yeung, G. E. Antonio, S. Y. S. Wong, T. C. Y. Kwok, J. Woo, and P. C. Leung, “Vertebral marrow fat content and diffusion and perfusion indexes in women with varying bone density: MR evaluation”, en, *Radiology*, vol. 241, no. 3, pp. 831–838, Dec. 2006.

[100] C. M. Costelloe, V. Kundra, J. Ma, B. A. Chasen, E. M. Rohren, R. L. Bassett Jr, and J. E. Madewell, "Fast dixon whole-body MRI for detecting distant cancer metastasis: A preliminary clinical study", en, *J. Magn. Reson. Imaging*, vol. 35, no. 2, pp. 399–408, Feb. 2012.

[101] D. Franz, D. Weidlich, F. Freitag, C. Holzapfel, T. Drabsch, T. Baum, H. Eggers, A. Witte, E. J. Rummeny, H. Hauner, and D. C. Karampinos, "Association of proton density fat fraction in adipose tissue with imaging-based and anthropometric obesity markers in adults", en, *Int. J. Obes.*, vol. 42, no. 2, pp. 175–182, Feb. 2018.

[102] V. Stahl, F. Maier, M. T. Freitag, R. O. Floca, M. C. Berger, R. Umathum, M. Berriel Diaz, S. Herzig, M.-A. Weber, A. Dimitrakopoulou-Strauss, K. Rink, P. Bachert, M. E. Ladd, and A. M. Nagel, "In vivo assessment of cold stimulation effects on the fat fraction of brown adipose tissue using DIXON MRI", en, *J. Magn. Reson. Imaging*, vol. 45, no. 2, pp. 369–380, Feb. 2017.

[103] C. Caussy, S. B. Reeder, C. B. Sirlin, and R. Loomba, "Noninvasive, quantitative assessment of liver fat by MRI-PDFF as an endpoint in NASH trials", en, *Hepatology*, vol. 68, no. 2, pp. 763–772, Aug. 2018.

[104] J. Gu, S. Liu, S. Du, Q. Zhang, J. Xiao, Q. Dong, and Y. Xin, "Diagnostic value of MRI-PDFF for hepatic steatosis in patients with non-alcoholic fatty liver disease: A meta-analysis", en, *Eur. Radiol.*, vol. 29, no. 7, pp. 3564–3573, Jul. 2019.

[105] M. A. Fischer, C. W. A. Pfirrmann, N. Espinosa, D. A. Raptis, and F. M. Buck, "Dixon-based MRI for assessment of muscle-fat content in phantoms, healthy volunteers and patients with achillodynia: Comparison to visual assessment of calf muscle quality", en, *Eur. Radiol.*, vol. 24, no. 6, pp. 1366–1375, Jun. 2014.

[106] J.-P. Kühn, D. Hernando, A. Muñoz del Rio, M. Evert, S. Kannengiesser, H. Völzke, B. Mensel, R. Puls, N. Hosten, and S. B. Reeder, "Effect of multipeak spectral modeling of fat for liver iron and fat quantification: Correlation of biopsy with MR imaging results", en, *Radiology*, vol. 265, no. 1, pp. 133–142, Oct. 2012.

[107] D. Hernando, Y. S. Levin, C. B. Sirlin, and S. B. Reeder, "Quantification of liver iron with MRI: State of the art and remaining challenges", en, *J. Magn. Reson. Imaging*, vol. 40, no. 5, pp. 1003–1021, Nov. 2014.

[108] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin, “Adaptive reconstruction of phased array MR imagery”, en, *Magn. Reson. Med.*, vol. 43, no. 5, pp. 682–690, May 2000.

[109] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, “The NMR phased array”, en, *Magn. Reson. Med.*, vol. 16, no. 2, pp. 192–225, Nov. 1990.