

Optimisation of Phase Data Processing for Susceptibility Reconstruction in Magnetic Resonance Imaging

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften der RWTH Aachen
University zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften
genehmigte Dissertation

vorgelegt von

Diplom-Physiker
Johannes Lindemeyer

aus Marburg

Berichter:

Universitätsprofessor Dr. Nadim Jon Shah
Universitätsprofessor Dr. Achim Stahl

Tag der mündlichen Prüfung: 3. Februar 2015

Diese Dissertation ist auf den Internetseiten der Hochschulbibliothek online verfügbar.

Abstract

Nuclear magnetic resonance (NMR) is exquisitely sensitive to variations of the static magnetic field. Nevertheless, historically mainstream magnetic resonance imaging techniques have focused on the magnitude of the complex NMR signal only. Hence, valuable information about the magnetic field contained within the signal phase was not considered. Only during the last decade has the phase been increasingly employed to obtain structural characteristics complementary to that included in the signal magnitude. Phase is influenced by the physical properties of the imaged object such as chemical shifts and in particular magnetic susceptibility and its spatial distribution. To efficiently exploit the phase information, a number of processing steps have to be performed. The recorded phase is ambiguous and has to be unwrapped, data from multiple receive channels must be combined, fieldmaps have to be obtained and field shifts of external origin must be removed in order to estimate the susceptibility distribution in tissue. In this thesis, the workflow for phase imaging is described step by step. Methods for enhancing each processing step as well as for optimising data acquisition are investigated and the development of several new evaluation techniques is described.

The key algorithms developed in this work bear the acronyms URSULA and MUBAFIRE. URSULA combines one of the most robust known spatial phase unwrapping strategies with volume compartmentalisation, allowing for reliable and fast phase unwrapping of large data arrays as acquired at ultra-high field strength. MUBAFIRE corrects for background fields that originate from sources residing outside of the volume of interest. Its great performance is due to the application of several sequential and complementary background-correction strategies, each preserving physical validity of the solution and addressing different characteristics of field distortions.

The novel algorithms combined with methods adopted from the literature allow for the calculation of detailed field and susceptibility distributions resulting in image contrast that is distinct from that found in magnitude images. The applicability of the established workflow is verified in several *post mortem* brain measurements and in studies on healthy volunteers as well as on patients with brain tumours or Parkinson's disease. In particular, the challenges of performing *post mortem* and *in vivo* imaging on a whole-body 9.4 T scanner - at present, the highest magnetic field available for human phase imaging worldwide - are met by employing the technical innovations developed within the scope of this thesis.

Zusammenfassung

Die Kernspinresonanz (NMR) ist hochempfindlich gegenüber Variationen des statischen Magnetfeldes. In der Vergangenheit basierten jedoch die meisten Magnetresonanz-Bildgebungsverfahren auf der Magnitude des komplexwertigen NMR-Signals. Somit wurde den wertvollen, in der Phase des Signals enthaltenen Informationen über das Magnetfeld wenig Beachtung geschenkt. Erst in den letzten Jahren fand die Phase zunehmend Anwendung bei der Bestimmung struktureller, zur Magnitude komplementärer Kontrastcharakteristika. Die Phase wird von physikalischen Eigenschaften des gemessenen Objektes, wie der chemischen Verschiebung und insbesondere der magnetischen Suszeptibilität sowie deren räumlicher Verteilung, beeinflusst. Um die in der Phase enthaltenen Informationen effizient nutzen zu können, bedarf es einer Reihe von Verarbeitungsschritten. Die gemessene Phase ist mehrdeutig und ihr tatsächlicher Verlauf wird mittels *Unwrapping* rekonstruiert. Daten aus Mehrkanal-Messungen sind zu kombinieren, die magnetische Feldverteilung ist zu berechnen und Feldverschiebungen äußeren Ursprungs sind zu kompensieren, bevor die zugrundeliegende Suszeptibilitätsverteilung im Gewebe abgeschätzt werden kann. Diese Arbeit beschreibt schrittweise den Bearbeitungsablauf der Phasen-Bildgebung und untersucht Methoden zur optimierten Datengewinnung sowie zur Durchführung der einzelnen Arbeitsschritte. Darüber hinaus werden mehrere neuartige Verarbeitungstechniken vorgestellt.

Die bedeutendsten Entwicklungen tragen die Namen URSULA und MUBAFIRE. URSULA kombiniert eine der zuverlässigsten bekannten räumlichen Unwrapping-Methoden mit einer künstlichen Bereichseinteilung der Daten. Dies ermöglicht eine robuste und zeiteffiziente Berechnung der Unwrapping-Ergebnisse von großen Phasen-Datensätzen wie sie etwa in der Ultra-Hochfeld-NMR-Bildgebung aufgenommen werden. MUBAFIRE kompensiert Hintergrundfelder mit Ursprung außerhalb der betrachteten Region. Die Leistungsfähigkeit dieses Algorithmus begründet sich in der Anwendung mehrerer aufeinanderfolgender und zudem komplementärer Filter zur Hintergrundkorrektur. Jeder dieser Filter behandelt Feldverzerrungen unterschiedlicher Charakteristika und bewahrt dabei die physikalische Gültigkeit der Korrektur.

Die entwickelten Algorithmen ermöglichen in Verbindung mit den angewandten, in der Fachliteratur beschriebenen Methoden die Berechnung hochauflösender Feld- und Suszeptibilitätsverteilungen. Die so entstehenden Kontraste unterscheiden sich deutlich vom Magnituden-Kontrast. Die Anwendbarkeit des entwickelten Bearbeitungsablaufs wird durch Messungen an *post mortem* Hirnen und im Rahmen mehrerer Studien an gesunden Probanden sowie an Tumor- und Parkinson-Patienten validiert. Insbesondere aber werden die in dieser Arbeit vorgestellten Entwicklungen den großen Herausforderungen, welche die Phasenbildgebung an einem 9.4 T Ganzkörper-Tomograph – derzeitig, die höchste weltweit verfügbare magnetische Feldstärke für Humanbildgebung – hinsichtlich Akquisition und Datenverarbeitung stellt, gerecht.

Contents

1	Introduction	13
2	Basics of Magnetic Resonance Imaging	17
2.1	Notation	17
2.2	Fourier Transform	17
2.3	Magnetisation	18
2.4	Resonance and Precession	19
2.5	Imaging – Spatial Encoding	23
2.6	General Pulse Types	27
2.7	Basic Pulse Sequences	28
2.8	Spoiling	33
2.9	MR Hardware	34
2.10	Noise	36
3	Phase	37
3.1	Magnitude and Phase	37
3.2	Off-Resonance in MRI Signal Equations	38
3.3	Sources for Off-Resonance	39
3.4	Aliasing and Wraps	41
3.5	Noise	44
3.6	Maximum Phase Accuracy	47
3.7	Non-linear Phase Evolution	48
3.8	Masks	49
3.9	Phase Measurement Example	49

4	Unwrapping the Phase	51
4.1	Basic Considerations	51
4.2	Unwrapping the Phase in Time Domain	52
4.3	Unwrapping the Phase in Spatial Domain	53
4.4	Simple Flood Fill Unwrap	55
4.5	Multi-Path Phase Gradient Integration	55
4.6	PRELUDE	61
4.7	Quality-Based Unwrapping	63
4.8	Sliced Unwrapping	70
4.9	Automated Phase Masking	72
5	Unwrap of Subdivided Large Arrays	75
5.1	Motivation	75
5.2	Subdividing the Array	76
5.3	Unwrapping the Subsets	77
5.4	Reassembling the Array	78
5.5	Implementation	80
5.6	Validation and Optimisation	80
5.7	Pitfalls	84
5.8	Outlook	85
5.9	Application Example	86
5.10	Overview	88
6	Fieldmaps	89
6.1	Single echo	89
6.2	Double Echo	90
6.3	Multiple Echo	92
6.4	Multiple Echo Differences	93
6.5	Shimming	95
6.6	Optimal Fieldmap Estimation	96
6.7	Spatial Distortions	104
6.8	Fieldmap Estimation Assistant	106

7	Magnetic Susceptibility	109
7.1	Susceptibility Weighted Imaging	109
7.2	Physical Background	110
7.3	Models	113
7.4	Reconstructing Susceptibility Distributions	115
7.5	Implementation	124
7.6	Optimising the Parameters	128
7.7	Value Range of Magnetic Susceptibility in the Human Body	135
7.8	Application Example	136
8	Removing Background Fields	139
8.1	Principle	140
8.2	Rudimentary Filters	140
8.3	Polynomial Fitting	143
8.4	Spherical Harmonics	145
8.5	External Susceptibility Distribution	149
8.6	Conclusions	152
9	Multistage Background Field Removal	153
9.1	Motivation	153
9.2	Theory	154
9.3	Material and Methods	157
9.4	Results	161
9.5	Discussion	163
9.6	Conclusions	165
9.7	Effects on Susceptibility Reconstruction	165
9.8	Implementation	166
10	Multiple Rx Recombination	169
10.1	Conventions	169
10.2	Optimal Reconstruction	170
10.3	Complex Sum	170
10.4	Sum of Squares	170
10.5	Smoothed Magnitude Weighting with Constant Phase Correction	171

10.6	Spatially Varying Phase Offsets	174
10.7	Downsampled Offset Correction	176
10.8	Per-Channel Field Estimation	177
10.9	Comparing the Methods	179
10.10	Reconstruction for a 9.4T Example	179
11	Applications	185
11.1	General Workflow for Phase Imaging	185
11.2	Subjects and Specimen	187
11.3	<i>Post Mortem</i> Brain at 9.4T	187
11.4	Visualising Brain Tumours	190
11.5	Susceptibility Reconstruction for a Study of Parkinson's Disease	195
11.6	<i>In Vivo</i> Phase and Susceptibility Imaging at 9.4T	201
11.7	Cerebellum Phase Images at 9.4T	207
11.8	Field in Parallel Fibre Structures	212
11.9	High-Resolution Phase Imaging of the Central Brain at 3T	217
11.10	Fast Susceptibility Imaging	222
12	Summary	229
13	Acknowledgements	233
14	References	235
A	Formulae	247
B	Phase Imaging – Images From a Sample Case	249
C	Field Susceptibility Estimator (FSE) – Function Overview	255
D	MULTistage BACKground Field REMoval (MUBAFIRE) – Function Overview	261

List of Abbreviations

BFR	Background Field Removal
BG	<i>basal ganglia</i>
BOLD	Blood-Oxygen Level Dependence
CG	Conjugate Gradients
CMPLX	Complex Sum Recombination
COSMOS	Calculation Of Susceptibility through Multiple Orientation Sampling
CN	<i>caudate nuclei</i>
CSI	Chemical Shift Imaging
DFT	Discrete Fourier Transform
DN	<i>dentate nuclei</i>
DIPF	Dipole Filter
DownCor	Downsampled Offset Correction
DTI	Diffusion Tensor Imaging
EPI	Echo Planar Imaging
FID	Free Induction Decay
FEST	Fieldmap ESTimation assistant
fMRI	functional MRI
FOV	Field Of View
FSE	Field Susceptibility Estimator
gid	Group ID
GF	Gaussian Filter
GRE	Gradient Recalled Echo
GUI	Graphical User Interface
LC	Local Coherence
LV	Local Variance
MPPGI	Multi-Path Phase Gradient Integration
MRI	Magnetic Resonance Imaging
MUBAFIRE	MULTistage BAcground Field REmoval
NMR	Nuclear Magnetic Resonance
PAL	<i>globus pallidus</i>

PET	Positron Emission Tomography
PGW	Phase Gradient Weighting
POLF	Polynomial Filter
PUT	<i>putamen</i>
P-URSULA	Parallel URSULA
QBU	Quality-Based Unwrapping
QSM	Quantitative Susceptibility Mapping
RF	Radio Frequency
RN	<i>red nuclei</i>
ROI	Region Of Interest
Rx	receive
SCPG	Shift-Corrected Phase Gradients
SD	Second Difference
SE	Spin Echo
SMW	Smoothed Magnitude Weighting
SN	<i>substantia nigra</i>
SNR	Signal-to-Noise Ratio
SPHINX	SPherical Harmonic INhomogeneity-suppressing eXpansion
SSQR	Sum-of-SQuaRes
STI	Susceptibility Tensor Imaging
STN	<i>subthalamic nuclei</i>
SWI	Susceptibility Weighted Imaging
Tx	transmit
URSULA	UnwRap of SUBdivided Large Arrays
VOI	Volume Of Interest
WLC	Weighted Local Coherence

More than three decades ago Bottomley et al. [1983] designed the first 1.5 T device for Magnetic Resonance Imaging (MRI) on humans. Today, MRI has become a well-established modality for non-invasive spatial characterisation of water in different environments. Standard MRI is based on the property of the hydrogen nuclei – the protons – to orient their magnetic moment in an external magnetic field. Large bodies of water in an external magnetic field create a substantial macroscopic magnetisation. The response of the magnetisation to radio waves of a suitable frequency is responsible for signal generation. Furthermore, spatially and temporally varying magnetic fields are employed to encode the information contained in a sample for imaging [e.g. Lauterbur, 1973; Kumar et al., 1975]. MRI has established itself as a fundamental medical imaging discipline [e.g. Henderson, 1983]. The contrast observed in MRI is influenced by the proton density, relaxation parameters of the tissue magnetisation or by the diffusion and flow of liquids [e.g. Haacke et al., 1999]. The influence of these parameters on image contrast can be enforced or attenuated in an extremely flexible way by the acquisition method. However, the majority of the MRI applications are traditionally based on the information contained in the magnitude of the complex-valued signal acquired in MRI, the phase being routinely discarded. Isolated examples of phase imaging reflect mainly pathologies with large deposits of material with high magnetic susceptibility (e.g. calcium). Only recently has the phase of the MR images been investigated as an important source of contrast, acting as a carrier of information on local deviations of the static magnetic field from its average value [e.g. Li et al., 1996; Duyn et al., 2007]. Within the last decades, the sources of phase and field variations were investigated more closely. Besides distortions introduced by imperfections of the static main field, influences are attributed to chemical shifts and, more importantly, to magnetic susceptibility differences within matter [e.g. Lodes et al., 1989; Lai et al., 1996; Reichenbach and Haacke, 2001]. Susceptibility Weighted Imaging (SWI), the first actual phase imaging application, employs the phase contrast generated by local perturbations to enhance the contrast in magnitude images [Haacke et al., 2004]. In the past ten years, methods for the reconstruction of the underlying distribution of magnetic susceptibility have evolved, introducing a novel contrast to MRI – Quantitative Susceptibility Mapping [e.g. Salomir et al., 2003; Marques and Bowtell, 2005; Koch et al., 2006; de Rochefort et al., 2010b; Wharton et al., 2010; Liu, 2010; Schweser et al., 2011b].

The inference of magnetic susceptibility distributions within a measurement volume implies certain requirements and preparatory steps. A model describing the field shifts induced by a susceptibility distribution has to be defined and algorithms for reconstruction, based on this model, have to be implemented. The actual field distribution must be derived from an appropriately configured measurement. This step includes three central topics of phase imaging: acquisition, unwrapping and fieldmap calculation. The acquisition has to be optimised for the best signal quality achievable within the time and resolution constraints [e.g. Deistung et al., 2008]. From measured data only the principal value of the phase is determined (ranging from $-\pi$ to π), whereas the true phase accumulation has an unconfined angular range. Thus, phase unwrapping – reconstruction of the true phase based on measured phase – is probably the most essential step when preparing the phase data and has been approached in various ways [see e.g. Jenkinson, 2003; Abdul-Rahman et al., 2007]. The calculation of fieldmaps based on the phase is strongly intermingled with the two other steps as it depends on the acquisition protocol and the unwrapping strategy has to be chosen according to the characteristics of the data. However, the estimation of fieldmaps that properly represent the actual static magnetic field experienced by the subject does not yet suffice for a convenient reconstruction of susceptibility. The currently employed models relate susceptibility to the field distortion generated as a consequence thereof [e.g. Salomir et al., 2003; Marques and Bowtell, 2005; Koch et al., 2006]. Artefacts and perturbations of external source corrupt the observed static field and lead to dramatic error propagation. Such influences have to be removed by performing *a priori* background field correction. Externally introduced perturbations originate in vague sources, yet they have typical characteristics that can be exploited by customised background correction algorithms [e.g. Abduljalil et al., 2003; Rauscher et al., 2005; Shmueli et al., 2009; Hirsch, 2009; de Rochefort et al., 2010b; Liu et al., 2011; Schweser et al., 2011b]. Finally, phase imaging at ultra-high fields, such as 9.4 T, was found to require a further preparation step. The automatic (manufacturer-provided) recombination of phase data acquired with multi-channel receive array coils leads to valid phase data at clinical field strengths up to 3 T, but tends to generate unusable results for higher field strength. This is mainly due to differing sensitivity profiles and destructive interferences perceived by the individual receive channels [e.g. Bernstein et al., 1994; Hammond et al., 2008; Ros et al., 2008; Robinson et al., 2011]. Hence, signal recombination has to be assessed in post-processing, considering all separate channels.

The current research work has emerged from the requirement to establish susceptibility reconstruction as a tool at the host institute. Since susceptibility reconstruction is only the last step of the long phase processing chain, a closer investigation of the preceding steps was required and this formed the main focus of the thesis work.

In order to provide a logical description of phase imaging, this thesis is arranged in a consecutive order of the involved processing steps. Even though the actual research work started rather at the end of the processing chain, the logical order was chosen. Basic considerations of MRI and phase imaging are presented as an introduction to the topic in Chpt. 2 and Chpt. 3. Subsequently, in Chpt. 4, phase unwrapping methods are discussed and investigated. Classic unwrapping techniques and conventional tools are described as well as advanced, quality-based approaches and custom developments. Chpt. 5 introduces a novel unwrapping algorithm, URSULA, that unifies one of the most robust spatial phase unwrapping strategies known with an artificial volume compartmentalisation and naturally allows for parallelisation capabilities – e.g. for cluster computing. The estimation of fieldmaps and the optimisation of the acquisition parametrisation are dealt

with in Chpt. 6. The theoretical background and the implemented algorithms for reconstruction of magnetic susceptibility distributions based on fieldmaps are described in Chpt. 7. The implemented algorithms from literature include closed-form and minimisation approaches. These approaches are validated and parameter optimisation is attempted. This chapter is followed by the investigation of background field influences observed in MRI fieldmaps (see Chpt. 8), including a detailed assessment of several background removal strategies. A novel approach, MUBAFIRE, is introduced, validated and qualitatively and quantitatively compared to other algorithms in Chpt. 9. MUBAFIRE is a multi-step technique that combines several models for background field correction in a sequential filter chain, whilst preserving the information required for susceptibility reconstruction. The last processing step addressed in this thesis is the recombination of phase data from multiple receive channels. Although this step is actually the very first in a computing chain, multichannel recombination is only necessary at high field strength, so the according chapter was placed in the end. Finally, a variety of applications for the phase imaging process including *post mortem* and *in vivo* measurements at 3 T and at 9.4 T are presented in Chpt. 11. Quantitative evaluation means and various measurement parametrisations are discussed and presented. Although all technical chapters contain examples for the application of the described methods, this last main chapter outlines the benefits and substantiates the applicability of the established phase imaging workflow with respect to the presented measurements.

Following the logic taken herein, this thesis was designed to serve as a guide to phase, field and susceptibility imaging. All presented methods – regardless of being custom-designed or inspired by literature – were specifically implemented and tested, if not stated otherwise, and the best performers are assembled into the standard phase processing method at the host institute.

Basics of Magnetic Resonance Imaging

This chapter shall give a rough introduction to Nuclear Magnetic Resonance (NMR) and more precisely to the acquisition techniques employed in MRI. Although the achievements discussed later do not necessarily require a description of the underlying quantum mechanics, origins of the NMR signal will be described to motivate further explanations.

2.1 Notation

The mathematical notation employed in this thesis is as follows. A scalar is written, a , a vector, $\vec{a}$, and a matrix or operator is written $\mathbf{A}$. Usually, $\vec{r}$ is a coordinate in continuous space, $\mathbb{R}^3$. The Cartesian axes are addressed by the variables x , y and z . The components of a vector, $\vec{a}$, along certain axis, such as the Cartesian ones, are: a_x , a_y and a_z . A unit vector is denoted by a hat, $\hat{a}$, and the Cartesian unit vectors are $\hat{x}$, $\hat{y}$ and $\hat{z}$. The components of the spatial wave vector, $\vec{k}$, acquire a subscript of the corresponding spatial domain axis, such as k_x for the x axis. In discrete space the coordinate or position counters are usually called i , j and k , oriented along x -, y - and z -axis, respectively. They are addressed as a tuple: $[i, j, k]$ or $[ijk]$ (if unambiguous). Functions of continuous space have the form $f(x, y, z)$, $f(\vec{r})$, $\mathbf{F}(x, y, z)$ or $\mathbf{F}(\vec{r})$, while functions of the rasterised image carry the index tuple $f[ijk]$. The real and imaginary part of the complex domain, $\mathbb{C}$, are addressed by the symbols $\Re$ and $\Im$. For a complex number, c , they are $\Re_c$ and $\Im_c$. The exponential function is written as e^x in equations and sometimes as $\exp(x)$ in text. Open interval boundaries are denoted by a reversed bracket, e.g. $[-\pi, \pi[$.

2.2 Fourier Transform

An essential mathematical tool in MRI is the *Fourier transform*. The Fourier transform is an operation that connects $\mathbb{R}^3$ to the wave vector representation and describes their relation. It transforms the familiar spatial representation e.g. in Cartesian coordinates, $\vec{r} = (x, y, z)$, to the frequency domain: $\vec{k} = (k_x, k_y, k_z)$.

Jean Baptiste Joseph Fourier described the relations – or rather transforms – between a function in the spatial domain, $f(\vec{r})$, and in wave vector space, $\hat{f}(\vec{k})$, as:

$$\hat{f}(\vec{k}) = \mathcal{F}(f(\vec{r})) = \int_{\mathbb{R}^3} f(\vec{r}) \cdot e^{-2\pi i \vec{k} \cdot \vec{r}} d^3r \quad (2.1)$$

$$f(\vec{r}) = \mathcal{F}^{-1}(\hat{f}(\vec{k})) = \int_{\mathbb{R}^3} \hat{f}(\vec{k}) \cdot e^{2\pi i \vec{k} \cdot \vec{r}} d^3r \quad (2.2)$$

The MRI signals are recorded in discretised samples, thus the continuous form of the Fourier transform does not apply here. The signals are instead computed using the Discrete Fourier Transform (DFT). For a three-dimensional dataset with $n_x \times n_y \times n_z$ samples, the 3D DFT is defined as:

$$S[k_x, k_y, k_z] = \sum_{x=1}^{n_x-1} \sum_{y=1}^{n_y-1} \sum_{z=1}^{n_z-1} S[x, y, z] \cdot e^{-i2\pi(k_x \frac{x}{n_x} + k_y \frac{y}{n_y} + k_z \frac{z}{n_z})} \quad (2.3)$$

By convention, only the inverse DFT is normalised by $1/(n_x n_y n_z)$:

$$S[x, y, z] = \frac{1}{n_x n_y n_z} \sum_{k_x=1}^{n_x-1} \sum_{k_y=1}^{n_y-1} \sum_{k_z=1}^{n_z-1} S[k_x, k_y, k_z] \cdot e^{i2\pi(x \frac{k_x}{n_x} + y \frac{k_y}{n_y} + z \frac{k_z}{n_z})} \quad (2.4)$$

[see e.g. Bracewell, 1986, p. 358f] – beware that the normalisation coefficients can be applied to forward or backward transform as long as they are consistently used. In approximation, the structure and laws for treatment of the Fourier transform are also properties of the DFT. A fundamental law for discretised measurements is the *Nyquist-Shannon* sampling theorem, which originally states:

„[...] if $f(t)$ contains no frequencies higher than W , it cannot change to a substantially new value in a time less than one-half cycle of the highest frequency, that is, $1/2W$.“
Shannon [1949]

For a digitised image, the highest spatial frequency observable thus has a wavelength of two voxels¹. In imaging this means that for an objective resolution of Δx a maximum spatial frequency of $k_{\max} = \frac{1}{2\Delta x}$ is necessary. Vice versa, the required stepping in frequency domain is implied by the dimensioning of the observed object, $x_{\max}$. This means: $\Delta k = \frac{1}{2x_{\max}}$.

A more detailed description of the Fourier transform can be found in Bracewell [1986] and its applications in MRI are thoroughly discussed in Haacke et al. [1999].

2.3 Magnetisation

The *nuclear magnetic moment*, $\vec{m}_{\text{nucleus}}$, is a consequence of the intrinsic charge and angular momentum properties of a nucleus. The hydrogen (H_1^1) nucleus consists of one proton and has a spin quantum number of 1/2. Being exposed to a static magnetic field the spin orientation

¹A voxel – a volume element – is the 3D equivalent to a pixel – a picture element – which is the basic unit of a rasterised image.

of hydrogen splits into two eigenstates, $+1/2$ and $-1/2$. The underlying effect was initially discovered in the context of spectral line splitting in atoms by Zeeman [1882], and is hence called *Zeeman splitting*. The states can be pictured as an alignment of the nuclear magnetic moment with or opposing to the static magnetic field. The Energy gap between both states is:

$$\Delta E = \hbar\gamma B_0 \quad , \quad (2.5)$$

where $\hbar$ is the *Planck* constant, γ is the gyromagnetic ratio of the observed nucleus and B_0 resembles the strength of the static magnetic field. The population of both eigenstates is given by *Boltzmann* statistics. Inserting Eq. 2.5 into the *Boltzmann* distribution, one gets:

$$\frac{N_{+1/2}}{N_{-1/2}} = e^{-\frac{\Delta E}{k_B T}} = e^{-\frac{\hbar\gamma B_0}{k_B T}} \quad (2.6)$$

[see Haacke et al., 1999, p. 89, Eq. 6.24] (k_B is the *Boltzmann* constant). For hydrogen protons ($\gamma = 42.576 \text{ MHz/T}$) at a static field strength of $B_0 = 1 \text{ T}$ and human body temperature, $T \approx 37^\circ\text{C} \approx 310.15 \text{ K}$, the population difference of the lower energy state ($N_{+1/2}$) and the higher one ($N_{-1/2}$) is about $6.6 \cdot 10^{-6}$. For 10^6 spins in the higher energy state, statistically 6 or 7 more spins populate the lower energy state. This slight imbalance generates an observable macroscopic magnetisation. Beware, that the effective magnetisation increases with the static magnetic field strength, B_0 . Higher field strength such as 9.4 T is thus employed e.g. to gain stronger signals.

2.4 Resonance and Precession

In this eigenstate quantum physics representation a spin can undergo a transition from the low energy state, $N_{+1/2}$, to the high energy state, $N_{-1/2}$. Ascending in energy is caused by the absorption of a photon with frequency ω_L , matching the energy gap (Eq. 2.5):

$$\hbar\omega_L = \Delta E = \hbar\gamma B_0 \quad . \quad (2.7)$$

Descending in energy triggers emission of a photon of the same frequency, ω_L . Further, Eq. 2.7 implies

$$\omega_L = \gamma B_0 \quad , \quad (2.8)$$

the so-called *Larmor Frequency* for a nucleus with characteristic γ at a static field strength of B_0 . More commonly, ω_L is known as the *resonance frequency*.

This frequency has an additional interpretation. In the most accepted model, the nuclear magnetic moment, m_{nucleus} , of a proton is understood to precess with the Larmor frequency, ω_L , about an axis parallel to the static magnetic field, $\vec{B}_0$. The polarity of the precession axis correlates with the polarity of the quantum state ($\pm 1/2$) [Graaf, 2007, p. 7f]. This phenomenon is called *Larmor precession*. Besides being a principle of nuclear physics, Larmor precession of magnetisation is a consequence of the Bloch equation introduced further below (Eq. 2.11). In MRI experiments the amount of spins (further on a synonym for nuclear magnetic moments) that are simultaneously measured is enormous. One observes a superposition of many spins contributing to a macroscopic

magnetisation vector:

$$\vec{M} = \sum_{\forall \text{ nuclei}} \vec{m}_{\text{nucleus}} \quad (2.9)$$

This interpretation preserves the conclusion, that the ensemble generates a resulting magnetisation along the direction of the static main field. In equilibrium the precession of the individual spins is incoherent. Hence, no effective magnetisation can evolve orthogonally to the direction of $\vec{B}_0$. Fig. 2.1 illustrates the model.




Figure 2.1: In the absence of a magnetic field, the orientation and the precession states of the nuclei are undefined (see left box) – strictly speaking no precession is encountered. When exposed to a static magnetic field, B_0 , Zeeman splitting occurs. Picturing this: the precession axis is oriented either parallel (lower energy state, spin $+1/2$) or anti-parallel (higher energy state, spin $-1/2$) with regard to the static field axis (second box). The precession states of the nuclei remain incoherent. The macroscopic magnetisation is thus only generated by the slight imbalance between the energy states (big arrow on the right).

As a convention the direction parallel to $\vec{B}_0$ will be called z direction and $\vec{M}_z$ the z or *longitudinal* magnetisation, accordingly. In equilibrium, magnetisation points in z direction and is assigned the equilibrium value M_0 . The x - y plane is often called *transverse plane* and M_{xy} is the projection of $\vec{M}$ into the plane or the *transverse magnetisation*.

Further on, the quantum mechanical description will be dropped and replaced by the classical model, describing only transformations of the effective magnetisation vector, $\vec{M}$. Two frames of reference are considered in MRI. The *laboratory system* is the frame in which $\vec{M}$ can be observed to precess with ω_L . The *rotating system* pictures the observer rotating “on” the magnetisation vector at ω_L , so that $\vec{M}$ stands still. Using the rotating system simplifies the descriptions of most processes.

2.4.1 Excitation

In short, *excitation* means to willingly influence the current state of $\vec{M}$. In the most simple case, a Radio Frequency (RF) pulse with ω_L is irradiated towards the sample. The magnetisation vector is *rotated* out of its equilibrium state (parallel to $\vec{B}_0$) towards the transverse plane (as in Fig. 2.2) in the rotating frame.




Figure 2.2: Different magnetisation states: A) Equilibrium state – magnetisation is oriented parallel to the field, $M_z = M_0$, $M_{xy} = 0$, the effective magnetisation shows no precession. B) Intermediate state after excitation about angle α – magnetisation has transverse (M_{xy}) and longitudinal (M_z) component. The vector precesses about the z-axis observed in the laboratory system. C) Maximum transverse magnetisation state ($\alpha = 90^\circ$) – ideally, no longitudinal magnetisation remains ($M_z = 0$), while transverse magnetisation is at its maximum ($M_{xy} = M_0$). Precession occurs in the laboratory system. In principle, any excitation angle is possible, including z-inversion.

The resulting angle between the excited and the equilibrium state is called the *flip angle*, α . The process can, for small tip angles, be described by:

$$\begin{aligned} M_{xy,exc} &= M_0 \cdot \sin \alpha \\ M_{z,exc} &= M_0 \cdot \cos \alpha \quad , \end{aligned} \quad (2.10)$$

with $M_{xy,exc}$ being the transverse magnetisation and $M_{z,exc}$ being the longitudinal magnetisation after excitation. Excitation generates coherent precession of the spin ensemble, leading to the actual observable, oscillating transverse magnetisation.

2.4.2 Free Induction Decay

After excitation, the magnetisation vector, $\vec{M}$, seeks to return to equilibrium, the lowest energy state of the system. This cannot happen instantaneously, but strongly depends on the electromagnetic properties and states of the environment of the single spins. The process of M_{xy} returning to the equilibrium state is called Free Induction Decay (FID).

Important: While the excited sample generates electromagnetic radiation at ω_L , the equilibrium state does **not**. So, only the transverse magnetisation M_{xy} can be measured directly.

2.4.3 Relaxation

The two decay processes involved in the FID were first theoretically described by Bloch [1946]. The *Bloch equation* is:

$$\frac{d\vec{M}}{dt} = \gamma \vec{M} \times \vec{B} - \frac{M_x \hat{x} + M_y \hat{y}}{T_2} - \frac{(M_z - M_0) \hat{z}}{T_1} - [\text{additional terms}] \quad , \quad (2.11)$$

with the constants T_1 and T_2 . Assuming that the magnetisation points along the x axis after excitation, a solution for this differential equation is a precession about the z axis, a decay in the

xy plane and a decoupled inverse decay along the z axis:

$$\begin{aligned} M_x(t) &= M_{xy,exc} \cdot \cos(\gamma B \cdot t) \cdot e^{-t/T_2} \\ M_y(t) &= -M_{xy,exc} \cdot \sin(\gamma B \cdot t) \cdot e^{-t/T_2} \\ M_z(t) &= M_0 - (M_0 - M_{z,exc}) \cdot e^{-t/T_1} \end{aligned} \quad (2.12)$$

[e.g. Haacke et al., 1999, p. 58], where T_1 and T_2 are exponential decay constants, describing the so-called *relaxation* of $\vec{M}$ towards equilibrium. T_2 is called *transverse relaxation*, while T_1 is the *longitudinal relaxation*. The sine and cosine terms describe exactly the oscillation at the Larmor frequency, $\omega_L = \gamma B$.

Important: In most biological tissue as well as in most liquids T_1 is significantly longer and thus relaxing slower than T_2 .

Fig. 2.3 shows illustrative examples of signal excitation and decay.




Figure 2.3: Excitation and Decay of the magnetisation vector. The upper row shows the laboratory frame, whilst the lower row shows the demodulated, rotating frame of reference. In the left column excitation is illustrated, whilst the second column shows signal decay. The third and the last column contain graphs for the transverse (M_{xy}) and the longitudinal (M_z) magnetisation in dependence of time after excitation. Note that M_z is not affected by a transition between the frames of reference. The illustrated simulations were optimised for visualisation and do not reflect true decay constants or frequencies.

In the rotating frame of reference ($x, y, z \rightarrow x', y', z$) the solution of *Bloch's* equation is simplified to:

$$\begin{aligned} M_{x'}(t) &= M_{xy,exc} \cdot e^{-t/T_2} \\ M_{y'}(t) &= M_{xy,exc} \cdot e^{-t/T_2} \\ M_z(t) &= M_0 - (M_0 - M_{z,exc}) \cdot e^{-t/T_1} \end{aligned} \quad (2.13)$$

The rotating frame of reference is gained by demodulating the received RF signal with the Larmor frequency, ω_L .

2.4.4 Complex Description

For the NMR signal the evolution of magnetisation in the transverse plane (M_{xy}) plays a predominant role. In order to achieve a more elegant description of the signals, the following convention is chosen:

The xy plane will further on be represented by the complex plane, with the real axis along x and the imaginary axis along y. The transverse magnetisation will thus be described as: $M_{xy} = M_x + iM_y = \sqrt{M_x^2 + M_y^2} \cdot e^{i\angle(M_x, M_y)}$, where $\angle(M_x, M_y)$ is the complex angle.

Using the complex notation, the magnetisation state can be written in a tuple of the form

$$\vec{M} = (M_{xy}, M_z) = (M_x + iM_y, M_z) \quad . \quad (2.14)$$

Eq. 2.12 can now be expressed as:

$$M_{xy}(t) = M_{xy,exc} \cdot e^{-i\omega_L t} \cdot e^{-t/T_2} \quad (2.15)$$

$$M_z(t) = M_0 - (M_0 - M_{z,exc}) \cdot e^{-t/T_1} \quad (2.16)$$

In Eq. 2.13, the resonance term $e^{-i\omega_L t}$ is dropped.

2.5 Imaging – Spatial Encoding

While classical NMR investigates the sample as a whole, identifying elements and chemical compounds, in MRI one seeks to acquire spatially encoded information of the sample, utilising only one certain nucleus, e.g. the proton, H_1^1 . Pioneer work in the field of MR *imaging* was published by Lauterbur [1973] and later by [Kumar et al., 1975]. Encoding the sample is performed by introducing linear *gradients* superimposed to B_0 . In their most general form, the applied field gradient is a tensor:

$$\mathbf{G} = \begin{pmatrix} g_{xx} & g_{xy} & g_{xz} \\ g_{yx} & g_{yy} & g_{yz} \\ g_{zx} & g_{zy} & g_{zz} \end{pmatrix} \quad . \quad (2.17)$$

In MRI usually only the z-component of the static magnetic field is considered, since $B_z \gg B_x, B_y$. Furthermore, the gradients are (ideally) linear. Thus, only the terms of the third row remain, and the gradients can be written in a simplified form – as vectors, $\vec{G} \vec{r}$. Gradients induce a shift of the local resonance frequency:

$$\omega_L = \gamma \cdot (B_0 + \vec{G} \vec{r}), \quad (2.18)$$

(see [Blümich, 2000, p. 19ff]). A typical MRI system comes with separately controllable gradients for x, y and z direction. The Cartesian gradients are usually centred about the *isocentre*² of the static field and their amplitude is measured in T/m. Gradient pulses are characterised by the product of their amplitude and their duration, the *gradient moment*.

²The position where B_0 shows greatest homogeneity – usually near the centre of the magnet bore.

For illustrative purposes homogeneous 90° excitation will be assumed in the following, so that after excitation the affected areas are in transverse state along the positive x axis: $M_{xy} = M_0$ and $M_z = 0$. Further, relaxation between excitation and the encoding steps will be neglected.

2.5.1 Cartesian Imaging

In *Cartesian Imaging* gradients are utilised in an orthogonal arrangement. Due to the linear nature of the three gradients, an arbitrary orientation of the frame of reference can be chosen by superposition of x , y and z fields. In the following, the respective gradients will thus be named after their purpose only. The signal equations in this section are described in the rotating frame, this means they all miss the resonance term, $e^{-i\omega_L t}$, describing precession about the static field axis:


$$M_{xy}(\vec{r}) \cdot e^{-i\omega_L t} \quad (\text{lab. frame}) \quad \longrightarrow \quad M_{xy}(\vec{r}) \quad (\text{rot. frame}) \quad (2.19)$$

2.5.1.1 2D Imaging

2D-imaging relies on a technique called *slice selection*. A linear gradient, the so-called slice selection gradient, $\vec{G}_s$, is activated, during excitation. An RF pulse is applied at the desired resonance frequency, ω_{slc} , which excites the magnetisation of the sample only where the Larmor frequency is matched:

$$(M_{xy}, M_z) = \begin{cases} (M_0, 0), & \gamma \vec{G}_s \vec{r} = \omega_{slc} \\ (0, M_0), & \text{otherwise} \end{cases} \quad (2.20)$$

Figure 2.4: Principle of slice selection. The desired position s_{slc} is addressed by choosing the corresponding gradient, G_s , and resonance frequency, $\omega_{slc} = \gamma G_s s$. A certain slice thickness Δs is selected by the frequency range, $\Delta\omega$ applied. After excitation, the spins in the selected slice are in their transverse state, while ideally all other spins are in equilibrium.



In fact, excitation is not performed using only a single, distinct frequency, but with a frequency bandwidth, $\Delta\omega$. This has several reasons. First, even when using the highest accuracy available, there will always be uncertainties arising from the RF wave generation and the deviations in the linear gradient shapes. And secondly, the signal strength emitted from the sample depends on the excited volume. The volume of an infinitesimally thin slice vanishes and so does the generated signal intensity. Hence, a slice should always cover a certain range, Δs , along the slice selection

direction, $\vec{G}_s$. The variable Δs is called the *slice thickness*. The excitation intensity profile along this direction is called *slice profile*. Pulses are designed with a certain bandwidth to match a chosen slice thickness.

Phase Encoding: After excitation, a second gradient pulse, the *phase encoding gradient*, is applied in a direction orthogonal to $\vec{G}_s$. Before, all excited spins were precessing at Larmor frequency. Now, applying $\vec{G}_p$ over a certain time, τ_p , causes spins at higher field strength to precess faster, and spins at lower field strength to precess slower. This leads to linear phase dispersion along the direction of $\vec{G}_p$, encoding the spins according to their position. After phase encoding is finished, $\vec{G}_p$ is switched off, again. The process can be described as:

$$M_{xy}(\vec{r}) = M_0(\vec{r}) \cdot e^{-i\varphi_p(\vec{r})} = M_0(\vec{r}) \cdot e^{-i\gamma \vec{G}_p \vec{r} \tau_p} \quad , \quad (2.21)$$

with the encoding time, τ_p . The phase dispersion is thus directly linked to the applied gradient moment, $(\vec{G}_p \vec{r})\tau_p$.




Figure 2.5: The principle of phase encoding: By applying a linear gradient to the excited slice (left hand side) along the phase encoding direction, magnetisation is spatially encoded with a starting phase (right hand side). After switching off the gradient, the magnetisation ideally precesses at Larmor frequency everywhere, and keeps its phase encoding.

Frequency encoding Finally, the signal of the sample is read out. But instead of just recording the signal, a third gradient, $\vec{G}_a$ – the *readout gradient* – is applied during acquisition. $\vec{G}_a$ encodes the signal in the remaining direction by frequency shifting: spins at higher field exhibit higher Larmor precession than spins at lower field (see Fig. 2.6).




Figure 2.6: The principle of frequency encoding: The previously encoded phase is still present (left hand side). Depending on the position along the readout (frequency encoding) axis, the readout gradient, G_a , induces different resonance frequencies and thus different phase spans overcoated during the readout block. This is visualised in the right image. Beware that the phase spans right of the neutral position have opposite polarity than the ones on the left.

The transverse magnetisation – still neglecting relaxation – now has the form:

$$M_{xy}(\vec{r}, t) = M_0(\vec{r}) \cdot e^{-i(\omega(\vec{r})t + \varphi_p(\vec{r}))} = M_0(\vec{r}) \cdot e^{-i\gamma(\vec{G}_a \vec{r} t + \vec{G}_p \vec{r} \tau_p)} \quad , \quad (2.22)$$

where t is the time evolution during readout. As a convention, one defines: $\vec{k}_a = \frac{\gamma}{2\pi} \vec{G}_a t = k_a \hat{a}$ and $\vec{k}_p = \frac{\gamma}{2\pi} \vec{G}_p \tau_p = k_p \hat{p}$ ($\hat{a}$ and $\hat{p}$ are unit vectors in the orthogonal directions of the gradients). The emitted signal is the sum of all spins – the resulting magnetisation:

$$S(\vec{k}) = \int_{\mathbb{R}^2} M_0(\vec{r}) \cdot e^{-2\pi i(\vec{k}_a \vec{r} + \vec{k}_p \vec{r})} d^2 r, \quad (2.23)$$

which can be rewritten to:

$$S(k_a, k_p) = \iint M_0(\vec{r}) \cdot e^{-2\pi i(k_a r_a + k_p r_p)} dr_p dr_a \quad (2.24)$$

($r_a = \vec{r} \hat{a}$, $r_p = \vec{r} \hat{p}$). This depicts a Fourier transform of the spatial domain (r_a, r_p) to the frequency domain (k_a, k_p). During readout one records the signal continuously over time – so S is known for a chosen k_a range with the accuracy of the sampling rate. Repeating the measurement multiple times with varying phase-encode gradients, G_p , leads to a series of k_p . Possessing a complete set of $S(k_a, k_p)$ parametrisations, one can transform (2.24) into spatial domain by Fourier transform as in Section 2.1 with two dimensions. This loosely outlines the principle of 2D imaging – single slices or volumes, subdivided in consecutive slices, can be acquired.

2.5.1.2 3D Imaging

Whilst in 2D imaging single slices are excited, true 3D imaging omits the slice selection gradient, $\vec{G}_s$, and excites the entire volume at once. The first phase encode axis from 2D imaging will now be renamed as k_{p1} . Furthermore, a second phase encoding direction, k_{p2} , is introduced. The direction of k_{p2} is arranged orthogonally to the ones of k_{p1} and k_a . The encoding principle is shown in Fig. 2.7 and an example of a measurement sequence is illustrated in Fig. 2.9.

The readout is performed as in 2D imaging. The generated signal equation – again under negligence of relaxation effects – is:

$$S(k_a, k_{p1}, k_{p2}) = \iiint M_0(\vec{r}) \cdot e^{-2\pi i(k_a r_a + k_{p1} r_{p1} + k_{p2} r_{p2})} dr_{p2} dr_{p1} dr_a. \quad (2.25)$$

With a three-dimensional inverse Fourier transform, the spatial distribution of the signal can be reconstructed. The exponent contains an explicit form of the dot product, as $\hat{a}$, $\hat{p}_1$ and $\hat{p}_2$ represent an orthonormal base of $\mathbb{R}^3$. A typical form of writing the exponent is: $i \vec{k} \vec{r}$, which resembles Eq. 2.1. To conclude, the two phase encoding gradients and the frequency encoding have the form:

$$\begin{aligned} \vec{k}_{p1} &= \frac{\gamma}{2\pi} \vec{G}_{p1} \tau_{p1} = k_{p1} \hat{p}_1, & \vec{k}_{p2} &= \frac{\gamma}{2\pi} \vec{G}_{p2} \tau_{p2} = k_{p2} \hat{p}_2, & \vec{k}_a &= \frac{\gamma}{2\pi} \vec{G}_{p1} t = k_a \hat{a}, \\ & & & & & \vec{k}_{p1} \perp \vec{k}_{p2} \perp \vec{k}_a \perp \vec{k}_{p1} \end{aligned} \quad (2.26)$$

The signal acquired in 3D imaging is stronger and less influenced by noise than 2D signals, since it originates in the entire volume instead of only single slices. On the other hand, in 3D imaging each repetition has to wait for the magnetisation to recover from the previous excitation, while for

2D imaging during the recovering process of one slice, several other slices can already be acquired. Both methods suffer from patient motion in different ways.




Figure 2.7: Cartesian acquisition scheme in k -space for 3D sequence. Arrows indicate readouts. Their position stands for phase encoding.

2.5.2 k -space

The frequency domain is by convention called k -space in MRI. The employed acquisition schemes are called *sequences*. For developing sequence concepts, the description of the encoding principle in k -space is an elementary tool. The most basic element in k -space is the *line*. Such a line is swept whenever a readout takes place. In Cartesian imaging, a range of wave vectors from $-k_{a,\max}$ to $+k_{a,\max}$ is recorded. The location of the line is determined by phase encoding $-k_p$ in the 2D case or k_{p1} and k_{p2} in the 3D case. Fig. 2.7 illustrates the k -space acquisition scheme for 3D imaging.

2.6 General Pulse Types

In MRI one differentiates between two RF pulse types: *slice-selective* and *non-selective* pulses. A third kind, situated somewhere in-between these two types, is the *slab-selective* pulse.

Slice-selective The slice-selective pulse was already implied in Section 2.5.1.1. It is played out during switched-on slice selection gradient, $\vec{G}_s$. Depending on the desired slice thickness, Δs , a frequency range of $\Delta\omega = \gamma G_s \Delta s$ has to be synthesised (see Fig. 2.4). Considering Fourier transform, the perfect design for exciting the slice profile of a rectangle function in frequency domain would be the according sinc function ($\text{sinc} = \sin(x)/x$) in time (see Fig. 2.8A), [Bernstein et al., 2004, p. 13]. Time constraints prevent the excitation of a slice with a perfectly sharp profile, as it would take infinitely long. In practice only the central lobe and a few side lobes of the sinc function are used for excitation. A slight smear-out of the ideal profile shape is accepted as a consequence of time-efficient imaging.

Non-selective Non-selective RF pulses do not employ a simultaneous gradient pulse. The most primitive form of a non-selective pulse is continuous irradiation of a wave at the Larmor frequency over a certain time span (see Fig. 2.8C). The flip angle, α , is proportional to the irradiation time and to the amplitude of the applied pulse [Bernstein et al., 2004, p. 36, Eq. 2.1]. Such pulses excite the entire volume and are used mainly in 3D imaging sequences.

Slab-selective A slab-selective pulse is a slice-selective pulse with large slice thickness (see Fig. 2.8B). It is employed for 3D imaging and designed for homogeneous excitation of a volume with finite dimensioning in slicing direction. This can be helpful for focusing on specific regions of the subject and excluding the rest, which would otherwise contribute to the signal.

The three pulse types are sketched in Fig. 2.8. More extensive descriptions of pulse design principles can be found in Bernstein et al. [2004].




Figure 2.8: Basic RF pulse types. In the top row, schematic RF pulse shapes are illustrated. The middle row shows excitation profiles (ordinate) in dependence of the position along the slice selection axis (abscissa). A) slice selection with narrow slice thickness – the profile exhibits excitation imperfections. B) slab-selection – the excited area looks like a very thick slice, but is treated as a 3D volume in imaging. In C), non-selective excitation is illustrated – excitation occurs everywhere. For visualisation, a square cross-section was chosen, although in reality excitation is only confined by the spatial coverage of the exciting coil.

2.7 Basic Pulse Sequences

In the following, two of the most basic pulse sequence layouts will be outlined. These sequences are often used in medical imaging on the one hand and depict the basis for many advanced sequences on the other. Since this thesis mainly deals with 3D Imaging, we will stick to the 3D versions of the imaging sequences.

2.7.1 Gradient Echo

The gradient echo sequence, also called Gradient Recalled Echo (GRE), has a layout closely following the scheme for basic Cartesian imaging, described in Section 2.5.1.1 and Section 2.5.1.2. GRE is the pulse sequence used for the acquisition of most of the data presented within this work. Its structure and properties will be outlined in the following. A more detailed discussion of the principle can be found in Haase et al. [1986].

2.7.1.1 Concept

The simplistic encoding principle and acquisition described before does not embody a real imaging sequence. The reason for this is, that with every gradient applied to the sample, the local magnetisation vectors inside the sample start to disperse. For complete sampling of the frequency domain, the k -space centre must not be neglected. Switching phase encode and readout gradients on and off only allows one to sample the positive axes in k -space. Therefore, it is necessary to introduce a so-called *prephaser* gradient pulse of negative amplitude, to move through k -space in a way that centres the acquisitions around zero. Fig. 2.9 shows a *sequence diagram*, illustrating the events occurring during sequence execution, and a graph outlining the generated k -space trajectory. Note that the readout gradient, and thus its trajectory, is always applied with the same amplitude and duration. Albeit, the gradient moments of the phase encode gradient pulses, G_{p1} and G_{p2} , are adapted to the current phase encoding step, moving the start of the acquisition trajectory to the desired position in k -space. All gradient combinations are applied to sample the full range of spatial frequencies in p_1 and p_2 direction.

2.7.1.2 Parameters

The gradient echo sequence is characterised by standard parameters. The flip angle of the excitation pulse, α , is adjustable to an arbitrary value, affecting the magnitude of the generated transverse magnetisation (or signal intensity). The *echo time*, TE , defines the time interval between the peak of the excitation pulse and the centre of the acquisition block (*ADC*). TE is used to control the image contrast, evolving through the different decay constants of the individual regions in a volume. The *repetition time*, TR , defines the duration of one complete sequence block from excitation to the end. It restricts the time for the magnetisation of the sample to recover and has thus also an effect on the observed image contrast. In the GRE sequence, the predominant contrast for signal decay is the T_2^* contrast (see Section 2.7.2). [e.g. Haacke et al., 1999, Chpt. 15 and 18]

2.7.1.3 Receive Bandwidth

An important parameter for signal acquisition is the *bandwidth* of the readout block. It can be defined in Hz/pixel and implies the slope of the readout gradient, G_a , meaning the maximum frequency difference between spins in a single voxel. The bandwidth is inversely proportional to the length of the readout block and to the noise variance introduced to the sample [Haacke et al., 1999, p. 334, Eq. 15.7; p. 340, Eq. 15.27].




Figure 2.9: The left hand side shows the sequence diagram of the 3D GRE sequence; In the top left, a schematic RF pulse is illustrated; TE is the echo time from excitation to the middle of the acquisition block; TR is the repetition time for the entire block; On the right hand side an example for phase encode and readout trajectory in k -space is illustrated (grey); The coloured projection lines visualise the gradients applied for positioning in k -space. Beware that the readout “pre-phaser” is included in the phase encoding block in the right hand side plot.

2.7.1.4 Multiple Echo Mode

The sequence described so far consists of one acquisition block per excitation. Practically, it is possible to record more than one k -space trajectory for one encoding step without applying another excitation pulse. To illustrate this, we will again consider what happens to the phase of the transverse magnetisation during the readout block. Each gradient pulse encodes the magnetisation in the applied direction. The acquisition block starts with the *prephaser* gradient pulse causing a negative (or positive) starting phase for the magnetisation in the sample along the readout direction (see Fig. 2.10). Then, during readout the magnetisation of all voxels passes through zero phase and ends up in a dephased state opposed to its starting phase. This is achieved by setting the integral of the prephaser pulse to half that of the readout gradient pulse. The illustration shows that one can repeat this procedure multiple times. Adding more readout blocks to the sequence allows for recording of the signal evolution over time. The prephaser gradient pulses for the second and all following acquisition blocks need to have the same moment as the readout gradient pulse³ in order to entirely rewind in k -space. Fig. 2.11 illustrates this method.

³This is twice the moment of the first prephaser pulse.




Figure 2.10: Prephasing and phase evolution during readout for spins along the readout gradient direction.




Figure 2.11: Adding multiple readout blocks to the GRE sequence turns it into a multiple echo GRE. The maximum number of acquisition blocks depends on the intensity of the remaining signal. The echotimes are called TE_1 , TE_2 and so forth.

The described multiple echo GRE sequence uses *monopolar readout gradient pulses*. Alternatively, the readout gradient pulses can be altered in direction – this saves the prephaser pulses and allows one to reduce spacing between the echo times. Fig. 2.12 illustrates this technique. However, bipolar readout will be rarely used within this work, as it can introduce phase offsets between even and odd echoes.




Figure 2.12: Bipolar version of readout gradient – the readout direction changes in every ADC block

2.7.2 Spin Echo

The spin echo sequence is the second basic pulse sequence that is, besides GRE, most frequently used in MRI and classical NMR [Hahn, 1950]. It uses a similar scheme as the GRE sequence, but incorporates an additional RF pulse.

2.7.2.1 Principle




Figure 2.13: Sequence diagram of the spin echo sequence. Beware, that the refocussing pulse is inserted exactly at half of the echo time, TE .

Fig. 2.13 outlines the scheme of this sequence. The excitation, the phase encoding block and the readout block are similar to the GRE sequence. The prephaser gradient pulse of the readout block is commonly separated and applied before the second RF pulse with opposing polarity. The new RF pulse is a so-called *refocussing pulse*. The effect of the pulse is illustrated in Fig. 2.14.




Figure 2.14: From left to right the figure shows: spins after excitation, dephasing, application of refocussing pulse, refocussing of the spins, spins in focused state at TE .

A crucial difference between GRE and Spin Echo (SE) is the observed transverse relaxation. The physical transverse relaxation parameter is T_2 . Local field inhomogeneities cause some isochromates to precess slightly faster or slower what reduces the overall signal in superposition. The reduced relaxation time is called T_2^* with $1/T_2^* = 1/T_2 + 1/T_2'$, where T_2' is the contribution attributed to the dephasing isochromates. While the GRE cannot compensate for this and thus exhibits T_2^* contrast, the spin echo sequence reduces this effect by flipping the magnetisation around 180° , leading to a mirroring effect for the orientations of the contained spins. The faster spins are then behind the in-phase state. The opposite applies for spins being slow before the pulse. Positioning the refocussing pulse in the centre between excitation and acquisition – the echo time, TE – allows the spins to refocus at the centre of the acquisition block.

2.7.2.2 Parameters

The parametrisation of the spin echo sequence is roughly similar to the GRE sequence. The minimum TE is higher for spin echo, because the refocussing pulse has to be inserted between the encoding and the readout block. This affects TR in two ways. The first factor is the duration of the RF pulse. Nevertheless, the crucial factor is that spin echo sequences usually employ 90° excitation, leading to a significantly longer time required by the magnetisation to recover.

2.8 Spoiling

After acquisition one usually wants to dephase all remaining signal sources for the followup acquisition block, this is called *spoiling*. There are two prominent techniques existing, that will be described in the following.

2.8.1 Gradient Spoiling

For the *gradient spoiling* technique an additional gradient pulse is positioned at the end of the acquisition block. The gradient can be applied in one or several gradient axis. It is designed to deliver a high gradient moment that will dephase all excited isochromates by an amount far greater than can be rephased by phase or frequency encoding. Dephasing ideally results in purely longitudinal magnetisation and so does not influence the next excitation and acquisition. Fig. 2.15 illustrates gradient spoiling with the example of a GRE sequence.





Figure 2.15: GRE sequence with a gradient spoiling block at the end. In this case all three axis are switched on for spoiling. Note that TR is extended accordingly, compared to Fig. 2.9

2.8.2 RF Spoiling

RF Spoiling is a technique in which the phase (*not the flip angle!*) of the excitation pulse and the receive (Rx) oscillator is different for each acquisition step. This method prevents magnetisation

imbalances which can emerge due to residual transverse magnetisation from the preceding excitation, when a low repetition time and a high flip angle coincide. Fig. 2.16 illustrates the spoiling principle. RF spoiling is commonly used for GRE sequences of short TR.

Figure 2.16: Rotation of the excitation direction in the transverse plane per excitation/acquisition block.



2.9 MR Hardware

The basic components of an MRI scanner are in short: the magnet, the gradient coils and the RF coils. Further crucial components are the shim coils, various amplifiers and the signal acquisition electronics (often: ADCs⁴). Fig. 2.17 gives a general overview on the architecture of an MRI system. More details can be found in literature [e.g. Haacke et al., 1999, Chpt. 27].

2.9.1 Static Field

The static magnetic field is generated by a solenoid or bipolar/c-shaped magnet. Its bore or opening has to be large enough to accommodate the object to be measured. In the case of measurements presented within this work, this would be a human. Not only has the patient to fit into the magnet, moreover, the static field has to be designed with sufficient homogeneity to allow for imaging. Common fields vary within a *ppm* of their average field strength over a typical field of view (e.g. $500 \times 500 \times 400$ mm). There are three types of magnets: Permanent, resistive and superconducting magnets. While permanent and resistive magnets are mainly employed for low-field scanners with large openings, capable of measuring large volumes, the most commonly used type is the superconducting magnet. Since the employed superconductors operate at temperatures $< 4.2^\circ\text{K}$, they are being cooled with liquid helium.

2.9.2 Gradient Coils

As described before, MRI requires three orthogonal and scalable linear gradient fields. These fields are generated by a set of three coils, that reside just inside the magnet bore. Gradient coils have two characteristic properties: the *maximum gradient strength* and the *maximum slew rate*. Gradient strength has units of T/m, as introduced before. The slew rate describes how fast the gradient strength can be changed and has units of $\text{T}/(\text{m} \cdot \text{s})$. Both parameters are limiting factors in sequence design (increasing them can speed up imaging).

2.9.3 RF Coils

For excitation and recording, appropriate receive and transmit coils are required – these are called *RF coils*. Both types exist in various forms, e.g. implemented in the scanner bore as a so-called

⁴ADC means analogue-digital converter




Figure 2.17: Scheme of a typical human MRI scanner. The upper left image shows a schematic frontal view while the right image shows a cross-section from the side. Further essential system components are outlined: Yellow components control RF transmission and receiving, green components control the gradients, and the blue box indicates shimming control. Shimming is indicated blue, since it corrects B_0 , although the shim coils reside close to the gradient coils. The illustration was partially inspired by [Ullisch, 2012, p. 29].

body coil, arranged in cylindrical or related form to fit around head or extremities (*birdcage coils* and *array coils*), or simple loops or arrays of loops for surface applications (*surface coils* or *surface arrays*). While birdcage and simple surface coils operate with a single channel, array coils are also known as *multichannel coils*

2.9.4 Shim Coils

Although the main magnet is extremely accurately manufactured and has fields precise to a *ppm*, the static field is perturbed, as soon as any object with magnetic susceptibility different from air is introduced. This has a negative impact on imaging, since the presented principles expect ideal conditions – meaning in this regard a flat static magnetic field. *Shim coils* attempt to compensate for such perturbations. Active shims consist of several coils, generating static fields with solid spherical harmonic form. Each coil can be adjusted to optimally compensate B_0 inhomogeneities. Common scanners possess about 8 shim coils, where the linear fields are generated by the gradient coils. Further information on shimming can be found in Section 6.5.

2.9.5 Signal Acquisition

The signal irradiated by the sample is recorded via transmit-and-receive RF coils or dedicated receive-only RF coils. Often, more than one Rx coil is used – in so-called multichannel or array coils. The signals recorded by each channel pass through a low noise pre-amplifier, before they reach further amplifier stages and the AD-converter. A peculiarity of NMR and MRI is, that the signal is recorded as real values, but transformed into a complex-valued signal [Haacke et al., 1999, p. 104ff]:

$$S(t) = \Re_S(t) + i\Im_S(t) \quad . \quad (2.27)$$

The complex representation of the signal is related to its circular polarisation inside the scanner. The recorded waveform is transformed into the complex signal by phase-sensitive demodulation.

2.10 Noise

MRI is generally subject to contamination by noise. Most importantly, the RF receive introduces errors in form of electronic noise from the receiver and amplifier circuits, where the signal is recorded and digitised. Assuming a reasonably well set up experiment, the electronic noise plays the major role of noise generation in the final image. *Haacke* described the noise composition with mainly, body-induced (σ_{body}), coil-induced (σ_{coil}) and electronically-induced noise ($\sigma_{electronics}$) as Gaussian (effectively and in parts) and gives an estimate for the effective noise in k -space: [Haacke et al., 1999, p. 334, Eq. 15.8]:

$$\sigma^2(\vec{k}) = \sigma_{body}^2(\vec{k}) + \sigma_{coil}^2(\vec{k}) + \sigma_{electronics}^2(\vec{k}) \quad (2.28)$$

In imaging noise is quantified by calculating the ratio between the average or the highest signal encountered in a volume and the average intensity in areas where noise dominates. The ratio is known as Signal-to-Noise Ratio (SNR). Noise in the spatial domain representation can be expected to be Gaussian-distributed, as long as the SNR is higher than 3 [see Gudbjartsson and Patz, 1995].

Besides noise, there are various other factors taking influence on MRI measurements. The recorded signal is falsified by imperfections in the static magnetic field, the gradients, unexpected tissue and material properties, but also by the signal shape, pulse design or the excitation and receive profiles, just to name a few.

The signal acquired in MRI measurements is complex valued (see Eq. 2.27). This implies that its representation in the spatial domain is, in general, also complex. There are two equivalent ways of expressing a complex signal, $S(\vec{r})$. Either in the Cartesian complex plane:

$$S(\vec{r}) = \Re_S(\vec{r}) + i\Im_S(\vec{r}) \quad (3.1)$$

or in polar representation:

$$S(\vec{r}) = A(\vec{r}) \cdot e^{i\varphi(\vec{r})} \quad (3.2)$$

Although the MRI hardware generates a signal of type (3.1), the more commonly used form is (3.2).




Figure 3.1: Describing the complex signal value – Polar form ($A \cdot e^{i\varphi}$) and Cartesian form ($\Re_S + i\Im_S$)

3.1 Magnitude and Phase

There are two reasons, why the form in Eq. 3.2 is more prominent. One is that most traditional MRI applications use the magnitude, $A(\vec{r})$, to generate images. The other, more conceptual one, is that $A(\vec{r})$ resembles a value proportional to the transverse magnetisation, M_{xy} , of the

sample. As will be described in the next section, the magnitude exhibits dependencies to most of the tissue-specific parameters, such as relaxation times, magnetisation and proton density. Thus, for a long time, the magnitude was the component which was understood to be the actual signal – the part bearing the information. Later, it was discovered that also the phase, $\varphi(\vec{r})$, carries important information on the measurement and on the observed sample [e.g. Reichenbach et al., 1997].

3.2 Off-Resonance in MRI Signal Equations

In an ideal MRI experiment with perfectly homogeneous static magnetic field, excitation and acquisition profiles, the acquired signal should be real and thus exhibit overall zero phase when observed in the rotating frame:

$$S(\vec{r}) = A(\vec{r}) \cdot e^{i0} = A(\vec{r}) = \Re_S(\vec{r}) \quad , \quad (3.3)$$

where $A(\vec{r})$ is in principle a contrast generated by the local proton density, ρ , the equilibrium magnetisation, M_0 , the relaxation times of the tissue, T_1 and T_2 (or T_2^*) and the sequence parameters, such as the flip angle, α , the repetition time, TR , and the echo time TE .¹ Eq. 3.3 would hence be more precisely expressed as:

$$A(\vec{r}) = A[\rho(\vec{r}), M_0(\vec{r}), T_1(\vec{r}), T_2(\vec{r}), \alpha, TR, TE, \dots] \quad . \quad (3.4)$$

These parameters mainly affect the magnitude, $A(\vec{r})$, and therefore can mostly be neglected when considering only phase. Spatially inhomogeneous excitation and receive profiles shall also be ignored for the moment to give a clean overview on the origin of the image phase in MRI. However, non-zero phase is not originating primarily in the parameters mentioned above. An obvious explanation for phase variation is that the effective resonance frequency varies throughout the observed sample. Considering Eq. 2.25 in the laboratory frame yields inclusion of the resonance term, $e^{-i\omega_L t}$:

$$S(k_a, k_{p1}, k_{p2}, t) = \iiint M_0(\vec{r}) \cdot e^{-2\pi i(k_a r_a + k_{p1} r_{p1} + k_{p2} r_{p2})} \cdot e^{-i\omega_L t} dr_{p2} dr_{p1} dr_a \quad . \quad (3.5)$$

Beware, that relaxation is still neglected in this description and an ideal 90° excitation of the equilibrium magnetisation, $M_0(\vec{r})$, is assumed. The Larmor frequency, ω_L , at the average static field strength, B_0 , will henceforth be named ω_0 . It will be assumed that the actual resonance frequency, $\omega(\vec{r})$, varies slightly with position in relation to ω_0 . Inserting the demodulating term, $e^{i\omega_0 t}$, to express the rotating frame leads to:

$$S(k_a, k_{p1}, k_{p2}, t) = \iiint M_0(\vec{r}) \cdot e^{-2\pi i(k_a r_a + k_{p1} r_{p1} + k_{p2} r_{p2})} \cdot \underbrace{e^{-i(\omega(\vec{r}) - \omega_0)t}}_{\text{off-resonance}} dr_{p2} dr_{p1} dr_a \quad . \quad (3.6)$$

The emerging off-resonance term was neglected before under the assumption that the actual resonance frequency would be $\omega = \gamma B_0 = \omega_0$, everywhere. In that case off-resonance is 0

¹Considering sequences other than the GRE sequence will result in even more parameters influencing the signal.

everywhere and the exponential term vanishes. Transforming Eq. 3.6 back to the spatial domain, leaves:

$$S(\vec{r}, t) = M_0(\vec{r}) \cdot e^{-i(\omega(\vec{r}) - \omega_0)t} , \quad (3.7)$$

where $(\omega(\vec{r}) - \omega_0) \cdot t$ is the phase that evolves from excitation to time t . Beware that phase encoding and readout also depend on time (see Section 6.7). As long as off-resonance is significantly lower than that induced by the field in-/decrease per voxel due to the gradients, off-resonance effects and gradients can be treated as approximately independent.

3.3 Sources for Off-Resonance

In the following potential factors causing the local resonance frequency to differ from the reference frequency, ω_0 , shall be described. Neglecting slight imperfections in the measurement electronics, there are several sources for frequency deviations with a physical reasoning.

3.3.1 Field Inhomogeneities

First of all, the measurement does not take place under ideal conditions, featuring a perfectly constant field $\gamma B_0(\vec{r}) = \text{const.}$ Experiments are subject to imperfections of the static field. The field differs from point to point leading to a spatial distribution of field shifts,

$$\omega(\vec{r}) = \gamma B(\vec{r}) = \gamma (B_0 + b_{\text{inh}}(\vec{r})) . \quad (3.8)$$

It is challenging to generate a static magnetic field that is constant within certain boundaries over a large area in space. Much effort is invested to homogenise the field inside the magnet bore (shimming, see Section 6.5), yet distortions are introduced by any component introduced to the field. Static field shifts, $b_{\text{inh}}(\vec{r})$, are an essential source for off-resonance, although their range is usually kept within *ppm* of B_0 .

3.3.2 Chemical Shifts

Apart from impurities of the static field, especially for protons, there is another significant source for frequency shifts. The hydrogen protons one ideally wants to observe, are the ones bound in water molecules. The body contains various chemical compounds, consisting partially of hydrogen – e.g. nearly all cellular carbon compounds and especially fat. Protons bound in large molecules can be partially shielded from the static magnetic field. The shielding effect mainly originates in the electron hull of the compounds and its constellation – so it directly links to the binding type and state of an atom. This means that the proton nucleus experiences an altered static field depending on the particular chemical compound. This effect is called *chemical shift*. A more detailed description of this topic can be found in [Haacke et al., 1999, Chpt. 17]. The chemical shift is measured as a relative frequency shift compared to water protons:

$$\sigma = \frac{\omega_0 - \omega_\sigma}{\omega_0} . \quad (3.9)$$

The units are commonly chosen to be parts per million (ppm). The chemical shift increases linearly with magnetic field strength. The most prominent shift in human MRI is the *water-fat-shift*. Fat has a chemical shift of about 3.35 ppm [Haacke et al., 1999, p. 422]. At 1 T the resonance frequency of fat is thus about 143 Hz lower than that of water. Chemical shifts can falsify measurements in MRI when not considered properly. On the other hand, chemical shifts are investigated in spectroscopic experiments and Chemical Shift Imaging (CSI) [e.g. Brateman, 1986; Brown et al., 1982].

In the following, chemical shifts will be denoted by b_σ . Considering (3.9), the frequency shift can be expressed as:

$$\omega_\sigma = \omega_0 \cdot (1 - \sigma) = \gamma(B_0 + b_\sigma) \quad , \quad (3.10)$$

where:

$$b_\sigma = -\sigma B_0 \quad . \quad (3.11)$$

3.3.3 Susceptibility-induced Field Shifts

Koch et al. [2006] published a comprehensive article concerning the role of magnetic susceptibility in shifts of the static magnetic field in NMR experiments. The crucial point of the article is, that magnetic susceptibility – in contrast to chemical shifts – generates field distortions with larger extent than the size of the actual structure causing the inhomogeneous susceptibility distribution. Thus, there is no linear or even direct relation, linking the frequency offset of a voxel to the susceptibility located in this position. The theoretical background on how magnetic susceptibility affects the static field and thus the phase is discussed in Chpt. 7. For now, the susceptibility-induced field shifts will be addressed with $b_\chi(\vec{r})$, as the magnetic susceptibility is denoted by the symbol χ .

Summing up the sources for deviations of the local resonance frequency from the system frequency, $\omega_0 = \gamma B_0$, the following expression evolves:

$$\omega(\vec{r}) = \gamma(B_0 + b(\vec{r})) = \gamma(B_0 + b_{\text{inh}}(\vec{r}) + b_\sigma(\vec{r}) + b_\chi(\vec{r})) \quad . \quad (3.12)$$

The observed frequency shifts are better understood as an effect of field shifts. Thus, γ will be omitted in later considerations, and *field shifts*, $b(\vec{r})$, and their origins will be described directly. The resonance *frequency shift* will be defined as:

$$\delta\omega(\vec{r}) = \omega(\vec{r}) - \omega_0 = \gamma(b_{\text{inh}}(\vec{r}) + b_\sigma(\vec{r}) + b_\chi(\vec{r})) \quad . \quad (3.13)$$

The general relation of field and frequency shifts is implied by:

$$\delta\omega(\vec{r}) = \gamma b(\vec{r}) \quad . \quad (3.14)$$

Summing this up, the signal for a voxel at position $\vec{r}$ at time t has the form²:

$$A(\vec{r}, t) \cdot e^{i\varphi(\vec{r}, t)} = A(\vec{r}, t) \cdot e^{i\delta\omega(\vec{r})t} \quad . \quad (3.15)$$

²The minus sign in the exponent is neglected from now on, relating positive phase with positive field shifts.

The actual values measured are just the complex components – this means, that one does not observe the frequency offset, $\delta\omega$, directly, but rather the *phase* of the complex signal, $\varphi(\vec{r})$.

3.4 Aliasing and Wraps

The observed phase has a limited range – depending on its definition this might be $[-\pi, \pi[$ or $[0, 2\pi[$. Within this chapter, the phase will be defined in balanced form, as $[-\pi, \pi[$.³ However, this means that frequency or field offsets are always mapped to this range.

The observed phase, φ , is a wrapped representation of the true phase, φ_{true} . Wrapping will be expressed using the *wrap operator*, $\mathcal{W}$, which extracts the division remainder of $(\frac{\varphi_{\text{true}} + \pi}{2\pi} - \pi)$:

$$\varphi = \mathcal{W}(\varphi_{\text{true}}), \quad \mathcal{W} : \mathbb{R} \rightarrow [-\pi, \pi[\quad (3.16)$$

Fig. 3.2 shows how an arbitrary one-dimensional phase-function is *wrapped* into the measurable interval.⁴




Figure 3.2: Example for phase evolution, shown as source (true) and wrapped (measured) graph.

3.4.1 Temporal Aliasing

The phase evolution of a single voxel after excitation will now be considered. The magnetisation in this voxel experiences a field offset, b , with respect to B_0 . According to Eq. 3.13 this leads to a frequency offset of $\delta\omega = \gamma b$. For a given time, t , after excitation this means for the signal of the voxel:

$$S(t) = A(t) \cdot e^{i\delta\omega t}, \quad (3.17)$$

and for the phase:

$$\varphi_{\text{true}}(t) = \delta\omega t. \quad (3.18)$$

It will be assumed that within the observed time frame the magnitude remains sufficiently strong to determine representative phase values. The true accumulated phase, φ_{true} , is projected on the

³Beware, that $+\pi$ is excluded, since $e^{i\pi} = e^{-i\pi}$. It might, anyway, be used occasionally as a phase value, when understanding of the context is simplified by doing so.

⁴In fact, the function shown is not typical for the temporal evolution, but rather for spatial phase characteristics – it was chosen for illustration of the wrapping process, only.

interval $[-\pi, \pi[$, resulting in the measured phase, φ . Measuring a phase evolution usually leads to discontinuities. These occur whenever the true phase passes through $\pi \pm n \cdot 2\pi$ ($n = 0, 1, 2, \dots$). The discontinuities are known as *aliasing* or *wrapping* of the phase.


Fig. 3.3 shows three different phase evolutions typical for the time domain. Multiple echo GRE sequences (see 2.7.1.4) are commonly employed to acquire information on the evolution of phase over time. Thus, the phase can only be measured for certain discretised time steps, e.g. implied by the GRE echo time spacing. A maximum phase gradient is implicitly defined, ensuring unambiguous determination of the local shift frequency, $\delta\omega$. Measuring multiple echoes with an echo time distance of $TE = \Delta t_{TE}$ leads (according to the *sampling theorem* mentioned in Section 2.2) to a maximum observable frequency of:

$$\delta\omega_{\max} = \frac{\pi}{\Delta t_{TE}} \quad (3.19)$$

The maximum observable field shift is thus:

$$b_{\max} = \frac{\delta\omega_{\max}}{\gamma} = \frac{\pi}{\gamma \Delta t_{TE}} \quad (3.20)$$

Figure 3.3: Different evolution graphs of phase in time. Graph A) and B) show evolutions with a positive and a negative slope ($\delta\omega$) which can be estimated correctly everywhere, while C) shows the case of an incorrectly estimated phase evolution – the measurement spacing is too wide to assume the correct slope.



Out of practical reasons, the echo time difference, Δt_{TE} , cannot be reduced infinitely. The readout duration increases inversely proportional to the bandwidth [e.g. Haacke et al., 1999, p. 238, Eq. 12.15], so it has a minimum, given by the measurement parametrisation. Besides that, keeping the bandwidth low enhances the SNR of magnitude and phase (see Section 6.6.2) – hence in practice reduction of Δt_{TE} is only possible within a very limited range [Haacke et al., 1999, p. 334, Eq. 15.7]. Recent literature [Dagher et al., 2014] shows attempts on circumventing the Nyquist limitation by optimised echo positioning, enabling field estimation in undersampled scenarios.

3.4.2 Spatial Aliasing

Phase values are not only acquired in time domain, but also in up to three spatial dimensions (e.g. for a 3D GRE measurement). One actually observes three additional and potentially dependent phase evolutions, simultaneously. Every line of voxels along one of the three dimensions shows

a different phase evolution and different wraps. Phase discontinuities can further occur between a voxel and all off its neighbours – meaning six connections. In an actual measurement the slopes observed in the spatial representation of the phase are commonly smoother than the ones observed in time (see the example phase measurement in Fig. 3.12). Yet, physically correct phase discontinuities occur and have to be taken into account when processing phase for any purpose.

3.4.2.1 Phase Gradients

For a rasterised phase image one can define phase gradients along the Cartesian directions as:

$$\begin{aligned}\partial_x \varphi &= \varphi_{i,j,k} - \varphi_{i-1,j,k}, & (0 \text{ for } i = 1) \\ \partial_y \varphi &= \varphi_{i,j,k} - \varphi_{i,j-1,k}, & (0 \text{ for } j = 1) \\ \partial_z \varphi &= \varphi_{i,j,k} - \varphi_{i,j,k-1}, & (0 \text{ for } k = 1)\end{aligned}\tag{3.21}$$

$$\tag{3.22}$$

($i, j, k = [2, \dots, n_{i,j,k}]$, the number of voxels in the particular dimension). Fig. 3.4 shows the spatial evolution of a phase measurement and its phase gradients. E.g. Bakker et al. [2008] used phase gradients for the illustration of susceptibility-induced field perturbations.





Figure 3.4: From left to right: Sample slice of phase data from measuring a fixed brain placed in a formalin-filled cylindrical container, phase gradient in x (horizontal) direction, phase gradient in y (vertical) direction, phase gradient in z (slicing) direction, gradient norm, $\sqrt{(\partial_x \varphi)^2 + (\partial_y \varphi)^2 + (\partial_z \varphi)^2}$. Wraps and noise become predominant in gradient contrast, whilst the contrast of local structures remains mostly below the strength of wrap gradients. Units are rad.

3.4.2.2 Fringe Lines and Surfaces


In two as well as in three dimensions a wrapped phase function, based on a uniformly continuous true phase function, will generate areas in which phase wraps occur – meaning, where the function spontaneously switches from $+\pi$ to $-\pi$ or vice versa. Here, the phase gradient shows local maxima or minima due to the $-\pi \rightarrow \pi$ or $\pi \rightarrow -\pi$ transition. These areas tend to form surfaces in space – denoted by *fringe lines* or *wrap lines* in 2D and *wrap surfaces* in 3D. Often they are simply

called *wraps*. Depending on the nature of the true phase, these surfaces can be either closed (e.g. in spherical form) or open, ending at the dataset boundaries. The minimum number of wrap surfaces is determined by the number of phase overflows. For a true phase range of $[-8\pi, 6\pi[$ at least 7 surfaces exist: $-7\pi, -5\pi, -3\pi, -\pi, \pi, 3\pi, 5\pi$. Wrap surfaces, belonging to one and the same transition, e.g. 3π , can emerge into one continuous surface or appear in a number of multiple disjoint surfaces. Fig. 3.5 illustrates a two-dimensional phase map.



3.4.2.3 Poles

The observed phase in MRI measurements does not necessarily have its origin in a uniformly continuous function. Moreover, the phase is biased by various influences, such that a smooth form of the signal with well-defined surfaces is not guaranteed by default. Especially for measurements with low SNR phase uncertainties tend to evolve. Very common artefacts in phase data are the so called *phase poles*. Within a square of four neighbouring voxels in a 2D phase image, a pole is observed, when the circular phase evolution over the voxels shows **only one** wrap. Fig. 3.6 illustrates this and Fig. 3.7 shows an example for a measured phase pole.



3.5 Noise

A crucial parameter for a significant measurement is the noise level. Thus, the characteristics and methods for quantification of phase noise will be discussed in short.




Figure 3.7: Sample slice from a phase measurement, showing a pronounced phase pole in its centre, indicating inconsistent phase data

3.5.1 Theoretical Relation

The noise in phase data can have numerous sources. Usually, suitable MRI data has an average magnitude SNR higher than 5. Under this premiss, a Gaussian noise distribution can be assumed for the complex signal [see Gudbjartsson and Patz, 1995]. The uncertainties, σ^2 , of the real and the imaginary part, $\Im_S$ and $\Re_S$, in the complex value affect magnitude, A , and phase, φ , as follows:

$$\begin{aligned}\sigma_A^2 &= \sigma^2 \\ \sigma_\varphi^2 &= \frac{\sigma^2}{A^2} = \frac{\sigma_A^2}{A^2}\end{aligned}\tag{3.23}$$

(as described in [Haacke et al., 1999, p. 374f]).

This means that the phase uncertainty grows with decreasing magnitude. This is self-evident due to the relation of phase and complex axes: $\varphi = \pm \tan^{-1}(\Im_S/\Re_S)$. The closer A , and thus $\Im_S$ and $\Re_S$, get to σ , the more uncertain becomes the $\tan^{-1}$ function. However, the noise calculation introduced above is fairly general, as it is simply relating magnitude variance to phase variance. It does not consider influences on phase signals that do not affect the magnitude in the same way. In actual measurements the relation of magnitude and phase variation is also influenced by the magnitude and phase contrast generated in the sample. Hence, if one intends to quantify the phase SNR or variation of an actual measurement, it is advantageous to estimate these values, based directly on the acquired phase data.


3.5.2 Practical Consideration

Straightforward methods exist for the determination of the SNR in magnitude images [e.g. Firbank et al., 2000]. One approach is the definition of signal and noise areas: One defines a region for which the SNR is to be determined and a reference region located outside the signal-emitting geometries. Both regions should consist of a statistically relevant number of voxels. The noise region should have low signal and exhibits mainly noise. The ratio of the average signal divided by the standard deviation of the noise region gives a measure for the SNR. Alternatively, the histogram of the measurement can be used to calculate the ratio between the signal peak and the

noise peak, leading to a rough estimate for the average SNR. Note that this simplistic approach is not valid for images reconstructed using so called “partial parallel imaging” techniques, causing the SNR to become spatially varying.

For phase data, SNR determination is somewhat more cumbersome. Phase exhibits all values between $-\pi$ and π . Further, areas containing technically no signal do usually provide completely random phase. Phase may contain gradients or arbitrarily shaped variations related to field or tissue structure properties. An example for a phase histogram is shown in Fig. 3.8

Figure 3.8: Histograms of magnitude and phase in (non-masked) 3D volume of GRE measurement. While the magnitude shows characteristic noise and signal peaks, the phase shows a rather random shape with slight centring tendency around zero. A phase peak near the centre is induced by a nearly homogeneous phase region.



The biggest problem in applying statistical methods on the phase, is the differentiation between signal and noise. The phase signal to be observed might have either strong phase contrast, exhibiting various local structures, or it might be a vast area of constant signal, e.g. an area with a constant field offset and just little variations. The latter case would be easy to distinguish from noise, whilst the former would probably not be separable.

One applicable method to get an estimate for phase SNR is the *phase variance*, described in the following section.

3.5.3 Measuring Phase Variance

The phase variance or its square root, the standard deviation, are means of estimating the variation of the phase within a certain neighbourhood. It can be analysed to compare similar datasets or to assess the phase quality within different regions of a single phase map [e.g. Witoszynskyj et al., 2009]. The variance can be calculated directly on the source data. It is usually determined for a cubical neighbourhood of each voxel (with dimensioning $\pm h$). Due to the discontinuous representation of the measured phase, one tends to calculate the standard deviation of the complex unity phase:

$$STD_h[ijk] = \text{std}(\forall_{|l-i| \leq h, |m-j| \leq h, |n-k| \leq h} \Phi[lmn]), \quad \Phi[lmn] = e^{i\varphi[lmn]}, \quad (3.24)$$

rather than of the wrapped phase, φ . The variance or standard deviation of φ is susceptible to phase wraps and automatically assigns a high value in their surroundings, whilst Φ is immune against wraps. Fig. 3.9 illustrates the properties of the variance calculated from wrapped and complex unity phase.

It should be noted, that the phase variance is a measure for steadiness of the phase, but can respond to true phase contrast of rapidly changing geometry in the same way as to phase noise




Figure 3.9: Simulated phase data. Top row: raw (wrapped) phase, φ , middle row: standard deviation of wrapped phase, bottom row: standard deviation of complex unity phase, Φ . From left to right: Phase noise with standard deviation 0, 0.2, 0.4, 0.6 and 0.8 rad added to φ . The right hand side of each image contains a repeating structure inducing true phase contrast. The standard deviation of φ recognises wraps as noise. Even in the absence of noise, the true contrast is recognised as a variation with the same strength as the actual noise with a level of 0.4 and higher.

(see Fig. 3.9). Thus, the phase variance has to be considered with care. In order to exclude phase gradients, it is useful to apply an *a priori* mean or median filter with a kernel width of several voxels, e.g.

$$\text{STD}_h[ijk] = \text{std} \left[\nabla_{|l-i| \leq h, |m-j| \leq h, |n-k| \leq h} \Phi[lmn] - \text{mean}(\nabla_{|o-l| \leq g, |p-m| \leq g, |q-n| \leq g} \Phi[opq]) \right], \quad (3.25)$$

(with g being the neighbourhood size), before determining the local standard deviation. Alternatively, a polynomial filter can be used to identify and remove slowly changing background variations, if no wraps occur in the observed region. Unlike the SNR the phase variance is an inverse measure for quality – the higher the variance or standard deviation, the worse is the phase quality. Phase variance is further discussed in Section 4.7.

3.6 Maximum Phase Accuracy

Even though the general relation between magnitude and phase was discussed before, there are more confinements to the accuracy of phase data, obtained within MRI measurements. Imaging data from human measurements, especially those of three-dimensional acquisitions, require vast amounts of hard disk space. To give an example, an imaging matrix of $320 \times 320 \times 256$ with 6 echoes is considered. If only the recorded data of one Rx channel (e.g. for a single-channel coil) is stored, including magnitude and phase, and *double precision* (4 bytes) is used for each dataset, 600 MB of storage space would be required. Thus, scanner manufacturers and researchers make

use of a truncated number range. For Siemens (Siemens AG, Erlangen, Germany) scanners, the default configuration is an integer range of $[-4096, 4096]$ for phase, and $[0, 4096]$ for magnitude. Although usage of full, double precision, raw data is in principle feasible, it is usually avoided due to the described size issues. In short, this means, common MRI datasets have a minimum uncertainty or maximum accuracy of:

$$\delta\varphi_{min} = \frac{360^\circ}{8192} \approx 0^\circ 2' 38'' = \frac{2\pi}{8192} \text{ rad} = 7.67 \cdot 10^{-4} \text{ rad} \quad . \quad (3.26)$$


3.7 Non-linear Phase Evolution

In this work, linear phase evolution over time is assumed (introduced in Eq. 3.15). Nevertheless, it has been reported, that phase data can exhibit non-linear phase evolution. Although, MRI measurements are intuitively pictured as a composition of intrinsically homogeneous voxels, one should be aware of the partial volume effect. Every voxel only shows an average of the magnitude or phase accumulated inside. This is still a simplified illustration of the problem. In fact, the discretised measurement result is heavily influenced by the point spread function, characteristic for the employed sequence. This does not only imply a “smear-out” effect, as in averaging over a certain window of a linear function, but can result in much more cumbersome superposition effects.

In general, a voxel cannot be expected to constitute of only one kind of material with a common magnetisation, M , common decay constants, T_1 , T_2 (or T_2^*) and a constant static magnetic field shift, b . A voxel may constitute of several so-called *pools*. It will be assumed now, that one certain voxel consists of two pools, p_1 (20% of the volume) and p_2 (80%). The pools may have the same magnetisation at $t = 0$, $S_0 = S_{0,1} = S_{0,2}$, and different T_2^* decay times, $(T_2^*)_1 > (T_2^*)_2$. Further, the field shifts experienced by the pools are $b_1 < b_2$. The signal evolution of the average signal of this voxel now has the form:

$$S(t) = S_0 \cdot \left(0.8 \cdot e^{-t/(T_2^*)_1} \cdot e^{i\gamma b_1 t} + 0.2 \cdot e^{-t/(T_2^*)_2} \cdot e^{i\gamma b_2 t} \right) \quad (3.27)$$

Fig. 3.10 shows an example for the phase evolution of the two pools, $\angle S(t)$.



Due to the differing exponential decay of both pools and the linear phase evolution of each pool, an average curve evolves, not even closely related to a linear phase evolution over the entire voxel. In fact, these effects are very weak, but offer new ways of understanding phase information.

Anyway, for the analysis of the phase evolution, as presented within this work, these effects will be neglected.

In literature, non-linear phase evolution is e.g. described in: Wharton and Bowtell [2012]; Chen et al. [2013]; Schweser et al. [2011a]; Wharton and Bowtell [2012].

3.8 Masks

In MRI it is common to define masks in order to identify the areas of a discretised measurement that contain valuable information. For measurements of the human head, e.g. such a mask usually segments the brain volume, while excluding all other sources of signal. For phase analysis, masks are of crucial importance, since regions of low signal exhibit phase values with a random range. These have to be excluded from observation.

Within this work, masks will be treated as a preceding function, $m[ijk]$. Depending on the index tuple the mask function takes the values:

$$m[ijk] = \begin{cases} 1, & [ijk] \text{ is a valid voxel} \\ 0, & [ijk] \text{ is an invalid voxel} \end{cases} \quad (3.28)$$

When applied to a phase map, the result is:

$$m[ijk] \cdot \varphi[ijk] = \begin{cases} \varphi[ijk], & m[ijk] = 1 \\ 0, & \text{elsewhere} \end{cases} \quad (3.29)$$

For an invalid voxel the result might alternatively be discarded, depending on the convention used in the particular situation. An example for a phase measurement with a confining brain mask is shown in Fig. 3.11. An automated method for mask estimation based on the phase is presented in Section 4.9.




Figure 3.11: Sample slice from phase data of 3D GRE. Left: the raw phase, Mid: brain confinement mask, Right: masked phase.

3.9 Phase Measurement Example

In order to illustrate the characteristics of a typical phase measurement, an example dataset originating from an MRI measurement is shown in Appendix B. Examples for phase evolution over time and space are illustrated in Fig. 3.12. Fig. 3.13 facilitates a comparison between magnitude and phase evolution over echo time in a sample slice. Fig. 3.14 shows the phase noise compared to the phase itself.




Figure 3.12: Phase evolution plots in spatial domain (along x-direction, indicated by a yellow line in the slice plot) and in time domain (indicated by a red dot in the slice plot and a red arrow in the slice stack). The lower graph shows echo numbers, indicating equally spaced echo times.




Figure 3.13: Phase (upper row) and magnitude (lower row) evolution with echo time (increasing from left to right). The phase value range is $[-\pi, \pi]$. The magnitude is shown in arbitrary grey scale, but consistently for all echoes.




Figure 3.14: Phase (upper row) and phase variance (lower row), both un-masked, over echo time (increasing from left to right). The value range is $[-\pi, \pi]$ for phase and $[0, 1]$ for the phase variance in greyscale colouring. Apparently, with increasing echo time local structures can easily be misinterpreted as noise due to the strong phase variation induced by these.

Unwrapping the Phase

The phase data acquired in MRI measurements represent only a truncated observation of the true phase evolution that takes place in the sample. The observed phase is confined to the circular phase range described by the wrap operator, $\mathcal{W}(\varphi) \rightarrow [-\pi, \pi[$, whereas the true phase evolution has no such confinements. In order to obtain the true evolution of phase, an essential intermediate step – *phase unwrapping* – has to be applied. *Unwrapping* stands for the process of estimating the true phase based on the truncated or wrapped phase observed in a measurement and has been investigated in various research areas such as radar inteferometry (e.g. Goldstein et al. [1988]). Different approaches on unwrapping phase data based on the previously introduced models for the evolution of phase in space and over time will be discussed in this chapter.

4.1 Basic Considerations

The unwrapping operation ideally imposes the inverse of the wrap function, $\mathcal{W}$ (see Eq. 3.16). It will be denoted by an *unwrap operator*, $\mathcal{U}$:

$$\mathcal{U} = \mathcal{W}^{-1}, \quad \mathcal{U} : [-\pi, \pi[\rightarrow \mathbb{R} . \quad (4.1)$$

The transformation can be described as:

$$\begin{aligned} \varphi_{\text{unw}}(\vec{r}, t) &= \mathcal{U}(\varphi_{\text{meas}}(\vec{r}, t)) \\ &= \varphi_{\text{meas}}(\vec{r}, t) + (2\pi) \cdot n(\vec{r}, t), \quad n \in [\dots, -2, -1, 0, 1, 2, \dots], \end{aligned} \quad (4.2)$$

where $n(\vec{r}, t)$ is an unknown function in space and time with integer co-domain, φ_{meas} is the measured and φ_{unw} the reconstructed or unwrapped phase. The function $n(\vec{r}, t)$ will be called the *wrap count* between φ_{meas} and φ_{unw} . Ideally, after a successful unwrap $\varphi_{\text{unw}}(\vec{r}, t) = \varphi_{\text{true}}(\vec{r}, t)$ (the true phase). Yet, even an ideal unwrap determines the wrap count and thus the unwrapped phase only up to a constant global offset, c :

$$\varphi_{\text{unw}}(\vec{r}, t) = \varphi_{\text{true}}(\vec{r}, t) + (2\pi) \cdot c , \quad (4.3)$$

In principle, c is part of the wrap count, $n(\vec{r}, t)$, but remains undetermined, since a specific point of reference between the true phase and the measured phase can hardly be defined.

4.2 Unwrapping the Phase in Time Domain

In phase imaging the focus is set on the evolution of the phase over time, which is related to the local shift of the static magnetic field (see Eq. 3.13). Practically, this indicates that instead of approaching the entire measurement volume at once, only the phase evolution over time needs to be unwrapped for the individual voxels. The problem is hence reduced to one-dimensional unwrapping. The *1D unwrap function*, $\mathcal{U}_1(\varphi_{\text{meas}}[i]) = \varphi_{\text{true}}[i] + (2\pi) \cdot n[i]$, can be described by the primitive approach:

$\mathcal{U}_1([-\pi, \pi] \rightarrow \mathbb{R})$:

$$\Delta n[i] = \text{round} \left[\frac{\varphi_{\text{meas}}[i] - \varphi_{\text{unw}}[i-1]}{2\pi} \right] \quad (4.4)$$

$$\varphi_{\text{unw}}[i] = \varphi_{\text{meas}}[i] + \Delta n[i] \cdot (2\pi) \quad (4.5)$$

This procedure has to be carried out iteratively over the counter, i , denoting the sequential values in the phase vector. In simple words:

The difference between the current and the previously unwrapped data point is calculated. If the difference exceeds $\pm\pi$, a multiple of 2π is added to or subtracted from the current point.

Fig. 4.1 illustrates the principle of one-dimensional unwrapping for a positive phase slope.




Figure 4.1: One-dimensional unwrapping for a positive phase slope. The measured data match the unwrap in the beginning, until the iterative unwrapping procedure encounters a phase difference larger than π . The red arrows illustrate the $n \cdot 2\pi$ adding operation, correcting the measured data.

As long as the Nyquist criterion (Section 2.2) is fulfilled everywhere – meaning, that the *true phase* contains no slope inducing a phase change of more than $\pm\pi$ between two data points (see Chpt. 3, Fig. 3.3) – this method works perfectly. Nevertheless, problems arise as soon as phase noise is considered (see Fig. 4.2).




Figure 4.2: Example for failing unwrap due to noisy phase. The true phase (green) roughly follows the linear evolution (green eye guide), but suffers from noise. The unwrap fails, when suddenly a measurement point is closer to the previous point than π , but should have been 2π -shifted, to match the true phase. All following points are corrupted due to this failure.

In that case more adept methods such as prospective regression fits, as e.g. described by Windischberger et al. [2004] or Hirsch [2009] have to be applied. In this method the wrap count function $n(t)$ is approximated by seeking for an optimal fit of the data points to a linear slope. It is based on the assumption, that phase evolution is approximately linear. Nevertheless, this method cannot completely circumvent the Nyquist limitation, so one-dimensional unwrapping can still lead to incorrect estimates for the phase evolution.

4.3 Unwrapping the Phase in Spatial Domain

Unwrapping in time domain seems straightforward, but in general unwrapping of the data in the spatial domain is inevitable. Whilst in spatial domain the phase changes between voxels are relatively small, the evolution in time domain may undergo rapid changes. This is not least due to the fact, that temporal sampling – the echo spacing in multiple echo GRE – has a lower limitation (see Eq. 3.20). A temporal unwrap applied to the time dimension of a 3D multiple echo dataset is likely to fail. Spatial phase wraps will occur in the unwrap result, in areas where the local field shift exceeds the maximum observable shift in time domain. Thus, commonly phase analysis requires unwrapping in the spatial domain. Fig. 4.3 illustrates this motivation.

Inside a volume all voxels have neighbours in three dimensions – direct neighbours at the surfaces and less proximate ones at the corners. Considering all of them, a voxel has six direct neighbours, 18 neighbours when including the edges and even 26 neighbours in the proximity when corners are included (see Fig. 4.4). In one dimension the unwrapping procedure is distinctly following the data vector – forward and backward processing are equivalent in their result. In three dimensions, one has to derive an optimal sequence in which the voxels are unwrapped against each other. Additionally, every voxel's phase has to fulfil certain edge conditions to *all* of its neighbours – that means, it might be necessary to address a voxel more than once while unwrapping. An essential advantage of spatial unwrapping is, that the limitation of the maximum observable frequency in time domain can be bypassed, if at least one region exists, in which it is fulfilled and all other relevant regions of the volume are connected to it by spatial phase relations that fulfil the sampling limits. In the following several approaches on two- and three-dimensional phase unwrapping will be outlined.




Figure 4.3: Phase evolution in space and time. A sample slice of a phase measurement at 9.4 T is shown in the left top. Below, the same slice is shown for the four recorded echo times. A red and a blue reference point are indicated in the slice plots, showing in the top image a spatial phase path along the vertical axis and indicating the time relation of the phase in the slice series below. The plots on the right hand side show, from top to bottom: Spatial and temporal phase evolution for the blue reference point, and spatial and temporal phase evolution for the red reference point. The spatial phase evolution appears smooth in both cases, exhibiting only few jumps where wraps occur. The temporal evolution is low for the blue reference, but for the red reference, it changes rapidly with echo time. This observation is supported by the slice series, showing multiple wrap surfaces emerging left of the red reference point with increasing echo count. The red reference is not valid for temporal unwrapping, since the Nyquist criterion is not fulfilled. The spatial phase relation and thus the spatial unwrap is valid in both cases.




Figure 4.4: Neighbourhood constellation for the six neighbours with common surface (left), and the 18 neighbours including those with common edges (right).




Figure 4.5: Flood fill unwrap scheme: A) unwrap all lines in one direction, B) unwrap the average offsets between the lines and C) unwrap the average offsets between slices.

4.4 Simple Flood Fill Unwrap

This technique is most intuitive for 2D and 3D unwrapping, but most prone to failure at the same time. All dimensions are addressed by three consecutive unwrapping steps (see Fig. 4.5):

- A Unwrap all voxels on lines along x direction separately ($n_y \cdot n_z$ unwraps) $\Rightarrow$ *unwrapped lines*
- B Determine the offsets between the lines from step A, neighbouring in y direction – comparable to n_z unwraps along $y \Rightarrow$ *unwrapped planes*
- C Determine the offsets between the planes from step B, neighbouring in z direction – comparable to one unwrap along $z \Rightarrow$ *unwrapped volume*

Flood filling operates on 2D or 3D data, although the main unwrap operation occurs only along one dimension, creating unwrapped lines. The other two steps are in principle offset-matching procedures. Thus, the method implies a high susceptibility to errors, as every error that is e.g. made in the line unwraps of step A will be propagated into the followup steps and finally falsify the entire volume. When further a mask is applied, this method is prone to fail, as it cannot account for data gaps that are regularly encountered when using arbitrarily shaped masks. An application example of the method is shown further below (see Fig. 4.8).

Summary:

- ▲ Straight-forward algorithm with low demands on computing operations
- ▼ Susceptible to any error accumulated during the unwrap process
- ▼ Not applicable for masked volumes

4.5 Multi-Path Phase Gradient Integration

Within the scope of this thesis an algorithm was developed, that makes use of the high computational speed of the flood fill method, while trying to at least partially circumvent its pitfalls. The technique is designed for 3D (spatial) unwrapping and is named Multi-Path Phase Gradient Integration (MPPGI).

4.5.1 Concept

The basic idea of MPPGI is to employ the phase gradient rather than the phase itself (as reported e.g. for a method by Strand and Tøft [2002]). The current algorithm attempts to remove the wraps within the gradient maps, where wraps generate only local peaks. The true phase can be reconstructed based on the corrected gradient maps via integration along multiple paths. The integration step implies that, unlike the flood fill technique, each voxel is individually addressed. To ensure maximum accuracy of the unwrap, the integrated gradients are only used to estimate a wrap count map. Fig. 4.6 outlines the workflow which is subsequently explained in detail.




Figure 4.6: Process Scheme of the MPPGI algorithm.

4.5.2 Descending to Phase Gradients and Improving Quality

The three Cartesian phase gradient maps, $g'_x = \partial_x \varphi$, $g'_y = \partial_y \varphi$ and $g'_z = \partial_z \varphi$, are calculated as in Eq. 3.21. Wrap surfaces are usually identified by voxel difference calculation – the phase gradient vector, $\vec{g}' = (g'_x, g'_y, g'_z)$, embodies this operation. All phase data that is compliant to Nyquist's sampling theorem and is not subject to severe noise should exhibit phase gradients with an absolute value of $|g'_{x(y,z)}| < \pi$. Wraps find expression in gradient values outside of these limits. In affected voxels, the correct gradient can be determined by adding or subtracting 2π . This leads to the following assignment:

$$\begin{aligned}
 -\pi \leq g'_{x(y,z)}[ijk] < \pi &\longrightarrow g_{x(y,z)}[ijk] = g'_{x(y,z)}[ijk] \\
 g'_{x(y,z)}[ijk] \geq \pi &\longrightarrow g_{x(y,z)}[ijk] = g'_{x(y,z)}[ijk] - 2\pi \\
 g'_{x(y,z)}[ijk] < -\pi &\longrightarrow g_{x(y,z)}[ijk] = g'_{x(y,z)}[ijk] + 2\pi
 \end{aligned} \quad (4.6)$$

where $[ijk]$ stands for the relative position of the voxel, and $x(y, z)$ indicates the direction of the respective gradient. The corrected phase gradient maps will be called Shift-Corrected Phase Gradients (SCPG), $\vec{g} = (g_x, g_y, g_z)$. Bakker et al. [2008] used such gradients for imaging.

4.5.3 Phase Gradient Smoothing

Eventual noise within the phase gradient maps can be mitigated by an optional step – a gentle smoothing filter. The filter will not change the nature of the phase shape, nor will it change its value range. Local noise accumulation could otherwise disturb the integration in the follow-up step. More precisely, a Gaussian filter with a standard deviation, σ_{smooth} , of about a voxel width is applied to the SCPG:

$$g_{\text{smooth},x(y,z)} = g_{x(y,z)}[ijk] * \left(\frac{e^{-\frac{i^2+j^2+k^2}{2\sigma_{\text{smooth}}^2}}}{\sqrt{2\pi}^3 \sigma_{\text{smooth}}} \right) . \quad (4.7)$$

4.5.4 Integrating Phase Gradients

A point of origin, $\vec{r}_0 = (x_0, y_0, z_0) = [i_0, j_0, k_0]$, is defined. In order to determine the unwrapped phase value in position $[ijk]$ based on the phase gradient maps, one has to integrate the phase gradients along a path from the origin towards that position:

$$\varphi_{\text{int}}[ijk] = \int_{S_{[ijk]}} \vec{g} \cdot \vec{r}(S) \, ds . \quad (4.8)$$

The choice of the path within this integral is ideally arbitrary (neglecting noise or signal voids). In MPPGI, the integration is performed along the Cartesian axes $\hat{x}$, $\hat{y}$ and $\hat{z}$. The reason for this is that the integral along each segment can be calculated by simply summing up the phase gradient values from the origin up to the target coordinate. If one chooses to integrate over x , then over y and finally over z , the path integration takes the following form (in continuous and discretised notation):


$$\begin{aligned} \varphi_{\text{int}}(x, y, z) &= \int_{x_0}^x \vec{g}(x', y_0, z_0) \hat{x} \, dx' + \int_{y_0}^y \vec{g}(x, y', z_0) \hat{y} \, dy' + \int_{z_0}^z \vec{g}(x, y, z') \hat{z} \, dz' \\ \varphi_{\text{int}}[ijk] &= \sum_{m=i_0}^i g_x[m, j_0, k_0] + \sum_{m=j_0}^j g_y[i, m, k_0] + \sum_{m=k_0}^k g_z[i, j, m] . \end{aligned} \quad (4.9)$$

This operation is remotely related to a flood fill unwrap. In contrast to global offset points for line and slice matching, here every voxel is addressed over one particular unwrap path. Furthermore, one is free of choice concerning the integration path. Clinging to the Cartesian axis allows for choosing between six different (partly overlapping) permutations: $(\int_x \int_y \int_z)$, $(\int_x \int_z \int_y)$, $(\int_y \int_x \int_z)$, $(\int_y \int_z \int_x)$, $(\int_z \int_x \int_y)$ and $(\int_z \int_y \int_x)$.

In Section 4.4 it was already indicated that following a single path can easily propagate errors – e.g. due to phase noise, signal voids, etc. To avoid this, integration over three independent permutations is performed (see Fig. 4.7). The median of the path integrations is calculated to ensure better stability of the result. It serves as the phase integration result – loosely speaking:

$$\varphi_{\text{int}}[ijk] = \text{median} \left(\int_x \int_y \int_z g; \int_y \int_z \int_x g; \int_z \int_x \int_y g \right) . \quad (4.10)$$

Figure 4.7: Three independent integration paths, following the Cartesian axis, x , y and z . The integrated phase is determined by the median of these three integrations.



The described technique can be implemented in a rather performant way. Most data processing languages such as the utilised tool, MATLAB® (The MathWorks Inc., Natick, USA), offer a cumulative sum (*cumsum*) operation on multidimensional arrays. With this operation the cumulative sum of the elements in all data vectors of an array can be rapidly calculated in the required direction. The integration of all voxels along one of the integration paths thus only requires three cumsum operations. This means that 9 operations are needed to integrate the array over three paths and 18 operations suffice for six paths.

4.5.5 Improving the Integration

The integration result, φ_{int} , can optionally be slightly smoothed in the same way as the SCPG. This potentially suppresses error propagation from the integration. The next step will point out, that gentle smoothing of the integration will not change the unwrap result dramatically. The same filter type as the one described in Eq. 4.7 is applied here.

4.5.6 Wrap Count Estimation

The integrated phase, φ_{int} , is not the final unwrap result, yet. Due to uncertainties in the unwrap path, including potential noise, and the optional smoothing operations, the integrated phase follows the true phase curve only roughly. Hence, φ_{int} is just used to estimate the wrap count in each voxel. The integration induces a global phase offset, φ_c . This offset is corrected for by evaluating:

$$\varphi_c = \text{mean}(\varphi_{\text{int}}[n] - \varphi_{\text{meas}}[n]), [n] \in \text{local neighbourhood of a reference point, } \vec{r}_0 \quad (4.11)$$

The wrap count is then estimated by evaluating:

$$n[ijk] = \text{round} \left[\frac{\varphi_{\text{int}}[ijk] - \varphi_{\text{meas}}[ijk] - \varphi_c}{2\pi} \right] \quad (4.12)$$

in the entire volume. The unwrapping process is finalised by adding the estimated 2π -multiples to the measured phase:

$$\varphi_{\text{unw}}[ijk] = \varphi_{\text{meas}}[ijk] + 2\pi \cdot n[ijk] \quad (4.13)$$

4.5.7 Implementation

The MPPGI was implemented in MATLAB[®] (The MathWorks Inc., Natick, USA). The phase gradient calculations are performed by quick circular array shift and adding operations. As indicated above, integrations are performed by assimilated cumsum operations. All other operations have a straight-forward implementation, following the workflow described above. Gaussian smoothing is computed with the filter introduced by Hirsch [2009] (see also Section 8.2).

4.5.8 Comparing Multi-Path and Flood Fill

In order to outline the advantages of the MPPGI with regard to simple flood filling, the algorithms are tested on a synthetic dataset using different noise levels. The accuracy of MPPGI is further assessed for different parametrisations.

4.5.8.1 Material and Methods

The simulated phase map of size $128 \times 128 \times 128$ contains several polynomial distortions up to 3rd order. Further contrast is introduced by adding spherical bodies of different sizes to the phase. Local contrast is generated by a small grid of spheres with small diameter aligned in a dense neighbourhood. The phase map, φ_{sim} , generated this way is corrupted by Gaussian noise with $\sigma_{\text{noise}} = [0^\circ, 11^\circ, 23^\circ, 34^\circ, 46^\circ, 57^\circ]$: $\varphi_{\text{true}} = \varphi_{\text{sim}} + \text{noise}_{\text{Gauss}}(\sigma_{\text{noise}})$. The observed phase, φ_{raw} , is the result of the wrap function (Eq. 3.16): $\varphi_{\text{raw}} = \mathcal{W}(\varphi_{\text{true}})$.

The flood fill unwrapper and the MPPGI are applied on φ_{raw} . MPPGI is applied using different smoothing factors, $\sigma_{\text{smooth}} = [0, 1, 2, 3]$ voxels, used for gradient smoothing and integration smoothing. In all unwrap results the offsets towards φ_{true} are corrected for by determining the phase difference in a central reference point. To evaluate the unwrap quality, the $\mathbb{L}_2$ -norm of the point-wise difference between the corrected unwrap results, φ_{unw} , and the true phase, φ_{true} , is calculated and rounded to the closest integer value:

$$\mathbb{L}_2(\delta\varphi) = \text{round} \left[\sqrt{\sum_{ijk} (\varphi_{\text{unw}}[ijk] - \varphi_{\text{true}}[ijk])^2} \right] . \quad (4.14)$$

4.5.8.2 Results and Discussion

Fig. 4.8 gives an overview on the results of the simulation and the applied unwrapping algorithms. Obviously, even for high noise levels, the un-smoothed MPPGI performs better than the simple flood fill algorithm. Flood filling is very sensitive to phase errors – this is well observable, e.g. for noise level $\sigma_{\text{noise}} = 34^\circ$ even though only one sample slice of the volume is shown. The optimal smoothing level for the MPPGI varies with the noise level. For low noise ($\leq 23^\circ$) smoothing should be disabled, while for high noise level higher smoothing is required, to obtain a phase map that roughly matches the original one. Interestingly, for $\sigma_{\text{noise}} = 34^\circ$ even smoothing with $\sigma_{\text{smooth}} = 1$ is optimal, while higher smoothing decreases the unwrap quality. Higher smoothing




Figure 4.8: Result table of the flood fill (second row) and the MPPGI (third to sixth row) unwrappers applied on raw phase (top row). The MPPGI was applied using several smoothing factors, σ_{smooth} . From left to right, the noise level increases from $\sigma_{\text{noise}} = 0^\circ$ up to $\sigma_{\text{noise}} = 57^\circ$ phase uncertainty. The yellow numbers represent the $\mathbb{L}_2$ -norm error of the unwrap result compared to the true phase, φ_{true} (bottom row). Yellow borders mark the best result at a given noise level (column), choosing always the most simple method to achieve it.

levels destroy local phase contrast robustness and thus decrease the unwrap accuracy. Applying high smoothing levels thus only makes sense, if the noise level is already too high to distinguish between noise and local contrast. Whilst the flood fill unwrap is performed in less than a second for the simulated $128 \times 128 \times 128$ array, the MPPGI requires about 5 seconds for the same dataset. This is still a very performant result.

The advantages and disadvantages of MPPGI shall be pointed out in short:

- ▲ Relatively fast computing, since the number of operations is small and all operations can be performed on the entire array
- ▲ Small areas of inconsistent phase can be compensated for by the smoothing operations
- ▲ Smoothing of low σ_{smooth} does not necessarily falsify phase values. The integrated phase is only used to determine the wrap count, n , to be applied to the measured phase
- ▼ Requires several arrays in the size of the source array to compute in the phase gradient domain $\rightarrow$ memory demands
- ▼ Error correction by smoothing is limited and can only compensate for very small inconsistencies. When compensating for higher noise levels, smoothing misses very small phase structures $\rightarrow$ loss in local contrast
- ▼ Due to the algorithm layout (integration paths), only convex, ideally rectangular masks are unwrapped correctly

4.6 PRELUDE

One of the most popular unwrap algorithms, freely available on the web, is *PRELUDE*, introduced by Jenkinson [2003]. It is part of the *FSL*¹ toolbox for analysis of MRI data. *PRELUDE* uses a “[...] region merging approach to optimize a cost function that penalizes phase differences across boundaries.” Jenkinson [2003]. The phase range of $[-\pi; \pi]$ is divided in n (8 by default) partitions. From this partitioning a set of disjoint regions is defined. The regions are expected to exhibit integer offsets – global wrap counts – with regard to each other. The core algorithm of *PRELUDE* tries to determine these offsets by minimising the phase discontinuities at the region surfaces through variations of the integer offset function. This remarkable approach works very well on 2D and 3D data, and the author, *Mark Jenkinson*, claims that it is expandable to n -dimensional problem sets. Nevertheless, the 3D mode of *PRELUDE* is less performant than the 2D mode, especially when applied to volumes of high matrix size (such as 256^3 and above) and high phase complexity. Especially noisy phase maps lead to very high computing times, as the number of disjoint regions increases dramatically. In contrast to the MPPGI method, the concept of *PRELUDE* is well suited for arbitrarily shaped brain masks. In order to illustrate the virtues of *PRELUDE*, another simulation was set up, testing the 3D mode of the programme.

¹<http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/>

4.6.1 Material and Methods

A simulation is performed as in the MPPGI case (see Section 4.5.8.1). Additionally, a validity mask is chosen. The mask consists of two separate spherical bodies connected by a small sphere in their centre. Again, noise is simulated with $\sigma_{\text{noise}} = [0^\circ, 11^\circ, 23^\circ, 34^\circ, 46^\circ, 57^\circ]$. Besides the $\mathbb{L}_2$ -norm, the required computing time is recorded.

4.6.2 Results and Discussion

Fig. 4.9 shows the results of the simulations computed with PRELUDE. It is apparent, that PRELUDE is very precise in unwrapping phase with high noise level – errors in the unwrap are not encountered below $\sigma_{\text{noise}} = 46^\circ$. The absolute error level cannot be compared directly to the MPPGI, as PRELUDE computes a smaller, but more demanding mask. On the other hand, although the volume (of originally $128 \times 128 \times 128$ voxels) is strongly confined by the mask, the computing time necessary for solving the unwrap problem grows dramatically. PRELUDE requires almost 20 minutes for the highest noise level. This increase of time demands is due to the phase partitions created. The more noise splits up previously contiguous regions, the more regions are created and have to be matched in the end. Finally, PRELUDE works well and efficiently for phase data with low and high noise level. For higher noise level, the robustness of the result is at the expense of enormous computing times.

Remarks, Advantages and Disadvantages of PRELUDE:

- ▲ Very robust approach, usually leading to consistent results, when applied in 3D mode
- ▲ Works on any kind of mask
- ▲ Offers 2D and 3D mode
- ▲ For bulky masks and calm phase, 2D mode with slice matching offers a very quick unwrap alternative
- ▼ For noisy phase, the number of partitioned groups increases dramatically → high computing time.
- ▼ For bigger arrays, the number of matching conditions increases dramatically → high computing time.




Figure 4.9: Results of the unwrap of the simulated data, performed with PRELUDE (mid row). For comparison the true phase (lower row) and the raw, wrapped phase (top row) are shown. From left to right the noise levels, σ_{noise} , are increasing. The yellow numbers indicate the rounded $\mathbb{L}_2$ -norm error, while the red numbers indicate the required time for computing each result.

4.7 Quality-Based Unwrapping

A completely different approach to unwrapping is the determination of a certain unwrap order on a voxel-wise base. Every particular voxel is addressed and matched to previously unwrapped voxels in its neighbourhood at a certain point in the unwrap chain. The order is generated based on a phase quality measure (e.g. Cusack and Papadakis [2002]) and can have various forms, e.g. an *unwrap path*, a *region growing* method or a *grouping algorithm*.

Unwrap Path: Voxels are addressed from one or multiple reference points using specific paths on which the unwrapper propagates. The path to a voxel might be different from that for its neighbours and should always have best quality. The most simple, but not quality-based, example for path unwrapping is the MPPGI.

Region Growing: Starting from one or several locations, the surrounding are unwrapped iteratively until all voxels have been processed. This can be pictured by interpreting the quality map as a surface, guiding a rising water level through the evolution of places it should fill.

Grouping algorithm: Voxels are unwrapped against each other or against existing “unwrapped groups” just according to their phase or link quality. This approach is pursued in the algorithm by Abdul-Rahman presented later in this chapter (Section 4.7.2).

For all of these methods the quality measure decides over the propagation, and over success or failure. In general such unwrapping methods are summarised with the term Quality-Based Unwrapping (QBU). In order to be complete, a variety of quality measures will be introduced in

the following, although most will not find application in this work. A crucial advantage of QBU is that the algorithms are not bound to strict path integration pattern or volume regions. Hence, unwrapping of ambiguous scenarios including phase poles becomes feasible.

4.7.1 Quality Measures

Phase quality can be measured in various ways. Several phase quality measures were presented by Witoszynskij et al. [2009] – some of these measures are presented below, applied to the sample dataset illustrated in Fig. 4.10, acquired at a 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany).




Figure 4.10: Sample in vivo data, measured at 3 T. Upper row: magnitude, lower row: phase, every 10th slice is displayed.

Phase Gradient Moments: The gradient moments of phase can be considered as quality measures. The stronger phase changes in the surroundings of a voxel, the more likely are wraps and errors in local unwrapping. The SCPG defined in Eq. 4.6 can be used as a directional weighting function, or as a scalar phase gradient indicator:

$$W_{Gx(yz)} = g_x(yz) \quad (4.15)$$

$$W_G = \sqrt{g_x^2 + g_y^2 + g_z^2} \quad (4.16)$$

The phase gradients, $\partial_{x(y,z)}\varphi$ used to determine $g_{x(y,z)}$ should be calculated symmetrically around $[ijk]$ as in the following example of the x gradient:

$$\partial_{x,symm}[i,j,k] = \varphi[i+1,j,k] - \varphi[i-1,j,k] \quad (4.17)$$

Phase Gradient Weighting (PGW) is a means of describing local directional changes in phase. It is an inverse quality measure, as it is minimal for low phase variations. Fig. 4.11 shows PGW applied on the sample data.




Figure 4.11: PGW applied on the sample data

Local Coherence: Considering how much the phase varies throughout local neighbourhood of n^3 (n : odd) voxels, one can derive the following measure, the Local Coherence (LC), using the complex polar notation:

$$W_{LC}[ijk] = \frac{1}{n^3} \left| \sum_{[pqr]} e^{i\varphi_{[pqr]}} \right|, \quad \text{with } [pqr] \in [ijk] \pm (n-1)/2. \quad (4.18)$$

For a neighbourhood of voxels exhibiting the same phase, this measure reaches its maximum ($= 1$). The LC was introduced by Witoszynskyj et al. [2009]. Fig. 4.12 shows the LC measure applied on the sample data.



Figure 4.12: *LC applied on the sample data*

Weighted Local Coherence: Whilst LC summarises the normalised complex vectors of a neighbourhood, Weighted Local Coherence (WLC) calculates the *complex average*. The maximum now includes maximum phase coherence and maximum signal (magnitude). The expression for this measure is:

$$W_{WLC}[ijk] = \frac{1}{n^3} \left| \sum_{[pqr]} A_{[pqr]} e^{i\varphi_{[pqr]}} \right|, \quad \text{with } [pqr] \in [ijk] \pm (n-1)/2. \quad (4.19)$$

The WLC can optionally be normalised and was also introduced by Witoszynskyj et al. [2009]. Fig. 4.13 shows WLC applied on the sample data.




Figure 4.13: *WLC applied on the sample data*

Local Variance Another measure is the Local Variance (LV). It determines the variance of the neighbourhood with regard to the centre-voxel (see also Section 3.5.3):

$$W_{LV}[ijk] = \frac{1}{n^3} \left(\sum_{[pqr]} S_{[pqr]} S_{[pqr]}^\dagger \right) - \frac{1}{n^3} \left(\sum_{[pqr]} S_{[pqr]} \right) \cdot \left(\sum_{[pqr]} S_{[pqr]} \right)^\dagger, \quad \text{with } [pqr] \in [ijk] \pm (n-1)/2. \quad (4.20)$$

Here, S is the complex signal, and $\dagger$ symbolises complex conjugation. The LV is helpful for edge detection [Witoszynskyj et al., 2009]. Fig. 4.14 shows LV applied on the masked sample data.




Figure 4.14: LV applied on the masked sample data

Second Differences A quite robust quality measure is the Second Difference (SD) calculation. While the gradient itself, e.g. $\partial_x \varphi = \varphi[i,j,k] - \varphi[i-1,j,k]$, is a 1st order difference measure, SD calculates the change of the phase gradients. For the Cartesian directions the straight second differences are:

$$\begin{aligned} H[ijk] &= \mathcal{W}(\partial_x \varphi[i+1,j,k]) - \mathcal{W}(\partial_x \varphi[i,j,k]) \\ &= \mathcal{W}(\varphi[i+1,j,k] - \varphi[i,j,k]) - \mathcal{W}(\varphi[i,j,k] - \varphi[i-1,j,k]) \end{aligned} \quad (4.21)$$

$$\begin{aligned} V[ijk] &= \mathcal{W}(\partial_y \varphi[i,j+1,k]) - \mathcal{W}(\partial_y \varphi[i,j,k]) \\ &= \mathcal{W}(\varphi[i,j+1,k] - \varphi[i,j,k]) - \mathcal{W}(\varphi[i,j,k] - \varphi[i,j-1,k]) \end{aligned} \quad (4.22)$$

$$\begin{aligned} N[ijk] &= \mathcal{W}(\partial_z \varphi[i,j,k+1]) - \mathcal{W}(\partial_z \varphi[i,j,k]) \\ &= \mathcal{W}(\varphi[i,j,k+1] - \varphi[i,j,k]) - \mathcal{W}(\varphi[i,j,k] - \varphi[i,j,k-1]) \end{aligned} \quad (4.23)$$

(as in Abdul-Rahman et al. [2007]). The quality measure is then defined as:

$$W_{SD} = \sqrt{H^2 + V^2 + N^2} \quad (4.24)$$

A drastic change between the phase gradients “left” and “right” of a voxel indicates a barrier or unsteadiness of the phase function, potentially leading to unwrap failure. Thus, voxels with high SD are rated as low phase quality voxels, while low SD indicates a low phase gradient change and thus a good quality. Fig. 4.15 shows SD applied on the sample data.




Figure 4.15: SD applied on the sample data

Abdul-Rahman et al. also describe the second difference mixed terms, D_i , representing the diagonal phase gradients. Applied to the SD this means:

$$W_{SD} = \sqrt{H^2 + V^2 + N^2 + D_1^2 + D_2^2 + \dots} \quad (4.25)$$

However, practical experience has shown, that for most MRI-based unwrapping applications, these terms can be neglected.

4.7.2 The Abdul-Rahman Algorithm

Abdul-Rahman et al. [2007] published an article on their approach on quality based unwrapping. Unlike other quality-based algorithm this method does not employ a phase quality map directly, but derives a quality estimate for the links between adjacent voxels based on their phase quality.

The algorithm consists of two major steps, a preparation step estimating phase and link qualities, and a second iterative part that solves the unwrapping problem link by link. The principle of this approach is sketched in Fig. 4.16. The first part starts with the estimation of SD quality map (see Section 4.7.1). The map will be denoted by W in the following. A set of three link quality maps is derived from W by calculating the mean quality of adjacent voxels in all directions:

$$W_{ex}[ijk] = W[i,j,k] + W[i+1,j,k], \quad i = 1, \dots, n_x - 1 \quad (4.26)$$

$$W_{ey}[ijk] = W[i,j,k] + W[i,j+1,k], \quad j = 1, \dots, n_y - 1 \quad (4.27)$$

$$W_{ez}[ijk] = W[i,j,k] + W[i,j,k+1], \quad k = 1, \dots, n_z - 1 \quad (4.28)$$




Figure 4.16: Process scheme of the Abdhul-Rahman Algorithm. The top row outlines the steps for preparation and unwrapping, while the lower circle graph illustrates the iterative unwrapping process for the links, called in the last step of the chain, and representing the core part of the algorithm.

In the presented implementation an indexing array, $[W_e[ijk], i, j, k, o]$, is generated and sorted in the order of descending link quality, W_e . The array stores the link position $[ijk]$ and its orientation (o). Note that W_e contains link qualities for x-, y- and z-direction. In the second part, the

unwrap process addresses all links regardless of their orientation one at a time with descending link quality. The current top quality link indicates the two neighbouring voxels, a and b , with phases, φ_a and φ_b . These voxels are unwrapped against each other. This can lead to three basic situations: A) both voxels were not addressed before – they are singles, B) only one voxel is free, while the other already belongs to a group, g_a , C) both voxels belong to different groups, g_a and g_b . If both voxels already belong to the same group, g , the link is omitted.

For case A), the voxels are simply unwrapped against each other,

$$[\tilde{\varphi}_a, \tilde{\varphi}_b] = \mathcal{U}_1[\varphi_a, \varphi_b] = [\varphi_a, \varphi_b + 2\pi n_b] \quad (4.29)$$

and joined to a new group, g . In case B) the free voxel (e.g. voxel b) joins the group, g_a , after being unwrapped towards the group voxel, a . In case C) The smaller group, e.g. $\#g_b < \#g_a$, joins the bigger one, while “unwrapping” the group by finding the wrap distance, n_b , of the reference voxel towards the voxel of the larger group.² The determined wrap count is assigned to all voxels of group b .

The process is finished as soon as all links in the volume were addressed. The number of link solving operations is about three times the number of voxels in the volume. Furthermore, each voxel is addressed several times (six times for inner voxels³), as a voxel possesses six surface links.

4.7.3 Implementation

The algorithm, originally developed by *Abdul-Rahman* et al., was implemented in MATLAB® (The MathWorks Inc., Natick, USA). The phase gradients and the quality measure of the SD, are derived as described above, generating the indexing array. Grouping is controlled by a separate indexing array, assigning a unique Group ID (gid) for each voxel. The gid can either denote a certain group ($g \geq 1$) or no group ($g = -1$), indicating a free voxel. Unwrapping of free voxels is performed by simply adding the estimated wrap distance, n , to the particular voxel. Groups are shifted by adding n to all voxels carrying the particular group identifier. The implementation further offers to include a mask. All connections residing outside of the regions indicated valid by the mask are excluded from the list.

4.7.4 Example

The Abdul-Rahman algorithm is tested on a small structure exhibiting phase poles generated by a particular geometric shape of the underlying phase distribution. A small-sized geometry is chosen in order to challenge the algorithm with effects arising with low spatial sampling and high local phase contrast. The algorithm is compared to the 3D mode of PRELUDE.

²# symbolises the member-count operator, returning the number of voxels contained in a group

³Inner voxels are voxels not situated at the surface of the volume or the confining mask.

4.7.4.1 Material and Methods

A small synthetic phase map of size $32 \times 32 \times 32$ voxel is generated, inducing a central region smoothly connected to its surroundings only at one side. The rest of the surface of the central sphere shows a strong offset towards the exterior phase. Due to the smaller dimensioning of the sample, the noise levels are reduced compared to the examples shown before: $\sigma_{\text{noise}} = [0^\circ, 5.7^\circ, 11.5^\circ, 17.2^\circ, 23.0^\circ, 28.6^\circ]$. As before, computing time and the 2π -rounded $\mathbb{L}_2$ -norm are calculated.

4.7.4.2 Results and Discussion

Fig. 4.17 shows the results of the experiment. Except for the highest noise level, the numerical accuracy ($\mathbb{L}_2$ -norm) of the quality-based unwrap is better than the accuracy of PRELUDE [Jenkinson, 2003]. Visual inspection further shows that the results of the quality-based unwrap are much smoother than the PRELUDE results. Even for the highest noise level the geometry of the central ball is still preserved, although the surroundings are split into various regions. The two separate areas (the inner ball and the exterior) possess only one physical connection in the top left of the displayed slice. Yet, the modulated (measured) phase exhibits more points of potential connections, falling into the partitioning scheme of PRELUDE and causing it to fail. The quality estimation of the Abdul-Rahman algorithm, nevertheless, recognises the surfaces with incorrect connections to be of weaker quality than the true connections. The locally changing phase gradients across the sphere surface expose a lower quality than the gradients along the surface and thus prevent an incorrect unwrap of internal against the external voxels. Both regions are grouped and connected in the correct reference point. This strategy succeeds at least up to a certain noise level.

Unfortunately, the presented implementation of the quality-based algorithm is orders of magnitude slower than PRELUDE. Yet, the computation time depends mainly on the problem size and less on the noise-level as for PRELUDE. The algorithm will later on be denoted by “Abdul-Rahman algorithm” or simply by “QBU”.

According to the findings of Abdul-Rahman et al. [2007] and the experiences gathered within the scope of this work, the virtues of the algorithm can be condensed to:

- ▲ Extremely robust unwrap
- ▲ 2D and 3D mode possible (2D by ignoring links in one direction)
- ▲ Compatible with any kind of mask
- ▲ Phase inconsistencies, such as poles, can in many cases be resolved
- ▼ Extremely high computing time demands due to grouping, seek and array operations




Figure 4.17: Unwrapping results for the simulated data, performed with PRELUDE (second row) and with quality-based unwrapping (third row). For comparison the true phase (last row) and the raw, wrapped phase (top row) are shown. From left to right the noise level, σ_{noise} , is increasing. The yellow numbers indicate the rounded $\mathbb{L}_2$ -norm error, while the red numbers indicate the required time for computing each result.

4.8 Sliced Unwrapping

4.8.1 Principle

It was shown that 3D unwrapping algorithms may require high computing times. One means of speeding up volumetric unwrapping is *sliced unwrapping*, processing the volume in a slice-wise manner. Each slice of the volume is treated as an individual two-dimensional unwrapping problem. In comparison to an arbitrary unwrap technique for three dimensions, sliced unwrapping saves one degree of freedom regarding neighbour considerations. Hence, it is significantly more performant than any full 3D unwrap. A volume described by the positioning tuple, $[ijk]$, is split into a stack of slices, $[ij]$, with the slice positioning index k . A slice can be computed in less than $1/n_k$ (where n_k is the number of slices) of the computing time for the full volume, since besides the reduced voxel number, the surface conditions to the neighbours in z -direction are omitted. This allows for an altogether faster computation in a sequential approach. Nevertheless, it is crucial to match the mutual phase offsets between the unwrapped slices – this is achieved e.g. by the unwrap of a reference vector orthogonal to the unwrapped slices as shown for the flood fill unwrap in Fig. 4.5c. Relying on only one reference vector is a technique prone to errors, induced by potential noise influences. A preferable technique to match the slice offsets while reducing noise-induced errors is, to determine the phase offset between two slices in as many reference points as possible with

regard to the applied mask, m :

$$\Delta\varphi_{sl.unw}[ijk] = m[i,j,k]m[i,j,k-1](\varphi_{sl.unw}[i,j,k] - \varphi_{sl.unw}[i,j,k-1]) \quad , \quad (4.30)$$

where k is the index along the slicing direction. To enhance the stability of the phase difference, the median filter should be preferred over the simple mean filter. Thus, one can determine the wrap count between two slices, k and $k - 1$:

$$n_{k,k-1} = \text{round}\left(\frac{\text{median}(\Delta\varphi_{sl.unw}[ijk])}{2\pi}\right) \quad (4.31)$$

Several of the algorithms presented within this chapter are applicable in 2D and can thus be sped up by sliced unwrapping. In the end, the applicability of sliced unwrapping has to be determined regarding the observed dataset.

This principle has certain pitfalls. It completely ignores the phase evolution along the third dimension, which might lead to unresolved *phase poles*, incorrect unwraps of disjoint areas that might pose a straight forward unwrap in three dimensions. Furthermore, Nyquist errors solvable in three dimensions can become sources of artefacts in sliced unwrapping. Noise is harder to compensate for, as well. Common pitfalls are illustrated and further discussed in the following.

4.8.2 Pitfalls of Sliced Unwrapping

One is easily tempted to prefer a slice-based unwrapping algorithm due to the reduction in problem size and the drastic reduction of required computing time thereby induced. To emphasize the problems potentially arising when slice-based algorithms are used for unwrapping, three examples will be shown. The examples are visualised on a two-dimensional phase sample, using one-dimensional (line) unwrap, exploiting the nature of the pitfalls, likewise. Despite this form of illustration, the indicated pitfalls do not lose generality.

4.8.3 Field Gradients and Structural contrast

In phase images, and in particular in phase images of the human brain, one can find delineated regions in which the phase is shifted with regard to their surroundings. If this local contrast is now superimposed with linear field gradients of different strength inside and outside of the region, the result can be cumbersome to unwrap. Fig. 4.18 shows an example for such a situation. Regardless of the slicing direction of the unwrap, a false result emerges.

4.8.4 Disjoint Mask Segments

Typically, one cannot expect a brain mask to cover the entire measurement volume, nor to be convex, but rather to be relatively random in shape. Thus, treating slices in separate unwraps can easily lead to two or more disjoint segments inside one slice. Fig. 4.19 illustrates an example for a phase and a mask that are incorrectly unwrapped due to this shortcoming. It is impossible, to match neighbouring slices/lines correctly, without considering the mask shape and valid regions in three dimensions.




Figure 4.18: Two-dimensional example of shortcomings in slice-based unwrap; The local phase offset is superimposed with internal and external phase gradients. The wrapped phase is simulated and (line-based) unwrap is approached along x and y direction. Both fail and result in unwraps with incorrect regions. Units are rad.




Figure 4.19: Two-dimensional example of issues with disjoint mask segments in slice-based unwrap. The source phase data is artificially wrapped and then unwrapped along x and y direction. The result is displayed within the validity mask. The slice unwraps both generate stripe artefacts with offsets in the outer regions of the mask, that are no longer unified in sliced view. Units are rad.

4.8.5 Undersampling and Noise

Sometimes, phase data contain strong local gradients or noise. Since sliced unwrapping offers only limited options for choosing the optimal unwrap path, the chance to be forced to pick an adverse path is high. This makes sliced unwrapping much more vulnerable to noise and further leads to slice offset matching errors when realigning all slices in the volume. Volume unwrapping would on the other hand isolate the problematic region to an area with much smaller extend and find a consistent solution for the properly sampled regions of the volume. Fig. 4.20 shows how noise interferes with sliced unwrapping.

4.9 Automated Phase Masking

When working with anatomic images, especially images of the entire brain, a mask has to be generated defining the Volume Of Interest (VOI). Unwrapping relies strongly on the choice of a mask that contains only consistent phase data. For modalities based on magnitude data, there are several tools for the generation of masks. Measurement of phase data always includes acquiring magnitude data. In principle a mask can hence be generated utilising the magnitude. However, masks created from magnitude data suffer from inhomogeneities of the average signal intensity and tell little about the robustness of the phase information. Thus, estimation of a mask based on phase data is desirable. A strategy for mask estimation will be presented in the following.




Figure 4.20: Two-dimensional example of local noise influences on slice-based unwrap. The unwrap applied in x-direction does not search for the best access to the noisy area, but consecutively unwraps the voxels, leading to incorrectly unwrapped lines. This leads to a corrupt matching of line phases and falsifies the entire result (third image). The full volume unwrap (last image, computed with PRELUDE) experiences less problems with the estimation of phase values in the noisy region and further leaves the whole volume consistency intact. Units are rad.

4.9.1 Method

Since the mask should be related to the phase robustness or quality, it makes sense to employ a quality measure. A beneficial choice is e.g. the LC measure (Section 4.7.1, Witoszynskij et al. [2009]). The measure reflects how strongly the phase of voxels differs in a local neighbourhood. It is maximal, if all voxels have the same phase. For an MRI measurement the phase within regions with low signal is typically randomised, thus the LC will be extremely low. In regions with sufficient signal, the behaviour of the LC still depends on phase behaviour: For phase varying smoothly with regard to the local neighbourhood dimensioning the LC will be close to 1, whilst for neighbourhoods with rapid phase changes, the LC decreases towards 0. If a small neighbourhood, such as $3 \times 3 \times 3$ voxels, is chosen, these characteristics become quite useful. One wants to exclude noise, but also local variations that are likely to create phase sampling problems. Thus, exclusion of voxels with a low LC even inside a measured object, can be advantageous.

Due to statistics it is possible, that in a few neighbourhoods residing in noisy areas the LC is accidentally high. E.g. random coherences of a major fraction of these voxels can have this effect. In order to avoid accidental inclusions of voxels in the signal void, a gentle smoothing filter, such as a *Gaussian* can be applied. Finally the smoothed LC is thresholded in order to obtain a mask for the region of valid phase. A visualisation of the steps involved in this method is shown in Fig. 4.21. Additionally, it is advisable to identify the biggest continuous region in the resulting mask. This avoids the inclusion of isolated and falsely assigned regions.




Figure 4.21: The steps of the automatic phase masking algorithm. From left to right: Source phase, LC, Gaussian-smoothed LC, mask after thresholding.




Figure 4.22: *Resulting separation of the automatic phase masking algorithm. From left to right: Source phase, phase inside valid mask, phase outside valid mask.*

4.9.2 Parameters

For the automatic phase mask estimation, the neighbourhood size is set to the minimum ($3 \times 3 \times 3$ voxels, as indicated above). A Gaussian filter with a standard deviation of $\sigma_{LD} = 2$ voxels is applied and the threshold used for mask estimation is $\epsilon_{LD} = 0.65$. The standard deviation and the threshold are adjustable to fit the problem in an optimal way. Fig. 4.22 shows the separation of valid and invalid phase for the example phase data used in Fig. 4.21. In case of accelerated data acquisition (iPAT), occasionally areas of continuous phase can be found in the signal void. These hinder proper mask estimation and can partly be suppressed by more aggressive thresholding and smoothing.

Unwrap of Subdivided Large Arrays

Nowadays, high-resolution imaging is routinely used in ultra-high field MRI (7T and above), resulting in the acquisition of large data arrays. This dramatically increases the demands on the computing time for unwrapping algorithms. Recent applications [e.g. Wharton and Bowtell, 2010] require efficient algorithms that can provide unwrapped phase data on a time scale of seconds or minutes rather than hours. In order to improve on computing speed, a novel algorithm, named UnwRap of SUBdivided Large Arrays (URSULA) [Lindemeyer et al., 2013], was developed, making use of an artificial volume compartmentalisation and a computing parallelisation to optimise the unwrapping process. The presented algorithms are based on QBU, such as introduced in Section 4.7.2.

5.1 Motivation

The elementary reason for increased computing time demands of a spatial unwrapping problem is the problem size, which is related to the size of the acquired data array. For PRELUDE [Jenkinson, 2003], e.g. larger array dimensioning induces a significantly increased number of surface conditions to be solved, since local noise contributions split up otherwise continuous partitions and thus result in a higher number of phase partitions encountered. For the Abdul-Rahman algorithm, e.g. the number of surface unwraps to be performed is rising with $3 \cdot n^3$, where n is the edge length of a cubic data set. Furthermore, any computing operation performed on the entire array requires computing time proportional to n^3 . For the favourable *Abdul-Rahman* algorithm this means a complexity of $\mathcal{O}(n^6)$.

Correspondingly, a reduction of the problem size will dramatically enhance computing speed. One way to reduce the problem size is subdividing the phase volume into a number of disjoint subsets, being treated as isolated unwrapping problems. Sliced unwrapping (Section 4.8) is a special case of this approach, as it reduces the volumetric problem to a set of two-dimensional problems. Any subdivision will cause a particular loss of generality, because phase evolution is assumed to be consistent within each subset. E.g. in sliced unwrapping, the relations along the third dimension are entirely ignored in the intrinsic unwrap procedure. However, the use of three-dimensional

subsets rather than slices allows the unwrapping problems to preserve spatial consistency to a larger extent. The idea behind the URSULA approach can be condensed to:

The phase volume to be unwrapped is split into a number of disjoint, three-dimensional subsets. All subsets are unwrapped separately. The unwrapped subsets are then reassembled and matched to generate the global unwrap

The critical parts of the algorithm are the choice of the subdivision, the subset-unwrapper and a convenient algorithm to reassemble the unwrapped subsets. A process scheme of the proposed algorithm is outlined in Fig. 5.1.




Figure 5.1: Scheme of the URSULA algorithm. The source phase is split up into subsets. The subsets are unwrapped individually and reassembled in the result array. Offsets between the subsets are corrected for in order to create the unwrapped array. The data used for this illustration are shown in the application Section 5.9 at the end of this chapter.

5.2 Subdividing the Array

The proposed approach does not enforce a certain method to split up the phase volume. Anyway, it is advantageous to have the subsets easily addressable and arranged within the volume. Intuitively, one would further choose the subsets to have resemblant sizes and properties. As the subsets need to be matched after unwrapping, it is beneficial to arrange their surfaces orthogonally to the Cartesian indexing axis. These considerations yield a rectangular grid for volume subdivision. Let $\varphi[ijk]$ be a discretised phase volume of size $n_x \times n_y \times n_z$. The subdivision will be defined by the subset size $s_x \times s_y \times s_z$, leading to the following subset count:

$$c_x = \lceil n_x/s_x \rceil, \quad c_y = \lceil n_y/s_y \rceil, \quad c_z = \lceil n_z/s_z \rceil \quad (5.1)$$

(where $\lceil \cdot \rceil$ denotes rounding up the included term). It is necessary to round up the results, since for subset sizes that are not an integer fraction of the volume dimensioning, the remaining voxels belong to an additional subset. A voxel in the original volume is addressed as $\varphi[ijk]$, while a voxel in the subset $\lceil uvw \rceil$ ($u, v, w = 1, \dots, c_{x,y,z}$), located at position $\lceil pqr \rceil$ is denoted by:

$$\lceil uvw \rceil \varphi[pqr] = \varphi[(u-1)s_x+p, (v-1)s_y+q, (w-1)s_z+r] \quad (5.2)$$

Depending on the source array dimensions, n_x , n_y and n_z , and the subset sizes, s_x , s_y and s_z , the subsets of highest index might have a smaller voxel count than the predefinition due to the division remainder. Beside this deviation, all subsets can be treated in the same way. Usually a phase map is accompanied by a mask, $m[ijk]$, which has to be subdivided as well, in order to provide information about the validity of the subset contents. As above, the subset masks are addressed by:

$$\ulcorner_{uvw}\urcorner m[pqr] = m[(u-1)s_x+p, (v-1)s_y+q, (w-1)s_z+r] \quad (5.3)$$

With this description, all subsets and voxel locations are well-defined. There are six neighbours of each internal¹ subset, addressed by in- or decreasing **one** of the subset indices, $\ulcorner_{uvw}\urcorner$, by 1.

5.3 Unwrapping the Subsets

All individual subsets, $\ulcorner_{uvw}\urcorner$, are independently unwrapped according to their phase, $\ulcorner_{uvw}\urcorner\varphi$, and mask, $\ulcorner_{uvw}\urcorner m$. In principle, there are no restrictions to the unwrapping algorithm used. Any of the algorithms introduced in Chpt. 4 can be applied. As the algorithm of *Abdul-Rahman* generates the most robust results, the original version of URSULA uses this algorithm (described in Section 4.7.2). Subsets with an empty mask, $\forall[pqr] : \ulcorner_{uvw}\urcorner m[pqr] = 0$, are ignored and their result is set to zero:

$$\ulcorner_{uvw}\urcorner \varphi_{unw}[pqr] = 0, \quad \forall[pqr] \quad . \quad (5.4)$$

Unwrapping can be performed in two modes: *sequential* and *parallel* mode. The sequential mode is fit for stand-alone computers and processes the subsets one after another (sequentially). The parallel mode submits the subsets (including phase and mask) to a computing cluster, allowing for parallel computing of the individual unwrapping tasks.




Figure 5.2: The raw phase is first unwrapped, then the surfaces between neighbouring subsets (coloured in red, yellow and blue) are matched for multiples of 2π . This yields the offset-matched unwrap of the entire volume. Offset-matching has to follow a certain order determined by the mask coverage of the surfaces (see Fig. 5.4).

¹Internal subsets are all subsets that do not touch the volume surface. Their positioning indices are thus bigger than 1 and smaller than the highest possible indices, c_x , c_y and c_z .

5.4 Reassembling the Array


After all subsets were processed, the unwrapped phase data are inserted into the solution array $\varphi_{\text{unw}}[ijk]$. Each subset is unwrapped with respect to a different reference phase. This means that the individual subsets are likely to be mismatched by an intersubject specific phase offset. The offset is bound to multiples of 2π and can be understood as the subset wrap count, $\lceil uvw \rceil n$. Hence, the global unwrap, ϕ , has the form:

$$\lceil uvw \rceil \phi[pqr] = \lceil uvw \rceil m[pqr] \cdot \lceil uvw \rceil \varphi_{\text{unw}}[pqr] + \lceil uvw \rceil n \cdot 2\pi \quad (5.5)$$

The objective of the reassembling step is to determine the subset wrap count, $\lceil uvw \rceil n$ (see Fig. 5.2).

Quality-based Wrap Count Determination: A sequential approach on matching the subset offsets seems straightforward and promising. Unfortunately, this only works on rectangular or convex masks. When subsets of a masked volume are arranged in a sequential row including empty subsets, offset-matching is prone to fail as illustrated in Fig. 5.3. This is, since the separated subsets miss relation (a common surface) to the previously unwrapped ones. The phenomenon is comparable to the mask-induced pitfalls in sliced unwrapping as described in Section 4.8.4.

Figure 5.3: Problems in sequential addressing of the subsets: In case A) two consecutive subsets do not share a common surface and thus miss relation (the mask confinement is outlined in gray), whilst in case B) a valid link between the two subsets is determined, choosing a pathway over adjacent surfaces covered by the mask.



Therefore, an approach was developed to match subsets in a well-defined order. Inspired by the method used for phase unwrapping by Abdul-Rahman et al. [2007] a quality measure for subsets was defined. The measure does not encounter phase quality itself, but is weighted by the *mask surface coverage*, $W_{sc,x(y,z)}$ of the subset surfaces. The *mask surface coverage* is the number of valid adjacent voxels shared by the surfaces of two neighbouring subsets. For a surface in x direction (first index), and the subsets, $\lceil u,v,w \rceil$ and $\lceil u+1,v,w \rceil$ the mask surface coverage is defined as:

$$\lceil uvw \rceil W_{sc,x} = \sum_{q,r} (\lceil u,v,w \rceil m[s_x, q, r] \cdot \lceil u+1,v,w \rceil m[1, q, r]) \quad (5.6)$$

For the y and z direction, this means accordingly:

$$\lceil uvw \rceil W_{sc,y} = \sum_{p,r} (\lceil u,v,w \rceil m[p, s_y, r] \cdot \lceil u,v+1,w \rceil m[p, 1, r]) \quad (5.7)$$

$$\lceil uvw \rceil W_{sc,z} = \sum_{p,q} (\lceil u,v,w \rceil m[p, q, s_z] \cdot \lceil u,v,w+1 \rceil m[p, q, 1]) \quad (5.8)$$

Fig. 5.4 illustrates the offset-matching principle. As in the quality-based unwrapping algorithms, the links – meaning the surfaces – between subsets are computed with descending link quality. For

every surface the subset wrap count difference of the adjacent subsets is estimated by calculating the 2π -modulated phase difference at the mask-covered parts of the opposing surfaces – here exemplarily described for a surface in x direction:

$$\begin{aligned} \lceil u,v,w \rceil \Delta n = & \text{round} \left(\left[\sum_{q,r} (\lceil u,v,w \rceil m[s_x, q, r] \cdot \lceil u+1,v,w \rceil m[1, q, r]) \right. \right. \\ & \cdot (\lceil u,v,w \rceil \varphi[s_x, q, r] - \lceil u+1,v,w \rceil \varphi[1, q, r]) \left. \right] \\ & \cdot (2\pi \cdot \lceil uvw \rceil W_{sc,x})^{-1} \end{aligned} \quad (5.9)$$




Figure 5.4: Scheme of matching the subset offsets: A quality value is assigned to all surfaces, describing how much of the surface is covered by the mask. According to this order, the subsets are matched iteratively with descending quality. The edge colour is green for high quality and red for low quality. Green edges will be unwrapped first.

The calculation for surfaces in y and z direction works analogous. As in quality-based phase unwrapping, the algorithm matches and groups two adjacent subsets with regard to their type:

- If none of the subsets is assigned to a group, yet, they are assigned a new group identifier and one subset is shifted by their phase difference $\lceil uvw \rceil \Delta n$.
- If one subset belongs to a group, and the other is unassigned, the second subset joins the group, while being shifted.
- If both subsets belong to a group, the smaller group joins the bigger one, with every assigned subset being shifted simultaneously.
- If both subsets belong to the same group, they remain invariant.

As the number of subsets is naturally smaller than the number of voxels in the original phase volume, the matching process is performed orders of magnitude faster than a volume unwrap. When all surfaces have been matched, $\lceil uvw \rceil n$ is defined. The resulting phase volume

$$\lceil uvw \rceil m[pqr] \cdot \lceil uvw \rceil \varphi[pqr] + \lceil uvw \rceil n \cdot 2\pi = \lceil uvw \rceil \phi[pqr] \longrightarrow \phi[ijk] \quad (5.10)$$

resembles the unwrap of the initial phase $\varphi[ijk]$, potentially shifted by a global offset of $2\pi \cdot c$.

5.5 Implementation

The entire procedure was implemented in MATLAB[®] (The MathWorks Inc., Natick, USA). A target array with the dimensioning of the source phase array is initialised, collecting the results of all subset unwrapping steps in their appropriate position. As indicated above, the algorithm provides a *sequential* and a *parallel* mode. The sequential mode is called URSULA while the parallel mode is called Parallel URSULA (P-URSULA). URSULA applies the MATLAB[®] implementation of the previously described Abdul-Rahman algorithm sequentially on all subsets with non-zero mask and immediately writes the result back to the target array. P-URSULA exports the phase and mask of all subset as indexed files to a location that is accessible by the computing cluster. The unwrapping script is a standalone version of the Abdul-Rahman algorithm. All unwrapping jobs are launched on the cluster, and a standby loop, running on the host machine, sorts the indexed results into the target array and waits for the completion of all jobs before proceeding. Once, the result array is complete, quality-based wrap count determination is launched, matching the offsets inside the result array.

5.6 Validation and Optimisation

In the following, the performance and the accuracy of the unwrapping algorithm will be tested on synthetic phase data, of which the true (unwrapped) phase is known. It is expected that the accuracy depends on the subset size, since the restriction of the spatial dependencies could lead to a higher number of errors in phase estimation. It is further important to estimate which subset size yields minimal computing time. Test series were performed for URSULA and P-URSULA.

5.6.1 URSULA

The performance graph of URSULA, describing the relation of the computing time, t_c , to the subset size, $s_{x(y,z)}$, has two well-defined fix points. On the one hand, reducing the subset size to $s_x = s_y = s_z = 1$ isolates the single voxels and transforms the problem back to the actual Abdul-Rahman algorithm (ignoring the quality measure). Increasing the subset size to $s_{x(y,z)} = n_{x(y,z)}$ ends up in one single subset – the entire volume – likewise solved by the Abdul-Rahman algorithm. Between these two constraints a minimum of computing time is to be found, representing the point where the algorithm performs optimal. The position of the minimum may vary depending on problem size, the employed hardware and the implementation. A balance between the entire computing time required to unwrap the individual subsets and the time required for the offset-matching process defines the minimum.

A cubic array with a dimensioning of $64 \times 64 \times 64$ voxels was simulated with field characteristics as illustrated in Fig. 5.5. Gaussian noise was added with a standard deviation of $\sim 28^\circ$. Accordingly, the following subset sizes were chosen as test parameters for URSULA: $s_{x(y,z)} = [(1), 2, 4, \dots, 30, 32]$. The case, $s_{x(y,z)} = 1$ was benchmarked using only the Abdul-Rahman algorithm.

Computing took place on a quad-core CPU (2.20GHz) with 8 GB of physical memory. Fig. 5.6 shows the computing time results and the result error. The overall accuracy was evaluated as the




Figure 5.5: Synthetic phase data used for performance optimisation test of URSULA; Left: simulated true phase, middle: wrapped phase, right: unwrapped phase; Volume size $64 \times 64 \times 64$ voxel.




Figure 5.6: Left: Performance of the URSULA algorithm applied on $64 \times 64 \times 64$ voxel array with respect to different subset sizes; size 1 was computed by the actual Abdul-Rahman algorithm. Right: Accumulated result error in measures of 2π

sum of the absolute result errors divided by the wrap phase, 2π :

$$e_n = \frac{1}{2\pi} \sum_{ijk} |\varphi_{true}[ijk] - \varphi_{unw,n}[ijk]|, \quad (5.11)$$

where n is the subset size. Except for subset sizes 1 and 2, the average error attains values between 20 and 80. This can e.g. be interpreted as a number of up to 80 incorrectly unwrapped voxels, each missing the correct result by one wrap, which would imply a very good and robust result, as 80 voxels are about 0.03% of the volume and the introduced phase noise was comparably strong (see Fig. 5.5). Realistically, it is much more likely that single voxels show a high mismatch, accumulating unwrap errors. Thus, even less voxels might be falsified. The increasing error rate for low subset sizes is directly related to the quality measure. Phase quality is not considered in the offset-matching process, but only the mask coverage. Surprisingly, the QBU whole volume unwrap shows more errors than the subdivided URSULA result. This is potentially caused by special features in the geometry of the synthetic data and should not be used to derive a general relation. Furthermore, this result will change, if a different algorithm is used to unwrap the individual subsets. However, the URSULA algorithm seems robust for arbitrary subset sizes. The optimal subset size was found to be $s_{x(y,z)} = 8 = \sqrt{n_{x(y,z)}}$. The slight ramps found in Fig. 5.6 (right subfigure) right of the minimum originate in subset geometries featuring smaller subsets at the volume surfaces due to non-integer divisions of the volume size. The error for the minimum computing time (approx. 60 s) is about $36 \cdot (2\pi)$, representing also a local error minimum. Even

with only 8 subsets, $s_{x(y,z)} = n_{x(y,z)}/2 = 32$, the computing time is reduced to almost a quarter of the time required for the entire volume.

A second study was performed to investigate the susceptibility of URSULA to errors in comparison to the standard Abdul-Rahman QBU algorithm. A numerical phantom with dimensioning of $48 \times 48 \times 48$ voxels was simulated and contaminated by Gaussian phase noise with standard deviation, $\sigma = [0^\circ, 11^\circ, 23^\circ, 34^\circ, 46^\circ, 57^\circ]^2$. Furthermore, a spherical mask was applied to the volume to challenge the reassembling step of URSULA. The simulated phase was processed using both algorithms and the hardware configuration described above. Fig. 5.7 shows the model for all simulated noise levels and the unwrapping results of QBU and of URSULA for a subset dimensioning of 4, 8 and 16 voxels edge length in comparison to the true phase. This time, the




Figure 5.7: Unwrapping errors of URSULA and Abdul-Rahman QBU. The bottom row shows the true phase, whilst the wrapped phase is plotted in the top row. The second to fifth row illustrate the unwrap outcome of the algorithms. Yellow numbers indicate the rounded $\mathbb{L}^2$ -norm difference between true and reconstructed phase. Red numbers indicate the required computing time. The yellow arrows indicate errors induced by subdivision.

²Degrees are used instead of radians to allow for better understanding of the noise influence.

result error was estimated as the rounded³ $\mathbb{L}_2$ -norm between true and reconstructed phase as in Eq. 4.14,

$$\mathbb{L}_2(\delta\varphi) = \text{round} \left[\sqrt{\sum_{ijk} (\varphi_{\text{unw}}[ijk] - \varphi_{\text{true}}[ijk])^2} \right]. \quad (5.12)$$

The results clearly show that the computing time of URSULA only depends on the subset dimensioning, but not on the noise level. A subset size of 8 seems ideal regarding time efficiency ($\sqrt{n_{x(y,z)}} = \sqrt{48} \approx 7$), underpinning the former finding. Up to a noise level of 23° the error remains low for all configurations. For higher noise levels the result error of URSULA rises quicker than for QBU. Especially for small subset dimensioning ($4 \times 4 \times 4$) the error increases dramatically, particularly from $\sigma = 57^\circ$. For the subset sizes 4 and 8, errors related to the subdivision and missing relations become apparent, while the result error of the subset size 16 is only slightly higher than the QBU error. Although subdivision-induced errors emerge for high noise level, the URSULA algorithm is fast and reliable. One should further consider that the noise levels triggering errors, are already relatively high and might only occur in regions of a measurement, where the signal is extremely low.

5.6.2 P-URSULA

Optimisation of the parallel mode is more cumbersome than for the sequential mode. One needs to consider many factors that take influence on parallel computing, e.g. data preparation time, read/write operations on files and scripts, refresh times, computing speed of the involved cluster nodes, cluster load induced by competitive jobs, time for reassembling computed subsets – in short: computing *overhead* that strongly depends on the employed computing cluster.

In order to sketch the optimisation process, in the following a synthetic dataset of $96 \times 96 \times 96$ voxels is unwrapped on the local institute cluster (constituting of 40 nodes, 2.0- 2.4GHz QuadCore CPUs) using different subset sizes. The simulated phantom was again falsified with Gaussian noise. An illustration of the phantom can be found in Fig. 5.8.





Figure 5.8: Sample slice of the numerical phantom used for performance optimisation test of P-URSULA; Left: simulated true phase, mid: wrapped phase, right: reconstructed phase; Volume size $96 \times 96 \times 96$ voxels.

A plot of the required computing time is shown in Fig. 5.9. The overall accuracy was not assessed again, since the decomposing and offset-matching method is identical for URSULA and P-URSULA. The optimal subset size, judging by the minimum in computing time, is to be found around 32 voxels, which is a third of the volume size, $s_{\text{opt}} = n_{x(y,z)}/3$. But also the surrounding

³Potential unwrapping errors are multiples of 2π , thus phase differences below 1 rad are of no relevance.

subset sizes, such as 24, 28 or 36 voxels show good results regarding performance and differ only slightly from the minimum value.

Figure 5.9: *Performance of the P-URSULA algorithm applied on $96 \times 96 \times 96$ voxel array with respect to different subset sizes*



At this point, it shall be remarked that the algorithm performance for the tested array size does not outperform the conventional URSULA algorithm significantly. This is due to the overhead generated by read/write operations, submitting and receiving of data. For arrays of larger extent parallelisation becomes more beneficial, since the actual unwrapping process dominates the computing time demands such that overhead plays a less predominant role.

5.7 Pitfalls

The URSULA method is accurate and fast. Nevertheless, it entails some pitfalls that should be considered, whenever the method is employed. This especially affects the choice of the subset size. Some of these pitfalls can be compensated for by careful configuration of the algorithm.

5.7.1 Disjoint Mask Segments

Whenever one subset carries a mask outlining two disjoint regions – or segments – the applied unwrapper is usually unaware of their mutual relation. For data behaving similar in both regions and exhibiting few to no phase wraps, the scenario might be safe. Albeit, if the disjoint segments differ by a non-zero average wrap count, $\bar{n}(\text{reg}_1, \text{reg}_2) \neq 0$, their relation gets lost. The wrap count is ignored, and the surfaces of the segment exhibit inconsistent offsets as in Fig. 5.10. Using a more advantageous subset geometry as in Fig. 5.11 can solve this shortcoming in a straightforward way.

5.7.2 Misinterpreted Regions

Occasionally, regions with well-defined surfaces range over several subsets. The optimal location⁴ to unwrap such a region against the rest of the map resides in one subset. As the unwrap algorithms computing the other subsets are not aware of this, they will try to find the best link inside their own subset. This can lead to inconsistent results, since the local unwrap ignores the actual wrap count between the region and its surroundings, as illustrated in Fig. 5.12.

⁴The link with highest quality




Figure 5.10: The effect of masking combined with unfavourable subset dimensioning: On the left the wrapped data and next to it the masked wrapped data are shown. The selected URSULA subsets lead to an incorrect result (third image) in the lower left segment, if compared to the true phase (right image).




Figure 5.11: Changing the subset size to improve on disjoint mask segments: On the left the wrapped data and next to it the masked wrapped data are shown. The URSULA subsets are now dimensioned to facilitate a proper result (third image), that agrees with the original phase (right image).




Figure 5.12: The effect of large regions, spanning a larger area than covered by one subset and possessing phase properties different from their surroundings. On the left: The simulated data in raw and secondly in wrapped form. The URSULA unwrap in the quadrants on the right fails, since it has insufficient information to estimate the true wrap count between inner and outer areas. The whole volume unwrap, however, solves the problem correctly (right image).

5.8 Outlook

Although URSULA is already a full-grown tool for large array unwrapping, it still bears potential for certain improvements. The most important features to be enhanced are:

- **Computing Speed:** URSULA, P-URSULA and the Abdul-Rahman algorithm were written in a simple MATLAB[®] (The MathWorks Inc., Natick, USA) implementation that still offers some opportunities for improvement. Besides, MATLAB[®] features some speed-ups for array computing, but an implementation in e.g. C++ could potentially further improve performance.

- **Quality:** As mentioned above, disjoint segments inside single subsets might be incorrectly unwrapped and attached. An algorithm identifying and tagging disjoint voxel groups inside the individual subsets could compensate for this shortcoming. The disjoint regions would then be individually matched to the adjacent subsets instead of assigning only a global phase offset per subset.

5.9 Application Example

The unwrap of an *in vivo* phase dataset from a 9.4 T human scanner (Siemens AG, Erlangen, Germany), also included in the first report of the algorithm in [Lindemeyer et al., 2013], is demonstrated. A standard multiple echo GRE sequence was applied including:

Sequence Parameters					
seq	GRE	exc	slab-selective	Tx/Rx	8 ch array (1ch Tx, 8 ch Rx)
matrix	$[308 \times 448 \times 80]$	FOV	$[154 \times 224 \times 40]$ mm	res.	$[0.5 \times 0.5 \times 0.5]$ mm
TR	38 ms	BW	235 Hz/Px	α	48°
TE	$[3.93, 9.68, 15.43, 21.18, 26.93]$ ms				

The acquired channels were recombined using an offset-corrected technique that is introduced in Section 10.6. The recombined phase is shown in Fig. 5.13. It contains numerous phase wraps and some areas with high local noise level. This is caused by destructive interference of the transmit fields resulting in localized signal dropoff. Nevertheless, the phase does not contain significant poles that would hinder successful unwrapping. The phase of the fourth echo was unwrapped using P-URSULA and a subset size of $45 \times 45 \times 45$ voxels. The unwrap results show a robust and highly continuous phase estimate as illustrated in Fig. 5.14. The unwrapped phase of a central line along x-direction is shown in Fig. 5.15 in comparison to the raw phase. The application reveals another pitfall of URSULA: In the slices containing subset edges, some lines of incorrectly assigned phase can be found. These indicate failure of the QBU algorithm at the edges of the subset cubes, presumably induced by local noise. Fig. 5.16 illustrates these effects. A compensation for such artefacts is relatively straightforward. Increasing the subset size by one voxel in each direction⁵, causing the subsets to overlap, and assigning only the inner unwrapped region avoids any surface or edge errors in the resulting volume. These effects are probably not encountered in whole-volume QBU unwraps, as usually the brain mask does not reach the volume surfaces. Nevertheless, this inconsistency does not affect URSULA itself, as it is caused by the utilised subset unwrapping algorithm that can easily be exchanged.

The presented example was additionally computed with (full 3D) PRELUDE [Jenkinson, 2003] for comparison. An earlier echo with fewer wraps had to be chosen, since otherwise the computing time would have exceeded a day. PRELUDE finishes after 169 min and 23 sec, while P-URSULA takes only 25 min and 15 sec using the above parametrisation. This could even be reduced to 12 min and 8 sec by adjusting the subset size to $48 \times 32 \times 16$. The computing time required by P-URSULA is only related to the problem (or array) size, thus it does not change significantly for phase data of higher complexity, such as the presented.

⁵Meaning positive and negative direction for all axis




Figure 5.13: Sample slices of the fourth echo of the acquired phase data (every 8th slice is shown).




Figure 5.14: Unwrapped phase data, after applying P-URSULA with subset size $45 \times 45 \times 45$ (every 8th slice is shown, same slices as in Fig. 5.13).






Figure 5.15: Left: Positioning of the line plot. Right: Line plot along x direction showing wrapped (red) and P-URSULA unwrapped (blue) phase.

Figure 5.16: *The illustrated slice is positioned at the subset surface position, in the centre of the volume (45th slice). It shows occasional line artefacts originating from errors in the subset unwraps, indicated by red arrows.*



5.10 Overview

- ▲ Robust unwrap
- ▲ Speed significantly higher than for Abdul-Rahman (QBU) algorithm and PRELUDE
- ▲ Speed is related to array size rather than to complexity
- ▲ Can be used in cluster-based parallelisation
- ▲ Works on any kind of contiguous mask
- ▲ Phase inconsistencies, such as poles, can in many cases be resolved (see QBU)
- ▲ Can be configured to employ any desired unwrap algorithm for subsets
- ▼ Potential errors in disjoint mask segments *inside subsets*
- ▼ Potential errors in misinterpreted regions expanding over multiple subsets

An approach that also employs image compartmentalisation was presented in Strand et al. [1999] and Strand and Taxt [2002], albeit operating only in 2D and using region growing to connect the compartments rather than making use of quality-based matching as in URSULA.

Following the description and discussion of phase imaging and phase unwrapping of MRI data, the focus now returns to the origin of phase contrast. In Chpt. 3, the demodulated resonance frequency, $\delta\omega$, and the causative local field shifts, b , were introduced (see Eq. 3.14). A dataset describing the local static magnetic field shift for a measurement is called *fieldmap*. This chapter describes methods to derive the fieldmap and related variables. Furthermore, the properties of fieldmaps are characterised.

6.1 Single echo

The most simple way to calculate a fieldmap is using the phase data of a single echo GRE measurement. With Eq. 3.14 and Eq. 4.1, the field shift, b , is related to phase and time by:

$$b(\vec{r}) = \frac{\delta\omega(\vec{r})}{\gamma} \quad (6.1)$$

$$= \gamma^{-1} \cdot \frac{\Delta\varphi(\vec{r})}{\Delta t} \quad (6.2)$$

$$= \gamma^{-1} \cdot \frac{\mathcal{U}(\Delta\varphi_{\text{meas}}(\vec{r}))}{\Delta t} \quad (6.3)$$

where γ is the gyromagnetic ratio, $\mathcal{U}$ symbolises ideal spatial unwrapping of the measured phase difference, $\Delta\varphi_{\text{meas}}(\vec{r})$, observed over the time span, Δt . For a single echo acquisition with the phase distribution, $\varphi_{TE}(\vec{r})$, measured at t_{TE} , one can approximate:

$$b(\vec{r}) = \gamma^{-1} \cdot \frac{\mathcal{U}(\varphi_{TE}(\vec{r}))}{t_{TE}} \quad (6.4)$$

under the assumption, that $\varphi_{TE}(\vec{r})$ at t_{TE} represents the entire accumulated phase since excitation – with other words: $\varphi(\vec{r})_{t=0} \stackrel{!}{=} 0, \forall \vec{r}$. In fact, the effective phase at excitation time is not necessarily zero. Depending on the sequence and hardware used, phase can exhibit a constant offset, φ_0 , such that:

$$\varphi(\vec{r}) = \gamma \cdot b(\vec{r}) \cdot t_{TE} + \varphi_0 \quad (6.5)$$

In order to derive $b(\vec{r})$, the offset φ_0 has to be known. For increasing field strength the zero phase can even exploit spatially varying characteristics, $\varphi_0(\vec{r})$ (e.g. Robinson et al. [2011]). These originate e.g. in inhomogeneous excitation of the sample or distortions in the signal receive process. With this consideration, Eq. 6.3 becomes:

$$b(\vec{r}) = \gamma^{-1} \cdot \frac{\mathcal{U}(\varphi_{\text{meas}}(\vec{r}) - \varphi_0(\vec{r}))}{t_{TE}} . \quad (6.6)$$

Practically, there are no means to measure $\varphi_0(\vec{r})$ *directly* and thus, single echo phase acquisition is generally not an optimal method for deriving a fieldmap. The unknown zero phase offset map can falsify phase and field observations, e.g. by generating phase poles or long range inhomogeneities that have no origin in the true field distribution. Fig. 6.1 shows examples for the phase evolution in measurements at 3 T and 9.4 T. The offset shown in the figure is estimated by the multiple-echo method discussed further below. The necessity of slope estimation instead of single echo phase imaging becomes obvious in the 9.4 T case. Also for 3 T, where only a relatively constant global phase offset is encountered, the resulting field would otherwise be falsified. Still, a constant offset can be accepted in contrast to the effects arising at ultra-high fields such as 9.4 T.

Errors: The uncertainty, $b \pm \varepsilon_b$, implied by single echo fieldmap estimation abides:

$$\varepsilon_{b,\text{single}} = \gamma^{-1} t_{TE}^{-1} \cdot \sqrt{\varepsilon_{\varphi}^2 + \varepsilon_{\varphi_0}^2} \approx \sqrt{2} \cdot \gamma^{-1} t_{TE}^{-1} \cdot \varepsilon_{\varphi} , \quad (6.7)$$

with ε_{φ} being the phase uncertainty at echo time, t_{TE} , and ε_{φ_0} being the phase uncertainty at $t = 0$. The uncertainty arising after excitation cannot be measured directly, as discussed above. The uncertainties may be expected in approximately the same range, as long as no significant offset in the phase map at zero time is encountered. If the latter is the case, Eq. 6.7 is no longer valid, since the offset is not a Gaussian distributed uncertainty.

6.2 Double Echo

Determining the phase offset, $\varphi_0(\vec{r})$, can be avoided by gathering information about the true evolution of phase over two points in time rather than just taking a “snapshot” at one echo time, t_{TE} . Using a double-echo GRE sequence as introduced in Section 2.7.1.4 allows one to determine the phase for two different echo times, t_{TE1} and t_{TE2} [e.g. Jezzard and Balaban, 1995]. This leads to a more adequate fulfilment of Eq. 6.3, since a true phase difference can be calculated:

$$b(\vec{r}) = \gamma^{-1} \cdot \frac{\mathcal{U}(\mathcal{W}(\varphi_{2,\text{meas}}(\vec{r}) - \varphi_{1,\text{meas}}(\vec{r})))}{t_{TE2} - t_{TE1}} . \quad (6.8)$$

The wrap operator, $\mathcal{W}$, is used here, since $\varphi_2(\vec{r}) - \varphi_1(\vec{r})$ may induce phase values in a range of up to $[-2\pi, 2\pi[$ that have to be wrapped back into standard phase range before applying the actual unwrapping process.




Figure 6.1: Sample phase data from post mortem measurements at 3 T and 9.4 T. The plots on the left show the phase evolution beginning from the estimated offset at excitation time over the first two acquired echoes. The blue eye guide symbolises the phase slope or the field value, $b = \delta\varphi/\delta t$, whilst the dotted green lines indicate the field slope that would be estimated for single-echo acquisition. The red line on the y-axis indicates the offset at excitation time. The slice plots show the zero offset and the first two echo phase maps. The reference points for the phase evolution graphs on the left are indicated in red. In the 3 T case (upper half), mostly constant phase offsets at excitation time are encountered, while the 9.4 T case shows strong spatially varying offsets that need to be considered. Even in the 3 T case the constant offset would lead to irritations and a global field offset, if only one echo was considered. Beware, that the presented 9.4 T data is not suited for spatial unwrapping.

Errors: The result uncertainty in this case is the same as in Eq. 6.7 (with $\Delta t_{TE} = t_{TE2} - t_{TE1}$):

$$\varepsilon_{b,\text{double}} = \gamma^{-1}(\Delta t_{TE})^{-1} \cdot \sqrt{\varepsilon_{\varphi_2}^2 + \varepsilon_{\varphi_1}^2} \approx \sqrt{2} \cdot \gamma^{-1}(\Delta t_{TE})^{-1} \cdot \varepsilon_{\varphi} \quad , \quad (6.9)$$

Echo Mode: It is crucial to acquire both echoes in the same readout direction (*monopolar readout*, see Section 2.7.1.4), since the MRI system can induce an indeterminate phase shift between the even and odd echoes (recorded forwards and backwards, with inverted gradient amplitudes). This offset would result in an additional constant phase offset in all phase values acquired.

Wrap Count: For the single echo and the double echo method an additional pitfall exists. In general, an unknown number of wraps might have occurred in time domain before the first echo is recorded, but also between the first and the second echo. This basically means that an estimated field value, b , is ambiguous to the field values:

$$b_n = b + \gamma^{-1} \frac{2\pi \cdot n}{t_{TE2} - t_{TE1}}, \quad n = \pm[1, 2, 3, \dots] \quad . \quad (6.10)$$

This ambiguity is avoided by choosing a proper reference point for unwrapping, where the phase changes only slightly in space and over time (see Section 6.3.2). Sufficient *shimming* can also

prohibit incorrect field estimations (see Section 6.5). Nevertheless, this effect generates a constant offset throughout the fieldmap, which is part of the unwrapping uncertainty.

Fieldmap Example: In Fig. 6.2 an example for a fieldmap derived from a double echo *in vivo* measurement acquired at a 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany) is illustrated. The fieldmap was derived from the echo difference using the P-URSULA algorithm. Further examples for estimated fieldmaps will be shown in later chapters.




Figure 6.2: Fieldmap calculation on double-echo *in vivo* phase measurement as 3 T. The phase difference between second and first echo ($t_{TE} = [9.68, 31.49]$ ms) is calculated and unwrapped using P-URSULA.

6.3 Multiple Echo

A more adept technique to estimate fieldmaps is the use of multiple (more than two) echoes [e.g. Windischberger et al., 2004]. Here, the field is estimated by linear regression. This is expressed as:

$$b(\vec{r}) = \gamma^{-1} \cdot \frac{\sum_{i=1}^n (t_i - \bar{t}) \cdot (\varphi_i(\vec{r}) - \bar{\varphi}(\vec{r}))}{\sum_{i=1}^n (t_i - \bar{t})^2}, \quad (6.11)$$

(e.g. [Hirsch, 2009, p. 42, Eq. 3.1]) where $\bar{t} = 1/n \cdot \sum_i t_i$ is the mean echo time and $\bar{\varphi}(\vec{r}) = 1/n \cdot \sum_i \varphi(\vec{r}, t_i)$ is the mean value of the *true* phase. Nevertheless, before the regression can take place, the phase has to be prepared using spatial or temporal unwrapping.

Errors: The error, ε_b , accumulated in linear regression of n echoes can roughly be estimated as:

$$\varepsilon_{b,\text{multiple}} \approx \gamma^{-1} \cdot \frac{\varepsilon_\varphi}{\sqrt{\sum_{i=1}^n (t_i - \bar{t})^2}} \quad (6.12)$$

The positioning of the echo times is essential for narrow error margins – wide echo spacing increases the $\mathbb{L}_2$ norm in the denominator and thus reduces the result error.

6.3.1 Unwrapping in Time Domain

If the true phase between two consecutive echoes changes less than π throughout the entire volume marked for imaging¹, linear regression can be performed directly after applying one-dimensional

¹i.e. the time domain fulfils *Nyquist's* criterion

unwrapping along time domain:

$$\varphi(\vec{r}, t) = \mathcal{U}_t(\varphi_{\text{meas}}(\vec{r}, t)) \quad (6.13)$$

Unfortunately, this is not the common case (see Chpt. 4, Fig. 4.3). This procedure is only applicable for measurements of extraordinarily well shimmed subjects containing no strong local field shifts and using reasonably short echo time spacing.

6.3.2 Unwrapping in Spatial Domain

Usually, each echo has to be unwrapped in space, separately:

$$\varphi_{\text{unw}}(\vec{r}, t_i) = \mathcal{U}_3(\varphi_{\text{meas}}(\vec{r}, t_i)) \quad , \quad (6.14)$$

where $\mathcal{U}_3$ indicates spatial unwrapping. Unlike the double echo method, the unknown wrap count between proximate echoes becomes fatal in the linear regression operation. In order to apply a proper regression, the actual evolution of the phase has to be reconstructed throughout all echoes after spatial unwrapping was performed. As the true wrap count between two echoes is unknown, this is not a primitive task. A trick is utilised to estimate a valid wrap count everywhere. A spatial reference position, $\vec{r}_0$, is identified, possessing a phase change over time, that is low enough to fulfil *Nyquist's* sampling law. Usually such a reference point can be found in regions, where also the spatial phase gradients are smooth throughout the echoes. For MRI measurements the image centre normally offers good conditions, as it is usually well shimmed and close to the isocentre of the magnet. At the determined reference point, unwrap in time domain is performed:

$$\varphi_{\text{off}}(t) = \mathcal{U}_t(\varphi_{\text{meas}}(\vec{r}_0, t)) \quad . \quad (6.15)$$

The offset vector, $\varphi_{\text{off}}(t)$, contains the true phase evolution in $\vec{r}_0$, connecting the separately unwrapped echo volumes. Although the phase volumes lose their relation in the spatial unwraps, they can be reordered in time in relation to the reference point:

$$\varphi(\vec{r}, t) = \varphi_{\text{unw}}(\vec{r}, t) - \varphi_{\text{unw}}(\vec{r}_0, t) + \varphi_{\text{off}}(t) \quad , \quad (6.16)$$

as illustrated in Fig. 6.3.

If all echoes are properly arranged with regard to each other, linear regression as in Eq. 6.11 can be performed to estimate $b(\vec{r})$.

6.4 Multiple Echo Differences

An alternative to directly performing linear regression on (unwrapped) multiple echo data is to calculate phase difference maps first. This approach is chosen, if the acquired phase data are not consistent and e.g. exhibit phase poles that prevent successful spatial unwrapping. Examples for this are shown above (Fig. 6.1) and in Chpt. 11, Section 11.5. The phase difference will in most




Figure 6.3: Evolution of phase on the example of one data line (along position axis) with evolving phase over the echo times; left: recorded signal, middle: separately unwrapped echoes – chaotic evolution in time, right: reordered echoes show linear temporal evolution; units: radians.

cases eliminate phase poles to the cost of a small region with higher noise level where the poles resided. Difference calculation can be performed using the first echo as a reference:

$$\varphi'(\vec{r}, t_i) = \mathcal{W}[\varphi(\vec{r}, t_i) - \varphi(\vec{r}, t_1)] \quad (6.17)$$

$$t' = t_i - t_1, \quad (6.18)$$

keeping the number of echoes constant. This technique will be called *division by 1st echo*, since this voxel-wise operation is usually performed in complex polar notation:

$$\varphi'(\vec{r}, t_i) = \angle \left(e^{i\varphi(\vec{r}, t_i)} / e^{i\varphi(\vec{r}, t_1)} \right). \quad (6.19)$$

The new phase dataset, $\varphi'(\vec{r}, t_i)$, and the echo times, t'_i , can then be analysed by linear regression as for a standard multiple-echo dataset. A second approach is the calculation of *multiple echo differences*. Unlike the division by the 1st echo, this method calculates differences between echo pairs:

$$\Delta\varphi_i(\vec{r}) = \mathcal{W}[\varphi(\vec{r}, t_{i+d}) - \varphi(\vec{r}, t_i)] \quad (6.20)$$

$$\Delta t_i = t_{i+d} - t_i, \quad (6.21)$$

where d is an integer smaller than the echo count, n_{TE} . After difference calculation, there are $n_{TE} - d$ separate datasets remaining. These are spatially unwrapped and potential phase offsets, n_i , in form of integer wrap counts are corrected for. The average of the unwrapped and corrected data is used to determine the fieldmap:

$$b(\vec{r}) = (\gamma \cdot (n_{TE} - d))^{-1} \cdot \sum_{i=1}^{n_{TE}-d} \frac{\mathcal{U}_3(\Delta\varphi_i(\vec{r})) - 2\pi \cdot n_i}{\Delta t_i} \quad (6.22)$$

The uncertainty arising with this method can be approximated as that of the double echo acquisition (Eq. 6.9) decreased by the factor $1/\sqrt{n_{TE} - d}$ (the number of echo differences calculated). This only holds true under the premiss that the phase uncertainty is constant over all echo times. Fig. 6.4 sketches the procedure for fieldmap calculation based on multiple echo differences.




Figure 6.4: Principle of multiple echo difference calculation, used for data with inconsistent phase. The recorded data have equally-spaced echo times, allowing for simple averaging. Phase poles are eliminated and the resulting fieldmap is the mean value of all unwrapped, offset-corrected phase differences calculated.

6.5 Shimming

The shim coils (e.g. Gruetter [1992]), mentioned in Section 2.9, are a crucial component in acquiring robust phase data and fieldmaps. They are engineered to compensate for harmonic field distortions, usually up to an order of two. For each spherical harmonic function, a set of coils exists, that can be controlled separately with regard to its field strength. As the major distortions inside the static field tend to have harmonic characteristics, spherical harmonics appear as an optimal choice for this task (see e.g. Zhao et al. [2005]). In Cartesian coordinates, the field characteristics, $B_{\text{shim},j}$, of the shim magnets up to second order can (under negligence of normalisation) be described by the following terms:

$$B_{\text{shim},j=[1:9]} : \quad 1; \quad x; \quad y; \quad z; \quad -x^2 - y^2 + 2z^2; \quad yz; \quad zx; \quad xy; \quad x^2 - y^2 \quad . \quad (6.23)$$

A parameter, A_j , defines the weighting of each term in the superposition:

$$B_{\text{shim}} = \sum_j A_j B_{\text{shim},j} \quad . \quad (6.24)$$

Shimming is usually optimised only on the selected Field Of View (FOV), to induce best field

distortion compensation in this area. This can be performed automatically, using a fast phase imaging sequence based on which the parameters are estimated, or manually by defining the individual shim currents. The shim currents ideally have a linear relation to the coefficients, A_j .

6.6 Optimal Fieldmap Estimation

The basis for a solid fieldmap estimation is a low noise level of the acquired phase data. Therefore, the influence of parameters such as the *echo time*, the *bandwidth* and the *number of acquired echoes* will be discussed in the following.

6.6.1 Echo Time

Whenever the magnitude of an MRI signal is observed or involved in contrast generation, some basic relations between the image SNR and the echo time, t_{TE} , are obvious. First, it is known that the magnitude – representing the transverse magnetisation of the sample – decays in an exponential manner² with T_2^* for GRE-based sequences. Since noise in general experiences no decay, one could roughly expect the SNR to decrease in the same way as the signal magnitude – exponentially – with increasing echo time.

As indicated before (see Section 3.5), the phase is much more cumbersome to handle, when estimating its SNR. In general, two effects compete in being responsible for the phase contrast level. The first one is the linear relation of phase SNR to magnitude SNR (as described in Section 3.5), leading to increasing uncertainty for late echo times, related to the exponential decay. The second effect is the enrichment in contrast with time after excitation. It was assumed that phase is ideally flat at excitation time ($\varphi(\vec{r}, t = 0) = 0, \forall \vec{r}$). The model used in this thesis furthermore assumes that phase is linearly changing with time (Eq. 6.3). The slope is dependent on the local field shift, b . This implies that especially voxels possessing low field shifts will show little phase variations for early echo times. These initial phase variations may range below the minimum phase uncertainty (Section 3.6) and yield extremely low SNR, whilst very late echo times yield high SNR even for low field shifts, according to this second argumentation. In the following, an estimation for the echo time with optimal phase contrast will be presented.

6.6.1.1 Material and Methods

Measurements of a *post mortem* brain, residing in a formalin-filled acrylic cylinder, were carried out at a 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany) using a single-slice GRE sequence. Shimming was optimised on the measurement volume in five iterations. The sequence was applied including the following parameters:

Sequence Parameters					
seq	GRE	exc	single slice	Tx/Rx	Body / 12 ch array
matrix	320×320	FOV	$[192 \times 192 \times 6]$ mm	res.	$[0.6 \times 0.6 \times 6]$ mm
TR	100 ms	α	25°	BW	120 Hz/Px
TE	[6.16, 10, 15, 20, 25, 30, 35, 40, 45, 50, 55, 60, 70, 80] ms				

²The fact that the decay is not necessarily solely exponential is neglected here.

The measured slice was positioned in transverse orientation near the central brain regions. A mask is generated using automated phase masking (see Section 4.9) and the phase inside the masked region is unwrapped using PRELUDE [Jenkinson, 2003]. A rectangular region within the mask, containing signal everywhere, is selected for numerical comparison of the measurements.

The local phase variation (see Section 3.5.3) is assessed by applying a median filter with a 3×3 neighbourhood to the phase and subtracting the result:

$$\varphi_{\text{local}} = \varphi - \text{median}_{3 \times 3}(\varphi) \quad . \quad (6.25)$$

The standard deviation of φ_{local} is a measure for the phase uncertainty:

$$\delta\varphi = \text{std}(\varphi - \text{median}_{3 \times 3}(\varphi)) \quad . \quad (6.26)$$

In order to estimate the SNR, the maximum value range, φ_{MinMax} of the unwrapped phase is determined. The SNR is then defined by:

$$\text{SNR} = \frac{\varphi_{\text{MinMax}}}{\delta\varphi} \quad (6.27)$$

As a reference, the average T_2^* value of the region is calculated by transforming the magnitude data to logarithmic scale and performing a linear regression fit in time domain to every voxel. The maximum of the T_2^* distribution is chosen as the representative value.

6.6.1.2 Results

An overview of the raw phase distribution, the unwrapped phase and the median-corrected local phase variation is illustrated in Fig. 6.5. Fig. 6.6 shows the dependence of φ_{MinMax} , $\delta\varphi$ and the SNR on echo time. The maximum SNR is found for an echo time of 35 ms. The T_2^* distribution of all voxels contained in the mask is shown in the histogram in Fig. 6.7. The maximum abundance is located at $T_2^* \approx 30$ ms.




Figure 6.5: The top row shows the raw phase, the middle row shows the unwrapped phase and the lower row illustrates the local phase variance after median-correction. From left to right: [6.16, 10, 15, 20, 25, 30, 35, 40, 45, 50, 55, 60, 70, 80] ms. Units are degrees ($^\circ$).




Figure 6.6: From left to right: Value range of unwrapped phase, φ_{MinMax} , standard deviation of local (median-corrected) phase, $\delta\varphi$, and calculated SNR with respect to the echo time. Units are degree ($^\circ$) for first and second plot




Figure 6.7: Histogram plot of the distribution of T_2^* values in the selected region of the brain.

6.6.1.3 Discussion and Conclusions

The value range increases linearly – this is to be expected, as phase linearly evolves over time. The standard deviation starts with a small slope, but ascends more rapidly after passing $t_{TE} = T_2^*$. Finally, the SNR increases between lower echo times and the average T_2^* value of the sample, shows a maximum around T_2^* and decreases for late echo times. It is interesting that the maximum SNR and the average T_2^* value of the measured tissue coincide. Consequently, for e.g. double-echo measurements the second echo should be positioned somewhere near T_2^* , whilst setting the first echo time as early as possible. This will maximise the time and thus the phase evolution between both echoes. Further, the contrast in the second echo is maximised due to its positioning. Deistung et al. [2008] report related findings in a different context. Since the primary use of an early echo is to serve as an offset-correction reference, its local contrast is only of little importance.

6.6.2 Bandwidth

The bandwidth of a measurement is known to be closely related to the noise variance, mainly with regard to magnitude information [Haacke et al., 1999, p. 334, Eq. 15.7]. Since the noise in phase data correlates to the magnitude noise (see Section 3.5), this suggests that a low bandwidth is also beneficial for the phase.

6.6.2.1 Material and Methods

A cylindrical phantom, filled with water and containing an air-filled tube close to its centre (sketched in Fig. 6.8), was used for reference measurements. Since air-filled cavities are known to induce strong field distortions in their surroundings, the described tube was used to generate spatially varying phase contrast. The measurement was performed on a 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany) using the following parameters:

Sequence Parameters					
seq	GRE	exc	slab	Tx/Rx	Body / 12 ch array
matrix	$128 \times 128 \times 96$	FOV	$[153.6 \times 153.6 \times 115.2]$ mm	res.	$[1.2 \times 1.2 \times 1.2]$ mm
TR	32 ms	α	9°	TE	$[9.27; 23.19]$ ms
BW	$[80, 100, 140, 220, 400, 800, 1220]$ Hz/Px				




Figure 6.8: Phantom filled with water, holding an air-filled tube in its centre that generates field distortions in its surroundings.

A three-dimensional region of interest (situated in the air-tube's stray field in the water) is selected and the phase difference of both echoes is calculated and spatially unwrapped using the *Abdul-Rahman* algorithm:

$$\Delta\varphi(\vec{r}) = \mathcal{U}[\varphi(\vec{r}, TE_2) - \varphi(\vec{r}, TE_1)] \quad (6.28)$$

The following parameters are extracted from the phase difference:

- $\overline{\Delta\varphi}$, the average ROI phase
- φ_{MinMax} , the value range of the slightly smoothed³ phase difference (as in the previous experiment, Section 6.6.1.1)
- $\delta\varphi = \text{std}(\Delta\varphi_{\text{local}})$, the standard deviation of the median-corrected phase difference, $\Delta\varphi_{\text{local}}$.
- The SNR of the measurement,

$$\text{SNR} = \frac{\varphi_{\text{MinMax}}}{\delta\varphi}$$

6.6.2.2 Results

An overview including sample slices for the phase of the measured volume, the unwrapped ROI phase and the local (median-corrected) phase is shown in Fig. 6.9. The resulting graphs for average phase, $\overline{\Delta\varphi}$, value range, φ_{MinMax} , the phase variation, $\delta\varphi$, and the SNR are plotted in Fig. 6.10.




Figure 6.9: Overview, showing sample slices of the raw phase (top row), the unwrapped phase inside the ROI (middle row) and the local phase variation after median-correction (bottom row). The applied bandwidth values are (from left to right): [80, 100, 140, 220, 400, 800, 1220] Hz/Px. The position of the ROI is indicated in the sample slice. Units are deg ($^{\circ}$)




Figure 6.10: Top left: Average phase value in the ROI; top right: value range of ROI; bottom left: Standard deviation of median-corrected ROI and bottom right: SNR.

Obviously, the phase variation increases dramatically with increasing bandwidth, whilst the SNR decreases. The average phase and the phase value range also increase with the bandwidth.

6.6.2.3 Conclusions

The phase variation and further the SNR show the expected behaviour. It becomes apparent that low bandwidth values such as 100 Hz/Px or even 80 Hz/Px significantly improve the resulting phase SNR. Surprisingly, the value range and the average phase increase with the bandwidth, in particular in the lower part of the bandwidth scale. At least the value range should theoretically

³Smoothed by a Gaussian filter of $\sigma = 1$ voxel

not change significantly, since the echo time parametrisation remains constant. The origin of this shift is not clear. A low bandwidth might generate a smear-out effect over time. However, the increase of the average phase might also originate in unwanted side-effects, taking influence on the observed phase characteristics. A low bandwidth should nevertheless be preferred in order to avoid loss of SNR and to minimise the influence of additional factors.

6.6.3 Number of Echoes

Although the acquisition of a higher number of echoes for the estimation of a fieldmap should statistically improve the stability of the result, it was shown, that the SNR of phase changes with the echo time chosen for recording. It was further discussed, that acquiring the phase for two or more echoes has additional advantages such as the compensation for constant and spatially varying phase offsets. In principle, it makes sense, to acquire as many echoes as possible to ensure statistical robustness of the fieldmap derived by linear regression. Unfortunately, for *in vivo* observations measurement time is strongly restricted, not only by the optimal echo time, but also by the maximum duration for measurements. And thus TR is also restricted. If a measurement was optimised for one or two echoes, the introduction of more echoes enforces shorter acquisition blocks meaning higher recording bandwidths, as the readout time is inversely proportional to the bandwidth [e.g. Haacke et al., 1999, p. 341, Eq. 15.30]. This again, decreases the phase SNR of the individual echoes as discussed in Section 6.6.2. The effect of different echo counts in a fixed repetition time is assessed in a small measurement series.

6.6.3.1 Material and Methods

Measurements were performed on the phantom introduced in Section 6.6.1.1. A full 3D GRE measurement with a non-selective pulse was applied using 2, 4, 6 and 8 echoes. Each time the echo time points were fit into the available time frame within $TR = 35$ ms, meaning that an optimal echo positioning for minimum possible bandwidth was chosen, to maximise the result quality:

Sequence Parameters					
seq	GRE	exc	non-selective	Tx/Rx	Body / 12 ch array
matrix	$192 \times 192 \times 144$	FOV	$[192 \times 192 \times 144]$ mm	res.	$[1.0 \times 1.0 \times 1.0]$ mm
TR	35 ms	α	24°	RF-spoil	on
TE	[7.93, 26.39] ms			BW	70 Hz/Px
	[4.36, 12.68, 21.00, 29.32] ms				140 Hz/Px
	[2.95, 8.51, 14.07, 19.63, 25.19, 30.75] ms				230 Hz/Px
	[2.21, 6.33, 10.45, 14.57, 18.69, 22.81, 26.93, 31.05] ms				350 Hz/Px

For all measurements, fieldmaps are calculated using spatial unwraps of each echo and a linear regression fit in time domain. A mask with continuous signal is selected for analysis. The $3 \times 3 \times 3$ neighbourhood average is subtracted from the fieldmap and the standard deviation of the result, $\delta b = \text{std}(b - \text{mean}_{3 \times 3 \times 3}(b))$, is calculated. Further, the difference between maximum and minimum, $(\Delta b)_{\text{MinMax}}$, of the fieldmap divided by the standard deviation is used to estimate an average SNR (as in the previous sections).

6.6.3.2 Results

Fig. 6.11 shows the fieldmaps and the local standard deviation for a sample slice of each measurement. Fig. 6.12 shows how the average standard deviation and the SNR throughout the masked region change with echo count. The mean standard deviation of the signal area remains in a very limited range and varies only slightly. A minimum standard deviation and at the same time maximum SNR can be found for 6 echoes, although the values differ only slightly from the other echo configurations (especially 4 echoes).




Figure 6.11: Overview, showing in the top row the calculated fieldmap and in the lower row the $3 \times 3 \times 3$ neighbourhood-corrected field map. From left to right the measurements for 2, 4, 6 and 8 echoes are shown. Visually no significant change can be observed in the noise level.




Figure 6.12: Left: Average local standard deviation for different number of echoes acquired within $TR = 35$ ms. Right: SNR. An echo count of 4 or 6 seems to be optimal, but the SNR for 2 or 8 echoes is only slightly lower.

6.6.3.3 Conclusions

Apparently, there is no great disadvantage in recording more than two echoes, even while being bound to lower bandwidths. The optimum, found for 6 echoes, shows only slight improvement compared to the other configurations. Recording with other echo counts can be performed without a huge loss of accuracy. Higher echo counts than eight or ten would not be advisable for the employed TR of 35 ms, as the bandwidth has to be increased too significantly, leading to a more dramatic loss of phase contrast and detail (see Section 6.6.2). Furthermore, this example only discusses a very special case of TR -confinement and object geometry. When other confinements and different objects (e.g. exhibiting higher local contrast) are observed, the optimal number of echoes might be different or more distinct.

6.6.4 Interfering Compounds with Chemical Shifts

Another consideration in choosing appropriate acquisition echo times, is the chemical shift discussed in Section 3.3.2. Human tissue consists mainly of water ($\sim 70\%$). Nevertheless, a significant part of the other compounds such as lipids also contain hydrogen. Due to the shielding effect induced by the bindings in the molecule, the resonance frequency of these protons is lower than that of free water. In human tissue, most fat has an offset of about $\sigma_{fat} = 3.35$ ppm with regard to water [Haacke et al., 1999, p. 422]:

$$\omega_{fat} = \omega_{water} - \sigma\omega_{water} \quad (6.29)$$

For example, at 3 T, this implies an off-resonance of:

$$\delta\omega_{wf} = -\sigma \cdot 42.576 \frac{\text{MHz}}{\text{T}} \cdot 3 \text{ T} \approx -429 \text{ Hz} \quad (6.30)$$

Practically, this means that after excitation water and fat mutually dephase. They reunite to the next *in-phase* state after 2.33 ms (then, after 4.67 ms and so on). Although, in the human head most fat is situated in the surroundings of the brain, also a fair amount of lipids may reside in cells, fluids or blood inside the brain. These lipid components contribute to the signal and thus also to the phase. Assuming a two-compartment model for a sample voxel, the phase of the fat compartment will contribute to the resulting phase observed in the voxel:

$$\varphi_{obs}(t) = \angle \left(A_{water}(t) \cdot e^{i \cdot \varphi_{water}(t)} + A_{fat}(t) \cdot e^{i \cdot \varphi_{fat}(t)} \right) \quad (6.31)$$

$$= \angle \left(A_{water}(t) \cdot e^{i \cdot \varphi_{water}(t)} + A_{fat}(t) \cdot e^{i \cdot \delta\omega_{wf} \cdot t} \right) \quad (6.32)$$

If the echo times for observation are chosen in an adverse way, e.g. $\delta\varphi_{wf} \cdot t_{TE1} = \pi/2$, and $\delta\varphi_{wf} \cdot t_{TE2} = -\pi/2$, the observed phase evolution would not have the same slope as the actual water phase (see 2nd and 4th echo in Fig. 6.13).




Figure 6.13: Schematic illustration of the phase effects induced by the water-fat shift. The upper row shows the phase of water (blue), fat (orange) and the combined signal (black). The graph illustrates the phase evolution and its states at the measurement points. Only the exact *in-phase* and *opposing-phase* states have the same phase as the true water signal.

Only the signal of free water is included into the model and analysis presented within this work. Chemical shifts – or namely: contributions from hydrogen protons precessing at other frequencies than that of free water – are not quantified. Therefore it is advantageous to avoid any influence from off-resonant protons. The recommendation shall thus be to choose echo times with identical water-fat phase offsets. This can be the *in-phase* situation, but since usually only the phase gradient is analysed over time, an arbitrary – but constant – offset between the signals is sufficient.

If one chooses exactly the in-phase and *opposing-phase* echo times, it is possible to achieve even half of the echo time spacing. For 3 T the temporal distance of the in-phase states is just 2.33 ms. A more dense echo positioning makes not much sense, since one wants to keep the bandwidth low (inducing long readout blocks). Further, if the in-phase and opposing-phase states are not matched accurately, one risks inducing offsets between even and odd signals as in the example above (see Fig. 6.13). Choosing an echo spacing matching the in-phase state timing is more reliable.

6.6.5 Summary

As in all imaging modalities that can be used for multiple purposes, it remains only to say, that the measurement parameters should be fit best to the given confinements. Roughly, optimal phase contrast is achieved with a low bandwidth and at least one echo acquired close to the average T_2^* of the observed region. When observing mixed signals, such as that of water and fat in human tissue, the echo times should be chosen according to the cycle of the in-phase state of the main components, if possible. Single-echo recordings are not recommended as they might promote unknown phase offsets from the excitation or the receive side. Double-echo scans are popular, as they allow for a very low bandwidth associated with low repetition times. Nevertheless, multiple echo acquisition, such as 4, 6, 8 or more are also feasible and might even improve the SNR of the derived fieldmap as long as the measurement parameters are kept in a reasonable range.

6.7 Spatial Distortions

While great interest is attributed to the static magnetic field and its local behaviour, it is at the same time a potential source for image distortions. Imaging is based on field gradients, and in particular the linearity of the readout gradient is crucial for acquiring a proper representation of the spatial domain.⁴ Eq. 2.22 describes the relation of r_a , the coordinate in readout direction and the encoding readout gradient, G_a . Usually, r_a is encoded with:

$$b_{G_a}(\vec{r}) = G_a r_a = (G_a \hat{a}) \vec{r} \quad . \quad (6.33)$$

Unfortunately, only in an ideal acquisition one can expect the magnetic field to be a constant superimposed with the linear gradient, G_a . If an arbitrary, spatially varying field shift, $b(\vec{r})$, is introduced the equation has to be modified:

$$(G_a \hat{a}) \vec{r}_{\text{obs}} = (G_a \hat{a}) \vec{r} + b(\vec{r}) \quad . \quad (6.34)$$

Literally this means, the observed position, $\vec{r}_{\text{obs}}$, is a shifted version of the real world position, $\vec{r}$. The shift occurs along the readout direction, $\hat{a}$. Eq. 6.34 can be rephrased as:

$$\hat{a} \vec{r}_{\text{obs}} = \hat{a} \vec{r} + \left(\frac{b}{G_a} \right) \quad . \quad (6.35)$$

⁴Hardware imperfections of the actual gradient linearity can be accounted for by the scanner software as these are independent of the subject.

For imaging parametrisation gradients can be defined over the readout bandwidth – the frequency range per picture element, Hz/Px. With the spatial position in the discretised image, measured in Px, the shifting term takes effect relative to the voxel width:

$$\Delta r_a = \left(\frac{b}{G_a} \right) [Px] \quad (6.36)$$




Figure 6.14: Schematic illustration of distortions induced by field offsets in connection with low receive bandwidth. A) and C) show the true situation, A) represents a greyscale fieldmap and C) a simple grid. B) and D) show the observed fieldmap and the grid distorted by observation, as in Eq. 6.35. The readout direction is indicated by the black arrow.

Regarding magnitude images this problem is e.g. described by [Haacke et al., 1999, p. 421ff, Eq. 17.5 and p. 570-580] and further by [e.g. Jezzard and Balaban, 1995] for Echo Planar Imaging (EPI). Ideally, this distortion has to be corrected for. As it was pointed out before, phase measurements provide optimal contrast for low bandwidths, such as ~ 100 Hz/Px. Any field shift, b , that approaches 10^2 Hz induces a significant spatial shift of the concerned voxel.

Fortunately, shifts usually remain below one or two voxels throughout the volume of a typical measurement. Yet, the discussion suggests to keep the static field homogeneous up to a high degree, using intensive active shimming (see Section 6.5). This will minimise the influence of field-induced distortions in the acquisition. When interpreting anatomic images generated with GRE sequences at low bandwidth, one should always consider the background fieldmap before discussing relative positions within the dataset. Local field contrast at 3 T scanners is usually in a range of less than ± 10 Hz, so the local surroundings are not significantly displaced. Mainly long range distortions generated by the background field are a potential source of distortions. This influence increases heavily for ultra-high field, thus B_0 distortion correction might become a serious topic in the near future.

that induce a high number of wraps for late echo times and potential poles. Generally, the delta mode offers a means to circumvent non-consistent phase data, since the phase difference is usually consistent and compatible with spatial unwrapping.

6.8.2 Unwrap Mode

The unwrap mode chooses between three major modes. Temporal mode performs a one-dimensional unwrap of phase in time domain. The other two modes are sliced and volume unwrap, both operating in spatial domain. While sliced unwrapping is always performed with the FSL tool PRELUDE [Jenkinson, 2003], volume unwrap can optionally be performed either with PRELUDE or with URSULA/P-URSULA.

The unwrap mode decides further over the mask preparation. FEST makes sure that the mask only contains one continuous region to be unwrapped – if several disjoint regions are present, FEST chooses the biggest one. In slice mode, an additional check is performed, excluding disjoint regions in every single slice (since sliced unwrapping is not able to guess the relation of multiple regions). This can lead to exclusion of regions that were in principle part of the continuous mask in three dimensions.

6.8.3 Echo Separation

FEST offers the option to compute even and odd echoes separately. This option should be chosen, when a GRE measurement with bipolar readout was acquired, since positive and negative readout lobes might induce a constant phase offset between even and odd echoes [see e.g. Scheffler, 2004]. If echo separation is performed, implicit wider echo spacing must be taken into account, since only every second echo is considered in each computing part. This is less important for the spatial unwrap modes, but can cause problems for time unwrapping due to undersampling issues.

6.8.4 Field Calculation

Depending on the previously selected modes, the fieldmap is calculated from prepared phase and echo time data.

- **Temporal Unwrap:** Linear regression in time domain is performed.
- **Spatial Unwrap:** The phase offsets of the spatially unwrapped echoes are reordered using a reference point (see Section 6.3.2). Following this, linear regression in time domain is performed.
- **Delta Spatial Unwrap:** The average of the spatially unwrapped echo-differences is calculated.
- **Echo Separation:** If even/odd echoes were separated, each of the two sets is computed by one of the choices above. Finally, the average is calculated.

6.8.5 Offset Estimation and Smoothing

In all modes the offset, $\varphi_0(\vec{r}) = \varphi(\vec{r}, t = 0)$ can optionally be returned. For the linear regression modes, the offset is a straightforward result of the computation. For the delta mode, the offset is given by:

$$\varphi_0(\vec{r}) = \varphi(\vec{r}, t_1) - t_1 \cdot \gamma b(\vec{r}) \quad , \quad (6.37)$$

(where $b(\vec{r})$ is the estimated local field value). This is potentially not the true offset at zero time and might still contain a non-zero wrap count ($n \cdot 2\pi$). This is due to the unknown relative wrap count of the first echo phase. Optionally, the offset map can be used to automatically mask voxels, in which the offset exhibits a strong deviation from zero.

If the first echo phase of a dataset contains a high noise level, a Gaussian smoothing operation can be applied to it. This preserves smooth phase variations, whilst suppressing noise. Local detail is usually hardly evolved in early echo phase.

Magnetic Susceptibility

Despite early but isolated reports regarding the usefulness of MR phase images for the characterisation of pathology (e.g. Lai et al. [1996] and Reichenbach et al. [1997]), its potential is only beginning to gain recognition in clinical research via susceptibility weighted imaging (SWI, Haacke et al. [2004]) and more recently Quantitative Susceptibility Mapping (QSM) [de Rochefort et al., 2010b]. In 1996 Schenck et al. published a comprehensive article concerning the role of magnetic susceptibility in NMR and MRI [Schenck, 1996]. Both local phase and local magnetic susceptibility distributions have been shown to exhibit excellent tissue contrast in the human brain, based on its microscopic structure (e.g. Langkammer et al. [2012]; Duyn et al. [2007]). Today, besides chemical shift effects, magnetic susceptibility is being held primarily responsible for local phase shifts inside the brain [Schenck, 1996; Duyn et al., 2007].

7.1 Susceptibility Weighted Imaging

SWI was first introduced by Haacke et al. [2004], whilst its underlying principle was described by Reichenbach and Haacke [2001]. It offers a fast and useful way of including additional phase-based contrast to T_2^* -weighted magnitude images. The technique implies an aggressive filtering of the phase in its complex representation. Wraps and background fields are removed while promoting only highly localised phase contrast. After application of thresholding, the resulting local phase contrast is multiplied with the magnitude image. This leads to contrast enhancement of certain regions, such as blood vessels, haemorrhages or calcifications [e.g. Wu et al., 2010; Mittal et al., 2009].

The phase contrast within these regions was attributed to susceptibility variations, distorting the magnetic field in their close surroundings. However, SWI does not yet exploit the full potential of phase information regarding electromagnetic properties of tissue. A more complete description of phase contrast, focusing solely on the phase signal, is nowadays offered by QSM.

7.2 Physical Background

7.2.1 Defining Magnetic Susceptibility

Let $\vec{H}$ be the magnetic field. Then, the magnetic flux density, $\vec{B}$, in linear media can be expressed as [e.g. Feynman et al., 2010, p. 725, Eq. 36.12]:

$$\vec{B} = \mu_0 \mu_r \cdot \vec{H} \quad , \quad (7.1)$$

where μ_0 is the vacuum permeability and μ_r is the relative permeability. This rather compact definition introduces $\vec{B}$ as a scaled version of $\vec{H}$. The relative magnetic permeability is characteristic of the material observed. The magnetic properties of a material can alternatively be characterised by the so-called magnetic susceptibility, χ , defined by:

$$\chi = \mu_r - 1 \quad \Leftrightarrow \quad \mu_r = 1 + \chi \quad , \quad (7.2)$$

and thus:

$$\vec{B} = \mu_0(1 + \chi) \cdot \vec{H} \quad . \quad (7.3)$$

Within this work the phrase *susceptibility* stands for the *magnetic volume susceptibility*. Alternative definitions like the mass susceptibility, χ_m , or the relative susceptibility, χ_r , are not considered here.

A more general description than a linear relation between $\vec{H}$ and $\vec{B}$ – and a fundamental definition of magneto-statics in matter – is achieved by introducing the magnetisation $\vec{M}$:

$$\vec{B} = \mu_0 \cdot (\vec{H} + \vec{M}) \quad \Rightarrow \quad \vec{M} = \chi \vec{H} \quad . \quad (7.4)$$

Eq. 7.3 can alternatively be expressed by using a tensor-valued χ_{uv} :

$$\vec{B} = B_u = \mu_0 \cdot (1 + \underbrace{\chi_{uv}}_{=M_u}) H_v \quad . \quad (7.5)$$

Beware, that magnetisation is now expressed as a transform of the static field, $\vec{H}$:

$$\vec{M} = M_u = \chi_{uv} H_v \quad . \quad (7.6)$$

Throughout this thesis, the phrase *magnetic field* simplistically denotes the flux density, $\vec{B}$.

7.2.2 Spatial Effects of Susceptibility

Unfortunately, the relation between the field and the underlying susceptibility distribution only takes the simple form, described in Eq. 7.3, when $\chi = \text{const.}$, everywhere. Usually, this is not the case, but one rather expects a spatial distribution of magnetic susceptibility in the observed sample: $\chi = \chi(\vec{r})$.

In order to understand the link between magnetic susceptibility and field shifts of objects inserted into a static magnetic field, one has to consider how the underlying magnetisation takes influence on the static field. This context was already examined by Deville and Bernier [1979] and more recently by Marques and Bowtell [2005] (see also [Jackson, 1999, Eq. 5.64]). The applied model is based on the *Lorentz Sphere* approximation. This approach, originally derived in electrostatics by [Lorentz, 1916, Note 54], is applied at the level of a single nucleus. The observed nucleus is positioned inside an imaginary sphere. The field, experienced by the sphere is approximately constant on the length scale of the sphere, since the neighbouring molecules, generating macroscopic field distortions in their superposition, reside at a large distance. The field on the outside of the imaginary sphere generated by its magnetic moment has dipole characteristics, while the field inside is constant (e.g. Bhagwandien et al. [1994] or [Haacke et al., 1999, p. 749ff]). With this way of interpreting matter in magneto-statics, the fields resulting from arbitrary susceptibility distributions can be interpreted in a much simpler way, than solving the exact problem analytically at nuclear level. Parts of the derivations shown below can be found in Cheng et al. [2009]. Since the underlying theory is only briefly outlined within the scope this thesis, the latter reference as well as Lorentz [1916] are suggested for the interested reader.

Jackson [1999] as well as Marques and Bowtell [2005] postulate the following relation between magnetisation, $\vec{\mathbf{M}}(\vec{\mathbf{r}})$, and resulting flux density, $\Delta\vec{\mathbf{B}}(\vec{\mathbf{r}})$:

$$\Delta\vec{\mathbf{B}}(\vec{\mathbf{r}}) = \frac{\mu_0}{4\pi} \int_{V'} \left(\frac{3\vec{\mathbf{M}}(\vec{\mathbf{r}}') \cdot (\vec{\mathbf{r}} - \vec{\mathbf{r}}')}{|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|^5} \cdot (\vec{\mathbf{r}} - \vec{\mathbf{r}}') - \frac{\vec{\mathbf{M}}(\vec{\mathbf{r}}')}{|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|^3} \right) d^3r' \quad (7.7)$$

As described in Chpt. 2, in MRI experiments, only the static field along the principal field axis – the z axis – is considered. Further, the magnetisation, $\vec{\mathbf{M}}$, will be approximated to also occur only along z axis: $\vec{\mathbf{M}} = M_z \hat{\mathbf{z}}$ (as in linear media). Since Eq. 7.7 does not contain vector products, the Cartesian components can be observed individually, leading to the following expression for the field shifts induced:

$$\Delta B_z(\vec{\mathbf{r}}) = \frac{\mu_0}{4\pi} \int_{V'} \left(\frac{3M_z(\vec{\mathbf{r}}') \hat{\mathbf{z}} \cdot (\vec{\mathbf{r}} - \vec{\mathbf{r}}')}{|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|^5} \cdot (\vec{\mathbf{r}} - \vec{\mathbf{r}}') \cdot \hat{\mathbf{z}} - \frac{M_z(\vec{\mathbf{r}}')}{|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|^3} \right) d^3r' \quad (7.8)$$

$$= \frac{\mu_0}{4\pi} \int_{V'} \left(\frac{3M_z(\vec{\mathbf{r}}') \cdot (z - z')^2}{|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|^5} - \frac{M_z(\vec{\mathbf{r}}')}{|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|^3} \right) d^3r' \quad (7.9)$$

Proposing the Green function:

$$d(\vec{\mathbf{r}}) = \frac{1}{4\pi} \frac{3z^2 - r^2}{r^5} \quad (7.10)$$

$$= \frac{1}{4\pi} \frac{3\cos^2\vartheta - 1}{r^3} \quad (7.11)$$

where $\vartheta = \cos^{-1}(z/r)$, Eq. 7.9 can be rewritten into:

$$\Delta B_z(\vec{\mathbf{r}}) = \mu_0 \int_{V'} M_z(\vec{\mathbf{r}}') d(\vec{\mathbf{r}} - \vec{\mathbf{r}}') d^3r' \quad (7.12)$$

Apparently, this is a convolution of the local magnetisation distribution, $M(\vec{\mathbf{r}})$, with the Green function, $d(\vec{\mathbf{r}})$. Inserting Eq. 7.4 into Eq. 7.12 leads to the following expression for deriving

magnetic field shifts, $b := \Delta B_z$, based on the underlying susceptibility distribution, χ :

$$b = \mu_0 H_z \cdot (\chi * d), \quad d = \frac{1}{4\pi} \frac{3 \cos^2 \vartheta - 1}{r^3}, \quad (7.13)$$

where $*$ stands for the convolution operation ($(f * g)(\vec{r}) = \int f(\vec{r}')g(\vec{r} - \vec{r}') d\vec{r}'$). The Green function, d , resembles the field of a dipole, thus it is commonly called *dipole kernel*. The induced shifts, b , apparently increase linearly with field strength, $B_0 = \mu_0 H_z$, yielding:

$$b = B_0 \cdot (\chi * d). \quad (7.14)$$

The identity introduced with Eq. 7.14 is the basis for most current susceptibility reconstruction methods. The effect of spatially varying susceptibility on the static magnetic field is illustrated in Fig. 7.1 for a spherical body. More generally, bodies of arbitrary shape that possess a susceptibility other than their surroundings generate orientation-dependent shifts of the static magnetic field according to Eq. 7.14. These shifts decrease with $1/r^3$ with distance to the generating structure.




Figure 7.1: Field shifts induced by an isotropically dimensioned spherical body of positive susceptibility with respect to its surroundings, introduced to a static magnetic field, B_0 . The susceptibility contrast (top left) is distinct, whilst the fieldmap shows orientation-dependent distortions with respect to the direction of B_0 . Units are arbitrary.

7.2.3 Susceptibility in k-space

It is advantageous to carry out the calculation of field shifts, induced by a certain susceptibility distribution, in k -space. This is mainly because the convolution operation is simplified to a point-wise multiplication, when transforming to the frequency domain [Bronstein et al., 2005, p. 752, Eq. 15.92]:

$$\mathcal{F}(\chi * d) = \mathcal{F}(\chi) \cdot \mathcal{F}(d) \quad (7.15)$$

In numerical applications the transformation of χ is simply performed by a DFT. Hence, the transformation of the dipole kernel, d , should ideally be available in an analytic form. Deville and Bernier [1979] derived an expression for the magnetisation-induced field in frequency domain.

Later, Marques and Bowtell [2005] further investigated the relation and found that the k -space representation of the dipole kernel, d , can be described by the following plain expression:

$$\mathcal{F}(d) = \frac{1}{3} - \frac{k_z^2}{\mathbf{k}^2}, \quad \mathbf{k}^2 = k_x^2 + k_y^2 + k_z^2 \quad . \quad (7.16)$$

Applying Fourier transform on Eq. 7.14 yields:

$$\mathcal{F}(b) = B_0 \cdot \mathcal{F}(\chi) \cdot \left(\frac{1}{3} - \frac{k_z^2}{\mathbf{k}^2} \right) \quad . \quad (7.17)$$

Usually, one is interested in the induced spatial domain field shifts. In rasterised data, the dipole kernel, $\mathcal{F}(d)$, can be discretised beforehand either by direct calculation in k -space or by transforming the kernel from spatial domain into frequency domain. The spatial representation of the field shifts can be determined by applying inverse Fourier transformation to the result of Eq. 7.17:

$$b = B_0 \cdot \mathcal{F}^{-1} [\mathcal{F}(\chi) \cdot \mathcal{F}(d)] \quad . \quad (7.18)$$

The Fourier transform of the dipole kernel, $\mathcal{F}(d)$, is radially constant and its orientation varies only with the angle towards the z -axis, the direction of the magnetic field. The point-wise multiplication makes computation of the equation straightforward, including only three Fourier transformations. The transformed dipole kernel will further on be denoted in a short form:

$$\hat{d} = \mathcal{F}(d) \quad (7.19)$$

7.3 Models

Neglecting externally induced static field variations, the acquired phase signal in MRI is a result of the electromagnetic properties of the tissue and their distribution. Several models were introduced describing these sources.

7.3.1 Scalar Susceptibility

The most common model for interpreting phase images of human tissue is that of *scalar susceptibility*. This is due to the fact that the structure and orientation of the components of tissue show a certain level of arbitrariness, and due to the superposition of the signal from many cells in a voxel. Thus, the model of a scalar susceptibility, as in Eq. 7.14, is still the most frequently employed model of phase contrast generation in human tissue [e.g. de Rochefort et al., 2008; Wharton et al., 2010; Bilgic et al., 2012]:

$$\varphi(\vec{r}, t) = \gamma b_\chi(\vec{r}) \cdot t \quad (7.20)$$

$$= \gamma B_0 \cdot [\chi * d](\vec{r}) \cdot t \quad (7.21)$$

7.3.2 Susceptibility and Chemical Shift

Human tissue does not consist only of water, but rather of various carbon-based molecules, such as proteins, lipids and many other compounds, causing chemical shifts as described in Section 3.3.2. A more complete model including the off-resonance generated by chemical shift contributions is expressed as follows:

$$\varphi(\vec{r}, t) = \gamma(b_\chi + b_\sigma) \cdot t \quad (7.22)$$

$$= \gamma B_0 \cdot [(\chi * d)(\vec{r}) - \sigma] \cdot t \quad (7.23)$$

[e.g. de Rochefort et al., 2008]. Beware that the effect of both influences on the phase is linearly equally dependent on the main static field strength. Although this model offers a more extensive description of phase contrast than susceptibility-only models, it is not yet commonly used for reconstructions. On the one hand chemical shifts are expected to play a minor roll in brain tissue, and on the other hand Eq. 7.23 shows, that conventional phase contrast cannot distinguish between both sources.

An indication, that chemical shifts actually influence phase contrast of the brain at ultra-high field is shown in Chpt. 11, section Section 11.7. Sommer et al. [2012] attempted the reconstruction of susceptibility and chemical shifts from phase data based on *a priori* knowledge of the observed geometry.

7.3.3 Anisotropic Susceptibility

In the last years, a few groups (e.g. Liu [2010]; Wharton and Bowtell [2012]) developed a new model for susceptibility inside the human brain. They focus on certain structures such as long range white matter fibres, exhibiting a high degree of geometrical organisation. According to the authors this organisation leads to local anisotropy of the macroscopically observed susceptibility. The anisotropy manifests itself in differently pronounced field shifts induced for different orientations of the structure towards the axis of the static magnetic field, $\vec{B}_0$, that cannot be explained only by the outer geometry of the object. Anisotropic susceptibility in these particular regions has been observed by many groups, yet it is still attributed to different sources. An anisotropic material susceptibility as in Eq. 7.6,

$$\chi_{uv} = \begin{pmatrix} \chi_{11} & \chi_{12} & \chi_{13} \\ \chi_{21} & \chi_{22} & \chi_{23} \\ \chi_{31} & \chi_{32} & \chi_{33} \end{pmatrix}, \quad (7.24)$$

is assumed. A reconstruction technique based on this model, introduced by Liu [2010], is hence called Susceptibility Tensor Imaging (STI). Another model for anisotropic susceptibility claims that the microscopic geometry of the tissue induces geometric field effects, having the macroscopic appearance of a susceptibility tensor [Wharton and Bowtell, 2012]. All models result in the same statement: certain isolated brain regions may generate phase shifts that depend on the orientation of the tissue towards the static field, suggesting that the apparent susceptibility varies with direction. The application of STI is beyond the scope of this thesis and was thus only discussed for completeness.

7.4 Reconstructing Susceptibility Distributions

All strategies for estimating susceptibility distributions in soft tissue are based on the *a priori* estimation of fieldmaps, and thus on phase imaging. Single-echo GRE measurements are not useful for this aim, since fieldmaps derived of such data may contain deviations from the true field due to unconsidered phase offsets (see Section 6.1). Fieldmaps used for susceptibility reconstruction should at least rely on double-echo, or offset-corrected phase data. Furthermore it is evident, that the estimation of susceptibility requires knowledge of the magnetic field in all spatial dimensions. Ideally, an isotropic sampling of space should be attempted. This is best achieved using an isotropically sampled, full 3D GRE sequence with slab- or non-selective excitation. Several strategies for the reconstruction of tissue magnetic susceptibility were introduced during the last years. Two main categories can be defined, single and multiple orientation measurements. Furthermore, reconstruction can be performed directly or by using a minimisation approach. All methods presented below use the model of scalar susceptibility (Section 7.3.1), and will be shown in application examples using optimised parametrisation. Parameter optimisation will be discussed later on in Section 7.6. A comprehensive discussion of the current reconstruction techniques and applications for susceptibility imaging in MRI can be found in Reichenbach [2012].

7.4.1 Direct Inversion

The most intuitive technique for estimating the magnetic susceptibility is the *direct inversion*. This approach, introduced by de Rochefort et al. [2008] and Shmueli et al. [2009], seeks to invert Eq. 7.18. This is in principle done by a simple mathematical operation:

$$\text{invert: } b_{[ijk]} = \mathcal{F}^{-1} \left[\mathcal{F}(\chi_{[ijk]}) \cdot \hat{d} \right] \quad (7.25)$$

$$\Rightarrow \chi_{[ijk]} = \mathcal{F}^{-1}(\mathcal{F}(b_{[ijk]})/\hat{d}) \quad , \quad (7.26)$$

where $\cdot$ and $/$ denote point-wise multiplication and division. Although this inversion seems straightforward, it bears a crucial pitfall. The dipole kernel, $\hat{d}$, contains conical surfaces with zero-values. These can be determined by recalling the dipole definition:

$$\hat{d} = \left(\frac{1}{3} - \frac{k_z^2}{\vec{k}^2} \right) \quad (7.27)$$

$$\text{zero surfaces: } 3k_z^2 = \vec{k}^2 \quad (7.28)$$

This implies *zero cones* at an *angle* of 54.7° (see Fig. 7.2a). When in Eq. 7.26, $\mathcal{F}(b_{[ijk]})$ is point-wise divided by the dipole kernel, *division by zero* occurs on the cone surface, and k -space voxels residing in the close proximity are amplified, leading to strong promotion of noise. This means, the inverse problem is ill-posed or under-determined, prohibiting the direct application of Eq. 7.26. Shmueli et al. [2009] circumvent this issue by masking the conical surfaces and their proximity, yielding a valid inverse kernel $\hat{d}^{-1}$:

$$\hat{d}_{[ijk]}^{-1} = \begin{cases} 1/\hat{d}_{[ijk]}, & |\hat{d}_{[ijk]}| \geq 1/t_{\text{cone}} \\ c_{\text{cone}}, & |\hat{d}_{[ijk]}| < 1/t_{\text{cone}} \end{cases} \quad (7.29)$$




Figure 7.2: Frequency domain representation of the dipole kernel, $\hat{d}$ (a), and truncated inverse dipole kernels, $\hat{d}^{-1}$, with constant value filling (b) and with zero filling (c). The left colour bar belongs only to sub-plot (a).

One possibility is, to have the truncated area zero-filled, $c_{cone} = 0$. Alternatively a fill value of $c_{cone} = \pm t_{cone}$ for the thresholded areas can be assigned. Both alternatives are illustrated in Fig. 7.2. The methods allow for a moderate estimate of the underlying susceptibility distribution, but are strongly susceptible to noise and commonly exhibit strong streak artefacts. The threshold, t_{cone} , can be optimised on the individual data. While aggressive thresholding (low t_{cone}) creates results with a smoothed, low contrast. Too weak thresholding (high t_{cone}) keeps the result falsified by streak artefacts.

Three numerical phantoms with a dimensioning of $128 \times 128 \times 128$ voxels, shown in Fig. 7.3, are designed to test and visualise the algorithms. The zero-filled and solid-filled inversion are tested on the phantoms. Fig. 7.4 and Fig. 7.5 show the according results. Both algorithms were applied with optimised parametrisation (see Section 7.6). The results illustrate the functionality of the algorithm.




Figure 7.3: Top row: Numerical susceptibility phantoms (representative central slice), Middle row: Fieldmaps simulated using Eq. 7.25 with double-sized padding and B_0 along y -axis. Bottom: Artificial Gaussian noise introduced to fieldmaps. The noise-distorted maps are used to test the algorithms.




Figure 7.4: Reconstruction using zero-filled direct inversion (lower row) compared to simulated phantoms (upper row). Artefacts of large extent and noise-induced streak artefacts are quite apparent. Yet, the phantoms are approximately correctly reconstructed.




Figure 7.5: Reconstruction using solid-filled direct inversion (lower row) compared to simulated phantoms (upper row). The result also shows streak artefacts, but long range streaks seem smoother than for zero-filling.

Advantages and disadvantages of the direct inversion are:

- ▲ Only two Fourier transforms and a point-wise division operation → Extremely high computing speed.
- ▲ Moderate qualitative result achieved at low computing cost.
- ▼ Result is naturally falsified by streak artefacts.
- ▼ Structures with high detail are usually not clearly outlined between artefacts and noise.

7.4.2 Minimisation-based Techniques

In 2010, de Rochefort et al. [2010b] presented an approach that strongly reduces the noise of the estimated susceptibility map. Rather than just inverting Eq. 7.18, the authors suggest to utilise minimisation for solving Eq. 7.14 in a forward approach.

7.4.2.1 Principle

Minimisation is performed using the $\mathbb{L}_2$ -norm of the problem formulation:

$$\min_{\chi} \|W[b - B_0(\chi * d)]\|_2^2 \quad . \quad (7.30)$$

The weighting matrix, W , allows one to restrict the minimisation to a target area such as a brain mask for an *in vivo* measurement. The algorithm used for minimisation has to find the optimal solution whilst ignoring noise, since noise induces local minima. de Rochefort et al. [2010b] suggests to use a *conjugate gradient* solver to approach this problem. Within the scope of this thesis such a solver was custom-built to fit the demands of the minimisation problem. It will be described in Section 7.5.1. The approach in Eq. 7.30 only slightly improves the determination of the underlying susceptibility, χ . In the very same publication, de Rochefort et al. [2010b] further present regularisation approaches to ensure the stability and to reduce the noise level in the solution.

7.4.2.2 Tikhonov Regularisation

A well-known method to regularise a minimisation problem, is the *Tikhonov Regularisation*. This method adds another term, containing the $\mathbb{L}_2$ -norm of the minimisation target variable, χ :


$$\min_{\chi} \|W[b - B_0(\chi * d)]\|_2^2 + \lambda \|\chi\|_2^2 \quad . \quad (7.31)$$

This term can be used to suppress noise, as it is sensible to solution values of high amplitude. The regularisation parameter, λ , controls the intensity of the regularisation. Additionally, the regularisation can also be restricted to certain areas, or enforced in regions, where low or zero susceptibility is expected. This is done by introducing a Tikhonov weighting matrix, W_T :

$$\min_{\chi} \|W[b - B_0(\chi * d)]\|_2^2 + \lambda \|W_T \chi\|_2^2 \quad . \quad (7.32)$$

Tikhonov regularisation, including the regularisation parameter, λ , and the weighting matrix, W_T , allows for restriction of the value range of data globally or in selected regions. Under the premiss of a valid fieldmap, de Rochefort et al. [2010b] propose to set W_T to 1, where susceptibility is known to be constant such as the brain exterior. The inverse of a brain mask can be used as W_T for that purpose. An application of unweighted Tikhonov regularisation on the numerical phantoms (introduced in Fig. 7.3) is shown in Fig. 7.6. Tikhonov regularisation significantly reduces the local noise level. Yet, the long-range streak artefacts are still present, but seem slightly smoother. Since no weighting matrix, W_T , is applied, the regularisation also affects the value range of the result, such that the estimated susceptibility is slightly lower than the true model.

Figure 7.6: Reconstruction using Tikhonov regularised minimisation (lower row) compared to the simulated phantoms (upper row). No specific weighting matrix, W_T , is applied. The noise level of the result is apparently lower than for the direct inversion approaches shown in Fig. 7.4 and Fig. 7.5. Some artefacts such as under- or overemphasised areas as well as soft long-range streaks remain. The overall susceptibility is lower than in the true phantom.



Advantages and disadvantages of the Tikhonov-regularised minimisation are:

- ▲ Suppression of local noise and peaks
- ▲ Useful for definition of regions, where $\chi \stackrel{!}{=} 0$
- ▼ Underestimating intensity
- ▼ Smooth streak artefacts remaining

7.4.2.3 Gradient Regularisation

Another way to regularise the minimisation problem is to add a term that evaluates the gradients of χ :

$$\min_{\chi} \|W[b - B_0(\chi * d)]\|_2^2 + \mu \|\mathbf{G}\chi\|_2^2, \quad (7.33)$$

where μ is the regularisation parameter and $\mathbf{G}$ is the gradient operator:

$$\mathbf{G} = \partial_x + \partial_y + \partial_z \quad (7.34)$$

$$\partial_x \chi[ijk] = \chi[i,j,k] - \chi[i-1,j,k] \quad (7.35)$$

$$\partial_y \chi[ijk] = \chi[i,j,k] - \chi[i,j-1,k] \quad (7.36)$$

$$\partial_z \chi[ijk] = \chi[i,j,k] - \chi[i,j,k-1] \quad (7.37)$$

As shown for Tikhonov regularisation, the term can additionally be accompanied by a gradient weighting matrix, W_G :

$$\min_{\chi} \|W[b - B_0(\chi * d)]\|_2^2 + \mu \|W_G \mathbf{G}\chi\|_2^2, \quad (7.38)$$

Fig. 7.7 shows the application of unweighted gradient regularisation. The result shows less noise and less artefacts than the Tikhonov case and the value range is close to the simulated map. Yet, a few smooth long-range streaks remain and the result is slightly blurred.




Figure 7.7: Reconstruction using gradient-regularised minimisation (without spatial priors) in the lower row compared to simulated phantoms (upper row). The noise level of the result is even lower than for the Tikhonov regularisation (Fig. 7.6). Surprisingly the long range streak artefacts are well compensated and only slight irritations remain. The local signal is clear and just suffers from a slight blurring effect.

Many groups [e.g. de Rochefort et al., 2010b; Wharton and Bowtell, 2010; Liu et al., 2012] nowadays use so-called *spatial priors* to estimate gradient weighting maps, restricting the solution

to smooth changes in regions where the spatial prior is smooth and allowing for stronger gradients where the spatial prior shows the same behaviour. Usually, the gradients of the magnitude, cropped to a certain value range, are used to estimate these prior maps.

Fig. 7.8 shows the application of spatial priors in gradient-regularised minimisation. The priors were derived from the gradient of the simulated source susceptibility map, $g_{x(y,z)} = \partial_{x(y,z)} \chi[ijk]$. The weighting matrix was calculated as a boolean array, containing 0 where gradients in the particular direction occur and 1 elsewhere to restrict variation within constant areas. This ideal spatial prior setting obviously works very well – the minimisation results differ from the simulation source only by a low local noise level. Two aspects have to be considered in actual measurements. Firstly, the spatial priors can only be estimated e.g. by using the signal magnitude, being subject to noise and distortions itself. And secondly, all structures not visible in the signal magnitude are automatically suppressed.




Figure 7.8: *Reconstruction using gradient-regularised minimisation without (top row) and with (bottom row) spatial priors. The spatial-prior regularisation enables a reconstruction with almost no artefacts, just containing a noise uncertainty level related to that of the fieldmaps (Fig. 7.3).*

The advantages and disadvantages of gradient-regularised minimisation can be summarised to:

- ▲ Strong noise suppression
- ▲ Promotes areas of large extent and homogeneous susceptibility contrast
- ▲ Spatial priors enable better estimation of the susceptibility within a known geometry
- ▼ Without using spatial priors, local detail gets blurred with increasing gradient weighting
- ▼ When using spatial priors, geometries not detected in magnitude are suppressed automatically

Summing up both regularisation techniques yields the following minimisation term:

$$\min_{\chi} \|W[b - B_0(\chi * d)]\|_2^2 + \lambda \|W_T \chi\|_2^2 + \mu \|W_G \mathbf{G} \chi\|_2^2 . \quad (7.39)$$

Recalling the Fourier domain relations, as introduced in Eq. 7.18, and changing slightly the no-

tation, Eq. 7.39 can be expressed as:

$$\min_{\chi} \|W[b - C\chi]\|_2^2 + \lambda \|W_T\chi\|_2^2 + \mu \|W_G\mathbf{G}\chi\|_2^2, \quad (7.40)$$

where $C = B_0\mathcal{F}^{-1}\hat{d}\mathcal{F}$. This expression means: transforming χ to frequency domain, multiplying it with the frequency domain representation of the dipole kernel and transforming back to spatial domain. Minimisation is performed iteratively until the solution deceeds a certain residual limit, defined *a priori*. For computation, usually the factor B_0 is omitted during minimisation, leading to a χ -estimate scaled to the B_0 -fold of the true result. This factor is included by division, when minimisation is finished.

7.4.3 Explicit Regularised Inversion

Bilgic et al. [2013] recently presented a novel method to achieve the robustness of the solution provided by regularised minimisation with a single-term closed-form solution. *Bilgic* based his approach on the gradient-regularised minimisation, Eq. 7.38. He points out that the gradient operators, $G_{x,y,z}$, can be expressed in k -space as:

$$G_{x,y,z} = \mathcal{F}^{-1}E_{x,y,z}\mathcal{F} \quad (7.41)$$

$$\text{where: } E_{x,y,z}[ijk] = 1 - e^{-2\pi i k_{x,y,z}[ijk]/N_{x,y,z}}, \quad (7.42)$$

where $N_{x,y,z}$ is the dimensioning of the measurement volume in voxels. This leads to a closed-form solution for the susceptibility:


$$\chi = \mathcal{F}^{-1} \left[\hat{d} \cdot \left(\hat{d}^2 + \mu \cdot (E_x^2 + E_y^2 + E_z^2) \right)^{-1} \cdot \mathcal{F}(b) \right]. \quad (7.43)$$

In contrast to Eq. 7.26, the matrix in brackets can be inverted, since it does not contain zero cones. *Bilgic*'s approach reduces the iterative minimisation solution to simple algebraic matrix operations. The author claims, that the solution equals the final result that is approximated by the minimisation approach. Using the same argumentation, one can also include Tikhonov regularisation into the closed-form solution. The $\mathbb{L}_2$ -norm is conserved under Fourier transform. Thus, besides a potential scaling factor the regularisation parameter, λ , can simply be included into Eq. 7.43:

$$\chi = \mathcal{F}^{-1} \left[\hat{d} \cdot \left(\hat{d}^2 + \lambda + \mu \cdot (E_x^2 + E_y^2 + E_z^2) \right)^{-1} \cdot \mathcal{F}(b) \right]. \quad (7.44)$$

Fig. 7.9 shows a comparison between the gradient-regularised minimisation and the same parametrisation in *Bilgic*'s approach. The result differences are marginal, which supports the proposed method and the correct implementation. Nevertheless, it was observed that for problem sets, in which the source field shifts and the target area for the underlying susceptibility distribution are not identical, the method tends to fail. It is thus advisable to make use of the method only for direct reconstructions in a common source and target area or mask. The publication of *Bilgic* is quite recent, so the presented technique is only included in a few of the susceptibility applications shown within this work.

Figure 7.9: Comparing the minimisation approach (upper row) with Bilgic's approach (mid row). Both are applied with the same optimised parametrisation of gradient regularisation, but without spatial priors. The bottom row shows the difference between the mid and top row in ten-fold amplification of the greyscale colour map. The variations of the difference map is on a low relative scale.




7.4.4 Multiple Orientation Measurements

The pitfalls implied by the ill-posedness of the inversion problem affect all reconstruction techniques presented above. The algorithms differ in their ability to cope with the noise-induced stripe artefacts, arising from the underdetermined problem. Liu et al. [2009] proposed a method to circumvent these pitfalls by collecting the information missing for inversion rather than trying to improve on the lack of knowledge. The problem with direct inversion arises with the zero cones (Eq. 7.28), rendering the division of $\mathcal{F}(b)$ by the dipole kernel impossible. The authors present an approach relying on multiple measurements, using each time a different orientation of the measured object with respect to the static magnetic field axis. The dipole kernel rotated over the angle θ_x about the x axis can be described as:

$$\hat{d}_{\theta_x} = \left(\frac{1}{3} - \frac{(k_z \cdot \cos \theta_x + k_y \cdot \sin \theta_x)^2}{\vec{k}^2} \right) \quad (7.45)$$

If one proposes two measurements, for $\theta_{x,1} = -30^\circ$ and for $\theta_{x,2} = +30^\circ$, the resulting kernels and the coverage of the result space have the form illustrated in Fig. 7.10.

Figure 7.10: Upper row: cross section of the truncated dipole kernel in k -space viewed in the y - z plane ($t_{\text{cone}} = 10$) for (a) $\theta_x = -30^\circ$ and (b) $+30^\circ$. Lower row: Number of truncated ($= 0$) points along x direction for (c) $\theta_x = -30^\circ$, (d) $+30^\circ$ and (e) the number of points not covered by at least one of the two kernels along x . The range of the lower row colour bar reaches from 0 to half of the volume width (voxel count in x direction).



The truncated areas around the cones are mutually compensated in most parts. Just in areas where the two kernels show intersections of the truncation, zero kernel values remain. Using more

than two measurements can compensate for these regions of intersection, if the orientation angles are chosen wisely, as Liu et al. [2009] state. The relation between field shifts and susceptibility can be expressed as:

$$\mathcal{F}^{-1} \left[\begin{pmatrix} \hat{d}_{\theta_{x,1}} \\ \hat{d}_{\theta_{x,2}} \\ \vdots \\ \hat{d}_{\theta_{x,n}} \end{pmatrix} \cdot \mathcal{F}(\chi) \right] = \begin{pmatrix} b_{\theta_{x,1}} \\ b_{\theta_{x,2}} \\ \vdots \\ b_{\theta_{x,n}} \end{pmatrix} . \quad (7.46)$$

This expression can be inverted – depending on the choice of the $\theta_{x,i}$ – either with only little truncation, or if the measurement directions cover all voxels, without any truncation. Liu et al. [2009] further show that the smallest possible set of measurements covering all voxels in frequency domain includes three measurements at -60° , 0° and $+60^\circ$ (see Fig. 7.11). However, this angular set-up is not really applicable in measurements, since a tilting of 60° does not cover the natural movement range of a human neck inside an RF coil.




Figure 7.11: *Three-angle approach. Upper row: cross section of the truncated dipole kernel in k -space viewed in the y - z plane ($t_{\text{cone}} = 10$) for (a) $\theta_x = -60^\circ$ (b) 0° and (c) 60° . Lower row: Number of truncated ($= 0$) points along x direction for (d) $\theta_x = -60^\circ$, (e) 0° , (f) 60° and (g) the number of points not covered by at least one of the three kernels along x . The range of the lower row colour bar reaches from 0 to half of the volume width (voxel count in x direction). No uncovered voxels remain in this measurement approach.*

The method – called Calculation Of Susceptibility through Multiple Orientation Sampling (COSMOS) – was further pursued and enhanced e.g. by Wharton et al. [2010], who showed that rotation about two different axis, but with smaller angles, also allows for robust direct inversions. Fig. 7.15 shows the implemented algorithm applied on the simulated synthetic data and using the Tikhonov-regularised minimisation. Already the $\pm 20^\circ$ angle set allows for a good approximation of the true susceptibility map. The recommended set with -60° , 0° , 60° works almost ideally. An alternative set uses $\pm 20^\circ$ tilting about the x -axis and $+20^\circ$ tilting about the y -axis. The configurations proposed in the literature were simulated – the field sets are shown in Fig. 7.12, Fig. 7.13 and Fig. 7.14 and the reconstruction result is shown in Fig. 7.15.

Figure 7.12: Fieldmaps for COSMOS angle set 1, using -20° and 20° rotation about x-axis; 0° is aligned along vertical image axis.




Figure 7.13: Fieldmaps for COSMOS angle set 2, using -60° , 0° and 60° rotation about x-axis; 0° is aligned along vertical image axis.




Figure 7.14: Fieldmaps for COSMOS angle set 3, using -20° and 20° rotation about x-axis and 20° about y-axis (in image plane); 0° is aligned along vertical image axis





Figure 7.15: Results of the COSMOS acquisition schemes, computed with unweighted, Tikhonov-regularised minimisation (top row) – from left to right: Simulated map, single angle reconstruction (0°), angle set 1, 2 and 3. The lower row shows the difference to the simulated map.

7.5 Implementation

7.5.1 Conjugate Gradients

The minimisation problem presented in Eq. 7.40 has to be solved in an iterative manner. Focussing on the actual minimisation objective function and ignoring the regularisation terms for now, the problem has the general form:

$$\min_{\chi} \|A\chi - \beta\|_2^2, \quad (7.47)$$

with the linear function

$$A = WC = W(\mathcal{F}^{-1}\hat{d}\mathcal{F}) \quad (7.48)$$

the parametrisation, χ , and the target vector,

$$\beta = Wb \quad (7.49)$$

All operations involved into the dipole convolution – here denoted by A – are linear. This is obvious, as the point-wise multiplication with the dipole kernel is a linear transformation as well as the DFT (see definition in Section 2.2). Due to this linearity, Eq. 7.47 can be written in a form, where all voxels of the observed volume are arranged in a single column vector:

$$\min_{\vec{\chi}} = \|A\vec{\chi} - \vec{\beta}\|_2^2 . \quad (7.50)$$

This is just an alternative notation, making the following descriptions easier. The dipole kernel matrix, $\hat{d}$, as well as the weighting matrix, W , contain only diagonal elements, while the Fourier transforms, $\mathcal{F}$ and $\mathcal{F}^{-1}$ are more sophisticated functions, relating all voxels.

A prominent technique in approaching such minimisation problems is the method of Conjugate Gradients (CG). This approach was suggested for the present problem by de Rochefort et al. [2010b]. Shewchuk [1994] published a comprehensive document, explaining the theoretical basis as well as applications of the CG method. Inspired by this publication, the method shall be introduced shortly in association with the presented problem. The reference given above is recommended for the interested reader. The $\mathbb{L}_2$ -norm expression in Eq. 7.47 is best approached by describing its derivative:

$$A^T A \vec{\chi} = A^T \vec{\beta} , \quad (7.51)$$

incorporating a positive-(semi-)definite, symmetric objective function, $A^T A$. Since the presented problem includes Fourier transformation, it certainly operates on complex values. Thus, the simple transpose will be replaced by complex conjugated transpose, indicated by the symbol $\dagger$.

This equation type ($A\vec{\chi} = \vec{b}$) is the typical starting point for a CG process. CG is related to minimisation with the *method of the steepest descend*. The latter technique tries to find the shortest way to the global minimum by following the direction of the steepest gradient of the minimisation function. CG utilises a more evolved descending technique. The iterative minimisation directions are kept mutually orthogonal to all previously applied descends.¹ This is achieved by a *Gram-Schmidt* based procedure, applied simultaneously with minimisation. The function “memorises” the previously applied directions, excluding them from future descends.

The iterative CG is best described using the involved operations. Beginning with a starting parametrisation, the descending direction, $\vec{d}$, is iteratively adapted to the steepest direction, preserving orthonormality to the previous directions. The step factor, α , determines how far along $\vec{d}$ descending occurs. The process sets the following starting parameters – the initial descending direction, $\vec{d}_0$, and the initial minimisation residuum, $\vec{r}_0$:

$$\vec{d}_0 := \vec{r}_0 := \beta - A\vec{\chi} \quad (7.52)$$

For the i -th step, α_i is determined by:

$$\alpha_i = \frac{\vec{r}_i^\dagger \vec{r}_i}{\vec{d}_i^\dagger A \vec{d}_i} \quad (7.53)$$

¹Orthogonality here means A -orthogonality – two parametrisations $\vec{d}_i$ and $\vec{d}_j$ are A -orthogonal, if $\vec{d}_i^\dagger A \vec{d}_j = 0$.

Proceeding from the i -th to the $(i + 1)$ -th step is performed by assigning the following values:

$$\text{solution guess: } \vec{\chi}_{i+1} = \vec{\chi}_i + \alpha_i \vec{d}_i \quad (7.54)$$

$$\text{residuum: } \vec{r}_{i+1} = \vec{r}_i - \alpha_i A \vec{d}_i \quad (7.55)$$

$$\text{descending direction factor: } \epsilon_{i+1} = \frac{\vec{r}_{i+1}^\dagger \vec{r}_{i+1}}{\vec{r}_i^\dagger \vec{r}_i} \quad (7.56)$$

$$\text{new descending direction: } \vec{d}_{i+1} = \vec{r}_{i+1} + \epsilon_{i+1} \vec{d}_i \quad (7.57)$$

This roughly describes the iterative CG process. The minimisation process can be stopped by a predefined maximum iteration count, $i_{\max}$, a residual limit, $r_{\min} = \sqrt{\vec{r}^\dagger \vec{r}}$, or a result change limit, $(\Delta\chi)_{\min} = \sqrt{(\vec{\chi}_i - \vec{\chi}_{i-1})^\dagger (\vec{\chi}_i - \vec{\chi}_{i-1})}$. The algorithm developed within the scope of this thesis further features a residual reset, allowing the user to define a certain iteration count, i_{reset} , after which residual and directional vectors, as well as the coefficients, α and ϵ , are reset to their initial values whilst keeping the current minimisation result. This allows minimisation inaccuracies to be compensated for by abandoning all previous minimisation directions and restarting the method. The parameters, $i_{\max}$, $r_{\min}$, $(\Delta\chi)_{\min}$ and i_{reset} , are summarised as CGparams in the function description in Appendix Section C.

The actual transformation matrix, $A^\dagger A$, shall now be described for the problem, Eq. 7.51. Inserting Eq. 7.48 and Eq. 7.49 into the derivative of the unregularised convolution equation leads to:

$$C^\dagger W^\dagger W C \vec{\chi} = C^\dagger W^\dagger W \vec{b} \quad (7.58)$$

Including the Tikhonov regularisation term, $\lambda \|W_T \vec{\chi}\|_2^2$, with the derivative, $\lambda W_T^\dagger W_T \vec{\chi}$, one gets:

$$(C^\dagger W^\dagger W C + \lambda W_T^\dagger W_T) \vec{\chi} = C^\dagger W^\dagger W \vec{b} \quad (7.59)$$

Similarly, for gradient regularisation the derivative, $\mu \sum G_{x(y,z)}^\dagger W_{G_{x(y,z)}}^\dagger W_{G_{x(y,z)}} G_{x(y,z)} \vec{\chi}$ leads to the combined minimisation equation:

$$\begin{aligned} (C^\dagger W^\dagger W C + \lambda W_T^\dagger W_T \\ + \mu \sum G_{x(y,z)}^\dagger W_{G_{x(y,z)}}^\dagger W_{G_{x(y,z)}} G_{x(y,z)}) \vec{\chi} = C^\dagger W^\dagger W \vec{b} \end{aligned} \quad (7.60)$$

$$A^\dagger A \vec{\chi} = A^\dagger \vec{\beta} \quad (7.61)$$

Since all weighting matrices contain only diagonal elements, W , W_T and $W_{G_{x(y,z)}}$ are identical to their conjugate transpose. Further,

$$C = \mathcal{F}^{-1} \hat{d} \mathcal{F} \quad (7.62)$$

$$\begin{aligned} C^\dagger &= (\mathcal{F}^{-1} \hat{d} \mathcal{F})^\dagger = (\mathcal{F})^\dagger (\hat{d})^\dagger (\mathcal{F}^{-1})^\dagger \\ &= \mathcal{F}^{-1} \hat{d} \mathcal{F} = C \end{aligned} \quad (7.63)$$

The gradient operators, $G_{x(y,z)}$, contain only the values 1 and -1 on two diagonal axis, in a direct neighbourhood or with one element in-between – depending on the definition. Assuming a symmetric gradient matrix, complex transposing just changes their polarity.

The entire algorithm described above was newly implemented in MATLAB[®] (The MathWorks Inc., Natick, USA) in order to be best adaptable to the susceptibility reconstruction problem. The CG solver offered by the MATLAB[®] minimisation toolbox could not be configured to deal with the array sizes involved and to operate on the implicit minimisation functions in a satisfying way. The implementation includes the objective function and constant described in Eq. 7.60.

7.5.2 Field Susceptibility Estimator

The developed evaluation framework is called FSE. It is written in MATLAB[®] (The MathWorks Inc., Natick, USA) and consists of two major components: The complete command-line function set and a simplistic Graphical User Interface (GUI). The function set splits into fundamental functions, direct inversion functions and the minimisation set. Fig. 7.16 outlines the FSE framework. A description of the individual function can be found in Appendix C.




Figure 7.16: Scheme of the FSE framework. The functions on the left hand side are interface functions representing the user interface. The functions on the right hand side embody the core functionalities, called by the interface functions.

7.6 Optimising the Parameters

Implementation of the methods for susceptibility reconstruction includes investigation of the adjustable parameters involved. Fig. 7.17 shows examples for three different truncation parametrisation for the zero-filled direct inversion method. Fig. 7.18 shows three regularisation parametrisations for Tikhonov- and gradient-based minimisation. It is apparent, that the choice of the parametrisation is essential for successful reconstruction. Thus, the robustness of the reconstruction results was investigated with regard to parametrisation and noise level for the direct inversion method as well as for the minimisation method. In both cases, optimisation was performed on a numerical phantom.

Figure 7.17: Different truncation parametrisations for zero-filled direct inversion. From left to right: $t_{\text{cone}} = [50, 10, 4]$. Only the image in the middle shows a compromise of noise level and contrast.





Figure 7.18: Different regularisation parameters, $\lambda = [0.001, 0.05, 0.5]$ for Tikhonov-regularised reconstruction (top row) and $\mu = [0.001, 0.03, 0.3]$ for gradient-regularised reconstruction (lower row). Low regularisation leads to strong noise amplification, whilst high regularisation leads to a loss in contrast in the Tikhonov case and to blurring in the gradient case.



7.6.1 Direct inversion

As previously described, direct inversion can be performed using zero-filling or solid-filling of the truncated areas. In both cases, the adjustable parameter is the truncation threshold, t_{cone} . A guideline for picking the optimal truncation parameter to achieve good result accuracy with respect to the noise level of the source map shall be determined in the following.

7.6.1.1 Material and Methods

Basis for this optimisation was a numerical susceptibility phantom of size $100 \times 100 \times 100$ voxels. The phantom consists of 250 spherical bodies with random position, a radius between 1 and 10 voxels and random pseudo-susceptibility values between -1 and 1 . The volume is initially zero-filled and the spherical bodies are added to the map. Intersections are intended rather than avoided in order to generate a complex structure serving as the source susceptibility distribution, χ_{true} . The according fieldmap, b , was simulated, inserting χ_{true} into Eq. 7.25 (with a dipole

kernel, discretised in k -space). Double-size zero-padding of the Fourier transform was used to avoid folding artefacts. A mask for the interior, m , was defined:

$$m[ijk] = \begin{cases} \text{true}, & \chi_{\text{true}}[ijk] \neq 0 \\ \text{false}, & \chi_{\text{true}}[ijk] = 0 \end{cases}, \quad (7.64)$$

marking the areas filled with spherical bodies. The complementary mask, $m_{\text{ref}}[ijk] = \text{not}(m[ijk])$ defines the exterior regions. Noise was finally added to the fieldmap, $b[ijk]$, using a zero-centred Gaussian noise distribution with a standard deviation chosen relative to the mean absolute field value:

$$b_{\text{sim}}[ijk] = b[ijk] + \delta(\sigma_{\text{noise}})[ijk], \quad \sigma_{\text{noise}} = (p/100) \cdot \overline{|b|}, \quad (7.65)$$

where p defines a percentage of $\overline{|b|}$. Fig. 7.19 shows the characteristics of the simulated susceptibility phantom, the generated fieldmaps and the mask. The noise percentages, $p = [0, 5, 10, \dots, 90, 95, 100]$, were simulated and analysed with zero-filled as well as with solid-filled direct inversion kernels using truncation values of $t_{\text{cone}} = 10^\nu$, with $\nu = [0, 0.1, 0.2, \dots, 2.8, 2.9, 3.0]$.




Figure 7.19: Simulation of arbitrary susceptibility distribution consisting of spherical bodies (A), the fieldmap derived by convolution with the dipole kernel without noise (B) and with 100% noise (C). The interior mask, m , is visualised in (D). The static field direction is horizontal.

In order to evaluate the result validity, two values were determined: The resulting relative susceptibility error and the relative SNR. The relative error was determined as the $\mathbb{L}_2$ -norm difference of the reconstructed map, χ_{rec} , to the simulated one, χ_{true} , normalised by the voxel count, 100^3 :

$$e_\chi = \frac{1}{100^3} \cdot \|\chi_{\text{rec}} - \chi_{\text{true}}\|_2. \quad (7.66)$$

The relative SNR in the reconstruction was defined as the ratio between the average absolute susceptibility inside $m[ijk]$ and the average absolute susceptibility in the exterior, $m_{\text{ref}}[ijk]$:

$$\text{rel. SNR} = \frac{\frac{1}{n_m} \cdot \sum_{ijk} |m[ijk]\chi[ijk]|}{\frac{1}{n_{m_{\text{ref}}}} \cdot \sum_{ijk} |m_{\text{ref}}[ijk]\chi[ijk]|}, \quad (7.67)$$

with n_m and $n_{m_{\text{ref}}}$ being the number of voxels inside of the masks.




Figure 7.20: Results for the optimisation of zero-filled direct inversion. The surface plots show the relation between noise level and truncation value. The white line in the surface plots serves as an eye guide to the minimal or maximal values. The third graph shows the relation between SNR, result error and noise level. The last plot combines the minima or maxima found for relative error and relative SNR.

7.6.1.2 Results and Discussion

The resulting values for SNR and susceptibility error are arranged in surface plots in Fig. 7.20 and Fig. 7.21. The white line represents an eye guide marking the respective minimum (result error) or maximum (SNR) value. The result error shows a clear tendency for the optimal truncation level with increasing noise level. The SNR surface plot shows the same tendency, but then a second maximum in the top right corner. This rise originates in strong noise amplification within the inner areas and can easily be excluded from the evaluation by considering the high error level in that particular area, visible in the error plot. Even when sticking to the ideal truncation, the relative SNR becomes relatively low, when the noise level rises above 30%. For the zero-filled inversion the result error suggests a truncation of $t_{\text{cone}} = 10^{0.6} \approx 4$ for high noise levels, while the SNR-related curve suggests, $10^1 = 10$. In the solid-filled case, the result error suggests $t_{\text{cone}} = 10^{0.4} \approx 2.5$ and the SNR implies $t_{\text{cone}} = 10^{0.7} \approx 5.0$. Naturally, for low noise levels, the truncation can be




Figure 7.21: Results for the optimisation of solid-filled direct inversion. Graphs have the same specifications as Fig. 7.20.

drastically reduced. The SNR is comparable in both methods.

In MRI experiments, the fieldmaps used for susceptibility analysis are typically subject to a significant noise level. The determined graph allows one to make an optimal *a priori* choice for the truncation of the applied inversion algorithm. Nevertheless, this can only serve as a guideline. Phase noise experienced in measurements is usually not just equally or Gaussian-distributed – unexpected falsifying gradients or local spikes can be observed, so the actual truncation has to be adapted to the individual measurement.

7.6.2 Minimisation and Explicit Regularised Inversion

The regularisation parameters for the minimisation approaches also require optimisation. Customisation of the weighting maps will not be considered at this point.

7.6.2.1 Material and Methods

The numerical phantom introduced in the previous section is reused here to simulate susceptibility-induced field shifts. The same noise levels from 0 – 100% of the mean absolute field are simulated. Instead of the direct inversion, the Tikhonov-regularised minimisation and the gradient-regularised minimisation techniques are applied here using solid-filled weighting matrices, W_T and $W_{Gx(.y,z)}$. A wide range of regularisation parameters is tested using 20 iterations:

$$\lambda = 10^{v_\lambda}, \quad v_\lambda = [-5, -4.9, \dots, -0.1, 0.0] \quad (7.68)$$

$$\mu = 10^{v_\mu}, \quad v_\mu = [-5, -4.9, \dots, -0.1, 0.0] \quad (7.69)$$

$$(7.70)$$

The interior and exterior masks, m and m_{ref} , are defined as before. SNR and susceptibility error are also evaluated as described above.

7.6.2.2 Results and Discussion

The resulting values for SNR and susceptibility error are arranged in surface plots in Fig. 7.22 and Fig. 7.23. The white line indicates an eye guide for the minimum or maximum achievable value. A clear minimum or respectively maximum graph can be determined for both gradient and Tikhonov regularisation at all noise levels. Unlike the direct inversion, the SNR graph shows a distinct maximum and no side-maxima, which is due to the strong restriction of result values for high regularisation parameters. For increasing noise contribution ($\geq 40\%$) the result errors of both, Tikhonov and gradient regularisation, suggest optimal regularisation values in the order of 10^{-2} . This roughly agrees with results from e.g. Bilgic et al. [2012]. The optimum for Tikhonov regularisation seems to be slightly higher than for gradient regularisation. Furthermore, in the Tikhonov approach the optimal parameterisations regarding SNR and result error diverge drastically above 50% noise level, whereas for gradient regularisation the optimum graphs are only slightly separated. While the minimisation of the Tikhonov-regularisation exhibits good SNR only at low noise levels ($\leq 25\%$), the gradient-regularisation seems to offer moderate SNR even at higher noise levels. The confinements of the area of maximal SNR are much more distinct for gradient-regularisation than for the Tikhonov approach. The relative error is also lower in the overall map for gradient regularisation. Gradient regularisation seems thus to be slightly preferable compared to the Tikhonov regularisation technique. Again, the reader may be reminded that actual MRI phase measurements are not only falsified by Gaussian noise. When outliers or spikes have to be suppressed, both gradient and Tikhonov regularisation can show good results. Also combined approaches can be applied to the fieldmap under investigation. In fact, most applications shown in this thesis use a combined minimisation including Tikhonov and gradient regularisation, since this configuration produced the most robust results on measured data.

7.6.3 Dipole Kernel and Fourier Transform

The deconvolution methods – regardless if choosing direct inversion or minimisation – base on a deconvolution in the discretised Fourier domain. This fact makes it necessary to include sampling considerations. In general, objects exhibiting susceptibility different from their surroundings




Figure 7.22: Results for the optimisation of Tikhonov-regularised minimisation. The surface plots show the relation between noise level and regularisation parameter, λ . The white line in the surface plots serves as an eye guide to the minimal or maximal values. The third graph shows the relation between SNR, result error and noise level. The last plot combines the minima or maxima found for relative error and relative SNR.

generate a field with an extent beyond the acquired FOV. For small objects on the length scale of several voxels or even slightly above, being positioned in a sufficient distance from the FOV boundaries, this effect can be neglected due to the small amplitude of the dipole field at a large distance.² Albeit, larger objects such as the human brain, projected on an array with only slightly larger dimensioning than the brain itself, will certainly cause significant field shifts outside the FOV. These distortions are missing information when trying to reconstruct the causative susceptibility. Additionally, the convolution of an object of this extend with the dipole field creates fold-ins in the result area. Thus, zero-padding is inevitable, if one tries to reconstruct bigger objects with partial field information. A thumb rule for reconstruction is that the extent of the object under investigation should not significantly exceed half the FOV range. Further, Bouwman

²The dipole field decays by $1/r^3$ with distance.




Figure 7.23: Results for the optimisation of gradient-regularised minimisation with parameter μ . Graphs have the same specifications as Fig. 7.22.

and Bakker [2012] recently published a promising approach on avoiding sampling effects without the need of zero-padding.

Another point of concern is the discretised dipole kernel. If discretisation takes place in Fourier domain, two conditions have to be fulfilled. Firstly, the aforementioned sampling considerations are fulfilled and secondly, no rotation of the kernel is applied. A rotated discretised kernel in k -space is not circularly continuous. Strong discontinuities are created between (e.g. the front and back) array surfaces.³ Discontinuities in k -space lead to strong ringing artefacts after Fourier transform to spatial domain. This problem arises due to the finiteness of the kernel. The highest frequencies represent voxel-by-voxel changes. These changes can only be observed along the discretised positions of the grid, and are not continuously resolvable in all directions. Thus, the true discretised k -space representation should exhibit rotation symmetry with a slight bending at the outer edges. Using a spatial representation of the dipole kernel and transforming it to Fourier

³In circular view.

domain leads to smooth surfaces and corners. Fig. 7.24 illustrates the effects of insufficient padding and rotation of the discretised kernel.




Figure 7.24: Effects of inappropriate dipole kernel choice. A) The top row shows the absolute value of the frequency domain representation of the kernels, the bottom shows the simulated field shifts using the susceptibility phantom introduced earlier. “freq.” stands for kernel definition in frequency domain, “spat.” for the Fourier-transformed spatial kernel. The discretised rotated kernel generates significant ringing artefacts, whereas the transformed spatial kernel does not. In B) the left fieldmap was simulated using a kernel size equal to the map size, in the right image zero padding to double image size was performed before convolution took place. The field characteristics differ significantly.

7.7 Value Range of Magnetic Susceptibility in the Human Body

In order to investigate magnetic susceptibility distributions in human tissue, one should be aware of the according value range. A list of reference values for the magnetic susceptibility of various elements and compounds can be found in [Schenck, 1996, p. 824]. While air is slightly paramagnetic, with $\chi_{air} \approx 0.36 \cdot 10^{-6}$, water is distinctly diamagnetic with $\chi_{water} \approx -9.05 \cdot 10^{-6}$ at 37°C (body temperature). Experiments performed in the context of this thesis showed, that water susceptibility may vary depending on the contained solvents.⁴ The value indicated by Schenck is interpreted as the “pure water” reference within this thesis. Human tissue shows a range of roughly -11 to $-7 \cdot 10^{-6}$ [Schenck, 1996]. Deoxygenated blood has a susceptibility of $-7.9 \cdot 10^{-6}$. Oxygenated and deoxygenated blood exhibit a susceptibility differences of about $\Delta\chi_{blood} = \chi_{blood,deox} - \chi_{blood,ox} \approx 0.28 \cdot 10^{-6}$ [Jain et al., 2012], leading to the Blood-Oxygen Level Dependence (BOLD) effect that is the basis for functional MRI (fMRI) [Ogawa et al., 1990]. High susceptibility contrast is thus induced by air cavities and blood vessels inside the head. Iron deposition in certain structures can generate a strong susceptibility contrast, even for low concentrations. This is caused by the high susceptibility value of the deposition form of iron in the human body, the *ferritin* molecule: $5.4 \cdot 10^{-4}$ [Schenck, 1996]. Such local concentrations are found in structures such as the *substantia nigra* (SN) and the *red nuclei* (RN) [Shmueli et al., 2009; Schäfer et al., 2009]. When addressing susceptibility values inside the human body, often times the unit *ppm* is used rather than the notation following the international system of units, 10^{-6} .

⁴The susceptibility of distilled water differs slightly from that of tap water.

7.8 Application Example

In order to demonstrate the implemented algorithms and their reliability, a susceptibility reconstruction example based on measured *post mortem* phase data will be presented. The data were acquired at 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany) with a single-channel transmit (Tx) body and a 12 ch Rx array coil. The measurement parametrisation included:

Sequence Parameters					
seq	GRE	exc	slab	Tx/Rx	Body / 12 ch array
matrix	$288 \times 384 \times 176$	FOV	$[172.8 \times 230.4 \times 105.6]$ mm	res.	$[0.6 \times 0.6 \times 0.6]$ mm
TR	46 ms	α	20°	RF-spoil	on
TE	[9.84, 22.14] ms			BW	110 Hz/Px

The uncombined data of the four physically separated channels were saved. Since phase offsets at excitation time are usually small at 3 T, the data were recombined using the SMW recombination (discussed later, in Section 10.5). A brain mask, m , was created with manual segmentation using ITK Snap [Yushkevich et al., 2006]. Background fields were corrected using MUBAFIRE and MUBAFIRE Local (also introduced later, in Chpt. 9) in standard configuration. For susceptibility reconstruction the following parametrisation with mixed regularisation including Tikhonov and gradient weighting was used:

FSE Parameters					
regmode	3	W	brain mask (m)	W_T	$1 + 10 \cdot m$
λ	0.035	μ	0.03	B_0	3.0 T
$W_{G_{x(y,z)}}$	m	iterations	30		

Fig. 7.25 shows the local fieldmap that was basis for the susceptibility reconstruction, and the resulting susceptibility distribution. The combination of Tikhonov and gradient regularisation suppresses stripe artefacts sufficiently. Yet, the local detail is relatively high, except for a certain noise level that is typical for a 0.6 mm measurement at 3 T. The zoomed regions in the yellow boxes show how effectively the reconstruction resolves spatial distortions into their underlying susceptibility distribution. The contrast in the cortical structures is mainly generated by the slight susceptibility difference between grey and white matter. The field distortions are best visible when the direction of the static field is parallel to the illustrated slice – as in the yellow boxes. Apparently, the true structure is hidden in the fieldmap. The reconstruction of susceptibility shows well outlined curvature of the gyri and a rich contrast between the grey matter structure and the white matter (appearing darker in the image). It shall also be remarked at this point that the illustrated reconstruction was obtained only from the field distortion map, without employing any kind of spatial prior (see Section 7.4.2.3).




Figure 7.25: The fieldmap shown in the orthogonal slices on the left hand side is used for susceptibility reconstruction. The susceptibility maps of the same slices are shown in the right half of the image. Additionally, a cube plot of the susceptibility map is shown in the middle. The yellow magnification boxes show, how the distorted field generated by grey/white matter susceptibility contrast is successfully employed to derive the underlying susceptibility geometry.

Removing Background Fields

Phase contrast offers outstanding insights into local tissue structure as reported particularly at high field strength (e.g. Wharton and Bowtell [2010]). Unfortunately, macroscopic field distortions are several orders of magnitude stronger and hence prevent a direct inspection of phase information relevant to brain structure. As indicated in Section 3.3, these distortions originate from a number of different sources. First, the static field (B_0) is to some extent inhomogeneous at the *ppm* level with a smooth, slowly varying behaviour. In addition, the intrinsic inhomogeneous magnetic susceptibility distribution in the human head and body as well as that caused by the RF coils and the patient table generate additional medium to long range field distortions relevant to MRI measurements. In the head regions with large susceptibility variations, most notably air cavities, can induce large field changes which cannot be fully removed by the commercially available shim systems, as reported e.g. by Schenck [1996]; Li et al. [1996]; Koch et al. [2006] and in [Marques, 2004, Chpt. 6]. In pursuit of the aim of investigating local phase and susceptibility contrast, the field distortions originating from outside the ROI need to be removed whilst preserving the local contrast. Background Field Removal (BFR) has been approached using numerous strategies. Thus far, the most prominent examples are high-pass or Gaussian filters [Abduljalil et al., 2003], polynomial filters [Shmueli et al., 2009], filters utilising dipole or susceptibility distributions [de Rochefort et al., 2010a; Liu et al., 2011] and, finally, filters based on the harmonic nature of background field distortions, such as SHARP [Schweser et al., 2011b].

In this chapter, the principle of BFR will be outlined. Most common algorithms, including several of those named above, will be introduced as well as our in-house development, SPHINX [Hirsch et al., 2009]. Application examples are shown, exploiting the virtues of the different filter techniques. A fieldmap with a dimensioning of $384 \times 288 \times 176$ voxels acquired from a *post mortem* measurement at a 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany) serves as a reference for all filter applications.¹ A direct comparison between several filters will be presented in Chpt. 9, where the multi-step filter MUBAFIRE, which was developed within the scope of this thesis, will be described in detail.

¹The precise measurement parametrisation is of no importance for the demonstrations shown within this chapter.

8.1 Principle

Based on the introduction on field shifts and their sources given in the previous chapters, deviations of the static field from its average value can be categorised into three main groups: B_0 inhomogeneities (also present in the absence of the sample), field distortions induced by sources outside the field of interest (e.g. from the body or the brain) and field distortions due to the structure of the object of interest (e.g. the brain) itself. We will refer to the first two factors as ‘field distortions’, b_{dist} , of external source, and to the last one as ‘internal field’, b_{int} , describing the phase contrast within the object. In approximation the field components can be understood as superimposed, thus the observed field appears to be their sum:


$$b := \Delta B = b_{\text{dist}} + b_{\text{int}} \quad . \quad (8.1)$$

This decoupled approximation is valid, since both variables, b_{int} and b_{dist} , are at the order of *ppm* of the static field B_0 , or lower. Phase imaging and moreover susceptibility reconstruction seek to gather information about local geometrical structures. One is interested in the internal perturbations, b_{int} , rather than in the field distortions, b_{dist} . Nevertheless, this information cannot be assessed in a direct manner, but yields the identification of external influences in the first place. Therefore, it is crucial to understand and utilise the nature of externally induced shifts, in order to successfully extract b_{int} . Once an estimate, $\tilde{b}_{\text{dist}}$, is obtained, Eq. 8.1 can be solved:

$$b_{\text{int}} = b - \tilde{b}_{\text{dist}} \quad (8.2)$$

Objective: A fundamental goal of BFR is to exclude external influences without corrupting the internally generated contrast. Inaccurate removal of field variations may harm the consistency of the internal field contrast and make the outcome useless for e.g. susceptibility reconstruction as illustrated in Fig. 8.1.

Figure 8.1: *Reconstruction of the susceptibility distribution (right hand side) based on the raw fieldmap (left hand side). The outcome contains severe gradients and local detail is hardly visible. Units are arbitrary.*



8.2 Rudimentary Filters

One characteristic of especially strong and extensive B_0 distortions is their smoothly varying nature. Although this property does not apply to all externally induced shifts, it characterises a major fraction of field distortions. A composition of slowly varying distortions and highly localised contrast suggests the utilisation of a *low-pass filter* in order to estimate b_{ext} . Low-pass filtering can be achieved with different methods, such as *mean* or *median filtering* or *Gaussian* filters. For mean and median calculation, analysis of a local neighbourhood, $N_{[ijk]}$, with extent e of every

voxel, $[ijk]$, is performed, where:

$$N_e[ijk] = [l, m, n] : |l - i| \leq e, |m - j| \leq e, |n - k| \leq e \quad . \quad (8.3)$$

The mean and median maps can thus be calculated:

$$\tilde{b}_{\text{ext}} = b_{\text{mean}}[ijk] = \text{mean}(b[N_e[ijk]]), \quad \tilde{b}_{\text{ext}} = b_{\text{median}}[ijk] = \text{median}(b[N_e[ijk]]) \quad . \quad (8.4)$$

Both measures have certain applications – e.g. the median shows a slightly higher immunity against outliers, while the mean map is smoother, even for small neighbourhoods. Nevertheless, both filters employ averaging. This removes contrast of all structures with an extent equal to or higher than the applied neighbourhood size, e .




Figure 8.2: Comparison of the rudimentary filters applied on the measurement example described above. The acquired fieldmap is shown on the left. The top row shows the true field and the estimated background, while the lower row shows the subtraction result, $b - \tilde{b}_{\text{dist}}$. The median filter (2nd column) shows a contrast level comparable to mean filtering (3rd column). Mean filtering shows stronger edge artefacts, but higher inner contrast. The best contrast is offered by the GF, showing smoother edges. Filter settings: Mean and median kernel width: 13 voxel, Gaussian $\sigma = 6$ voxel.

The Gaussian filter method estimates the background field by convolving b with a Gaussian kernel with a standard deviation of several voxels width (normalisation is omitted here):

$$\tilde{b}_{\text{dist}} = b * \left[\frac{1}{\sqrt{2\pi}^3 \sigma^3} \cdot e^{-\frac{r^2}{2\sigma^2}} \right] \quad , \quad . \quad (8.5)$$

(as before, $*$ symbolises three-dimensional convolution). Since a Gaussian function keeps Gaussian form under Fourier transformation, Eq. 8.5 embodies a strongly smoothed low-pass filter. Alternatively, applying the convolution to the complex signal can circumvent the unwrapping step

and is also known as *homodyne filtering* (e.g. [Reichenbach and Haacke, 2001], [Bernstein et al., 2004, p. 552], [Rauscher et al., 2005]):

$$S_{\text{local}} = \frac{S}{S_{\text{homodyne}}} = \frac{S}{S * \exp - \frac{r^2}{2\sigma^2}} \quad (8.6)$$

This is a fast and straightforward method which has been implemented in SWI post-processing (e.g. Shmueli et al. [2009]). Using an explicit convolution with limited kernel size, a neighbourhood, $N_e[ijk]$, as described above, and exclusion of non-mask voxels within a weighting step for each individual voxel as described by Hirsch [2009] can improve the filter performance at the edges of a mask. This technique will be used within this thesis as the standard Gaussian technique for BFR and will be abbreviated GF. The Gaussian filter family is yet another option to eliminate long-range background field variations using an averaging procedure. Fig. 8.2 shows the effects of mean, median and GF, applied on an phase data set example. Fig. 8.3 demonstrates mean filtering and GF with different neighbourhood sizes or kernel widths in application on the sample data. A line plot for the different settings is shown in Fig. 8.4. Fig. 8.5 shows the distribution of the field before and after background correction in the form of width-normalised histograms, using the same parametrisation. The field distributions found in the histograms and the findings from Fig. 8.3 indicate that lower filter widths result in narrower peaks, centring the distribution around zero, but remove image contrast. Strongly broadened peaks, such as for the raw distribution, are associated with significant field inhomogeneities in the maps. The loss in contrast with decreasing filter size is well observable in Fig. 8.4. Here, the minimal filter width eliminates almost all contrast.

Optimisation: A desirable filter setting should generate high contrast (of short to medium range variations) in the resulting maps, while keeping the distribution peak as narrow as possible.

The filter width (for median and mean filter) and the standard deviation used for the GF can be adapted to the particular fieldmap. Albeit, one has to avoid extreme filter values. Very low filter widths practically delete all contrast from the image, whereas large filter kernels generate strong edge artefacts and remove background variations less efficiently.

Pro and Contra of Rudimentary Filters:

- ▲ The rudimentary filters can be applied rapidly
- ▲ Filters are useful to identify the quality of local contrast or to search for hidden local structures
- ▼ Due to the high-pass characteristics, the filters remove a wide range of field variations that potentially carry information on internal structures.
- ▼ Except for the GF, the filters cannot be adapted to a mask
- ▼ Strong edge artefacts arise at the mask surfaces




Figure 8.3: Different parameter settings for mean (top row) and GF (bottom row) filter. The lowest setting removes nearly all contrast, whereas the highest setting induces strong edge distortions (even with GF), failing to remove the background field entirely.





Figure 8.4: Line plots for the filter widths shown in Fig. 8.3. Both, mean (top) and GF (bottom) reduce the overall contrast in the image with decreasing filter size or standard deviation. For minimal width the line plot shows a flat shape with very low contrast.

8.3 Polynomial Fitting

The rudimentary filters are quite aggressive in the reduction of contrast as they apply a frequency threshold without taking the origin of field variations into account. A more discriminating approach is the use of three-dimensional polynomials – the 3D Polynomial Filter (POLF). The use of polynomials suggests, that the externally induced field shifts are of a smooth, long ranging nature.

Figure 8.5: Histograms showing the effect of the rudimentary filter operations, mean (left) and GF (right), for the neighbourhood sizes or standard deviation values illustrated in Fig. 8.3. Units of y-axis are arbitrary (but identical to the other histograms shown within this chapter). The distributions become narrower with decreasing filter width.



Here, b_{dist} is approximated as:

$$\tilde{b}_{\text{dist}} = \sum_{p,q,r}^{(p+q+r) \leq n} \alpha_{pqr} \cdot x^p y^q z^r, \quad \text{with:} \quad \min_{\alpha_{pqr}} \left[\sum_{p,q,r}^{(p+q+r) \leq n} \alpha_{pqr} \cdot x^p y^q z^r - b \right] \quad (8.7)$$

where n is the maximum order, p , q and r are the orders of the polynomials along the Cartesian axis and α_{pqr} are the according coefficients. The position index $[ijk]$ was dropped in this equation for better understanding. In a least-squares minimisation the coefficients can be determined up to a preselected polynomial order, $(p, q, r) \leq n$. It is important, to apply the algorithm only inside the masked region in order to avoid falsification of the results by noise. For the implementation of the filter, the MATLAB[®] (The MathWorks Inc., Natick, USA) function, `polyfitn`, written by D'Errico [2006], was utilised. This package enables fitting of n -dimensional polynomials to coordinate tuples such as (x_s, y_s, z_s, b_s) , where b is the field value at the particular point, and $s = 1, \dots, n$ is the running index over the included voxels. The method was adapted with a wrapper script, to work with three-dimensional masked data, such that the mask voxels are assigned to valid data tuples:

$$m[ijk]b[ijk], \quad m[ijk] = 1 \longrightarrow b[s] \quad (8.8)$$

The direct specification of spatial positions allows the filter to be applied to arbitrarily-shaped masks, whilst excluding the rest of the voxels. An application of the POLF for several fitting orders is shown in Fig. 8.6. On a normal workstation computer it is unfortunately not possible to compute three-dimensional polynomials above a certain order, since the utilised method creates matrices for every polynomial term in the size of the fit mask, leading to vast memory demands. Fig. 8.7 illustrates line plots for the calculated fitting orders and Fig. 8.8 shows the distribution of the raw and corrected fields in width-normalised histograms. Although Fig. 8.6 still shows rich contrast for a filter order of 5, the peak observed in the according histogram is quite narrow and high. Yet, the width of the peak is wider and the effective peak is lower than for the rudimentary filters. This characteristic most likely originates in the higher contrast level achieved with POLF, becoming quite evident in the line plots shown in Fig. 8.7. Although long range variations are compensated for, most of the local contrast is preserved. The filter generates oscillation artefacts at the mask edges, since the fit accuracy decreases drastically with increasing distance from the origin. For a sufficient suppression of these artefacts a high polynomial order would be necessary. Very high orders would again eliminate local contrast in an attempt to properly fit the fieldmap under investigation. However, the POLF offers an elegant approach on BFR with moderate computational cost, even though not all polynomials are natural solutions for a current-free field

distribution.




Figure 8.6: *POLF applied with different orders. The top row shows a transverse slice, while the bottom row shows a three-dimensional cube with orthogonal cross-sections through the volume. The raw field exhibits a strong offset and is thus rendered dark. The accuracy of the polynomial fit increases with the order. Periodical artefacts at the mask edges (yellow line) are quite evident even at order 5.*

Pro and Contra of POLF:

- ▲ The filter can be applied rapidly and needs no extensive calculations
- ▲ Fitting can be confined to a mask
- ▲ The obtained contrast level is richer than for the rudimentary filters
- ▼ Generates periodic artefacts along the mask surfaces due to radially decreasing fit accuracy
- ▼ Except for the 0^{th} and 1^{st} order fits and some individual term combinations, the filter does not necessarily identify fields with physical reasoning
- ▼ With the utilised algorithm no order higher than 5 or 6 is feasible for a typical MRI dataset ($\geq 128^3$ voxels) at a workstation computer due to memory demands

The correction algorithms introduced above can highlight the effect of local field variations. Nevertheless, neither Gaussian, median, mean or POLF (above 1^{st} order), do necessarily produce physically valid solutions. Especially the rudimentary filters strongly restrict the field variations.

8.4 Spherical Harmonics

A filter, developed by a former member of the institute, *Sebastian Hirsch*, aims to find a solution with more physical validity. The SPherical Harmonic INhomogeneity-suppressing eXpansion (SPHINX) [Hirsch, 2009] is based on the same physical concept as shimming (e.g. Zhao





Figure 8.7: Line plots for the polynomial orders shown in Fig. 8.6. The overall contrast is preserved in its rough shape, while long range variations decrease with increasing fit order.

Figure 8.8: Histogram showing the effect of the POLF for the orders 1 – 5, as illustrated in Fig. 8.6. Units of the y-axis are arbitrary. The field distribution gets narrower with increasing polynomial order.



et al. [2005]) and utilises a basis built from spherical harmonics to describe the field generated by outside sources, since these are solutions of the MAXWELL equations (more precisely of the LAPLACE equation [e.g. Jackson, 1983, p. 117ff]).

8.4.1 Theory

Instead of using common polynomials, SPHINX uses the real regular solid spherical harmonic (SSH) functions, $\mathbf{v}_{lm}$, as a basis for approximating the background field. The SSHs can be expressed as:

$$\mathbf{v}_{lm}(\vec{\mathbf{r}}) = N_{lm} \cdot r^l \cdot Y_{lm}(\theta, \phi) \quad , \quad (8.9)$$

where l represents the order and $m \in \mathbb{N}$ is confined to $m \in [-l, \dots, -1, 0, 1, \dots, l]$. Further,

$$Y_{lm}(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-|m|)!}{(l+|m|)!}} \cdot P_l^m(\cos \theta) \cdot e^{im\phi} \quad (8.10)$$

are the simple spherical harmonic functions [Jackson, 1999, p. 108] in spherical coordinates $\mathbf{r} = (r, \theta, \phi)$. N_{lm} depicts the normalisation coefficient. The SSHs, $\mathbf{v}_{lm}$, are naturally orthonormal. Yet, for a non-spherical domain, such as a brain mask, their orthonormality is destroyed, so their discrete representation has to be orthonormalised using a GRAM-SCHMIDT-based procedure. In the following only the valid voxels as indicated by $m[ijk]$ are considered. The indexing tuple, $[ijk]$, will be omitted for better understanding. Nevertheless, the presented equations apply to any valid voxel and internal products summarise over the valid mask.

Orthonormality is achieved employing an iterative process:

$$\hat{\mathbf{u}}_{00} = \frac{\mathbf{v}_{00}}{|\mathbf{v}_{00}|} , \quad (8.11)$$

$$\begin{aligned} \hat{\mathbf{u}}_{lm} = \mathbf{v}_{lm} & - \sum_{j=0}^{l-1} \sum_{k=-j}^j \langle \hat{\mathbf{u}}_{jk} | \mathbf{v}_{lm} \rangle \cdot \hat{\mathbf{u}}_{jk} \\ & - \sum_{k=-l}^m \langle \hat{\mathbf{u}}_{lk} | \mathbf{v}_{lm} \rangle \cdot \hat{\mathbf{u}}_{lk} . \end{aligned} \quad (8.12)$$

The modified SSHs, $\hat{\mathbf{u}}_{lm}$, fulfil:

$$\langle \hat{\mathbf{u}}_{ln} | \hat{\mathbf{u}}_{mo} \rangle = \delta_{lm} \delta_{no} , \quad (8.13)$$

where $\langle | \rangle$ is the internal (voxelwise) product and δ_{ij} represents the *Kronecker* delta symbol. The field, b , inside the brain mask is now projected onto $\hat{\mathbf{u}}_{lm}$, yielding projection coefficients, $c_{lm} = \langle b | \hat{\mathbf{u}}_{lm} \rangle$. This allows for the formulation of an SSH-approximated expansion of the static field:

$$\tilde{b}_{\text{dist}} = \sum_{l=1}^n \sum_{m=-l}^l \underbrace{\langle b | \hat{\mathbf{u}}_{lm} \rangle}_{c_{lm}} \cdot \hat{\mathbf{u}}_{lm} . \quad (8.14)$$

Further information on the procedure, expressed in a slightly different notation, can be found in [Hirsch, 2009]. Fig. 8.9 shows the SPHINX filter applied to the sample data using filter orders between one and seven. The line plot in Fig. 8.10 illustrates how local contrast behaves for the different orders and Fig. 8.11 gives an overview on the field distributions in form of width-normalised histograms. The compensation for long range distortions is apparently better than with POLF, especially at higher fit orders. Certain ringing artefacts at the edges remain. Nevertheless, these are less pronounced than for POLF and can be further reduced by choosing higher fit orders. The histogram peak lies slightly below the peak for the POLF. This fact can probably again be attributed to better contrast preservation, since the visual results, shown in Fig. 8.9 and Fig. 8.10, indicate proper background correction with preservation of a rich contrast.




Figure 8.9: *SPHINX* applied with different orders. The top row shows a transverse slice, while the bottom row shows a three-dimensional cube with orthogonal cross-sections through the volume. As for the polynomial fit, the accuracy of the *SPHINX* fit increases with the order. Fit results at order 6 or 7 already show a visually sound result.





Figure 8.10: Line plots for the SPHINX orders shown in Fig. 8.9. SPHINX preserves the overall contrast in its rough shape better than POLF. Long range variations are compensated for well at high fit orders.

Figure 8.11: Histogram showing the effect of the SPHINX filter for the orders 1 – 7, as illustrated in Fig. 8.9. Units of the y-axis are arbitrary. As for the polynomial filter, the field distribution gets narrower with increasing order.



8.4.2 Implementation

The prototype of the SPHINX algorithm was implemented by *Sebastian Hirsch*. In the scope of this thesis, the algorithm was slightly improved and modified. Most importantly, some application-related features were introduced such as a customisable reference grid including scalable axes. Further, an approximative SPHINX approach, SPHINXapprox was designed. SPHINXapprox allows one to reuse spherical harmonic base functions computed for a mask with high similarity to the current one. The orthonormality is not perfect in that case, leading to potential fitting imperfections, but the method avoids the time-consuming orthonormalisation procedure and generates acceptable results. However, within the applications in this thesis SPHINXapprox was not employed in order to maintain maximum accuracy.

8.4.3 Practical Consideration

SPHINX can be applied with an arbitrary spherical harmonic order, l . In contrast to the polynomial filter, where the order is limited by memory demands, here the order is just confined by the desired correction quality of the result and the time limitations.

To summarise this method, a few key features and further properties found in practical applications are listed here:

- ▲ Fit results follow physical basis and do not falsify field variations
- ▲ Low ($2^{nd} - 4^{th}$) order fits can be computed rapidly. The approach is related to shimming, which offers even the opportunity to estimate shim values, or to analyse imperfections of the shims
- ▲ Fitting can be confined to a mask
- ▲ The fitting order can be increased infinitely and is primarily limited by time confinements
- ▼ When robust background correction with high-order SSHs is desired, computational demands increase dramatically (as usually orthonormal SSHs have not been pre-calculated)
- ▼ Even at high orders, variations originating in local disturbances such as air-filled cavities cannot be compensated for. Since SPHINX performs optimal on a volume influenced only by distant external sources, the physical principle is not designed to match arbitrary internally generated distortions

8.5 External Susceptibility Distribution

In Chpt. 7 it was shown that distortions of the static magnetic field originating in susceptibility changes can be approximated by convolving the underlying distribution of magnetic susceptibility, χ , with the field of a magnetic dipole, $d = (3 \cos^2 \vartheta - 1)/(4\pi r^3)$ (see Eq. 7.14). With this assumption, Eq. 8.1 can be expressed as:

$$b_{\text{dist}} = B_0 \cdot [(\chi_{\text{ext}} * d) + (\chi_{\text{int}} * d)] \quad (8.15)$$

(for small $\chi_{\text{int/ext}}$), where χ_{ext} is the susceptibility distribution outside the brain mask, m , and χ_{int} is the distribution inside. The minimisation algorithm for susceptibility reconstruction using Tikhonov regularisation has the general form (see Eq. 7.32):

$$\min_{\chi} \|W(b - B_0(\chi * d))\|_2^2 + \lambda \|W_T \cdot \chi\|_2^2 \quad (8.16)$$

It was further shown that with only slight modifications this approach may be used to estimate a pseudo-susceptibility distribution outside the brain mask that generates the non-intrinsic field shifts inside [de Rochefort et al., 2010b]. This is achieved by setting W_T to a high regularisation value inside the mask, favouring solutions with (approximately) zero susceptibility in this area. Further, W is set zero outside the mask to induce field compensation only on the inside, where




Figure 8.12: Schematic view of the DIPF: A pseudo-susceptibility distribution (second image) is estimated for the raw fieldmap (first image), residing only outside the brain mask. Using the pseudo-susceptibility, a background fieldmap is calculated by dipole convolution (third image). Subtracting the background field from the raw data leads to the background-corrected local fieldmap (last image).

the field is known. The external susceptibility distribution, χ_{ext} , and the distortions it induces, b_{dist} , can be estimated by minimising:

$$\min_{\chi_{\text{ext}}} \|W(b - B_0(\chi_{\text{ext}} * d))\|_2^2 + \lambda \|W_T \cdot \chi\|_2^2 \rightarrow b_{\text{dist}} \approx B_0(\chi_{\text{ext}} * d) \quad . \quad (8.17)$$

The background field determined by this DIPF is based on a dipole approximation of the static magnetic field (see Koch et al. [2006]) and thus automatically incorporates a certain physical reality. Fig. 8.12 outlines the steps for BFR via estimating the external pseudo-susceptibility distribution, χ_{ext} .

The elementary parameter changing the minimisation outcome is the iteration count, being naturally related to the fit accuracy. Fig. 8.13 shows background fields and correction results for different iteration counts, whilst in Fig. 8.14 line plots throughout the corrected fields for the different parametrisations are illustrated. Fig. 8.15 shows the corresponding width-normalised field distribution histograms. Evidently, a higher iteration count generates a result with higher focusing around zero. Yet, a high level of (local) contrast is preserved in all parametrisations, indicating that the physical basis of the field is not destroyed. Especially at the mask edges, that poses the biggest problem for the former filters, DIPF shows very robust results without giving rise to extensive artefacts. In contrast to the POLF that would potentially remove image contrast for very high fit orders, the DIPF approximates an optimal solution with external sources only. The change in the results becomes very low when approaching 50 iterations. Thus, it appears to be safe to postulate 50 – 100 iterations for a default correction procedure.

8.5.1 Implementation

For this kind of pseudo-susceptibility estimation the explicit regularised inversion discussed in Section 7.4.3 cannot be applied, as the source area containing the susceptibility distribution and the target area, containing the field distortions are not identical. Thus, a traditional minimisation as in Section 7.4.2 has to be employed. An approach using Tikhonov-weighting serves the problem




Figure 8.13: Results of the DIPF applied using different iteration counts for minimisation. The top row shows sample slices and the lower row volumetric plots of orthogonal slices of the background-corrected field. The improvements for each increase of the iteration count are evident.




Figure 8.14: Line plot example for the results of the DIPF using different iteration counts for minimisation. The centring around zero becomes stronger with increasing iteration count, though the local contrast widely remains.




Figure 8.15: Histogram showing the effect of the DIPF using different iteration counts for minimisation. Units of the y-axis are arbitrary. The field distribution becomes narrower with increasing iteration count. The peak for 50 iterations shows even greater amplitude than the highest fitting order of POLF (Fig. 8.8) and SPHINX (Fig. 8.11).

well. The Tikhonov weighting matrix, W_T , is set to the brain mask, m , as well as the result weighting matrix, W :

$$W_T = W := m \quad , \quad (8.18)$$

enforcing a zero solution inside the brain, but causing the evaluation of field shifts to take place only in the very same region. A high Tikhonov regularisation factor, $\lambda = 500$, is chosen (affecting only the brain mask), and the count of iteration steps is set to at least $n = 50$ or 100 , depending on the problem size and time constraints. Gradient regularisation is not employed. The background field is calculated by a simple discretised dipole convolution of the minimisation result. Optionally, the fieldmap can be zero-padded before computation, in order to increase the source area size for containing the pseudo-susceptibility and to avoid folding artefacts (see Section 7.6.3).

Pro and Contra of the DIPF:

- ▲ High fit accuracy for local distortions due to accurate susceptibility distribution
- ▲ Fitting can be confined to a mask
- ▲ The approach is physically sound and should not falsify the overall field characteristics → important, if susceptibility reconstruction is desired.
- ▼ High iteration count is necessary to compensate for long range fields, since minimisation for distant field shifts is less effective.
- ▼ Whenever the mask exterior available for carrying the pseudo-susceptibility is insufficient to express the structure causing the field shifts, compensation is not possible.

8.6 Conclusions

Various filters were demonstrated, each possessing different preferences and disadvantages. Rudimentary filters can be applied at low computation cost. Yet, they suffer from non-selective deletion of image contrast in form of high-pass filtering. The POLF preserves more local to medium-range contrast, but loses accuracy with distance to the centre. POLF furthermore potentially removes local contrast at high fit orders. SPHINX shows more robust results and relies on a physical basis – the spherical harmonic functions. This ensures that contrast is not removed in an uncontrolled manner, as for the former filters. The DIPF also has a physical basis – an external pseudo-susceptibility distribution. It generates robust fit results. Especially at the mask edges DIPF seems more accurate than SPHINX.

The parametrisation of the presented filters has to be adapted to the data under investigation. Field histograms and line plots support the determination of optimised parametrisations and exploit information about the field characteristics. The peak of the field distribution should be as narrow as possible, whilst the image contrast should be kept high.

Multistage Background Field Removal

In this chapter, a novel stepwise BFR algorithm is presented, addressing the versatile sources of static field distortions one by one in separate sequential steps. This MULTistage BAcground Field REmoval (MUBAFIRE) algorithm [Lindemeyer et al., 2011b, 2012] is characterised in detail and compared to the previously introduced approaches in application on *post mortem* and *in vivo* measurements. Moreover, it is tested on a simulated numerical sample to facilitate a comparison with other methods. Parts of the contents of the current and the previous chapter form a publication, which is currently in preparation, bearing the title: *Multistage Background Field Removal (MUBAFIRE) – Compensating for B_0 Distortions at Ultra-High Field*.

9.1 Motivation

The filters presented so far are all based on a particular method of determining and removing background fields. While POLF and SPHINX achieve this by fitting or projecting functions on the field inside the brain mask, the DIPF seeks for an iterative compensation for the field variations by creating a pseudo-susceptibility distribution. The novel MUBAFIRE filter aims at combining filters with different optimal domains of applicability such that the disadvantages of one filter are, to a large extent, covered by the advantages of other filtering steps. For example, the long-range, smooth filtering provided by low-order spherical harmonics is complemented by the strong local sensitivity of the DIPF. In order to visualise and prove this hypothesis, an example is presented below.

In Fig. 9.1 a simulated fieldmap is analysed with SPHINX and DIPF. The field was created using a simple spherical susceptibility phantom with a small cut-out on the mask edge. The susceptibility on the inside was assigned a value of 5 (in arbitrary units) while the exterior is zero-filled. The field generated by the phantom was calculated using dipole convolution as in Eq. 7.14. Only the small artificial imperfection generates inner distortions, since spherical bodies remain otherwise homogeneous when exposed to a static field. An artificial external field disturbance with the characteristics of a spherical harmonic of order $l = 2$, $m = 0$ was added to the field:

$$[r^l Y_l^m]_2^2 = [\text{norm}] \cdot r^2 \cdot (3 \cos^2 \theta - 1) \quad , \quad (9.1)$$




Figure 9.1: The numerical phantom (left) was transformed into a fieldmap by dipole convolution and falsified by the solid spherical harmonic of order $l = 2$, $m = 0$ (second image). Background correction with SPHINX (third image) and DIPF (last image) both show insufficient accuracy.

where r is the radial distance and θ is the inclination angle towards the z axis. The resulting field, b_{dist} , was analysed by SPHINX and the DIPF. Fig. 9.1 clearly shows that while SPHINX removes long range distortions introduced by the harmonic very well, the local distortions remain. On the other hand, DIPF removes the local distortions near the surface of the sphere almost perfectly. However, compensation of the harmonic influence is only partially feasible. The concept presented in this chapter combines the advantages of both filters, creating an algorithm sensible to short and long range field perturbations.

9.2 Theory

The MUBAFIRE filter consists of three essential filter cores: POLF, SPHINX and DIPF. The workflow of MUBAFIRE is illustrated in Fig. 9.3. Starting with the distorted fieldmap, b , the algorithm progresses through the following steps. First, a 1st order POLF is used to correct for linear field gradients and constant shifts (see Section 8.3):

$$b_{\text{POLF}} = \min_{\alpha_{ijk}} \left(b - \sum_{i,j,k=0}^{(i+j+k) \leq 1} \alpha_{ijk} \cdot x^i y^j z^k \right) \quad (9.2)$$

$$= b - b_{\text{off}} - G_x x - G_y y - G_z z \quad (9.3)$$

The fitting coefficients, α_{ijk} ($i, j, k \leq 1$) represent a constant field shift, b_{off} , and potentially remaining linear field gradients, $G_{x,y,z}$. Compensating for these will balance the phase distribution around zero – if quantitative analysis is desired, the constant offset correction, b_{off} , may be retained for later use. In order to preserve the physical consistency, no higher order polynomials are considered. This is especially important if susceptibility reconstruction is planned as a further step. The 0th and 1st order polynomials are entirely part of the spherical harmonics, thus they represent physical solutions to the static field problem.

Harmonic variations above 1st order are now addressed by SPHINX (see Section 8.4). The corrected field is generated after determining the orthonormalised function set, $\hat{\mathbf{u}}_{lm}$, as:

$$b_{\text{SPHINX}} \approx \sum_{l=1}^n \sum_{m=-l}^l \underbrace{\langle b_{\text{POLF}} | \hat{\mathbf{u}}_{lm} \rangle}_{c_{lm}} \cdot \hat{\mathbf{u}}_{lm} . \quad (9.4)$$

The highest order included in the fitting usually varies between four and six. Higher orders can also be included to the cost of processing time. From experience, an order of up to six is sufficient to prepare the fieldmap for the follow-up step. The applied filters primarily remove field shifts that are caused from outside of the FOV, since spherical harmonics yield smooth field solutions. Sources residing nearby or inside of the FOV must be addressed by an additional filter step. The initial removal of long range contributions and harmonics by POLF and SPHINX allows for a more precise and faster analysis regarding local distortions, since the strong weighting of major field slopes is no longer part of the solution.

External sources of long- and short-range dipole fields are determined by the DIPF (see Section 8.5), estimating the external pseudo-susceptibility distribution, χ_{ext} :

$$b_{\text{MUBAFIRE}} = b_{\text{DIPF}} = \chi_{\text{ext}} * d \quad (9.5)$$

$$\text{with } \chi_{\text{ext}} : \min_{\chi_{\text{ext}}} \|W_{\text{mask}}(b_{\text{SPHINX}} - B_0(\chi_{\text{ext}} * d))\|_2^2 + \lambda \|W_T \cdot \chi\|_2^2 \quad (9.6)$$

$$W_{\text{mask}} = W_T = m , \quad (9.7)$$

where $*$ denotes convolution, d is the dipole kernel (Eq. 7.11), W is the weighting matrix defining the minimisation region and W_T is the Tikhonov weighting matrix enforcing zero susceptibility inside the brain mask, m . The generated distortion map is then subtracted from the field inside the ROI. This concludes the basic MUBAFIRE workflow. Applying this scheme on the example given in Fig. 9.1 results in a smooth fit, as shown in Fig. 9.2. With careful observation, one recognises




Figure 9.2: The MUBAFIRE algorithm compensates for both (long range harmonic and local distortive) field disturbances (explanation for images 1-4, see Fig. 9.1).

that SPHINX as well as the DIPF applied solitarily suffer from that kind of field distortions that do not match their application scheme. Due to the cut-out imperfection, SPHINX does not perform a perfect fit to the sphere, although the introduced background function is of harmonic nature. In the same way, DIPF cannot compensate for the imperfection well enough, as the

harmonic distortion hinders an optimal solution. In sequential application, DIPF corrects for the fit inaccuracies of SPHINX. The combined algorithm is indeed able to generate an acceptable, flat fit result, as desired.




Figure 9.3: Schematic workflow of the MUBAFIRE algorithm, applied on a post mortem brain at 3 T. The top row shows the standard filter chain. POLF corrects for constant shifts and linear gradients, SPHINX compensates for harmonic contributions and, finally, DIPF takes care of disturbances with higher localisation. The lower row illustrates the optional MUBAFIRE Local extension: A thresholding step identifies outliers, gentle mask erosion excludes the outliers and their direct neighbours and a second DIPF corrects for field distortions generated by the newly masked regions.

Local Distortions: Occasionally some strong remnant field disturbances can be generated by small cavities, by blood vessels or by erroneous phase/field values that were not detected during data preparation and are thus still residing inside the brain mask. These field distortions have strong impact on e.g. susceptibility reconstruction, as especially the stray field of air cavities extends into surrounding regions. This leads to a strong local contrast that generates incorrect field values, hindering correct susceptibility reconstruction.¹ Such effects can be taken into account by an enhancement of MUBAFIRE, a local field filter. This filter, MUBAFIRE Local, consists of three additional substeps – thresholding, connecting/erosion and yet another DIPF. The filter principle is illustrated in the lower row of Fig. 9.3.

The MUBAFIRE Local filter chain expects that the base chain has been applied beforehand. This is crucial, since the first step – thresholding – needs all regions of the volume containing useful contrast to be centred about zero and to be describable in a well-determined field range. E.g. uncorrected field gradients would lead to the exclusion of vast areas. One may define the temporary fieldmap:

$$\tilde{b} = b_{\text{MUBAFIRE}} \quad . \quad (9.8)$$

¹This means, local phase gradients exceed the Nyquist condition, see Section 2.2.

Local spike distortions, as described before, are identified by excluding field shifts larger than a certain expected range, ξ , resulting in a modified mask,

$$m_\xi = m \wedge (|\tilde{b}| \leq \xi) \quad , \quad (9.9)$$

where m is the original brain mask and $\wedge$ is the logical and point-wise “and”-operator. The threshold is chosen with regard to the nature of the signal and the noise behaviour of the sample. Setting ξ to a multiple of the standard deviation, $\sigma_{\tilde{b}}$, of the fieldmap inside the mask,

$$\xi = n_\sigma \cdot \sigma_{\tilde{b}} \quad , \quad (9.10)$$

is an beneficial choice, as it reflects the width of the field distribution. The threshold has to be chosen to best fit the field under investigation. Typical values of n_σ lie between 5 and 15.² Often times, the direct surroundings of the masked areas and voxels are also affected by noise or incorrect phase. A customised erosion operation is used to connect standalone voxels in the exclusion of m_ξ , that do not share a common surface, but are arranged as diagonal neighbours, such as the indexes $[i,j,k]$ and $[i+1,j+1,k]$. The algorithm will in this case additionally exclude $[i+1,j,k]$ and $[i,j+1,k]$ to join the neighbours to one group. A typical six-neighbours erosion is applied thereafter. With this step, noisy phase in the proximity, especially between masked voxels, is excluded without eroding the mask too aggressively. The thresholding and erosion step is outlined in Fig. 9.4. Finally, the DIPF estimates susceptibility in the excluded areas and the field shifts generated by these. After subtraction of the determined local field shifts, ideally, the remnant field will be generated by the susceptibility distribution inside the corrected mask, m_ξ , only. The outcome will thus represent an optimal basis for further analysis of the field, e.g. for susceptibility reconstruction.




Figure 9.4: Erosion filters applied on a synthetic fieldmap with artificially introduced smeared-out, peak-like artefacts. The upper row shows the fieldmaps, while the lower row shows the corresponding masks. From left to right: original field, thresholded field, minimal (diagonal) erosion, six-neighbour erosion.

9.3 Material and Methods

Although the reasoning for the chosen filter chain is relatively clear, a validation in applications and a comparison is required. The MUBAFIRE and MUBAFIRE Local algorithms are in the following applied on fixed brain tissue measurements at 3 T and 9.4 T. Also, a human *in vivo* measurement at 9.4 T is analysed and the algorithm is tested on a simulated numerical sample to facilitate a comparison with other methods.

² n_σ defines the multiple of $\sigma_{\tilde{b}}$ to be used for thresholding.




Figure 9.5: Applying the MUBAFIRE (third image) and the MUBAFIRE Local algorithm (last image) to the simulated fieldmap (second image) generated from the synthetic susceptibility distribution in the first image. The strong local distortions were not excluded from the mask – MUBAFIRE ignores the influences, whilst MUBAFIRE Local recognises and excludes them by thresholding and successfully compensates for their influence on the surroundings.

9.3.1 Measurements

For the *post mortem* application a fixed female brain obtained from the brain donor programme of the University of Düsseldorf scanned in a custom-made cylindrical acrylic glass container filled with formalin was used. Measurements were performed a) on a 3 T TIM Trio scanner and b) on a 9.4 T human MR scanner, both: (Siemens AG, Erlangen, Germany). In case a), a twelve channel phased array Rx coil was used. A standard 3D gradient echo sequence with slab-selective excitation was applied.

The sequence parameters for the 3 T post mortem case included:

Sequence Parameters					
seq	GRE	exc	slab	Tx/Rx	Body / 12 ch array
matrix	$288 \times 384 \times 176$	FOV	$[172.8 \times 230.4 \times 105.6]$ mm	res.	$[0.6 \times 0.6 \times 0.6]$ mm
TR	46 ms	α	20°	RF-spoil	on
TE	$[9.84, 22.14]$ ms			BW	110 Hz/Px

The entire brain was covered by the FOV using transverse slicing. A brain mask was generated with bet2 [Smith, 2002]. In case b), an eight-channel-coil was used in single-channel transmit and eight-channel receive mode. In order to maximise the accuracy and spatial consistency of the phase data, the channel with highest SNR was used for further analysis. A slab-selective 3D gradient echo sequence was employed.

The sequence parameters for the 9.4 T post mortem case included:

Sequence Parameters					
seq	GRE	exc	slab-selective	Tx/Rx	8 ch array, 1 ch used
matrix	$480 \times 512 \times 160$	FOV	$[115.2 \times 122.9 \times 38.4]$ mm	res.	$[0.24 \times 0.24 \times 0.24]$ mm
TR	63 ms	α	25° (nominal)	RF-spoil	on
TE	16.38 ms	avgs	12	BW	40 Hz/Px

The measurement, acquired with sagittal slicing, covers parts of the central brain, the brain stem and the cerebellum. An ROI, in which the sensitivity of the coil was high, was chosen for further analysis. The *in vivo* measurement c) was performed with informed consent on a healthy volunteer at the 9.4 T human scanner described above.

The sequence parametrisation for the 9.4 T *in vivo* case included:

Sequence Parameters					
seq	GRE	exc	slab	Tx/Rx	8 ch array
matrix	$308 \times 448 \times 86$	FOV	$[159 \times 224 \times 43]$ mm	res.	$[0.5 \times 0.5 \times 0.5]$ mm
TR	38 ms	α	48° (nominal)	RF-spoil	on
TE	$[3.93, 9.68, 15.43, 21.18, 26.93]$ ms			BW	235 Hz/Px

The whole brain phase data from a) and the ROI phase images from b) and c) were unwrapped using PRELUDE [Jenkinson, 2003]. In case a) and c) fieldmaps were calculated by linear regression of the unwrapped phase, whilst case b) required only division by the echo time. An alternative processing workflow for measurement b) is further described in the application chapter (see Section 11.7). All presented algorithms were implemented in MATLAB® (The MathWorks Inc., Natick, USA).

9.3.2 Simulation

In order to investigate the accuracy of the MUBAFIRE algorithm quantitatively, a numerical model was used. The model is composed of field shifts induced by local susceptibility contrast superimposed with long range static field variations obtained in a phantom measurement. The actual distorted field, B_{true} , induced by an arbitrary susceptibility distribution, can be expressed as:

$$B_{\text{true}}(\mathbf{r}) - B_0 = B_0 \cdot (d * \chi)(\vec{\mathbf{r}}) \quad , \quad (9.11)$$

where d is the dipole kernel (see Eq. 7.11), as in [Salomir et al., 2003; Marques and Bowtell, 2005]. Considering long range spatial variations, the static field, B_0 , can be split into a constant component and a spatially varying part, $B_0 \rightarrow B_0(\vec{\mathbf{r}}) = B_{\text{base}} + b_{\text{inh}}(\vec{\mathbf{r}})$. Inserting this into Eq. 9.11 results in the expression:

$$\begin{aligned} b_{\text{sim}}(\vec{\mathbf{r}}) &= B_{\text{true}}(\vec{\mathbf{r}}) - B_{\text{base}} \\ &= (B_{\text{base}} + b_{\text{inh}}(\vec{\mathbf{r}})) \cdot (d * \chi)(\vec{\mathbf{r}}) + b_{\text{inh}}(\vec{\mathbf{r}}) \quad , \end{aligned} \quad (9.12)$$

where the left hand side stands for the simulated (observable) field shift. A brain confinement mask, m , and an internal and external susceptibility distribution,

$$\chi_{\text{int}} = m\chi = \chi_{\text{int}}(\vec{\mathbf{r}}) \quad , \quad (9.13)$$


$$\chi_{\text{ext}} = (\neg m)\chi = \text{const.}, \forall \vec{\mathbf{r}} \in (\neg m) \quad , \quad (9.14)$$

are artificially generated. The magnitude of a separate multiple echo gradient echo measurement applied on the *post mortem* brain described in Section 9.3.1 is used to estimate a T_2^* distribution. The distribution is Gaussian-smoothed, scaled and windowed to a susceptibility range roughly

related to that of human tissue, water and air (order of $\leq \pm 10^{-5}$). The constant susceptibility of the brain exterior is set higher than the average χ_{int} , such that distortions originating in the brain geometry evolve.

The real field inhomogeneity, $b_{\text{inh}}(\mathbf{r})$, in a water-filled spherical phantom was measured using a standard 3D gradient echo sequence, $80 \times 128 \times 128$ voxels at 1.5 mm isotropic voxel size, $\text{BW} = 610 \text{ Hz/Px}$, $\text{TR} = 50 \text{ ms}$ and $\text{TE} = [2.36, 5.02, 7.87, 10.72, 13.57, 25.00, 35.00, 37.85] \text{ ms}$. An area of continuous signal from around the centre of the sphere was extracted and stretched to the simulation FOV while keeping the aspect ratio in order to avoid modifying the physics of the problem. The value range of this external distortion map, b_{inh} , was scaled up to more than one order of magnitude above the internal field distortions ($\pm 2 \cdot 10^{-4}$). Nevertheless, exact imitation of the brain anatomy and its susceptibility values is not intended, since this simulation serves as a universal and independent proof of principle. The simulated susceptibility distribution, $[\chi_{\text{int}}, \chi_{\text{ext}}]$, and the distorted field, b_{inh} are inserted into Eq. 9.12 to obtain the simulated field, $b_{\text{sim}}(\vec{\mathbf{r}})$. b_{sim} is then analysed by all BFR algorithms. A scheme of the field simulation can be found in Fig. 9.6.

Figure 9.6: *Simulation scheme: A gradient echo measurement is used to estimate a T_2^* map; using the T_2^* anatomy and rescaling the values to an appropriate range, a susceptibility distribution is generated; dipole convolution leads to a local fieldmap; external fluctuations are introduced by adding a fieldmap derived from a sphere phantom measurement.*



As a reference, an internal fieldmap, $b_{\text{int}}(\vec{\mathbf{r}})$, is calculated by evaluating only the convolution term in Eq. 9.12 with $\chi_{\text{sim}}(\vec{\mathbf{r}}) = \chi_{\text{int}}(\vec{\mathbf{r}})$ inside the mask, and $\chi_{\text{sim}}(\vec{\mathbf{r}}) = \overline{\chi_{\text{int}}(\vec{\mathbf{r}})}$ (the average of χ_{int}) in the exterior. While exhibiting only local contrast, this fieldmap serves as a reference for quantifying the accuracy of the BFR algorithms:

$$b_{\text{int}} = (B_{\text{base}} + b_{\text{inh}}(\mathbf{r})) \cdot [(d * \chi_{\text{sim}})(\mathbf{r})] \quad (9.15)$$

9.3.3 Parameters

In order to quantify the robustness of the MUBAFIRE algorithm, it was compared to the rudimentary GF, the POLF, DIPF and SPHINX. All filters were applied on the designated regions in the prepared fieldmaps of the *post mortem* and *in vivo* measurements as well as on the simulated maps. The GF was applied with $\sigma = 6$ voxel widths, POLF was calculated up to 5th order. The DIPF was applied at full resolution with 50 iterations and Tikhonov regularisation with $\lambda = 500$.

An order of 10 was used for the standalone SPHINX correction. Finally, in MUBAFIRE, polynomial filtering was performed up to 1st order, SPHINX was processed up to 4th order and the DIPF was applied with 50 iterations and $\lambda = 500$ for external filtering as well as for local filtering. Thresholding was performed with $n_\sigma = 8$ for the *post mortem* data and with $n_\sigma = 14$ for the *in vivo* case. The simulated data were analysed accordingly, except for the GF being applied with $\sigma = 4$ voxels width due to lower resolution. Voxels exceeding the measurement range or giving erroneous information were not simulated, that is, MUBAFIRE Local was only applied on measured data, but not on the synthetic data.

9.4 Results

Table 9.1 shows the numerical results for mean value and standard deviation of the corrected field from the simulation and from all measurements. Both values were evaluated in a modified brain mask that was eroded once in order to exclude numerical edge effects. For the simulation, the table includes the $\mathbb{L}_2$ norm difference described below, in Eq. 9.16.

Correction Results for Simulation						
	GF	POLF	SPHINX	DIPF	MUBAFIRE	true
mean	0.0236	-0.0653	-0.0663	0.6772	0.0414	0.0144
std(b)	4.2368	2.4473	2.1526	4.2000	2.0585	2.0491
$\ \Delta b\ _2^2$	0.0199	0.0130	0.0103	0.0187	0.0088	0.0000
Correction Results for fixed brain at 3 T						
	GF	POLF	SPHINX	DIPF	MUBAFIRE	MUBAFIRE Local
mean	0.0303	0.0409	0.0136	-0.0484	0.0049	0.0050
std(b)	0.8997	1.2085	1.1095	1.1374	0.9348	0.9922
Correction results for cerebellum at 9.4 T						
	GF	POLF	SPHINX	DIPF	MUBAFIRE	MUBAFIRE Local
mean	0.0185	0.0055	-0.0154	0.0963	0.0337	0.0737
std(b)	5.7816	7.0856	6.8894	6.4352	6.3977	4.6326
Correction results for 9.4 T in vivo measurement						
	GF	POLF	SPHINX	DIPF	MUBAFIRE	MUBAFIRE Local
mean	0.0548	-0.0267	-0.0316	-0.2274	0.0243	0.0927
std(b)	10.2731	21.9982	12.8623	7.7240	6.9407	7.3318

Table 9.1: Results of the BFR algorithms for Simulation, 3 T, 9.4 T *post mortem* and 9.4 T *in vivo*; Units are Hz for the measurements and *a.u.* for the simulation.

The corrected fieldmaps from the simulated data are compared to the local contrast fieldmap, b_{ref} (see Eq. 9.16), visually and quantitatively. Fig. 9.7 shows plots of representative orthogonal slices for each algorithm, exhibiting the difference between the corrected fields and b_{ref} . The histogram in this figure shows the field distribution inside the mask. For quantitative analysis the $\mathbb{L}_2$ -norm residua of the difference between the corrected fields and b_{ref} were calculated:

$$\begin{aligned} \|\Delta b\|_2^2 &= \frac{1}{n_{\text{voxels}}} \cdot \sqrt{\sum_{\mathbf{r}} (b_{\text{corr}}(\mathbf{r}) - b_{\text{ref}}(\mathbf{r}))^2} \quad , \\ b_{\text{ref}} &= \overline{b_{\text{int}}} \quad , \end{aligned} \quad (9.16)$$

where $\overline{b_{\text{int}}}$ is the average of b_{int} over the modified brain mask. For all measurements a representative slice was chosen and plotted for each individual BFR algorithm, separately. This is shown in




Figure 9.7: BFR algorithms applied on the simulated map – (a) showing from left to right: GF, POLF, SPHINX, DIPF, MUBAFIRE correction and true simulated field, b_{ref} ; (b) shows the difference between corrected fields and b_{ref} , from left to right: GF, POLF, SPHINX, DIPF, MUBAFIRE. (a) and (b) show from top to bottom: transverse, coronal and sagittal slicing; (c) histograms of the field distribution; (d) MUBAFIRE-corrected slice with indicated line plot path; (e) line plot.

Figs. 9.8, 9.9 and 9.10. A histogram is plotted for each experiment, describing the field distribution inside the brain mask. In the *post mortem* cases at 3 T and 9.4 T, GF shows the narrowest peak followed by MUBAFIRE Local. For the *in vivo* measurement at 9.4 T, MUBAFIRE Local shows the narrowest field distribution, although its standard deviation is slightly larger compared to MUBAFIRE. All BFR algorithms are further compared using line plots, shown in Figs. 9.8 and 9.10. In particular, the benefits of the MUBAFIRE Local filter at 9.4 T *post mortem* and *in vivo* are outlined in Figs. 9.9 and 9.10. The line plot of MUBAFIRE at 9.4 T *post mortem* in Fig. 9.9 shows a strong spike. Judging from the slice image localising the line plot, the responsible structure appears to be a dipole-like distortion, and thus most likely resembles an air bubble inside the fixed brain tissue. Due to the amplitude of this distortion it can no longer be treated as regular contrast. The distortion as well as its influences on the surroundings are successfully removed by MUBAFIRE Local while keeping the underlying curve shape intact. The line plot at 3 T in Fig. 9.8 shows, that MUBAFIRE approximates the SPHINX solution in some parts, but favours DIPF in others.

Both, simulation and measurements show, that MUBAFIRE (Local) is rich in contrast and narrow regarding the distribution peak of the corrected field. Including visual and numerical evaluations, MUBAFIRE (Local) provides the most robust results, closely followed by DIPF and SPHINX.




Figure 9.8: Results of the BFR algorithms applied on measurement a) (post mortem at 3T). (a) from left to right: raw, GF, POLF, SPHINX, DIPF, MUBAFIRE, MUBAFIRE Local; (b) histogram comparison; (c) position of line plot in MUBAFIRE Local corrected slice; (d) line plot.




Figure 9.9: Results of the BFR algorithms applied on measurement b) (post mortem at 9.4T). (a) from left to right: raw, GF, POLF, SPHINX, DIPF, MUBAFIRE, MUBAFIRE Local; (b) histogram comparison; (c i-ii) position of the line plot in MUBAFIRE (c i) and MUBAFIRE Local (c ii) corrected slice; (d) line plot.

9.5 Discussion

The results of the correction algorithms applied on the simulated data, illustrated in Fig. 9.7, show a behaviour comparable to the case study examples. Obviously, MUBAFIRE exhibits the best approximation of the true field, closely followed by SPHINX and DIPF. This finding is substantiated by the line plot, showing a good agreement between true field and MUBAFIRE. Here, the edge effects of GF become apparent. The quantitative evaluation of the $\mathbb{L}_2$ difference indicates that SPHINX may already have a higher accuracy than the DIPF. This difference is presumably caused by edge artefacts (over-/underestimations) in the latter filter, observable in the difference maps shown in Fig. 9.7. Still, the accuracy is enhanced even when applying the DIPF additionally as in MUBAFIRE. MUBAFIRE obviously outperforms DIPF and SPHINX in precision.




Figure 9.10: Measurement c) (*in vivo* BFR results at 9.4 T). (a) from left to right: raw, GF, POLF, SPHINX, DIPF, MUBAFIRE, MUBAFIRE Local; (b) histogram overview; (c) shows the MUBAFIRE Local correction, (d i) line plot comparing the BFR algorithms and (d ii) line plot comparing MUBAFIRE and MUBAFIRE Local in magnified view (magnification is indicated in (d i) by yellow overlay).

It should also be noted, that the MUBAFIRE algorithm uses a significantly lower SSH order (4) than the stand-alone SPHINX, allowing the entire algorithm to be computed in significantly less time than the presented SPHINX of order 10.

The sample slices for the case studies illustrate the characteristics of the different algorithms. As expected, the GF produces significant artefacts at the mask edges. In most cases, GF generates the narrowest or second narrowest peak, but results in a relatively low image contrast compared to all other filters. The filter has little physical basis and cannot distinguish between long to medium size field distortions induced by outside sources and contrast originating from inside the ROI on a medium-size scale. POLF seems to be able to filter the largest field variations without inducing strong artefacts. Still, it exhibits periodic oscillations due to radially decreasing fit accuracy. SPHINX evidently results in significantly higher image contrast compared to GF and POLF. DIPF and MUBAFIRE (Local) visually show the greatest richness in overall contrast. Except for the GF, the filters DIPF, SPHINX and MUBAFIRE (Local) induce the narrowest peak for the field distribution while preserving image contrast. The standard deviation for MUBAFIRE Local resides in the same range as for MUBAFIRE, except for measurement b) (9.4 T *post mortem*), where it lies significantly below. It is apparent that compensation for sources excluded from the mask does not necessarily reduce the overall standard deviation of the field distribution. Yet, the quality of the field map in regions corrected for local disturbances is improved. Note, that the physical basis of the field constitution should ideally be preserved in SPHINX, DIPF and MUBAFIRE (Local), only. The line plot for case a) shows that MUBAFIRE seeks a compromise between spherical harmonic and dipole field approximation, as it supports both solutions in parts. The line plot of case b) (Fig. 9.9) shows that MUBAFIRE Local triggers significant improvement of local distortions – in this case most likely air bubbles – that prevent direct inspection of the underlying structure or exceed the thresholded value range. Also, in the *in vivo* experiment c), it was possible to successfully compensate for local distortions (Fig. 9.10). The numerical results support the experimental (visual) observations.

9.6 Conclusions

It can be concluded by referring to the experimental visual and numerical results from the sample cases as well as to the quantitative results obtained by simulation that the MUBAFIRE (Local) algorithm offers a robust means to correct fieldmaps for arbitrary background shifts. Externally and internally induced shifts can be addressed as well as static field variations. Especially at ultra-high field strengths, such as 9.4 T, field shifts induced by small vessels, air cavities or other material can become strong and even exceed the determinable field range. The latter is dictated by the Nyquist criterion, expecting local phase gradients to stay within a variation of $\pm\pi$ per voxel or per echo time (see Chpt. 3, Eq. 3.20). The MUBAFIRE (Local) algorithm developed in this work is geared to the needs of studies investigating phase contrast and field information. Studies further exploring magnetic susceptibility strongly benefit from consistent fieldmaps, those that MUBAFIRE can provide. Investigation of the algorithm in a susceptibility reconstruction framework is the next objective.

9.7 Effects on Susceptibility Reconstruction

The presented results are clear, particularly the simulated data, regarding the accuracy of MUBAFIRE (Local). Yet, the actual effect of different BFR algorithms on susceptibility reconstruction shall be illustrated in an example. The correction results for measurement a) are each evaluated by using the closed-form susceptibility solver (see Section 7.4.3) with Tikhonov- and gradient-regularised minimisation terms ($\lambda = 0.06$ and $\mu = 0.005$). A sagittal slice from the generated susceptibility maps for the uncorrected field and for each BFR algorithm is plotted in Fig. 9.11.





Figure 9.11: Susceptibility reconstruction applied to the raw field, and the correction results of GF, POLF, SPHINX, DIPF, MUBAFIRE and MUBAFIRE Local. The colour range is identical for all images and has arbitrary units. Yellow arrows indicate a gyrus of underestimated contrast. Green arrows indicate edge artefacts and over-/underestimated regions. Red arrows indicate the nuclei of the basal ganglia, corrupted by artefacts. Cyan arrows indicate internal regions of erroneous field values, corrupting their surroundings. All artefacts are widely eliminated either by MUBAFIRE or MUBAFIRE Local.

Applying the reconstruction to the raw data leads to a completely distorted image. The GF shows weak contrast compared to the other filters and the displayed structures are blurry or corrupted by artefacts. The POLF-corrected result shows stronger contrast, yet the map seems to suffer

from global inhomogeneities. The same applies for the map of the SPHINX filter. In the DIPF-corrected result the contrast is rich, but especially the mask edges, e.g. the cortical gyri, show overestimated values. MUBAFIRE compensates for most of the shortcomings of the other filters and MUBAFIRE Local, finally, removes even internal distortions originating in erroneous phase or field values and prevents propagation of these errors to their surroundings, successfully.

Figure 9.12: Schematic workflow of background field correction with the MUBAFIRE framework. The left hand side shows the interface functions, while the right hand side shows the separate BFR methods applied on the fieldmap, b , stored in the interface function. The green methods in the bottom manipulate the mask, m . The blocks possessing a yellow glow, are the ones used in the optimised MUBAFIRE workflow, presented above. Note that the implementations of SPHINX and GaussianIntelligentNormalisation are based on code from Sebastian Hirsch, [Hirsch, 2009]



9.8 Implementation

The developed MUBAFIRE algorithm consists of the workflow introduced above, but allows for a much higher degree of flexibility than the presented fixed order. Originally, MUBAFIRE evolved from the intention of creating a BFR framework, incorporating all previously presented filters. Similar to FSE, a handling interface was developed, including a configuration structure and a control function, addressing the different filter types. Fig. 9.12 shows the structure of the MUBAFIRE framework. A description of the core function can be found in Appendix D.

The advantages and disadvantages of MUBAFIRE and MUBAFIRE Local can be summarised to the following:

- ▲ Highest filter accuracy in comparison with other presented filter types
- ▲ All employed filters represent a physically valid solution for static field problems
- ▲ Ideal approach for local distortions as well as for long range field characteristics
- ▲ MUBAFIRE Local allows for compensation of strong local distortions or phase inconsistencies
- ▲ Only moderate computing time increase compared to SPHINX or DIPF standalone filters – SPHINX is used with low fitting order in the chain
- ▼ High computing complexity compared to rudimentary filters
- ▼ For the MUBAFIRE *Local* add-on thresholding has to be adapted manually according to the fieldmap under investigation

10

Multiple Rx Recombination

All considerations presented so far found on the availability of a consistent single or multiple echo phase dataset. In fact, modern MRI systems work with receive arrays rather than single-channel coils. Array coils possess a number of Rx channels, each potentially acquiring a full measurement dataset on their own. One particular Rx element exhibits a so-called *sensitivity profile*, that is the result of a complex interaction of the coil structure and the sample. Of great importance here is also the RF wavelength [Wiesinger et al., 2004]. In order to obtain a global dataset, the data acquired by all Rx elements have to be combined.

10.1 Conventions

The following considerations take place in the spatial domain. Let $\hat{S}(\vec{r}) = A(\vec{r})e^{i\varphi(\vec{r})}$ (with the amplitude, A , and the phase, φ) be the actual complex signal generated by the measured volume in spatial domain. Then, the signal recorded by one receive element is subject to the *receive profile* of the particular element, $\mathcal{R}(\vec{r}) = A_{\mathcal{R}}(\vec{r}) \cdot e^{i\varphi_{\mathcal{R}}(\vec{r})}$. A simplified view will be used within this work, where the receive profile just scales the true signal's magnitude and increases or decreases the phase. The recorded signal thus has the form:

$$S = \mathcal{R}(\vec{r}) \underbrace{\left[A(\vec{r}) \cdot e^{i\varphi(\vec{r})} \right]}_{=\hat{S}} = A_{\mathcal{R}}(\vec{r}) A(\vec{r}) \cdot e^{i(\varphi(\vec{r}) + \varphi_{\mathcal{R}}(\vec{r}))} \quad . \quad (10.1)$$

In this chapter, the spatial position, $\vec{r}$, is dropped wherever the operation described would apply to an arbitrary voxel.

10.2 Optimal Reconstruction

The task of guessing the true signal is theoretically a straightforward calculation. Neglecting noise, the true signal could be derived from any of the receive channels by inverting Eq. 10.1:

$$\hat{S} = \frac{S}{\mathcal{R}} = \frac{A \cdot e^{i\varphi}}{A_{\mathcal{R}} \cdot e^{i\varphi_{\mathcal{R}}}} \quad , \quad (10.2)$$

which is a point-wise division of the measured signal by the receive profile. Unfortunately, the receive profiles are unknown. Although some arbitrary characteristics – such as higher sensitivity for regions close to the receive element – can roughly be estimated, the true receive profile depends on the measurement set-up and is not predictable. Especially the phase profile is a troublesome parameter. Hence, the channels need to be recombined either in a way that avoids the calculation of receive profiles, or the actual profiles or parts of the profiles have to be estimated from the measurement. Noise correlation between different Rx channels is often encountered for low signals, but will be neglected here.

10.3 Complex Sum

The most straightforward way to combine multiple receive signals is forming the complex sum. Let the index $c \in [1, \dots, n]$ identify the n receive elements of a coil. Then, the complex sum of the signals is simply described by:

$$S = \sum_{c=1}^n S_c = \sum_{c=1}^n \mathcal{R}_c A \cdot e^{i\varphi} \quad (10.3)$$

$$= \left(\sum_{c=1}^n \mathcal{R}_c \right) \cdot \underbrace{A \cdot e^{i\varphi}}_{\hat{S}} \quad . \quad (10.4)$$

The resulting phase is:

$$\varphi_S = \angle \left[\left(\sum_{c=1}^n \mathcal{R}_c \right) \cdot A \cdot e^{i\varphi} \right] \quad . \quad (10.5)$$

This expression equals the true phase value, φ , if the receive profiles, $\mathcal{R}_c$, do not introduce any phase offset. In general, this very optimistic condition is not fulfilled. To sum this up, the complex sum method assumes the receive profile phases to be approximately identical, while the magnitudes of the profiles may differ without taking influence on the combined phase, φ_S . Fig. 10.1 shows an example for successful complex sum reconstruction.

10.4 Sum of Squares

A popular method dedicated to the fast extraction of the magnitude, A , is the so-called Sum-of-SQuaRes (SSQR) approach. It completely ignores the phase, calculating only the conjugate




Figure 10.1: Numerical example for the complex sum. The upper row shows the magnitude, the lower row shows the phase. For the artificial geometry (left), four channels were simulated, exhibiting different magnitude sensitivities, but similar phase. The complex sum is shown on the right, giving a good representation of the source data.

squares of the signal of each channel:

$$A_{\text{SSQR}} = \sqrt{\sum_{c=1}^n \langle S_c | S_c \rangle} \quad , \quad (10.6)$$

where $\langle A | B \rangle$ stands for the hermitian and point-wise product $A^\dagger \cdot B$. This estimate has little quantitative value, but allows for a decent estimation of the signal intensity distribution. It can further be utilised for mask generation. The magnitude SNR achieved with this method is very high compared to other methods [see Roemer et al., 1990].

Beware, that the receive profiles are ignored entirely in this approach. The phase of the receive profile is discarded by multiplication with the complex conjugate, while the magnitudes of the profiles are included into the sum:

$$A_{\text{SSQR}} = \sqrt{\sum_{c=1}^n S_c^* \mathcal{R}_c^* \mathcal{R}_c S_c} = \sqrt{\sum_{c=1}^n A_c^2 A_{\mathcal{R}_c}^2} \quad . \quad (10.7)$$

Fig. 10.2 shows an example. When talking about evaluation with SSQR recombination, one usually implies that the combined phase is, at the same time, evaluated with the complex sum.

10.5 Smoothed Magnitude Weighting with Constant Phase Correction

In general, one cannot expect the different receive channels to possess the same relative phase. This means, at least constant phase offsets have to be determined for each channel before using a method such as the *complex sum*. Hammond et al. [2008] introduced a technique that estimates and corrects for the individual receive profiles and removes constant phase shifts between the




Figure 10.2: Numerical example for the SSQR. Only the magnitude is shown. The simulation is identical to Fig. 10.1. The SSQR magnitude represents the source data well. The phase is discarded.

channels using a central reference point. However, this method employs complex summation and expects the image centre to exhibit representative phase values. The approach presented in the following determines mutual channel phase offsets in regions of high signal intensity and excludes values below a certain noise threshold from the summation, employing smoothed magnitude images as rough sensitivity estimates.

The resulting phase from the *complex sum method* is strongly influenced by the magnitude contrast. In regions with very low SNR, the magnitude reaches near the noise level. Here, the phase of the complex sum becomes relatively arbitrary due to the high noise contribution of the individual channels. For the magnitude the SSQR method offers an advantageous means to partially compensate for this shortcoming. The magnitude of a single-channel signal, A_c , takes the role of a weighting matrix for the contribution of the exponential term, $e^{i\varphi_c}$, to the complex sum. Replacing the magnitude by the modulus of the receive profile, A_{R_c} , should avoid noise propagation in regions with low signals. A proper weighting matrix considers only voxels with sufficient signal level for the averaging procedure over all channels. If, further, the phase information of each channel is corrected for a potential relative phase offset of the channel, $\Delta\varphi_{o,c}$, the result of the complex sum becomes more robust and shows higher statistical stability. The phase determined using a weighted, offset-corrected complex sum can be expressed as:

$$\varphi \approx \angle \left(\sum_{c=1}^n W_c \cdot e^{i(\varphi_c - \Delta\varphi_{o,c})} \right) , \quad (10.8)$$

with W_c being the weighting matrix for the channel, c . Usually, the magnitude of an early echo in a GRE measurement offers highest signal, but little contrast, and therefore can serve as a rough estimate of each channel's sensitivity.¹ Applying a gentle smoothing filter² to the magnitude avoids the promotion of dominant local contrast and noise into the estimated profile:

$$A_{R_c} := \text{smooth}(A_{c,1^{st} \text{ echo}}) \quad (10.9)$$

Normalised profiles, p_c , are defined by windowing the modulus of the receive profiles to the interval $[0, 1]$. The upper limit is given by the maximum value of the profile, $a_{\max} = \max(A_{R_c})$, whilst the lower value is chosen individually as a fraction of the mean value, e.g. $a_{\min} = 0.2 \cdot \text{mean}(A_{R_c})$. It

¹This is sufficient for phase recombination. When focusing on the combined magnitude, one prefers to use the magnitude divided by the SSQR recombination as an estimate for the channel receive profiles.

²Such as Gaussian smoothing.

defines the noise threshold for signal windowing:

$$p_c = \begin{cases} 0, & A_{\mathcal{R}_c} < a_{\min} \\ (A_{\mathcal{R}_c} - a_{\min}) / (a_{\max} - a_{\min}), & a_{\min} \leq A_{\mathcal{R}_c} \leq a_{\max} \\ 1, & a_{\max} < A_{\mathcal{R}_c} \end{cases} \quad (10.10)$$

The windowing procedure ensures that for each channel only data above the noise threshold is considered for the sum. A channel weighting factor, v_c , is determined using the signal averages of the individual channels. It defines the contribution of the channels to the sum:

$$v_c = \frac{\text{mean}(A_{\mathcal{R}_c})}{\max_c(\text{mean}(A_{\mathcal{R}_c}))} \quad (10.11)$$

The weighting matrices, W_c , can now be expressed using the normalised profiles, p_c and the channel weighting factors, v_c , and by normalising the matrix with the sum of all factors:

$$W_c(\vec{r}) = v_c \cdot \frac{p_c(\vec{r})}{\sum_c v_c \cdot p_c(\vec{r})} \quad (10.12)$$

In order to determine the phase offsets, a small volume, $\mathcal{V}_c$, of continuous signal is selected as a reference to ensure statistical stability. The cuboid area is positioned around the maximum of the summed weighting matrices for the previous and the current channel:

$$\vec{r}_{\mathcal{V}_c} : \max_{\vec{r}} [W_{c-1}(\vec{r}) + W_c(\vec{r})] \quad (10.13)$$

For the first channel, $\vec{r}_{\mathcal{V}_c}$ is set to the maximum position of its weighting matrix, W_1 . This choice for the reference points ensures that the reference area has strong SNR and a stable phase distribution. The reference area, $\mathcal{V}_c$, commonly has a dimensioning between 5^3 and 10^3 voxels. The phase inside $\mathcal{V}$ is unwrapped for all channels – this can be done by MPPGI (see Section 4.5), as the area is cuboid and the algorithm offers the shortest computing time:

$$\varphi_{\text{unw},c}[\text{ijk}] = \mathcal{U}(\varphi_c[\text{ijk}]), \quad [\text{ijk}] \in \mathcal{V}_c \quad (10.14)$$

The phase differences of the unwrapped reference regions for consecutive channels are determined:

$$\begin{aligned} \Delta\varphi_1 &= 0 \\ \Delta\varphi_{c>1} &= \frac{\sum_{[\text{ijk}] \in \mathcal{V}} (\varphi_{\text{unw},c}[\text{ijk}] - \varphi_{\text{unw},c-1}[\text{ijk}])}{n_{\mathcal{V}_c}}, \end{aligned} \quad (10.15)$$

where $n_{\mathcal{V}_c}$ is the number of voxels in $\mathcal{V}_c$. This results in the phase offsets with respect to the first channel:

$$\Delta\varphi_{o,c} = \sum_{i=1}^c \Delta\varphi_i \quad (10.16)$$

Thus, the weighting matrices and the phase offsets are well-defined and the combined phase can be determined:

$$\varphi_{\text{comb}}(\vec{r}) = \angle \left(\sum_{c=1}^n W_c(\vec{r}) \cdot e^{i(\varphi_c(\vec{r}) - \Delta\varphi_{o,c})} \right) = \angle \left(\frac{\sum_c v_c \cdot p_c(\vec{r}) \cdot e^{i(\varphi_c(\vec{r}) - \Delta\varphi_{o,c})}}{\sum_c v_c \cdot p_c(\vec{r})} \right) . \quad (10.17)$$

An improved estimate for the magnitude can also be derived, considering the calculated per-channel phase offsets:

$$A_{\text{comb}}(\vec{r}) = \sum_c A_c(\vec{r}) \cdot e^{i(\varphi_c(\vec{r}) - \Delta\varphi_{o,c})} . \quad (10.18)$$

The estimate, A_{comb} , has characteristics comparable to the sum-of-squares, since it ideally represents the sum of the magnitudes of all channels (assuming perfect correction for phase offsets). Of course, this estimate could also be achieved by calculating, $A_{\text{abs}} = \sum_c |A_c|$, directly, yet it gives a good indicator for the robustness of the recombined phase estimate. The presented method will further on be called Smoothed Magnitude Weighting (SMW). The method is suited for the recombination of datasets that exploit constant phase offsets for each channel. An illustration of the SMW scheme is shown in Fig. 10.3 and an example is illustrated in Fig. 10.4.




Figure 10.3: Scheme of the SMW method: The raw data (left) are used to calculate weighting profiles; reference points are determined for consecutive channels based on the weighting profiles, in order to determine mutual phase offsets; the corrected phase and the weighting profiles are finally combined.

10.6 Spatially Varying Phase Offsets

It was mentioned above that even the constant phase offsets for the individual channels are an idealistic assumption. In reality and especially for increasing field strength such as 9.4 T, the phase offset encountered between two channels develops a spatial distribution. This means that the receive profile no longer just exhibits a channel-dependent phase difference, $\Delta\varphi_{o,c}$, but a phase offset dependent on the relative location, $\varphi_{\mathcal{R}_c}(\vec{r})$. For high field strength, the wave length used for imaging is of the order of the head dimensioning. Hence interferences evolve, distorting the phase and the modulus of the individual receive profiles. As expressed in Eq. 10.1, each receive element, c , records the emitted signal from position $\vec{r}$ with a modified magnitude, $A_c(\vec{r})$, and




Figure 10.4: Numerical example for the SMW. The upper row shows the magnitude, the lower row shows the phase. For the artificial geometry (left), four channels were simulated, exhibiting different magnitude sensitivities and an individual constant phase offset (see channel phase plots). The complex sum is shown on the right hand side in comparison to the SMW. The complex sum fails and results in a magnitude image with strong inhomogeneities and a corrupted phase, showing a vertical fringe in its middle. Whereas the SMW generates a good estimate of the source magnitude and phase. The magnitude seems to be more homogeneous than in the examples in Fig. 10.1 and Fig. 10.2.

phase, $\varphi_c(\vec{r})$, both dependent on the position. Considering further phase offsets caused by inhomogeneous excitation³, φ_0 , yields the following model for a particular voxel:

$$A_{\text{obs}} \cdot e^{i\varphi_{\text{obs}}} = A_{\mathcal{R}_c} A \cdot e^{i(\varphi + \varphi_{\mathcal{R}_c} + \varphi_0)} \quad (10.19)$$

For phase imaging, the additional offsets have to be removed. A linear evolution of phase with time is assumed (see Eq. 3.7). When one is interested only in the evolution of phase, it is of no importance, if phase evolution is observed with respect to the excitation time or with respect to any other time point (echo time) after excitation. This implies, that any kind of voxel- or channel-dependent phase offsets can be removed by echo difference calculation, because the difference remains invariant when experiencing absolute phase offsets:

$$\frac{\Delta\varphi}{\Delta t} = \frac{\varphi_{2,\text{obs}} - \varphi_{1,\text{obs}}}{t_2 - t_1} = \frac{(\varphi_2 + \varphi_{\mathcal{R}_c} + \varphi_0) - (\varphi_1 + \varphi_{\mathcal{R}_c} + \varphi_0)}{t_2 - t_1} = \frac{\varphi_2 - \varphi_1}{t_2 - t_1} \quad (10.20)$$

where $\varphi_{*,\text{obs}}$ are the observed and φ_* are the true phase values for a voxel. The underlying principle was already discussed for fieldmap estimation in Section 6.2. While for fieldmap estimation the difference is calculated for the combined phase data, here the individual channels are corrected before the they are combined. Thus, if a double or multiple echo dataset was recorded, one can obtain the true phase evolution by performing the following calculations for the signals and echo

³Beware that excitation offsets are experienced simultaneously by all channels.

times of each channel:

$$\begin{aligned}
 S'_c(\vec{r}, TE_i) &= S_c(\vec{r}, TE_i)/S_c(\vec{r}, TE_1) \\
 &= A_c(\vec{r}, TE_i)e^{i\varphi_c(\vec{r}, TE_i)}/A_c(\vec{r}, TE_1)e^{i\varphi_c(\vec{r}, TE_1)} \\
 &= \frac{A_{c, TE=i}}{A_{c, TE=1}} \cdot e^{i(\varphi_c(\vec{r}, TE_i) - \varphi_c(\vec{r}, TE_1))}
 \end{aligned} \tag{10.21}$$

$$TE'_i = TE_i - TE_1 \tag{10.22}$$

(see also Section 6.4). The magnitude resulting from this calculation – a relative value with respect to the first echo – is not considered for phase imaging and can thus be discarded during calculation. The resulting phase can be enhanced in stability by slightly smoothing the complex signal of the first echo at TE_1 (using a Gaussian with a standard deviation of about one voxel width). Smoothing suppresses noise in the early echo, while preserving pronounced phase offsets – local phase contrast is much more emphasised for late echo times, as it increases (linearly) over time.

This principle is suitable for processing of ultra-high field data with inhomogeneous phase offsets. The offset-corrected phase data can be recombined using SMW. This approach is the basis to the method presented in the next section (Section 10.7) – an example is thus given in the very same section. For further reading, see also Bernstein et al. [1994] or Ros et al. [2008]. Furthermore, Chen et al. [2010] presented a recombination technique based on (smoothed) complex receive profiles that employs regularised recombination of multichannel data.

10.7 Downsampled Offset Correction

A related approach is the estimation of the actual offset of the phase at excitation time, $t = 0$. This requires linear regression fitting (as for fieldmap calculation), which usually implies that the phase data have to be spatially unwrapped. Considering n_c separate channels and n_e echoes, $n_c \times n_e$ unwraps have to be performed. Spatial unwrapping is very time-consuming, thus it is beneficial to temporarily reduce the data resolution. Downsampling is further reasonable, when zero offsets are assumed to contain relatively smooth variations throughout the sample. Each channel is resampled to 1/2 or 1/4 of the native resolution *a priori*. FEST is applied to the downsampled data. Depending on the characteristics of the phase, either unwrapping in time or unwrapping in the spatial domain, e.g. using URSULA, PRELUDE [Jenkinson, 2003] or QBU, is performed as described in Section 6.3.2. FEST uses the downsampled and unwrapped phase, $\varphi(\vec{s}, t)$, to calculate phase offset maps, $\varphi(\vec{s}, 0)$, for each channel by linear regression while discarding the calculated fieldmap. The phase offsets in native resolution are determined by interpolating the low-resolution offset map up to full resolution, $e^{i\varphi_c(\vec{s}, 0)} \rightarrow e^{i\varphi_c(\vec{r}, 0)}$.⁴ Subtracting the offset maps from each channel results in the offset-corrected phase:

$$\varphi'_c(\vec{r}, t) = \angle e^{i(\varphi_c(\vec{r}, t) - \varphi_c(\vec{r}, 0))} \tag{10.23}$$

The channels can then be combined e.g. by using the complex sum. Sometimes, constant offsets might occur between the channels, so the corrected data are ideally recombined using the SMW

⁴Linear interpolation is usually sufficient for this task.

introduced above (see Eq. 10.18 and Eq. 10.17). Fig. 10.5 shows a numerical example for the Downsampled Offset Correction (DownCor) method. A technique based on the same principle was introduced by [Robinson et al., 2011].




Figure 10.5: Numerical example for the DownCor method. The upper row shows the first echo, the second row shows the second echo phase. For the artificial geometry (left), four channels were simulated, exhibiting the same magnitude sensitivities as e.g. in Fig. 10.1, but a spatially varying phase offset was added, differing for each channel. The third row shows the $t = 0$ phase offsets estimated for the downsampled map (half resolution), whilst the lower row shows the true simulated phase offset maps. On the right hand side, the simple SMW algorithm is compared to an SMW with a priori offset-correction (DownCor). While SMW generates a strongly disturbed phase, the offset correction is able to restore the true phase evolution (in both echoes).

10.8 Per-Channel Field Estimation

Another form of offset-corrected recombination is the *per-channel field estimation*. The method determines individual fieldmaps from the data of each channel and estimates a weighted combination. Downsampling is not employed, since for this method as much phase information as possible shall be preserved and exploited. The smoothing and windowing operation from the SMW is used to obtain a phase validity mask, m_c , for the individual channels. A manually created mask can additionally be applied to define the valid volume for recombination. For each channel, then, a complete fieldmap is estimated inside the according validity mask (using FEST):

$$[\varphi_c(\vec{r}, t), m_c, TE_i] \longrightarrow b_c(\vec{r}) \quad (10.24)$$

The parametrisation of FEST depends on the type of the dataset and on the problem size. Thresholding for the per-channel SMW masks is performed more aggressively than for normal brain mask. This avoids unnecessarily long unwrap times due to inclusion of phase noise. The fields of the channels are finally combined using the channel weighting and the sensitivity profiles determined with SMW:

$$b(\vec{r}) = \frac{\sum_{c=1}^n v_c \cdot p_c(\vec{r}) \cdot b_c(\vec{r})}{\sum_{c=1}^n v_c \cdot p_c(\vec{r})} \quad (10.25)$$

The combined field represents the weighted average of the independently estimated fields of the individual channels. An example for the per-channel field estimation is shown in Fig. 10.6. In the example, the SMW recombination allows one to calculate a rough estimate for the true fieldmap. This is due to the fact, that echo-differences are preserved in all operations (see Eq. 10.20), even in the complex sum. Nevertheless, SMW promotes noise and erroneous field values in areas near phase poles. The per-channel field estimation method avoids this problem. The fields generated for each channel suffer much less from phase poles, since these are particularly generated during incorrect recombination of channel phase data. The presented approach is suited for and benefits from multiple echo acquisition. Robinson and Jovicich [2011] presented a double-echo recombination principle related to this technique.




Figure 10.6: Numerical example for per-channel field estimation. The upper row shows the magnitude, the second row shows the first and the third row the second echo phase. The fourth row shows the fieldmap estimated for each dataset, channel and method, using FEST with temporal unwrap. For the artificial geometry (left), four channels were simulated, using the same spatially varying phase profiles as in Fig. 10.5. The result of the SMW is shown in comparison with the per-channel field estimation. SMW fails dramatically in magnitude and phase recombination. Yet, the fieldmap created with the SMW data is partially correct, but includes noisy areas and streaks (red arrows) due to the incorrect phase recombination. The per-channel field estimation shows a consistent fieldmap, constituting of the fields from each channel in the areas where these are stable.

10.9 Comparing the Methods

When analysing phase data, one desires to employ a channel recombination method that generates robust results, but is also time-efficient. As long as the data fulfil the preconditions of e.g. the complex sum – meaning the absence of constant or spatially varying phase offsets in the channel data – one can use this method, since it is the quickest one. However, this is rarely the case. The SMW method is a quite robust choice that compensates for eventual phase offsets between the coil channels and offers a good compromise for everyday use. Further, SMW is utilised as part of the advanced methods for inhomogeneous phase offsets. When working with higher field strengths, it is advisable to pick a method such as DownCor or per-channel field estimation. Though both methods employ per-channel unwraps and estimate linear regression fits, DownCor is more time efficient, since it computes downsampled data. A downsampling factor of 1/4 of the source resolution implies massive speed-up for the unwrapping algorithm and thus for the entire offset estimation – the problem size is reduced to 1/64. Nevertheless, downsampling always leads to a loss in accuracy and to interpolation errors when resampling the offset data back to native resolution. When measurements contain phase contrast with high complexity, analysis at maximum accuracy is desired. Then, it is strongly advisable to choose the per-channel field estimation. Although the calculation of fieldmaps for each channel is highly time-consuming, one makes sure to keep the full resolution and all information contained in the measured dataset. In addition to that, novel techniques such as the URSULA algorithm presented in this thesis (see Chpt. 5) allow for significant speed-up and make the per-channel field estimation feasible within hours or minutes, even for high-resolution data.

10.10 Reconstruction for a 9.4T Example

In order to motivate the importance of recombination techniques considering spatially varying per-channel phase offsets, an example is presented in the following, illustrating a measurement at 9.4T. A *post mortem* brain was measured using a custom-built 16 channel array coil in parallel transmit mode. The protocol used for excitation optimisation and the measurement were planned and performed by *Dr. Michael Poole* and *Dr. Desmond Tse*⁵, who generously shared their measurement data. The applied sequence was a multiple echo GRE including the following parametrisation:

Sequence Parameters					
seq	GRE	exc	non-selective	Tx/Rx	16 ch array
matrix	$384 \times 312 \times 240$	FOV	$[265 \times 215 \times 166] \text{ mm}$	res.	$[0.69 \times 0.69 \times 0.69] \text{ mm}$
TR	35 ms	α	20° (nominal)	RF-spoil	on
TE	$[4.44, 11.1, 17.76, 24.42] \text{ ms}$			BW	270 Hz/Px
iPAT	2				

The channels were saved uncombined for post processing. Recombination was performed using the Complex Sum Recombination (CMPLX), the SSQR, simple SMW and DownCor. Further, *per-channel field estimation* was compared to the field calculated from the outcome of the DownCor algorithm. A mask was generated using ITK Snap [Yushkevich et al., 2006] to include only relevant regions in the results section.

⁵Both: Institute for Neuroscience and Medicine – 4 (INM-4), Forschungszentrum Jülich, Germany.




Figure 10.7: Sample slice of all acquired channels shown from left to right (channel 1 – 16). The top row shows the magnitude, while the bottom row shows the phase.

10.10.1 Results

The obtained multichannel data are shown for a sample slice in Fig. 10.7. For all following illustrations, the second echo of the recombined (phase and magnitude) data is shown. The CMPLX results are shown in Fig. 10.8. The resulting magnitude is very inhomogeneous, showing strong destructive interferences of the summarised complex channel signals. The phase also shows unordered characteristics, including numerous poles. Unwrapping of this phase, even of phase echo-differences would be difficult. The simple SSQR algorithm results in the magnitude data shown in Fig. 10.9. For completeness the CMPLX phase is shown here again, as it is usually calculated when using SSQR. The magnitude is relatively homogeneous, except for an obvious blackened area where the signal excitation seems to possess a minimum. The SMW recombination leads to slightly better results, shown in Fig. 10.10. The magnitude is more homogeneous than in the CMPLX case, but still significantly less homogeneous than for the magnitude-optimised SSQR. Yet, the channel-dependent sensitivity profiles lead to a much more robust phase estimate than in the previous methods. Still, several phase poles can be encountered, but the overall appearance of the phase data shows a higher degree of organisation and might even be usable for a fieldmap estimation. The DownCor algorithm offers the best results for magnitude and phase, shown in Fig. 10.11. The corrected phase data does not contain poles and appears to be smoother than in the other recombination attempts. The magnitude is as homogeneous as for the SSQR recombination. Finally the per-channel field estimation is compared to the field derived from the recombination result of the DownCor algorithm, as shown in Fig. 10.12. The background-corrected local field shows that both algorithms lead to approximately identical fieldmaps, except for a constant offset, visible in the raw fields. Fig. 10.13 illustrates the local variance of the phase (the LV measure from Section 4.7.1) in a $3 \times 3 \times 3$ neighbourhood for a sample slice. SMW apparently removes several artefacts and the overall variance level is slightly lower than for the simple CMPLX recombination. With DownCor most artefacts are removed and the background variance level appears significantly lower than for SMW. The remaining artefacts seem to origin in other sources than incorrect channel matching.

10.10.2 Discussion and Conclusion

The example clarifies that for multichannel data obtained at ultra-high field strength conventional recombination schemes do not provide a sufficient basis for further phase analysis. Even the SMW, designed for constant-offset-compensation, can only partly compensate for the spatially varying phase offsets encountered for each channel. DownCor and per-channel field estimation seem the

only applicable algorithms for such data. The choice of the utilised method may depend on the particular problem set. Per-channel field estimation might be preferred, as it computes data at full resolution, while DownCor requires down- and upscaling, inevitably leading to interpolation errors that could propagate into the offset estimation. On the other hand, DownCor requires less computational effort, since all unwrapping and estimation operations are performed at half or quarter resolution, whereas the inaccuracy introduced by this is still comparably low. However, the techniques are closely related and thus both fit for processing ultra-high field data.




Figure 10.8: Multi-slice view of the result obtained by CMPLX recombination. The magnitude is shown on the left, the phase is shown on the right hand side. Every 12th slice is included.




Figure 10.9: Multi-slice view of the magnitude obtained by SSQR recombination (left). The phase obtained from CMPLX is shown on the right hand side. Every 12th slice is included.




Figure 10.10: Multi-slice view of the result obtained by SMW recombination. The magnitude is shown on the left, the phase is shown on the right hand side. Every 12th slice is included.




Figure 10.11: Multi-slice view of the result obtained by DownCor recombination. The magnitude is shown on the left, the phase is shown on the right hand side. Every 12th slice is included.




Figure 10.12: Multi-slice view of the fieldmap generated by DownCor (left column) and per-channel field estimation (right column). The top row shows the raw fieldmap and the lower row shows the local field after subtracting the GF-filtered map with $\sigma = 6$ from the raw field. Every 12th slice is included.




Figure 10.13: Phase (upper row) and LV (lower row, as in Eq. 4.20) with a $3 \times 3 \times 3$ neighbourhood for a sample slice of the second echo – from left to right: CMPLX, SMW and DownCor. Yellow arrows indicate phase pole disturbances removed in SMW, red arrows indicate disturbances removed with DownCor and blue arrows show remaining phase corruptions.

11

Applications

The purpose of this chapter is to illustrate several applications of the methods discussed within this thesis. It includes applications for the custom-developed algorithms as well as for various implemented methods contributing to the established phase imaging framework. Needless to say, that the entire process of phase, field and susceptibility imaging is still under further development. Yet, this chapter functions as a *proof of principle* for the proposed methods. The applications include a variety of *post mortem* and *in vivo* measurements at 3 T and 9.4 T.

11.1 General Workflow for Phase Imaging

Before single applications will be described, a general discussion shall outline the formal workflow of phase imaging. Following the order of this thesis, the workflow starts with consistent phase data. Just in the very last chapter (Chpt. 10), it is shown that at ultra-high field a proper recombination of the Rx channel signals is inevitable when employing array coils. The actual workflow thus begins with this step, even though it is optional for measurements at lower field strength.

Fig. 11.1 illustrates a full processing scheme for phase imaging. Starting with the channel recombination, the phase and magnitude data is initially prepared using either SMW, DownCor or per-channel field estimation as described in Chpt. 10 (Fig. 11.2). The recombination provides an exclusion mask, applicable in brain masking, that delineates regions of insufficient signal. Regardless of the chosen recombination method, magnitude data of all channels is additionally recombined using the SSQR method and then used to segment a brain mask by employing Bet2 (FSL framework, Smith [2002]) or ITK Snap [Yushkevich et al., 2006]. The final mask is the conjunction of exclusion and segmented mask. The corrected phase data are unwrapped and transformed into a fieldmap using FEST. The unwrapping algorithms used within FEST were described in Chpt. 4 and Chpt. 5 and mainly include PRELUDE [Jenkinson, 2003], URSULA/P-URSULA and MPPGI. The estimated fieldmap is then corrected for background fields using MUBAFIRE. Although the parametrisation of this algorithm may vary slightly, the filter chain usually keeps the form described in Chpt. 9. Depending on the properties of the data under investigation, the MUBAFIRE Local




Figure 11.1: General phase imaging workflow. White boxes imply input or result data, whereas grey boxes symbolise operations and algorithms.




Figure 11.2: Phase imaging workflow for per-channel field estimation. White boxes imply input or result data, whilst grey boxes symbolise operations and algorithms.

filter is included into the background-correction step or applied as a separate filter. Fieldmaps are then either directly inspected or used for susceptibility reconstruction. Susceptibility is estimated using the FSE framework. Judging by observations made within the presented studies, combined

Tikhonov and gradient regularisation yields the best results for measured data.

The workflow illustrated in Fig. 11.1 applies for the case of a direct recombination of the measured magnitude and phase. The method of *per-channel field estimation*, however, estimates a fieldmap for all channels and thus unifies recombination with the FEST step that is otherwise performed separately. Here, the SSQR recombination is inevitable, since the field-based recombination method does not provide a combined magnitude map. The sum-of-squares magnitude can serve for brain-mask estimation. The general workflow for per-channel field recombination is illustrated in Fig. 11.2. Besides the application in the background correction steps and in the susceptibility estimation, the brain mask also serves as a confinement for the volume to be recombined in the *per-channel field estimation*. Unwrapping can be sped up significantly with strict mask confinement and thus exclusion of noisy areas. The rest of the framework is identical to Fig. 11.1. The particular configuration of the algorithms will be indicated for each application, individually. Further, differing evaluation procedures will be remarked.

11.2 Subjects and Specimen

All studies presented within this thesis were performed in compliance with all ethical requirements and standards. Volunteers and patients were only scanned with informed consent to the procedure and to anonymous usage of the acquired data. The *post mortem* brain measurements presented in this chapter and in former occasions were performed on brains acquired from the brain donor programme of the University of Düsseldorf, Germany.

11.3 Post Mortem Brain at 9.4T

In order to motivate the multichannel recombination techniques discussed in Chpt. 10, and at the same time to demonstrate the functionality of the phase imaging framework including susceptibility reconstruction, a *post mortem* brain measurement at 9.4 T was analysed. The primary objective was to determine the background-corrected fieldmap and to successfully apply susceptibility reconstruction to it. The actual evaluation was performed using *per-channel field estimation*. Besides, a comparison to results obtained with a simple SSQR recombination was facilitated.

11.3.1 Material and Methods

The data were obtained in a GRE measurement at a 9.4 T human scanner (Siemens AG, Erlangen, Germany) using the following parametrisation:

Sequence Parameters					
seq	GRE	exc	non-selective	Tx/Rx	8 ch array
matrix	$312 \times 384 \times 240$	FOV	$[215 \times 265 \times 166] \text{ mm}$	res.	$[0.69 \times 0.69 \times 0.69] \text{ mm}$
TR	35 ms	α	20° (nominal)	iPAT	2
BW	270 Hz/Px	TE	$[4.44, 11.1, 17.76, 24.42] \text{ ms}$		

Excitation was performed using parallel transmission with a customised phase scheme developed by Dr. Michael Poole and Dr. Desmond Tse¹. It included two acquisitions at two different,

¹Both: Forschungszentrum Jülich, Germany

optimised phase settings. Recombination was hence performed on 2×8 (real) channels, i.e. 16 virtual channels. A quick recombination was approached using the SSQR method, whereas advanced analysis was performed using per-channel field estimation. For both recombination techniques, FEST was configured to use URSULA as back-end. Unwrapping was applied on the first three echoes. A brain mask was derived using ITK Snap [Yushkevich et al., 2006]. Both recombined datasets were analysed with the following MUBAFIRE parametrisation:

MUBAFIRE Parameters				
Filters:	POLF $\rightarrow$ SPHINX $\rightarrow$ DIPF			
resample	1 (half)	POrder	1	SPorder 4
DIPF iterations	50			

Further, MUBAFIRE Local was applied using the parameters listed below:

MUBAFIRE Local Parameters				
Filters:	Thresholding $\rightarrow$ erode diag. neighbours $\rightarrow$ 6 neighbour erosion $\rightarrow$ 6 neighbour dilation $\rightarrow$ DIPF			
resample	0 (full)	DIPF iterations	50	THwidth 8

Finally, susceptibility was reconstructed using FSE with the following parametrisation:

FSE Parameters					
regmode	0	W	brain mask (m)	W_T	brain mask (m)
λ	0.03	μ	0.01	B_0	9.38 T
spat	true	iterations	50	reset	20
$W_{Gx(y,z)}$	brain mask (m)				

11.3.2 Results and Discussion

An overview on the data processing steps is illustrated in Fig. 11.3. The fieldmap, the MUBAFIRE correction and the susceptibility reconstruction can successfully be derived. The line plot shows, how the susceptibility reconstruction outlines the cortical gyri that were hidden in the distorted field, before. As later applications will show, the grey/white matter contrast observed in these gyri structures is much more pronounced for fixed brain than for *in vivo* measurements. The overall reconstruction of the brain is further illustrated in Fig. 11.4. The magnitude exhibits sharp edges, but also varies smoothly throughout the volume. Field and susceptibility show a weaker contrast, but a better homogeneity and better noise behaviour. Although the gyri structure is only visible as an “oscillating distortion” in the background-corrected fieldmap, the reconstructed susceptibility map shows precise gyri geometry and other structures. The visible anatomy appears to be related to that pronounced in the magnitude, though susceptibility is exhibiting a more homogeneous intensity distribution, whilst showing smoother edges.

Fig. 11.5 illustrates a comparison between reconstruction with the SSQR method and per-channel field estimation. Slices for the fieldmaps, the background-corrected field and the susceptibility are shown. The fieldmap obtained by the SSQR method is obviously subject to strong phase inconsistencies and thus shows various artefacts including local noise propagation and smooth inhomogeneities. The fieldmap from per-channel field estimation, however, shows homogeneous characteristics. The observed differences become even more pronounced in the background-corrected field. Background-correction based on the SSQR recombination is obviously not capable




Figure 11.3: Reconstruction process – from left to right: Fieldmap, MUBAFIRE correction, MUBAFIRE Local correction, susceptibility map and line plots of the local field and the susceptibility, indicated by yellow lines in the sample slice. Whilst the field shows long-range variations, the line plot of the susceptibility outlines distinct contrast for the gyrus (plateaus).





Figure 11.4: Magnitude (top), MUBAFIRE-corrected field (left) and reconstructed susceptibility (right) in volumetric plots showing orthogonal sample slices. The magnitude originates in SSQR recombination, whilst field and susceptibility are derived by per-channel field estimation.

of correcting for all external sources, just because the inhomogeneities partly originate in an incorrect recombination of the coil channels. Susceptibility variations in local contrast can roughly be estimated with the SSQR method. Yet, in comparison to the per-channel field estimation the SSQR-derived susceptibility distribution seems noisy and bears pronounced long-range variations.

Figure 11.5: The SSQR method (top row) shows artefacts in the fieldmap (left column), especially areas with strong noise amplification (red arrow). The background-corrected SSQR fieldmap (second column) exhibits strong inhomogeneities due to incorrect recombination (yellow arrows). Susceptibility (last column) inherits the inhomogeneities and areas of amplified noise. The results of the per-channel field estimation (lower row) are more robust.



11.3.3 Conclusions

The illustrated example gives good reasoning, why advanced phase recombination techniques for multichannel data at ultra-high field strength are inevitable. SSQR recombination is not applicable for reliable field mapping and susceptibility reconstruction in this context. The proposed reconstruction chain, including the algorithms presented in former chapters, performs well on the measured sample data of *post mortem* tissue. Distinct susceptibility contrast can be reconstructed using MUBAFIRE (Local) and FSE on the data obtained by per-channel field estimation. A great advantage can be found in the homogeneity of the field and susceptibility data in comparison to the magnitude that exhibits relatively inhomogeneous intensity.

11.4 Visualising Brain Tumours

Diagnosis, investigation and classification of brain tumours is a central field of research for medical imaging modalities such as MRI. It has been reported that also phase and susceptibility can provide structural or classifying information about tumours [e.g. Rauscher et al., 2005; Sehgal et al., 2006; Deistung et al., 2013]. Within the scope of this thesis, it was fortunately possible to obtain phase data from a cooperation study between the University Hospital of Cologne and the Jülich Research Centre. The objective of this collaboration was to acquire a broad range of imaging information, including PET Positron Emission Tomography (PET) [e.g. Weckesser et al., 2005] and several MRI modalities. A multi echo GRE sequence was part of the measurement, and could thus be used for phase imaging. Due to strong measurement time restrictions (the entire measurement protocol included e.g. several anatomic sequences as well as fMRI and Diffusion Tensor Imaging (DTI)) a sliced GRE had to be applied instead of a desirable 3D sequence. Furthermore, the applied GRE was configured for water content mapping, so the sequence parametrisation was unfortunately not optimised for phase imaging (see Section 6.6).

11.4.1 Material and Methods

Measurements took place on a 3 T MR-PET hybrid human scanner (Siemens AG, Erlangen, Germany). The sliced GRE sequence included the following parametrisation:

Sequence Parameters					
seq	GRE	exc	slice-selective	Tx/Rx	1 ch/8 ch array
matrix	$192 \times 144 \times 60$	FOV	$[220 \times 165 \times 135]$ mm	res.	$[1.145 \times 1.145 \times 2.25]$ mm
TR	10000 ms	α	90°	slc.thickness	1.5 mm
iPAT	2	BW	320 Hz/Px	Readout	bipolar
TE	[3.73, 7.74, 11.75, 15.76, 19.77, 23.78, 27.79, 31.8, 35.81, 39.82, 43.83, 47.84] ms				

The number of phase encode steps (144) may vary with the head volume of the subject. All scanned subjects were patients, diagnosed with different kinds of brain tumours. The scans were only acquired at 1.145 mm in-plane resolution, but still suffered from increased patient motion, attributable to the disease. Brain masks were derived after smoothing the magnitude at the first echo (by a Gaussian filter) and applying an adequate threshold. A standard deviation of 4 voxels was used in the smoothing operation, and the threshold used for excluding lower values was individually adapted to each dataset (order of $A = 700 - 800$, in the arbitrary units of the dataset). FEST had to be executed in echo-separation mode (separating even and odd echoes), since the acquisition was performed in bipolar mode. FEST performed satisfyingly in temporal unwrapping mode thanks to the dense echo spacing. Furthermore, automatic masking of voxels with an offset at $t = 0$ that significantly deviates from $\varphi_0 = 0$, was performed by FEST. In order to pass isotropic data to the susceptibility reconstruction programme, the derived fieldmaps were resampled using spline interpolation along z direction to resize the z voxel scaling from 2.25 mm to the in-plane resolution² of 1.145 mm. The same rescaling operation was applied to the masks with nearest-neighbour interpolation. The isotropically sampled field and mask were then background-corrected using MUBAFIRE with a minimal parametrisation including the following parameters:

MUBAFIRE Parameters					
Filters:	POLF $\rightarrow$ SPHINX $\rightarrow$ Thresholding $\rightarrow$ DIPF				
resample	0 (full)	POLorder	1	SPorder	4
DIPF iterations	100	THWidth	6		

Thresholding was chosen to be more aggressive, as the observed standard deviation in most maps still showed strong variations – beware that no DIPF is applied before thresholding. Finally, FSE was applied on the background-corrected fieldmaps, yielding the underlying susceptibility distribution with the following parametrisation:

FSE Parameters					
regmode	2	W	brain mask (m)	$W_{G_{x(y,z)}}$	brain mask (m)
λ	0.03	spatial dipole	true	B_0	3.0 T
μ	0.0005	iterations	50	reset	20

Needless to say, the reconstruction algorithm suffered from the under-sampled slicing (in z direction). Besides this issue, many of the patients already underwent surgery and thus had implanted

²Currently, this is no longer necessary.

ferromagnetic screws and other objects disturbing the magnetic field in a way that prevents proper susceptibility reconstruction in the vicinity of these implants. The contrasts illustrated for the measurement examples were coregistered with Flirt [Jenkinson and Smith, 2001; Jenkinson et al., 2002].

11.4.2 Results

An example of the workflow for phase imaging and susceptibility reconstruction algorithms applied on the sliced GRE data is illustrated in Fig. 11.6. Besides the expected problems in reconstruction




Figure 11.6: Workflow for phase, field and susceptibility imaging on the data acquired from the brain tumour patients. From left to right: Phase and Magnitude of the 8th echo, fieldmap, background-corrected fieldmap and susceptibility distribution. The strong edge artefacts visible in the local fieldmap and in the susceptibility map originate in the under-sampled z direction (slicing direction) and the interpolation to obtain isotropic images.

arising from the under-sampling and the interpolation of the slicing direction, another systematic artefact can be recognised. The phase shows slice-wise varying offsets. Unfortunately, this problem cannot be compensated for when calculating the fieldmap, but remains visible (see Fig. 11.7). No choice of the FEST modes is capable of eliminating these slice artefacts. Nevertheless, the




Figure 11.7: Phase-offset artefact between slices, illustrated on phase (left), background-corrected field (middle) and susceptibility reconstruction (right). The offsets are visible as horizontal lines. The artefact prohibits detailed reconstruction of susceptibility.

quality of the derived background-corrected fieldmaps suffices to generate susceptibility maps with significant contrast for several patients. Fig. 11.8, Fig. 11.9, Fig. 11.10 and Fig. 11.11 illustrate results for selected patients exhibiting different contrasts for the tumours. The figures are outlining the magnitude of the 6th echo, the fieldmaps and susceptibility distributions as well as the FET-PET images for the tumour region. The diagnosis data, gathered without relying on

susceptibility information, were generously provided by *Dr. Carolin Weiss*, Univeristätsklinikum Köln.




Figure 11.8: From left to right: Magnitude of 6th echo, field, susceptibility, PET activity. Diagnosis: Irregularly shaped, left precentral primary glioblastoma (WHO grade IV) with perifocal edema, mass effect and strong contrast enhancement at the tumour margins. In the susceptibility map the tumour surface appears brighter.




Figure 11.9: From left to right: Magnitude of 6th echo, field, susceptibility, PET activity. Diagnosis: Drop-shaped, right precentral primary glioblastoma (WHO grade IV) with atypically strong and rather homogeneous contrast enhancement, embedded necrotic zones and connection to the posterior pole of the lateral ventricle. The tumour is surrounded by a moderate perifocal edema. Although no bleeding is reported, the susceptibility distribution indicates high values for tumour and edema.




Figure 11.10: From left to right: Magnitude of 6th echo, field, susceptibility, PET activity. Diagnosis: Primary glioblastoma (WHO grade IV) of the right frontal lobe (precentrally located) with perifocal edema and mass effect lacking intense contrast enhancement. Large tumour bleeding originating from the frontally located tumour part. The susceptibility map outlines bleeding in the surface area and contrast between interior and exterior.




Figure 11.11: From left to right: Magnitude of 6th echo, field, susceptibility, PET activity. Diagnosis: Regularly shaped, round metastasis of an adenocarcinoma of the intestine, left precentrally located. The tumour shows reduced diffusion at the tumour margins and is surrounded by a remarkable perifocal edema. Here, susceptibility contrast outlines the tumour darker than the surroundings.

11.4.3 Discussion

The phase imaging workflow and the susceptibility reconstruction lead to expressive results. The major drawback concerning data acquisition are the slice offset artefacts, visible in phase, fieldmaps and background-corrected fieldmaps. The error can systematically be found in all measurements.

This suggests, that physiological noise such as body movement is an unlikely source. The assumption is that the artefacts result either in an inappropriate recombination of the channels or in a problem with the scanner imaging acceleration system, GRAPPA [Griswold et al., 2002]. Acceleration was used for acquisition, so it is indeed a potential source for the encountered artefacts. Nevertheless, compensation could not easily be accomplished. In case of a channel-recombination issue, it could be solved by post-processed channel recombination. Unfortunately, within this study the data were acquired only in combined mode, so this improvement is not applicable.

However, the information collected potentially provides a novel insight into the structure of various brain tumour types. The susceptibility contrast shows very distinct properties for the illustrated examples. Susceptibility reconstruction might evolve as an additional aid in diagnosing tumour tissue, haemorrhages and oedemas. The fact that despite the adverse measurement parameters the developed workflow for phase imaging leads to meaningful field and susceptibility maps confirms the robustness of the employed methods. Future studies on susceptibility contrast in brain tumours, e.g. on analysing its use for tumour classification, are a valuable perspective.

11.5 Susceptibility Reconstruction for a Study of Parkinson's Disease

In the scope of a study investigating MRI contrast mechanisms in Parkinson's disease (PD), carried out on a 3 T human scanner (Siemens AG, Erlangen, Germany) at the University Hospital of Aachen, the phase imaging workflow was tested on an extensive measurement series. The workflow, including unwrapping, fieldmap generation, background field correction and susceptibility reconstruction was customised for the data acquired within the study. As in the previous section, the measurements were not optimised for phase and susceptibility imaging, since this objective was added to the study at a later point of time.

11.5.1 Material and Methods

The measurement protocol included a 3D GRE sequence with the following parametrisation:

Sequence Parameters					
seq	GRE	exc	non-selective	Tx/Rx	1 ch/12 ch array
matrix	$192 \times 156 \times 128$	FOV	$192 \times 156 \times 128$ mm	res.	$1 \times 1 \times 1$ mm
TR	51 ms	α	8°	Flow Compensation	on
iPAT	2	BW	260 Hz/Px		
TE	[2.58, 6.57, 10.53, 14.52, 18.48, 22.47, 26.43, 30.42, 34.38, 38.37, 42.33, 46.32] ms				

Depending on the willingness of the subject³, the described GRE measurement was performed between one and three times in a row with identical positioning and parametrisation. The study included 30 Parkinson's patients as well as age- and gender-matched controls, of which 26 patients and 24 controls could be evaluated. The data from the receive array were recombined by the manufacturer-provided software (Adaptive Combine). As a consequence the recombined phase is not consistent everywhere. This is indicated by the presence of pronounced poles in the recombined

³His or her capacity to lie still for a long time.

phase data. Direct spatial unwrapping of the echoes is thus not feasible and the configuration for FEST has to be set to *delta phase mode* with half echo-count spacing ($nTE/2$). Although this circumvents most inconsistencies, the area around the pole still suffers occasionally from the acquisition mode, as an increased noise level and local field falsifications can be observed near the poles (see Fig. 11.12). Brain masks, generated using Bet2 (FSL framework, Smith [2002]), are used for unwrapping and as initial masks for the background correction.




Figure 11.12: The incorrectly recombined phase exhibits pronounced poles (A). The spatial unwrap of the separate echoes fails and generates yet other poles (B). When combining phase and time differences rather than only the phase itself for fieldmap generation (C), the result is robust, but includes local noise propagation in the close proximity of the pole. In the background-corrected fieldmap this effect becomes even more pronounced (D) whereas the rest of the map appears robust.

In order to enhance the statistical stability and to further reduce noise, the fieldmaps of all acquired measurements (up to three) are averaged. Global offsets between the fieldmaps, determined on the brain mask, m , are eliminated in the averaging process:

$$b_{comb}[ijk] = (1/3) \cdot \sum_{n=1}^3 b_n[ijk] - \overline{b_i(m)} - \overline{b_1(m)} \quad , \quad (11.1)$$

where (m) indicates summation over all true mask voxels, and $\overline{\dots}$ expresses the mean value. The derived fieldmaps are corrected for background variations using an integrated MUBAFIRE Local framework with the following parametrisation:

MUBAFIRE Parameters					
Filters:	POLF → SPHINX → DIPF → Thresholding → Erode diag. neighbours → 6 neighbour erosion → 6 neighbour dilation → DIPF				
resample	0 (full)	POLorder	1	SPorder	4
DIPF iterations	50	THWidth	18		

The MUBAFIRE Local part of the chain is necessary to remove artefacts generated by the non-consistently acquired phase near the phase poles (see Fig. 11.12). Susceptibility is reconstructed using the FSE framework with Tikhonov-regularised minimisation:

FSE Parameters					
regmode	1	W	brain mask (m)	W_T	brain mask (m)
λ	0.03	spatial dipole	true	B_0	3.0 T
iterations	50	reset	20		
padding	true	padsizes	0.5 · FOV		

The orientation data are extracted from all datasets in order to configure background correction and susceptibility reconstruction algorithms in each measurement with the true angle of the FOV towards B_0 . The entire processing scheme is illustrated in Fig. 11.13.




Figure 11.13: Processing scheme for the Parkinson's data. The essential difference to the standard procedure is the field averaging step.

For a number of cases an additional segmentation step was performed to facilitate comparison between the susceptibility of different anatomic regions. The *dentate nuclei* (DN), the RN and a part of the globus pallidus in the region of the *basal ganglia* (BG) were segmented manually, based on the reconstructed susceptibility maps using ITK Snap [Yushkevich et al., 2006]. The median value and the standard deviation of the susceptibility distribution inside the segmented regions for the selected patients and volunteers are compared with respect to the subject's age. Further, the Pearson correlation between the segmented areas and the age was examined for all subjects.

11.5.2 Results and Discussion

The majority of the reconstructed susceptibility maps shows sufficient contrast, especially within the central brain. An overview illustrating four sample datasets of controls and PD patients is shown in Fig. 11.14. Despite the suboptimal measurement protocol, the quality of the reconstructed susceptibility maps is relatively good. The level of stripe artefacts or blurring effects is low. A PD patient and a control are shown in orthogonal views in Fig. 11.15 and Fig. 11.16, outlining the robustness of the susceptibility maps in all directions.




Figure 11.14: Different stages of data processing for four PD patients and four controls. In each column one example is illustrated. The displayed slices are positioned in a similar region. The top row shows the 6th echo of the magnitude, the second row shows the 6th echo of the phase. Here the phase pole is well visible. The third row shows the field derived by phase difference calculation with applied brain mask, m . The fourth row shows the MUBAFIRE Local corrected field. In the last row the derived susceptibility distribution is illustrated. Recombination artefacts are exemplarily indicated by yellow arrows.

In some datasets strong motion artefacts, inconsistent phase or other sources for noise made further analysis impossible. Finally, 17 PD patients and 17 controls were chosen for the additional segmentation step. Segmentation is illustrated in Fig. 11.17 and the results are shown in Fig. 11.18. None of the regions shows a significant difference between PD patients and controls. Nevertheless, the RN and the DN show a tendency towards a reduced susceptibility in patients. The BG do not exploit any obvious difference between PD and controls. Based on this observation a fourth relation graph was generated, showing the difference of the median values for BG and RN for each subject. The separation in this graph is not clear, but indicates that PD subjects tend to show a slightly higher susceptibility difference between these regions. In the group of all subjects the correlation with the subject age is insignificant ($p \geq 0.2$) for RN, DN and BG-RN difference, but yielded a significant correlation for the BG with $p = 0.045$. The linear relation between BG susceptibility and age was estimated by a linear sum-of-squares fit as $8.3 \cdot 10^{-10} \text{ 1/year}$.




Figure 11.15: Background-corrected field (left side) and susceptibility (right side) in transverse (top) and sagittal (bottom) view for sample PD patient dataset. The nuclei of the basal ganglia are well pronounced.






Figure 11.16: Background-corrected field (left side) and susceptibility (right side) in transverse (top) and sagittal (bottom) view for sample control dataset. Local structures are pronounced as in Fig. 11.15

11.5.3 Conclusions and Outlook

The described study and the processing of the acquired data shows, that it is worthwhile to invest time in field and susceptibility reconstruction, even for data that was not specifically optimised for that purpose. Though in this measurement series e.g. the bandwidth is relatively high and no appropriate channel recombination technique was applied, the resulting signal could still be transformed into a widely robust fieldmap and finally into a susceptibility estimate for the volume.



Concerning the quantitative analysis, the results show a very slight tendency, but no statistical significance for differences between patients and controls. This could certainly be improved by performing measurements optimised for phase imaging and by employing a more robust segmentation algorithm. Segmentation in this study was not performed by a medical doctor, but by the author having only little medical background, so it is likely that the areas were not optimally covered. Yet, a correlation between the mean susceptibility in the BG and the subject age could be determined for joint patients and controls.

11.6 In Vivo Phase and Susceptibility Imaging at 9.4T

Ultra-high field imaging is still a highly experimental area of research. Scientists and engineers have approached successful image reconstruction by compensating for the additional pitfalls arising at field strengths such as 9.4 T. One crucial issue is, that the resonance frequency of water at 9.4 T is ~ 400 MHz. For technical RF applications this is a quite typical frequency. Yet, the primary disadvantage lies in the wavelength of radio waves with such frequency. The wavelength is about 0.75 m in vacuum and even less ($\sim 10^{-1}$ m) inside human tissue. This is the length scale of the human head. Thus, interference phenomena have to be considered at ultra-high field, whilst these were neglectable at lower field strength. The most dominant effects originating mainly in wave theory are inhomogeneous (magnitude and phase) excitation and the occurrence of spatially varying phase offsets on the receive side, as discussed in Chpt. 10. Phase imaging relies on properly estimated phase evolution, so an advanced channel recombination technique, such as the ones described in Section 10.7 or Section 10.8, is essential for successful processing.

11.6.1 Material and Methods

The measurement was prepared with *Dr. Michael Poole* and *Dr. Desmond Tse*⁴. A healthy volunteer of age 26 was scanned in a 9.4 T human scanner (Siemens AG, Erlangen, Germany) using an 8-channel Tx/Rx array coil. The utilised sequence was a standard GRE measurement using the following parametrisation:

Sequence Parameters					
seq	GRE	exc	non-selective	Tx/Rx	8 ch array, 7 ch Rx used
matrix	$216 \times 256 \times 160$	FOV	$[216 \times 256 \times 160]$ mm	res.	$[1 \times 1 \times 1]$ mm
TR	33 ms	α	20° (nominal)	RF-spoil	on
iPAT	2	Total t	8 min, 32 s	BW	260 Hz/Px
TE	$[4.65, 10.45, 16.25, 22.05]$ ms				

The acquired data were exemplarily recombined using simple CMPLX/SSQR recombination, in order to visualise the artefacts arising by using this technique. Further, the SSQR magnitude was used for creating segmentation masks with ITK Snap [Yushkevich et al., 2006]. For a proper phase combination, the DownCor technique, introduced in Section 10.7, was applied. The algorithm was parametrised as follows:

DownCor Parameters			
Resampling	1/2	Echoes used	4 (all)
FEST	$\Delta\varphi, n_{TE/2}, \text{PRELUDE}$	SMW smoothing	4
		SMW truncation	$0.3 \cdot \text{mean}$

Background correction was performed using MUBAFIRE with:

MUBAFIRE Parameters			
Filters:	POLF $\rightarrow$ SPHINX $\rightarrow$ DIPF		
resample	0 (full)	POOrder	1
DIPF iterations	100	SPOrder	4

⁴Both: Forschungszentrum Jülich, Germany

and afterwards MUBAFIRE Local, including the following parametrisation:

MUBAFIRE Local Parameters			
Filters:	Thresholding → Erode diag. neighbours → 6 neighbour erosion → 6 neighbour dilation → DIPF		
THwidth	6	DIPF iterations	100

From both results, b_{hom} and $b_{\text{hom,Local}}$, the susceptibility distribution was estimated using the FSE environment. The optimal parametrisation was found to be a combined Tikhonov- and gradient-regularised approach:

FSE Parameters					
regmode	2	W	brain mask (m)	W_T	1 everywhere
λ	0.06	μ	0.002	B_0	9.38 T
$W_{G_{x(y,z)}}$	m	iterations	25		

11.6.2 Results

From the eight Rx channels, one had to be excluded due to a high noise level, leaving seven channels for recombination. The steps of DownCor recombination are illustrated in Fig. 11.19. The overview image shows how essential an offset-sensitive recombination technique is for obtaining a proper phase result. The raw phase of every channel shows a variety of phase poles. The characteristics of the phase and the estimated offset maps change from channel to channel. The DownCor results of the channels 2 and 3 differ from the other channels at first glance. Yet, they only differ by a global offset that is compensated for by the SMW algorithm applied as the final DownCor recombination step. The CMPLX/SSQR recombination contains phase poles and characteristics that would hinder any unwrap attempt, whilst the DownCor appears robust. Fig. 11.20 shows a comparison of the LV measure (Section 4.7, Eq. 4.20) with neighbourhood size 1 calculated for CMPLX/SSQR and DownCor phase. Whilst the CMPLX/SSQR variance shows strong local artefacts and global imbalances, the contrast of the DownCor phase variance is predominantly related to anatomic structures.

The estimated fieldmap as well as the MUBAFIRE and MUBAFIRE Local corrections are illustrated in Fig. 11.21. Local contrast is enhanced with the MUBAFIRE algorithm, while long range variations are eliminated. MUBAFIRE Local additionally compensates for pronounced field disturbances within the central brain. Here, air cavities residing below the illustrated slice induce sharp, noisy field perturbations. These perturbations are part of the initial brain mask and could thus not be completely compensated for by the basic MUBAFIRE algorithm, but by MUBAFIRE Local. Fig. 11.22 shows a volumetric plot of the results of background field correction.

Occasionally, the MUBAFIRE Local algorithm excludes areas of too wide extent, and thus abandons valuable field information. Thus, for susceptibility reconstruction, both background correction results are considered. For illustrations, the best reconstruction result of a certain region is chosen. Fig. 11.23 shows the reconstructed susceptibility map in comparison to the MUBAFIRE Local result. The reconstructed susceptibility map shows rich contrast, especially in the central brain, e.g. for the basal ganglia. The utilised parametrisation for FSE seems to suppress artefacts efficiently. Grey/white matter contrast is clearly visible, even in the fieldmap. Surprisingly, the cortical gyri are not much more pronounced in the reconstruction than in the fieldmap. For




Figure 11.19: The mosaic on the left hand side shows the magnitude of the first echo (top row), the phase of the second echo (second row), the estimated phase offset at $t = 0$ (third row) and the offset-corrected phase (last row) for all channels included into the recombination. On the right hand side, the results of a simple CMPLX/SSQR and the DownCor recombination are illustrated. The red arrows indicate phase poles in the CMPLX/SSQR recombination hindering unwrapping.




Figure 11.20: Sample slice comparison of the phase variance inside the brain mask, employing the LV measure. The CMPLX/SSQR method (left) shows stronger disturbances (red arrows) than the DownCor phase (right).

the *post mortem* examples (see Section 11.3) the observed grey/white matter contrast in the susceptibility maps is distinctly higher. Finally, Fig. 11.24 shows a volumetric illustration of the local field and the reconstructed susceptibility and magnified examples for the susceptibility map in several brain regions that are known to possess significant susceptibility contrast.

11.6.3 Discussion and Conclusions

The phase and LV comparison prove that advanced recombination techniques such as DownCor are inevitable for ultra-high field. The rich contrast in the reconstructed susceptibility maps of the properly combined data verify the functionality of the employed processing steps. The structure inside the basal ganglia becomes very clear in the 3D plot and in the magnified view. The RN and the SN are obviously outlined from their surroundings. Yet, their reconstruction




Figure 11.21: Background correction of the measured field is shown for a sample slice. The raw fieldmap (left) is corrected with MUBAFIRE (middle) and finally with MUBAFIRE Local (right). Some highly localised distortions are compensated for in the latter step (red arrows). The results both show a strong overall contrast.




Figure 11.22: Three-dimensional plot of the background correction. The left image shows the raw field, and right on top of it the determined background contributions, including MUBAFIRE Local, are illustrated. The right image shows the result after full background-correction, exhibiting high image contrast.

suffered from surrounding blood vessels which are also visible in the magnified view. The high susceptibility in these vessels disturbs the local field and thus generates phase changes beyond




Figure 11.23: The top rows show the local field in relation to the susceptibility, using sample slices (every 12th slice) throughout the measured and masked volume. The lower images show one sample slice of the central brain comparing field (left) and susceptibility (right). The BG are pronounced in the reconstruction (yellow arrow). In the lower centre of these slices a masked region is visible, where the local field distortions exceed the value range that is accepted by MUBAFIRE Local (red arrows)

Nyquist confinements. The DN show very rich contrast. Even structures of the cerebellum with large spatial extent are visible as a corona in the transverse view. Showcase D) underpins that the *in vivo* susceptibility map of the cortical gyri does not yield more emphasised contrast than the local fieldmap.

The experiment shows that phase imaging at ultra-high field is an extremely promising field of research. Even with a very low repetition time and thus a quick measurement, very rich contrast can be achieved in form of phase and fieldmaps. The susceptibility reconstruction of these maps




Figure 11.24: Overview on the susceptibility reconstruction. The central volumetric images show the MUBAFIRE Local corrected fieldmap (left) and the susceptibility (right). The showcases illustrate local structures viewed in multiple orientations, including A) the BG (transverse, sagittal and frontal), B) the RN and the SN (transverse and frontal), C) the DN residing in cerebellum (transverse and frontal) and D) the structure of cortical gyri (sagittal). Case D) shows field (left) and susceptibility map (right). The colour ranges for the showcases were slightly adapted in each case to generate optimal contrast for the illustrated signals and can thus differ from the susceptibility range shown for the volumetric plot.

shows even richer contrast in regions such as the basal ganglia and the various nuclei residing inside the human brain. The identification and determination of haemorrhages inside brain tissue, findings of calcifications or any kind of tumour-like tissue mutation inducing field and phase contrast are potential future applications of ultra-high field phase imaging.

An interesting finding in this experiment is that the structure of cortical gyri in the human

brain seems very emphasised in background-corrected fieldmaps. Nevertheless, its contrast is not enhanced in susceptibility reconstruction as one would expect. Theoretically this can indicate, that the effects generating field shifts within these structures are not entirely originating in susceptibility properties of the tissue, but may additionally be related to chemical shifts that lead to a slight off-resonance (see Section 3.3.2).

11.7 Cerebellum Phase Images at 9.4T

The potential of ultra-high field phase imaging will be illustrated on a long-time measurement of a fixed *post mortem* brain. The fieldmap obtained from this measurement was already shown in Chpt. 9 to validate the reliability of the MUBAFIRE algorithm. In contrast to most of the other methods shown within this work, only a single echo is used for evaluation, although two echoes were recorded. The second echo contained an unacceptably high noise level and had to be abandoned. Due to the long duration of the measurement, the protocol could not be repeated at an arbitrary time.

11.7.1 Material and Methods

The measurement was performed at a 9.4 T human scanner (Siemens AG, Erlangen, Germany). *Dr. Ana-Maria Oros-Peusquens* deserves acknowledgements for designing the protocol and for planning and performing the entire measurement. A standard GRE sequence was applied with the following parametrisation:

Sequence Parameters					
seq	GRE	exc	slab-selective	Tx/Rx	8 ch array, 4 ch used
matrix	$480 \times 512 \times 160$	FOV	$115.2 \times 122.9 \times 38.4$ mm	res.	$0.24 \times 0.24 \times 0.24$ mm
TR	63 ms	α	25° (nominal)	RF-spoil	on
TE	16.38 ms	avgs	12	BW	40 Hz/Px

Although eight Rx channels were acquired, only one channel was used in the comparison shown in Chpt. 9. The reason for this lies in the high noise level of the second acquired echo. The noise level made the second echo unusable for evaluation and thus prevented the proper determination of an offset-corrected fieldmap. Recombination of several channels without offset-matching is not advisable (see Chpt. 10), since serious phase artefacts are likely to evolve. Nevertheless, in the following example a trick will be used to achieve a solution based on multiple receive channels.

The data contained in the ROI as defined in Chpt. 9 were extracted from all receive channels. The employed ROI differs from the previously introduced one only by one voxel and spans $174 \times 284 \times 110$ voxels, still covering mainly the cerebellum. Although at 9.4 T a multichannel recombination with constant offset correction is not advisable, it might still be applicable for a reasonably small ROI and for data of several channels that show similar sensitivity characteristics in this region. Under this premiss, the phase offset between these channels should behave smooth inside the ROI and recombination can be approached using only a constant per-channel offset. Four of the eight recorded channels, exhibiting sufficient magnitude signal within the defined ROI were chosen for further analysis, whilst the other channels were discarded. The remaining channels

possess relatively stable magnitude profiles inside the ROI and motivate the application of the SMW channel recombination technique (described in Section 10.5). The phase of the first echo of the recombined signal is unwrapped using P-URSULA with a subset size of $36 \times 36 \times 36$. The fieldmap is derived by scaling the unwrapped phase by the factor $10^{-3}/(2\pi \cdot t_{TE1}[ms])$. The MUBAFIRE framework is then employed for background field removal. The MUBAFIRE base workflow is applied using:

MUBAFIRE Parameters				
Filters:	POLF → SPHINX → DIPF			
resample	1 (half)	POrder	1	SPorder 4
DIPF iterations	50			

Distortions originating in incorrect phase or field values and preparation imperfections such as air bubbles are corrected for by a MUBAFIRE Local filter chain using the following parameters:

MUBAFIRE Local Parameters			
Filters:	Threshold → Erode diag. neighbours → 6 neighbour erosion → 6 neighbour dilation → DIPF		
THwidth	6	DIPF iterations	100

Susceptibility reconstruction was not pursued for this scenario, since the underlying fieldmap was not corrected for offsets and presumably contains arbitrarily shaped phase distortions as well as local shift artefacts.

11.7.2 Results

The sum-of-squares calculation as generated by the scanner is shown in Fig. 11.25. The high number of averages during measurement allows for a relatively good coverage of the measured volume with only few spots of low signal intensity. Obviously the signal level of the second echo is too low to be used for phase analysis. Only in the regions containing formalin, but no brain tissue, the signal is sufficient. The selected ROI, shown in the overview image, avoids areas of low signal or phase poles. The SMW recombination, applied to the extracted ROI data of channel 1, 3, 4 and 8, is illustrated in Fig. 11.26. The illustrations shows the offset-corrected phase maps and the smoothed magnitude maps used for phase recombination. The recombined phase maps show an overall good phase quality (except for the very edges of the ROI), sufficient for phase unwrapping. Fig. 11.27 features several sample slices for the recombined phase and magnitude data, underlining the robustness of the locally applied SMW recombination. Unwrapping the phase with P-URSULA took about ten minutes. The unwrapped phase is smooth throughout most parts of the ROI. Some areas show local inconsistencies which are in *post mortem* measurements usually attributed to preparation artefacts such as air bubbles. The fieldmap calculation only involved the multiplication of the unwrapped phase with a scaling factor, so no additional sources of errors had to be considered. The fieldmap and the background-corrected field are illustrated for a sample slice in Fig. 11.28.

Fig. 11.29 illustrates the background correction for orthogonal slices aligned as a block inside the ROI. The figure points out the strong reduction in value range of the field and the correction for local disturbances. To facilitate a comparison between magnitude and background-corrected field, a local region inside the cerebellar gyri structure was extracted. Besides a direct comparison of a sagittal sample slice (shown in Fig. 11.30), a line plot illustrates the contrast differences in Fig. 11.31.




Figure 11.25: Overview of the measurement, showing a sagittal sample slice for magnitude (left) and phase (right) for both recorded echo times (top: 1st echo, bottom: 2nd echo). The yellow box indicates the cross section of the ROI selected for analysis.




Figure 11.26: Runtime illustration of the SMW recombination method, showing in the top row a sample slice of the offset-corrected phase data of the selected channels and in the bottom row the smoothed magnitude used to recombine the phase properly.

11.7.3 Discussion and Conclusions

Judging by the sample slice shown in Fig. 11.26 and by the slices shown in Fig. 11.27 the offset-matching appears to be successful in the range of the ROI. Although under normal circumstances it is advisable to employ more adept methods for channel recombination at 9.4 T, the SMW




Figure 11.27: Series of sagittal sample slices through the SMW-recombined volume for magnitude (upper row) and phase (lower row).




Figure 11.28: Sagittal sample slice of the raw fieldmap (left), the MUBAFIRE correction (middle) and the additionally applied MUBAFIRE Local correction (right). Note, that the value range of the left image is indicated by the first colour index, whilst the middle and right images are described by the second colour index. The red arrows indicate shifts, induced by air bubbles that are removed by MUBAFIRE Local

method generates robust results when selecting a ROI of small extent and relatively constant signal strength for the considered Rx channels. The P-URSULA algorithm could successfully be applied to the sample, enabling a proper unwrap in reasonable computing time. The background field correction with the MUBAFIRE algorithms shows convincing results: A reduction of the result value range is achieved whilst preserving local contrast. The strong distortions visible in the cube top in Fig. 11.29 are most likely generated by air bubbles. Air bubbles cause strong, dipole-shaped fields with a wide extent. Besides masking the bubbles by thresholding, the influences




Figure 11.29: 3D illustration of the raw field (left) and the determined background contribution (top), being subtracted from the raw field in order to obtain only local variations (right).




Figure 11.30: Comparison of the magnitude (left) and the background-corrected field (right) contrast for a sagittal slice in a confined region inside the cerebellum. The local structures exhibit richer contrast in the fieldmap, indicating multiple layers. The yellow lines indicate the line plot position for Fig. 11.31.




Figure 11.31: Line plot illustrating the contrast properties of magnitude and local field. The magnitude has arbitrary units and is scaled to a range that allows comparison to the background-corrected field.

on the surrounding field characteristics are almost completely compensated for by MUBAFIRE Local. The presented study, featuring a nominal isotropic resolution of $240\ \mu\text{m}$, shows that high resolution phase imaging at ultra-high field has great potential for the visualisation of microscopic structures such as the cerebellar layers (shown in Fig. 11.29 and Fig. 11.28). In comparison with the magnitude acquired at the same resolution (Fig. 11.30 and Fig. 11.31), the phase

and the derived fieldmaps complement the contrast with additional information on microscopic structures. As indicated above, susceptibility reconstruction is not performed, since the observed phase does not enable derivation of an offset-corrected fieldmap. The contained phase shifts are not necessarily only attributable to susceptibility contrast. It is rather surprising how homogeneous the cerebellar layers appear in the illustrated fieldmaps. Similar to the 9.4 T *in vivo* study, shown in the last section, another source of phase contrast is potentially reflected in the fieldmaps. The homogeneous structure once more indicates, that weak chemical shifts might generate a contrast that adds up with that induced by structures with susceptibility contrast (see Chpt. 7, section Section 7.3.2). This should be considered in any future study that investigates tissue microstructure at ultra-high field.

11.8 Field in Parallel Fibre Structures

The presented reconstruction methods for magnetic susceptibility are based on a well-established model describing the observed field shifts as a convolution of χ with a dipole-shaped field (discussed in Chpt. 7, e.g. Eq. 7.21). The model works well for structures that are homogeneous within a voxel, but will most likely suffer from substructures with much smaller extent than a voxel, producing intrinsic field distortions. A prominent example for this are white matter fibre bundles that cannot be resolved with conventional in MRI [He and Yablonskiy, 2009; Lee et al., 2010]. Furthermore, the measured volume fraction (water) is only partially representative for the content of the imaged voxel. This study aimed to investigate the behaviour of field shifts and observed susceptibility in a diffusion phantom containing a bulk of parallel straight Dyneema[®] fibres. The phantom was originally designed by *Dr. Ezequiel Farrher* (Forschungszentrum Jülich, Germany). The presented findings were first reported in [Lindemeyer et al., 2011a].

11.8.1 Material and Methods

The phantom consists of a cylindrical, upright water container that holds the fibre model. To achieve relaxation times close to values inside the human brain, the phantom was filled with distilled water and doped with manganese sulphate in a concentration of approximately 3 ‰. The T_1 value of the phantom was determined by spectroscopy to be 780 ms. The fibres (consisting of Dyneema[®], diameter $16\ \mu\text{m}$) are densely wound around a perspex plate and compressed by two additional plates. The fibre bulk is aligned with the base and positioned at the centre of the phantom. The susceptibility of Dyneema[®] is known to be slightly different from that of water [Fieremans et al., 2009] and is thus optimally suited to our purpose. The geometry of the phantom is shown in the segmentation map derived from a GRE magnitude measurement, shown in Fig. 11.32. The phantom was measured with a twelve-channel Rx array coil at multiple orientations (see Fig. 11.33), where 0° indicates that the fibre direction is parallel to the direction of the static magnetic field, $\vec{B}_0$. All experiments were performed on a 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany). Rotation was applied about the upright axis of the phantom, using the angles $[-90^\circ, -67.5^\circ, -54.7^\circ, -45^\circ, -33.75^\circ, -22.5^\circ, 0^\circ, 22.5^\circ, 33.75^\circ, 45^\circ, 54.7^\circ, 67.5^\circ]$. At each position a multiple echo 3D GRE sequence was acquired using the following parametrisation:

Sequence Parameters					
seq	GRE	exc	non-selective	Tx/Rx	1 ch Tx, 12 ch Rx array
matrix	$174 \times 192 \times 144$	FOV	$[139.2 \times 153.6 \times 115.2]$ mm	res.	$[0.8 \times 0.8 \times 0.8]$ mm
TR	110 ms	α	30°	RF-spoil	on
Ellipt.	true	BW	500 Hz/Px		
TE	$[1.6 + n \cdot 3.38]$ ms (32 echoes)				




Figure 11.32: Segmentation of the phantom components based on the magnitude data of a GRE measurement. For analysis, the free water and the Dyneema[®] fibre area are of highest importance.





Figure 11.33: Phantom placement inside the Rx coil. Rotation took place along the symmetry axis of the cylinder. A rotation of 0° indicates that the fibres are oriented along the axis of the static field, B_0 .

Fieldmaps were calculated using FSE with unwrapping in time domain and linear regression on the phase images of the first six echoes (showing sufficient SNR). Masks were derived by thresholding the magnitude. The rotation of all measurements was aligned by inverse rotation about the true measurement angle. The aligned measurements were then transversally coregistered by minimising the sum of the point-wise difference between each mask and the mask for 0° . Two pairs of representative regions were selected to estimate an average difference between the fibre-area field, B_{Fibre} , and the free-water field, B_{Free} . Each pair was positioned symmetrically along the vertical axis with respect to the phantom centre (see Fig. 11.34). The free-water pair possessing a dimensioning of $43 \times 33 \times 16$ voxels was positioned in such a way that one region resides above and one region resides below the compressing perspex plates. The fibre-area pair possessing a dimensioning of $43 \times 33 \times 9$ voxels was positioned between the plates. Calculation of the average difference is the best approach to obtain a representative field shift. The application of a background field correction algorithm is not advisable, since it would potentially misinterpret the field generated by the fibres as an external influence. Due to the radial symmetry of the

Figure 11.34: ROIs for field difference calculation are highlighted in the fieldmap. The ROIs are positioned inside the fibre regions (central blocks) and above as well as below in the free water.



phantom and since it was kept in place for all angle measurements, the externally generated field inhomogeneities should remain constant, while the inner fields vary due to the influence of the rotated, non-symmetric geometry of the fibres. The susceptibility distribution inside the phantom was estimated using FSE with a simple Tikhonov-based approach as discussed in Section 7.4.2.

To interpret the field difference two simulations using a simple dipole convolution (as in Eq. 7.18) were performed, each based on a model that roughly matches the geometrical structure of our phantom. Fibres were simulated a) as a bulk and b) stacked with a chequerboard voxel-wise alternating cross-section. In the simulations constant (but arbitrary) fibre susceptibility was assigned and set negative compared to that of water, since Dyneema[®] shows this behaviour (determined by inspection and as in Fieremans et al. [2009]). For the bulk case, the entire ROIs in the fibre area were used to calculate the field difference. For the stacked simulation, only the field between the fibres – meaning every second voxel of the cross-section – was evaluated, since Dyneema[®] itself does not generate a signal. The numerical phantoms are sketched in Fig. 11.35.




Figure 11.35: Slices showing cross section of the numerical phantoms, left: bulk susceptibility, right: fibres stacked in chequerboard pattern.

11.8.2 Results

The fieldmaps show large variation with orientation inside the fibre area. Due to the phantom filling process, despite careful preparation and handling, a number of small air bubbles remained between the fibres. As the bubbles are sparsely distributed, the comparatively large ROIs compensate for the bubble influences when calculating the ROI average. Fig. 11.36 shows a central slice in the fibre area of each measurement aligned with the rotation plane and a vertical cross section. Fig. 11.37 shows a graph of the measured field difference between the reference regions as well as the standard deviation of the difference. Fig. 11.38 illustrates a comparison of simulation and measurement and Fig. 11.39 displays cross sections of the simulated fields. Due to the arbitrarily chosen susceptibility value in the simulation, the simulation graphs are scaled to match the field




Figure 11.36: Fieldmaps of the measurements for different angles. The left image shows a transverse cross section through one of the fibre planes, whilst the right image shows a vertical cross section through the phantom aligned with the fibres. Each illustration shows from left to right and from top to bottom the angles: $[-90^\circ, -67.5^\circ, -54.7^\circ, -45^\circ, -33.75^\circ, -22.5^\circ, 0^\circ, 22.5^\circ, 33.75^\circ, 45^\circ, 54.7^\circ, 67.5^\circ]$.




Figure 11.37: Average field difference between the fibre area, B_{Fibre} , and the free water, B_{Free} for all measured angles. The vertical error bars indicate the standard deviation of the calculated difference average.

amplitude of the measurement:

$$\Delta B_{\text{sim}} \rightarrow \alpha \Delta B_{\text{sim}} \quad (11.2)$$

$$\alpha = \frac{B_{\text{Fibre}} - B_{\text{Free}}}{B_{\text{Fibre,sim}} - B_{\text{Free,sim}}} \quad (11.3)$$

The scaling parameter, α , was determined for the graph of the chequerboard fibres, and the bulk graph was scaled with the same factor. Fig. 11.40 shows the susceptibility difference between the same ROIs in arbitrary units. Susceptibility appears to change with orientation and does not remain within the error margin.

11.8.3 Discussion and Conclusions

The simulation of the field caused by a homogeneous bulk susceptibility shows a completely different shape compared to the performed measurements. The simulation of the chequerboard-

Figure 11.38: Average field difference (blue) compared with the simulation of bulk (red plus) and chequerboard stacked susceptibility (red x) for all measured angles. The vertical error bars indicate the standard deviation of the calculated difference average.




Figure 11.39: Simulation of fields generated by bulk susceptibility (A and C) and chequerboard fibres (B and D). The top row shows a vertical cross section, whilst the lower row shows a transverse cross section through the fibre area. Units are arbitrary, but identical for greyscale colouring. The order of angles is the same as in Fig. 11.36.





Figure 11.40: Average susceptibility difference between the reference ROIs for all measured angles. The vertical error bars indicate the standard deviation of the calculated difference average.



stacked fibres on the other hand, exhibits an astonishing similarity when fitted to the measured results. Although the simulated model is very rudimentary and does not represent true fibre stacking and scales, it imitates the observed field effects convincingly well. It is obvious that macroscopic field measurements of areas containing parallel elements cannot be described by bulk susceptibility in the corresponding region. The findings rather suggest, that microscopic structures can have a major effect on the observed field shifts. Such structures might even induce the observation of seemingly anisotropic susceptibility, as shown in Fig. 11.40. This is an important

aspect for susceptibility studies, especially those concerning regions such as the corpus callosum, consisting of parallel nerve fibres. The apparently anisotropic susceptibility can be explained by the locally inappropriate spatial resolution which does not satisfy the underlying model – single fibres cannot be resolved. In continuous space the model would still solve the problem, but as soon as field shifts are generated by structures with an extent below the resolution of the measurement, the problem is insufficiently described. Thus, it is essential to be aware of structure-specific effects that might corrupt the consistency of the reconstruction whenever the assumed model is violated. Anisotropy of the susceptibility of Dyneema[®] itself or the included diffusing water volume might play an additional role, but these influences have not been quantified within this study. One can conclude regarding susceptibility reconstruction, that a hybrid model employing deconvolution methods as well as *a priori* knowledge on particular regions could be a promising approach.

11.9 High-Resolution Phase Imaging of the Central Brain at 3T

In the former sections it was pointed out that certain structures in central brain regions, such as the SN, the RN or the *globus pallidus* (PAL) expose strong susceptibility contrast with respect to their surroundings. Susceptibility mapping of e.g. the RN and the SN was already demonstrated e.g. in Section 11.6 and in literature [e.g. Wharton et al., 2010]. Yet, these regions are of small extent and are thus described by only a few voxels in measurements at 1.0 mm or 0.8 mm isotropic resolution. Higher resolution, allowing for a more precise description of the anatomy, would be beneficial for medical diagnosis or for planning of surgeries. Especially the *subthalamic nuclei* (STN), located close to the SN, are of great interest as shown in Schäfer et al. [2012] in a 7 T study.

Measurements at ultra-high field are not applicable in daily routine, since hospitals usually possess scanners of only up to 3 T field strength or less. Despite the limited potential of 3 T magnets regarding image resolution, this example shall exploit the potential of high-resolution GRE phase measurements at medium field strength, in this case a 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany). The measurement was performed on a patient of age 68.

11.9.1 Material and Methods

The measurement protocol was mainly designed by *Dr. Ana-Maria Oros-Peusquens*. The protocol makes use of slab-selection and saturation bands⁵, in order to confine the imaged region to the central brain. The readout direction was oriented along the inferior-to-superior axis. The GRE sequence is parametrised as follows:

Sequence Parameters					
seq	GRE	exc	slab-selective	Tx/Rx	1 ch Tx, 32 ch Rx array
matrix	448 × 154 × 112	FOV	[192.6 × 66.2 × 44.8] mm	res.	[0.43 × 0.43 × 0.40] mm
TR	64 ms	α	17°	RF-spoil	on
iPAT	2	Ellipt.	true	partial F.	6/8, 6/8
BW	120 Hz/Px	Total t	8 min, 2 s	Saturation	left and right
TE	[4.54, 13.59, 22.64, 31.69] ms				

⁵Pre-defined areas in the sample are saturated – meaning the temporary destruction of effective magnetisation – before the encoding and acquisition blocks of a sequence are applied

Due to the high channel count and in order to stay in the scope of clinical applicability, no offline channel recombination was employed. The signal was generated by the *Adaptive Combine* mode provided by the scanner manufacturer. The SNR and the image contrast are expected to be relatively low for an approximately isotropic resolution of 0.4 mm at this field strength. Thus, four consecutive measurements with identical parametrisation were acquired. The measurements are merged using off-line post-processing and averaging. The first post-processing step is the coregistration of the measurements. Although the measurements are acquired with roughly constant patient positioning they suffer from slight displacement due to involuntary patient motion between or during the individual measurements. Coregistration is performed using Flirt [Jenkinson and Smith, 2001; Jenkinson et al., 2002]. For phase imaging, coregistration is not only performed on the magnitude, but is also necessary for phase information. Since direct interpolation of the phase itself is not possible as it leads to interpolation of wrap surfaces, the operation has to be applied to the complex data. The coregistration matrices are therefore estimated for the magnitude data, before they are applied to the real and imaginary components of the complex data. These are finally converted back to polar representation.

For deriving joint fieldmaps and susceptibility estimations from the coregistered measurement series, two methods are suggested:

1. Complex data → Averaging → FEST → MUBAFIRE → FSE → Susceptibility
2. Complex data → FEST → MUBAFIRE → Averaging → FSE → Susceptibility

Whilst the first method embodies the traditional strategy of dealing with coregistered data, the second method accounts for pitfalls arising due to the sequential recording order. The first filter simply averages all coregistered acquisitions before performing the field and susceptibility analysis. The pitfalls addressed by the second filter are field changes caused by patient motion and potential gradient drift over time. Field changes do not influence magnitude imaging much. Yet, even small local field variations lead to a change in phase evolution. Fig. 11.41 shows how drift and motion generate field differences leading to varying phase values in the acquired data. Dephasing occurs especially for late echo times. Direct averaging of complex signals containing phase states that diverge with the echo time inevitably leads to signal loss, noise amplification and propagation of incorrect field values. Thus, the second method is proposed and applied in this study, compensating for such shortcomings. The field for every individual measurement, b_i , is calculated. The individual fieldmaps are corrected for background fields using MUBAFIRE → $b_{i,cor}$. The influences of gradient drifting and patient motion are of external origin and will thus be addressed and removed by the developed filter chain. Slight differences between the background correction results will be smoothed – or “averaged out” – by the final averaging step:

$$b = \frac{1}{4} \cdot \sum_{i=1}^4 b_{i,cor} \quad . \quad (11.4)$$

FEST is applied with linear regression and spatial unwrapping using PRELUDE [Jenkinson, 2003] in sliced mode. The following MUBAFIRE (Local) configuration is used for both methods:




Figure 11.41: Field and phase shifts occurring throughout the measurements. The left hand side shows the magnitude (top) and the phase (bottom) of the 4th echo for the four acquired measurements on one patient. The coloured boxes in the phase images indicate a local neighbourhood of $5 \times 5 \times 5$ voxels used to estimate a local mean phase value for each measurement and each echo time. The graph on the right shows the evolution of this mean phase with echo time. Obviously, not only a constant phase offset is encountered, but the slope – which is proportional to the local field shift – changes dramatically between the measurements.

MUBAFIRE Parameters					
Filters:	POLF → SPHINX → DIPF → Thresholding → Erode diag. neighbours → 6 neighbour erosion → 6 neighbour dilation → DIPF				
resample	0 (full)	POLorder	1	SPorder	4
THwidth	18	DIPF iterations	50		

The susceptibility distribution is finally reconstructed using the combined, background-corrected fieldmap. FSE employs the following configuration for hybrid Tikhonov and gradient regularisation:

FSE Parameters					
regmode	3	W	brain mask (m)	W_T	brain mask (m)
λ	0.02	μ	0.02	B₀	2.89 T
iterations	50	reset	20	W_{G_{x(y,z)}}	brain mask (m)
spat	true				

11.9.2 Results and Discussion

In Fig. 11.42 the data processing chain is illustrated for the measured sample data set. The background-corrected fieldmaps resulting from the measurement-averaging is smooth and stable enough to allow for susceptibility reconstruction. Though the calculated susceptibility map shows




Figure 11.42: Evaluation of the measured data. The top row shows the fieldmaps derived from the individual measurements. The left hand side of the bottom row shows the fieldmaps after MUBAFIRE correction. On the right hand side the average of the individual corrections and the reconstructed susceptibility map resulting from the latter are shown. The first measurement shows the highest visual noise level, while the other three measurements are comparably smooth. This is due to the fact that all measurements were coregistered to the first one.

oscillating artefacts, the area around the RN, SN and the STN is well-outlined, exhibiting the extent of the small-structured tissue formations. The origin of the oscillation artefacts found in field and susceptibility maps is not completely clear. Yet, automatic channel recombination as offered by the scanner was used to generate the data. Presumably, the involved processing steps operate on a slice-based principle, leading to potentially different results in phase estimation for adjacent slices. This effect can indeed be observed in the recorded data (see Fig. 11.43). A combination of slice-based processing and partially high local noise level is believed to lead to the observed oscillations. Nevertheless, the presented measurement protocol and evaluation method lead to good results regarding imaging of the central brain structures, the RN, the SN and most importantly the STN as shown in Fig. 11.44 and Fig. 11.45. The STN show distinct contrast in the susceptibility map while being hardly visible in the magnitude data.




Figure 11.43: Phase dataset of a single measurement viewed in a confined region with the slicing axis aligned in vertical orientation. The stacking shows slice-dependent phase offsets, indicated by red arrows. Unfortunately, these offsets can only be partly compensated for by fieldmap calculation.




Figure 11.44: The left hand side images show magnitude contrast of the 4th echo whilst the right hand side shows the reconstructed susceptibility distribution. Frontal and sagittal view are illustrated and their mutual position is indicated by the yellow lines. The region labelled in red depicts the SN, whilst the blue region marks the STN.




Figure 11.45: Transverse slices in the region of the RN, SN and STN (every second slice is shown, from left to right and from top to bottom). The left hand side shows the magnitude, whilst the right hand side shows the reconstructed susceptibility.

11.9.3 Conclusions

The presented method has great potential for imaging of an isolated region in the central brain. The protocol is still under development and will be further optimised in the scope of a more extensive study. The slice-dependent field variations, introduced by the automatic channel recombination of the scanner, still lead to suboptimal results and could be improved by offline channel recombination. Offline recombination of a high resolution dataset recorded with 32 individual receive channels is very demanding regarding computing time and resources and thus not suited for clinical application. Yet, the advantages of high resolution imaging of the central brain for medical imaging are already evident. The 9.4 T imaging applications show higher contrast and resolution capabilities, but in the near future no ultra-high field scanners will be available for conventional hospitals. Hence, optimisation of measurement protocols for common scanner types and field strength will still be of high importance. The data acquired on clinical scanners with a method such as the presented one can be a great tool for MRI-guided implantation of electrodes for deep brain stimulation, e.g. for treatment of Parkinson's disease, in the future.

11.10 Fast Susceptibility Imaging

Several of the studies discussed so far accompanied another major research topic and were thus bound to their measurement protocols, allowing for only slight modifications. The approach presented in the following attempts to find a fast imaging solution for moderate resolution and noise level, which can be included into existing protocols at low time cost. The algorithms for computing the results were chosen to be time-efficient where possible, to allow for prompt inspection of the results.

11.10.1 Material and Methods

The desired resolution for this approach is 0.8 mm isotropic with full brain coverage. The repetition time is kept at $TR = 30$ ms to allow phase contrast to evolve sufficiently. Typical T_2^* values of human brain tissue are slightly higher (above 40 ms), but choosing a longer TR would either prolong the measurement time significantly or allow only for partial brain coverage or lower resolution. To accelerate the measurement, elliptical scanning was applied. Acceleration with a GRAPPA factor of two [Griswold et al., 2002] and 6/8 partial Fourier in both phase encoding directions were applied. A healthy volunteer scan was performed with informed consent on a 3 T TIM Trio whole-body scanner (Siemens AG, Erlangen, Germany) including the following parametrisation:

Sequence Parameters					
seq	GRE	exc	slab-selective	Tx/Rx	1 ch Tx, 12 ch Rx array
matrix	$256 \times 232 \times 176$	FOV	$[204.8 \times 185.6 \times 140.8]$ mm	res.	$[0.8 \times 0.8 \times 0.8]$ mm
TR	30 ms	α	12°	RF-spoil	on
TE	$[5.00, 17.36]$ ms	BW	$[200, 60]$ Hz/Px	Total t	5 min, 52 s
iPAT	2	Ellipt.	true	partial F.	6/8, 6/8

The data were recorded in uncombined mode (separation in four physical channels) for post-processed recombination. *Per-channel field estimation* (see Section 10.8) was employed, resulting

in the combined fieldmap. The FET-based fieldmap estimation of the channel recombination used linear regression mode and P-URSULA with a subset size of $32 \times 32 \times 32$ voxels for unwrapping. The first echo was acquired with a higher bandwidth in order to save time for a low bandwidth for the second echo. The phase contrast of early echoes is usually not much evolved and is mainly used for offset-correction, so detail contrast is of little importance here. In order to reduce the noise introduced by the higher bandwidth, the first echo is smoothed with a Gaussian filter of 1 voxel kernel size (0.8 mm) before the fieldmaps are calculated.

Since the employed recombination method does not produce magnitude and phase images, the channels are additionally recombined using the SSQR to illustrate the measured data as shown in Fig. 11.46. Processing masks are generated by a quick automatic segmentation with ITK Snap [Yushkevich et al., 2006] based on the SSQR magnitude and phase inconsistencies are excluded using automatic phase masking (see Section 4.9) with a threshold of 0.68 and Gaussian smoothing of 3 voxels (2.4 mm) based on the SSQR phase. Background correction is performed with an integrated MUBAFIRE Local approach using the following parametrisation:

MUBAFIRE Parameters					
Filters:	POLF → SPHINX → DIPF → Thresholding → Erode diag. neighbours → 6 neighbour erosion → 6 neighbour dilation → DIPF				
resample	0 (full)	POLorder	1	SPorder	4
THwidth	10	DIPF iterations	50		

Susceptibility reconstruction is performed using FSE (see Chpt. 7) in a hybrid approach employing Tikhonov and gradient regularisation. The method has the following parametrisation:

FSE Parameters					
regmode	3	W	brain mask (<i>m</i>)	W_T	1 everywhere
λ	0.01	μ	0.01	B₀	2.89 T
W_{G_{x(y,z)}}	1 everywhere	iterations	50	reset	20

11.10.2 Results and Discussion

The magnitude images in Fig. 11.46 reflect that the employed echo timing resides below the average T_2^* values of the tissue observed. The magnitude contrast has only partly evolved up to the echo time of about 17 ms. Nevertheless, the phase shows distinct contrast for the later echo time. The noise level is moderate, recognisable by the clearly delineated wrap surfaces. The phase recorded for the first echo time is stronger corrupted by noise due to the higher recording bandwidth. The noise also originates in the early echo timing giving the phase contrast no time to develop (which is compensated for by the smoothing operation described above).

Fig. 11.47 shows the fieldmap derived by per-channel field estimation, the background-corrected field calculated by MUBAFIRE (Local) and finally the reconstructed susceptibility distribution. Due to the relatively low regularisation parameters, λ and μ , the noise level visible in the background-corrected field is only slightly smoothed in the susceptibility map. Depending on the purpose of the examination, the regularisation parameters can be modified to either try to extract more local detail (lower regularisation) or to achieve better noise reduction and greater smoothness in continuous regions (higher regularisation). Fig. 11.48 additionally illustrates three orthogonal plots of the reconstructed susceptibility map.




Figure 11.46: *Magnitude and phase data of the presented measurement as resulting from the SSQR recombination. Selected slices for magnitude (top row) and phase (bottom row) are shown. The left column shows the first echo, the right column shows the second echo.*

In order to estimate the mean susceptibility value and to evaluate the degree of homogeneity inside a number of structures inside the central brain, the following regions were segmented: RN, SN, PAL, *putamen* (PUT) and *caudate nuclei* (CN). Segmentation was performed manually using ITK Snap [Yushkevich et al., 2006]. For each region the mean susceptibility value was determined, representing the relative offset to the mean susceptibility of the whole volume. The standard deviation inside each region was determined, indicating the homogeneity and separability of the different regions. Fig. 11.49 shows the segmented regions, their mean susceptibility values and their variation in a bar plot. The following table gives an overview on the numerical results obtained from the segmentation:

		RN	SN	PAL	PUT	CN
Mean:	[ppm]	0.081	0.1833	0.139	0.060	0.060
Std:	[ppm]	0.029	0.051	0.035	0.034	0.022
Rel:	[%]	35.3	27.8	24.9	58.3	37.1

Susceptibility is relatively stable inside the different regions. This is indicated by the fact that the standard deviation is in an acceptable relative range. The determined values allow one to distinguish between value ranges for e.g. RN and SN or between PAL and e.g. the CN. CN and PUT show similar mean values, although the variation inside the CN is smaller than inside the PUT.

11.10.3 Conclusions

A fast phase imaging protocol, ready for use within clinical studies, was established. It allows for a whole-brain acquisition at 0.8 mm isotropic resolution within about 5 minutes and 52 seconds. Being designed for phase imaging, the measurement is suboptimally parametrised for further analysis of the magnitude signal, since the echo times remain clearly below the typical T_2^* values expected inside human tissue (about 50–60 ms). Nevertheless, the phase contrast accumulated for the early echo times of 5 and 17 ms shows sufficient contrast to derive a meaningful background-corrected fieldmap that enables the reconstruction of a proper susceptibility distribution. The reconstructed susceptibility introduces no streak artefacts, but propagates the noise level of the fieldmap. Quantification of susceptibility values for several structures in the central brain is possible and shows sufficient statistical stability. Fig. 11.47 further shows that the employed MUBAFIRE Local algorithm for exclusion of inappropriate field values efficiently removes a small subset-matching artefact generated by the P-URSULA unwrap method. Eventually, this fast imaging approach, consisting of the customized measurement protocol and the function chain used for evaluation, leads to convincing results. This is especially remarkable as the fast acquisition is situated on the brink of applicability for a 3 T scanner.




Figure 11.47: Selected slices of the fieldmaps from channel recombinations (top row) and of the background-corrected fieldmaps with local contrast (middle row) after MUBAFIRE Local correction and the reconstructed susceptibility (lower row) are shown. The red arrow indicates an artefact arising from an incorrectly matched subset of the P-URSULA unwrapping algorithm. The MUBAFIRE Local correction step recognises the contained field shifts as too high and excludes them from the data. The background-corrected fieldmaps (top row) are used to derive susceptibility maps (lower row) via FSE.




Figure 11.48: *Orthogonal slices of the calculated susceptibility map. The maps have smooth characteristics, exhibiting contrast in the central brain structures, but no significant streak artefacts.*




Figure 11.49: Segmentation of the reconstructed susceptibility map. Several brain regions are marked (left image), and their mean susceptibility (bars) as well as the standard deviation (red delimiters) are determined (right hand side plot).

12

Summary

In this thesis a comprehensive discussion of phase imaging in the context of MRI is presented, including phase unwrapping, fieldmap estimation, removal of contributions from fields other than those created by the tissue itself (background fields), susceptibility mapping, optimisation of phase acquisition and appropriate processing of data from multiple receive channels.

Differences and properties of several unwrapping algorithms, such as PRELUDE and a quality-based approach by Abdul-Rahman, are discussed and demonstrated. Two custom-built unwrapping approaches, MPPGI and URSULA are also introduced. MPPGI spatially unwraps phase data within cubical masks at high speed by integrating the wrap-corrected phase gradients while being tolerant to low noise levels. URSULA employs a QBU approach as a back-end, implemented for this purpose, and introduces an artificial volume compartmentalisation to enable time-efficient and highly robust unwrapping of large data arrays. A parallelisation extension, P-URSULA, is an additional feature of the algorithm. Numerical samples are used to assess accuracy and susceptibility to noise of URSULA and P-URSULA. Correct unwrapping is achieved for the majority of cases even for high noise levels. Various experiments are evaluated successfully using (P-)URSULA. It is possible to significantly reduce computational time in comparison to the classical Abdul-Rahman algorithm and PRELUDE.

Automated masking of phase data is presented in an approach that requires no further magnitude information, but evaluates the phase variation within a local neighbourhood. Measurement protocols are optimised for the acquisition of robust fieldmaps that are rich in contrast. Optimal parameters correspond to a monopolar multi-echo GRE acquisition with small bandwidth. Wide echo spacing is desirable to maximise the phase contrast. Whilst the first echo can be acquired early after excitation, the positioning of a second echo in the range of the tissue-dependent T_2^* time allows for maximum phase SNR. Besides the essential recording of two echoes, depending on the measurement constraints, the acquisition of four or six echoes, implying denser time sampling, is recommended. FEST, a framework for the calculation of fieldmaps based on MRI phase data, is introduced. It is designed to deal with various acquisition schemes and allows one to choose between several unwrapping and field estimation modes. In particular, the field estimation modes are linear regression, echo difference averaging and single echo mode.

The influence of magnetic susceptibility on the phase and particularly the induced orientation-dependent field distortions are discussed. Susceptibility reconstruction models and solver algorithms from literature, of which several were implemented within the scope of the thesis, are described, including direct inversion (k -space division), Tikhonov- and gradient-regularised minimisation as well as regularised closed-form solving. Multi-angle acquisition and reconstruction were also implemented. Optimisation maps were derived for direct inversion and for regularised minimisation to enable a problem-oriented configuration of the algorithm parametrisation. As expected, low noise levels demand only slight truncation or regularisation. For higher noise levels the optimised regularisation constants, λ and μ , approach a range of 10^{-2} to 10^{-1} , which is in agreement with literature values. Optimised truncation for direct inversion yields values in the range of 2 to 10, depending on the applied method. The minimisation-based reconstruction algorithms are the recommended choice for susceptibility reconstruction. They are successfully applied and validated on various case studies and series, obtaining qualitative and quantitative susceptibility distributions. The functionality of the involved algorithms is substantiated by the fact that the estimated maps show clearly delineated structures that describe the known anatomy and remove the gross orientation dependence of the contrast, which is visible in fieldmaps.

The necessity of *a priori* exclusion of externally induced distortions of the static field, in order to enable successful susceptibility reconstruction, is critically evaluated and several approaches are discussed. Background correction is performed by classical high-pass filtering (Gaussian, median and mean), by polynomial filters and by using spherical harmonic functions with SPHINX. Furthermore, the approximation via external pseudo-susceptibility distributions by the DIPF is discussed. Whilst the high-pass filters remove field contrast aggressively, SPHINX and the DIPF both preserve valuable, but different field information for susceptibility reconstruction. In this context, the novel MUBAFIRE filter chain is presented, comprising 1st order POLF, the SPHINX and the DIPF for physically consistent correction of background fields. This implies that the characteristics of the removed field as well as the remaining field have physical reasoning. Its performance is evaluated qualitatively and quantitatively by application on a numeric model and on *post mortem* and *in vivo* data of 3 T and 9.4 T measurements. Its accuracy clearly outperforms all traditional techniques and moreover shows clearly improved results in comparison with SPHINX and DIPF. The presented technique provides a highly optimised way to prepare MRI fieldmaps for the reconstruction of the underlying distribution of magnetic susceptibility.

In order to improve phase and susceptibility imaging with receive array coils, the recombination of multiple Rx channels was investigated. A custom development, SMW, performs a smoothed-magnitude-weighted complex summation that considers constant, channel-dependent phase offsets, in order to generate consistent phase data from multiple receive channels at field strength of up to 3 T. It is shown that for ultra-high field measurements spatially varying offsets have to be taken into account, and that traditional recombination is not applicable. For this purpose, two methods, DownCor and per-channel field estimation were designed. In both approaches, phase offsets are determined independently for each channel. DownCor estimates these offsets in a resampled low-resolution dataset to reduce computational time and corrects the acquired data with the interpolated offset estimate. Per-channel field estimation determines fieldmaps directly for each individual channel and combines them with SMW, rendering the calculation of phase offsets unnecessary. These techniques perform well in comparison to conventional recombination, which is demonstrated on a numerical sample and in application on ultra-high field data.

The algorithms developed and implemented in the framework of this thesis result in a thorough workflow for obtaining maps of the phase, the static magnetic field and the magnetic susceptibility distribution within an observed sample or subject. Several variations of the processing chain are applied to experimental data. The workflow is successfully employed on entire measurement series such as a study on 30 Parkinson's patients and age- and gender-matched controls as well as a hybrid MR-PET study on tumour patients which was part of presurgical planning. Notwithstanding the fact that the parameters of the acquisition were optimised for mapping of quantities other than phase (water content, T_2^*), serial evaluation could be performed, leading to robust results. The susceptibility maps of brain tumours show several interesting contrast features that could potentially serve as a novel tool for diagnostics. Although at this point in the data processing chain the differences between patients and controls are not obvious, the Parkinson's study underpins the reliability of the evaluation methods well. Due to constraints involved in data acquisition in a clinical environment, phase images from a multichannel coil array were recombined with manufacturer-provided software and show typical phase pole artefacts. This is compensated for by adapting the processing chain to employ echo-difference unwrapping. The derived susceptibility maps show clearly pronounced anatomy and no significant gradient artefacts. In order to provide valuable information on a region of the central brain, containing the *subthalamic nuclei*, a customized measurement protocol is presented. The protocol allows for 0.4 mm isotropic imaging at a clinical field strength of 3 T and is successfully demonstrated in a Parkinson's patient measurement. Additionally, a fast, whole-brain susceptibility imaging protocol was designed for 3 T, which can be included into existing study protocols with minimal time demands of only 5 minutes and 52 seconds at 0.8 mm isotropic resolution. Based on this protocol, structures of the central brain can be identified regarding form and susceptibility value range, as demonstrated for a volunteer measurement. Peculiarities of field shifts observed in media that contain a high level of organisation, such as parallel fibres, were investigated on a diffusion phantom. The experiment shows that knowledge about the underlying microstructure of tissue can be advantageous for the accurate evaluation of fieldmaps, since the observed field characteristics tend to deviate from the shifts predicted by the employed dipole model of susceptibility-induced perturbations.

The benefit of ultra-high field strength for phase imaging becomes evident in several examples. A high-resolution *in vivo* measurement, covering a central slab of the brain, is used to demonstrate the potential of URSULA and MUBAFIRE, and exhibits detailed and strong phase and field contrast throughout the volume. A whole-brain *in vivo* case with 1 mm isotropic resolution provides a broad phase contrast range, allowing for the reconstruction of susceptibility maps with distinctly outlined anatomy, i.e. of the *basal ganglia*, the *red nuclei*, the *substantia nigra* and the *dentate nuclei*. The *post mortem* measurement of the *cerebellum* enables unprecedented insights into tissue microstructure such as cerebellar layers, that are only partially visible in the magnitude, but revealed by the phase contrast at 0.24 mm isotropic resolution. The lack of pronounced orientation effects in the phase contrast of such a geometrically complicated structure is surprising and might indicate contributions of influences other than susceptibility contrast. The whole-brain *post mortem* measurement at 0.69 mm isotropic resolution exhibits resolution and contrast advantages of ultra-high field phase imaging unified in one experiment.

Further development of phase imaging for ultra-high field applications depicts a primary research objective for the future. With this thesis the basis for such research is substantiated by establishing the phase imaging workflow at 3 T and 9.4 T field strength at the host institute. In particu-

lar, the improved unwrapping performance provided by URSULA, source-sensitive correction for background fields by MUBAFIRE and the developed recombination strategies for multi-channel acquisition facilitate detailed investigations of the phase contrast in human tissues.

Acknowledgements

First of all, I want to thank my beloved wife, Steffi, my parents, my sister and my grandma, further our entire family and dear friends. These people and their love allow me to be the person I am today. I did not explicitly devote my thesis to either of them, since I know that they would find this pretty peculiar.

There are quite a bunch of people whom I want to thank and acknowledge and who took significant part in making this thesis possible, in the first place:

- I would particularly like to thank Prof. Dr. N. Jon Shah, for inviting me to the research position at his institute and for giving me the opportunity to become a member of this great group. I am most thankful for the freedom, support, advice and help he gave me in the last years, allowing me to develop my research ideas.
- Further, I want to gratefully acknowledge Dr. Ana-Maria Oros-Peusquens, who directly supervised me. She supported me and my ideas with overwhelming and constant enthusiasm. Thank you for all help, advice and criticism I received from you during the last years. Not least, thanks for providing me with numerous helpful suggestions and corrections on my thesis.
- My life at Forschungszentrum Jülich was immeasurably enriched by “the old PhD office” in which I was studying with my dear colleagues, Rüdiger, Daniel F., Zaheer, Daniel B. and Fabian (ordered by Cartesian distance between our desks). We had such an amicable atmosphere that one could easily forget about any other issues in life. Thanks to all of you!
- I am overly grateful for all help I received from current and former colleagues who advised me and corrected my thesis chapters – in alphabetical order of their first name: Dr. Christoph Weirich, Daniel Brenner, Dr. Daniel Fiege, Dr. Desmond Tse, Dr. Ezequiel Farrher, Dr. Ivan Maximov, Klaus Möllenhoff, Dr. Michael Poole, Rüdiger Stirnberg, Dr. Vincent Gras and Dr. Wieland Worthoff. Most of these people helped correcting several chapters. Beyond that, Michael deserves special thanks for giving me very extensive feedback on my initial chapter on MR theory. Wieland and Ezequiel would further drop in every so often for some corrections, or just to leave some helpful remarks concerning my current screen content.

- Planning, organising and finally performing measurements requires a lot of time invested beforehand. In that regard I want to thank several people. Dr. Ezequiel Farrher planned and built the straight-fibre diffusion phantom and the tube phantom. He further helped performing the measurements with the straight-fibre phantom. Dr. Desmond Tse and Dr. Michael Poole put tremendous effort into enabling the acquisition of proper phase data with the 9.4 T scanner. I want to thank Dr. Ana-Maria Oros-Peusquens for her help in planning sequence protocols such as the high resolution protocol at 3 T and performing the study on the Parkinson's patients. Further, for performing scans at 3 T and 9.4 T and for the acquisition of the measured fixed *post mortem* brains. Dr. Kathrin Reetz (University Hospital Aachen) deserves special thanks for her support in the Parkinson's study. Dr. Carolin Weiss organised the tumour patient study on behalf of Universitätsklinikum Köln, and supported me generously with diagnosis information on the illustrated cases.
- Thanks to Dorothe Krug for her support in many measurements and for the acquisition of volunteers. In this regard I also want to thank Barbara Elghahwagi and I owe thanks to all volunteers that were measured, contributing to what is shown within this thesis – needless to say, I cannot name them here.
- Most importantly, I want to gratefully acknowledge all of my colleagues for the great working atmosphere, their kindness, understanding and help. In particular, Arne, Daniel B. (for being a universal source for knowledge and for great evenings in places where the massive density oscillations of air originate predominantly in vibrating strings), Daniel F. (for being most inspiring in your professional purposefulness, and for sharing my sense of humour), Daniel P. (for all help and great garden parties, “so long...”), Desmond (for the most cheerful company in scientific and other life matters), Dorothe, Dorothee, Étienne, Ezequiel (for all your help and enchanting evenings of Hausmusik ¡Muchas gracias!), Fabian (for your genius in statistics, taking things with humour... and I'm sorry, Grizzly Adams lost), Frank, Ivan (for your Russian sense of humour), Jorge, Jörg M., Jörg F., Koko, Klaus (for helping out with IDEA(s) and reminders on “Latten et. al” or coffee and tea walks), Marcus (for being most helpful and cheerful and for delicious Danish hot dogs), Markus (for programming help and tea time – “zwei Meter, zwei Mark?”), Michael P., Michael S., Martina B., Martina R. (for your help and encouragement), Petra, Rüdiger (for your infectious perfectionism, inspiring illustrative talent and your universally positive attitude), Sandro (for inspiring me to cook pesto and for Italian vocabulary), Simon, Tony, Tracy, Vincent, Wieland (for your support, your expertise, “Tie-Day” and skiing company), Zaheer (for all you taught me about *in vivo* measurements, wonderful spicy Pakistani food and rides to work) and the “new PhD office”. Furthermore, I thank all other current or former members of our institute whom I could not name in person – you are a magnificent group!
- A thousand thanks are due the supervisor of my diploma thesis, Prof. Dr. Uwe Pietrzyk. It was him, who inspired me to study medical imaging physics, in the first place. He encouraged me to pursue my ideas and to never stop being curious.
- Not to forget, I couldn't have done all the work without finding balance in music. I thank my personal idol of classical music, Johann Sebastian Bach, who would compose such wonderful and stunningly complex pieces of art. Besides him Vocal Spectrum, Sting, Andy McKee, Dream Theater, Pink Floyd, Blind Guardian, Glenn Gould and Funny van Dannen helped me to survive particularly the last months of thesis writing time.

- Abdul-Rahman, H. S., Gdeisat, M. a., Burton, D. R., Lalor, M. J., Lilley, F., and Moore, C. J. (2007). Fast and robust three-dimensional best path phase unwrapping algorithm. *Applied Optics*, 46(26):6623–35.
- Abduljalil, A. M., Schmalbrock, P., Novak, V., and Chakeres, D. W. (2003). Enhanced gray and white matter contrast of phase susceptibility-weighted images in ultra-high-field magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, 18(3):284–90.
- Bakker, C. J. G., Leeuw, H. D., Vincken, K. L., Vonken, E.-j., and Hendrikse, J. (2008). Phase gradient mapping as an aid in the analysis of object-induced and system-related phase perturbations in MRI. *Physics in Medicine and Biology*, 53:N349–N358.
- Bernstein, M. a., Grgic, M., Brosnan, T. J., and Pelc, N. J. (1994). Reconstructions of phase contrast, phased array multicoil data. *Magnetic Resonance in Medicine*, 32(3):330–4.
- Bernstein, M. A., King, K. F., and Zhou, X. J. (2004). *Handbook of MRI Pulse Sequences*. Academic Press.
- Bhagwandien, R., Moerland, M. A., Bakker, C. J. G., Beersma, R., and Lagendijk, J. J. W. (1994). Numerical Analysis of the Magnetic Field for Arbitrary Magnetic Susceptibility Distributions in 3D. *Magnetic Resonance Imaging*, 12:102–107.
- Bilgic, B., Chatnuntawech, I., Fan, A. P., and Adalsteinsson, E. (2013). Regularized QSM in Seconds. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 0168.
- Bilgic, B., Pfefferbaum, A., Rohlfing, T., Sullivan, E. V., and Adalsteinsson, E. (2012). MRI estimates of brain iron concentration in normal aging using quantitative susceptibility mapping. *NeuroImage*, 59(3):2625–35.
- Bloch, F. (1946). Nuclear Induction. *Physical Review*, 70:460–473.
- Blümich, B. (2000). *NMR Imaging of Materials*. Clarendon Press, Oxford.

- Bottomley, P. A., Hart, H. R., Edelstein, W. A., Schenck, J. F., Smith, L. S., Leue, W. M., Mueller, O. M., and Redington, R. W. (1983). NMR Imaging/Spectroscopy System to Study Both Anatomy and Metabolism. *The Lancet*, July(30):273–274.
- Bouwman, J. G. and Bakker, C. J. G. (2012). Alias subtraction more efficient than conventional zero-padding in the Fourier-based calculation of the susceptibility induced perturbation of the magnetic field in MR. *Magnetic Resonance in Medicine*, 68(2):621–630.
- Bracewell, R. N. (1986). *The Fourier Transform and Its Applications*. McGraw-Hill Inc., 2nd edition.
- Brateman, L. (1986). Chemical shift imaging: A review. *American Journal of Roentgenology*, 146:971–980.
- Bronstein, I., Semendjajew, K., Musiol, G., and Mühlig, H. (2005). *Taschenbuch der Mathematik*. Verlag Harri Deutsch, nachdruck edition.
- Brown, T. R., Kincaid, B. M., and Ugurbil, K. (1982). NMR chemical shift imaging in three dimensions. *Proceedings of the National Academy of Sciences of the United States of America*, 79(11):3523–6.
- Chen, W. C., Foxley, S., and Miller, K. (2013). Investigating the orientation dependence of non-linear GRE phase evolution in White Matter using a high resolution geometric magnetic susceptibility WM model. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 21, page 0705.
- Chen, Z., Johnston, L. a., Kwon, D. H., Oh, S. H., Cho, Z.-H., and Egan, G. F. (2010). An optimised framework for reconstructing and processing MR phase images. *NeuroImage*, 49(2):1289–300.
- Cheng, Y.-C. N., Neelavalli, J., and Haacke, E. M. (2009). Limitations of calculating field distributions and magnetic susceptibilities in MRI using a Fourier based method. *Physics in Medicine and Biology*, 54(5):1169–89.
- Cusack, R. and Papadakis, N. (2002). New Robust 3-D Phase Unwrapping Algorithms: Application to Magnetic Field Mapping and Undistorting Echoplanar Images. *NeuroImage*, 16(3):754–764.
- Dagher, J., Reese, T., and Bilgin, A. (2014). High-resolution, large dynamic range field map estimation. *Magnetic Resonance in Medicine*, 71(1):105–17.
- de Rochefort, L., Brown, R., Prince, M. R., and Wang, Y. (2008). Quantitative MR susceptibility mapping using piece-wise constant regularized inversion of the magnetic field. *Magnetic Resonance in Medicine*, 60(4):1003–9.
- de Rochefort, L., Delzor, A., Guillermier, M., Houitte, D., Chaigneau, M., Déglon, N., Hantraye, P., and Lebon, V. (2010a). Quantitative Susceptibility Mapping In Vivo in the Rat Brain. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 18, page 4997.

-
- de Rochefort, L., Liu, T., Kressler, B., Liu, J., Spincemaille, P., Lebon, V., Wu, J., and Wang, Y. (2010b). Quantitative susceptibility map reconstruction from MR phase data using bayesian regularization: validation and application to brain imaging. *Magnetic Resonance in Medicine*, 63(1):194–206.
- Deistung, A., Rauscher, A., Sedlacik, J., Stadler, J., Witoszynskyj, S., and Reichenbach, J. R. (2008). Susceptibility weighted imaging at ultra high magnetic field strengths: theoretical considerations and experimental results. *Magnetic Resonance in Medicine*, 60(5):1155–68.
- Deistung, A., Schweser, F., Wiestler, B., Abello, M., Roethke, M., Sahm, F., Wick, W., Nagel, A. M., Heiland, S., Schlemmer, H.-P., Bendszus, M., Reichenbach, J. R., and Radbruch, A. (2013). Quantitative susceptibility mapping differentiates between blood depositions and calcifications in patients with glioblastoma. *PLOS ONE*, 8(3):e57924.
- D'Errico, J. (2006). polyfitn (MATLAB File Exchange, The MathWorks, Natick, USA).
- Deville, G. and Bernier, M. (1979). NMR multiple echoes observed in solid ^3He . *Physical Review B*, 19(11):5666–5687.
- Duyn, J. H., van Gelderen, P., Li, T.-Q., de Zwart, J. a., Koretsky, A. P., and Fukunaga, M. (2007). High-field MRI of brain cortical substructure based on signal phase. *Proceedings of the National Academy of Sciences of the United States of America*, 104(28):11796–801.
- Feynman, R. P., Leighton, R. B., and Sands, M. (2010). *Feynman-Vorlesungen über Physik, Band II*. Oldenbourg Verlag, München Wlen, 5. verbess edition.
- Fieremans, E., Deene, Y. D., Baete, S., and Lemahieu, I. (2009). Design of Anisotropic Diffusion Hardware Fiber Phantoms. *Diffusion Fundamentals*, 10:10.1–10.3.
- Firbank, M. J., Harrison, R. M., Williams, E. D., and Coulthard, A. (2000). Quality assurance for MRI: practical experience. *The British Journal of Radiology*, 73(868):376–83.
- Goldstein, R. M., Zebker, H. a., and Werner, C. L. (1988). Satellite radar interferometry: Two-dimensional phase unwrapping. *Radio Science*, 23(4):713.
- Graaf, R. A. D. (2007). *In Vivo NMR Spectroscopy – 2nd Edition, Principles and Techniques*. John Wiley & Sons, Ltd., Chichester.
- Griswold, M. a., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., and Haase, A. (2002). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6):1202–10.
- Gruetter, R. (1992). Fast, noniterative shimming of spatially localized signals. In vivo analysis of the magnetic field along axes. *Journal of Magnetic Resonance*, 96(2):323–334.
- Gudbjartsson, H. and Patz, S. (1995). The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34(6):910–914.
- Haacke, E. M., Brown, R. W., Thompson, M. R., and Venkatesan, M. R. (1999). *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. Wiley-Liss.

- Haacke, E. M., Xu, Y., Cheng, Y.-C. N., and Reichenbach, J. R. (2004). Susceptibility weighted imaging (SWI). *Magnetic Resonance in Medicine*, 52(3):612–8.
- Haase, A., Frahm, J., Matthaei, D., Hänicke, W., and Merboldt, K.-D. (1986). FLASH imaging: rapid NMR imaging using low flip-angle pulses. *Journal of Magnetic Resonance*, 67:258–266.
- Hahn, E. L. (1950). Spin Echoes. *Physical Review*, 80(4):580–594.
- Hammond, K. E., Lupo, J. M., Xu, D., Metcalf, M., Kelley, D. a. C., Pelletier, D., Chang, S. M., Mukherjee, P., Vigneron, D. B., and Nelson, S. J. (2008). Development of a robust method for generating 7.0 T multichannel phase images of the brain with application to normal volunteers and patients with neurological diseases. *NeuroImage*, 39(4):1682–92.
- He, X. and Yablonskiy, D. a. (2009). Biophysical mechanisms of phase contrast in gradient echo MRI. *Proceedings of the National Academy of Sciences of the United States of America*, 106(32):13558–63.
- Henderson, R. G. (1983). Nuclear magnetic resonance imaging: a review. *Journal of the Royal Society of Medicine*, 76(3):206–12.
- Hirsch, S. (2009). *Development of a Novel Method for B0 Inhomogeneity Suppression in MRI Field Maps Using Adapted Solid Spherical Harmonic Functions*. Diploma thesis, Johannes Gutenberg-Universität Mainz.
- Hirsch, S., Oros-Peusquens, A.-M., and Shah, N. J. (2009). Phase Contrast: Describing the Global Field Variations. *Conference Contribution, High Field Systems & Applications*, (31).
- Jackson, J. D. (1983). *Klassische Elektrodynamik*. Walter de Gruyter, Berlin - New York, 2nd edition.
- Jackson, J. D. (1999). *Classical Electrodynamics*. John Wiley & Sons Inc., third edit edition.
- Jain, V., Abdulmalik, O., Propert, K. J., and Wehrli, F. W. (2012). Investigating the magnetic susceptibility properties of fresh human blood for noninvasive oxygen saturation quantification. *Magnetic Resonance in Medicine*, 68(3):863–7.
- Jenkinson, M. (2003). Fast, automated, N-dimensional phase-unwrapping algorithm. *Magnetic Resonance in Medicine*, 49(1):193–7.
- Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *NeuroImage*, 17(2):825–841.
- Jenkinson, M. and Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2):143–156.
- Jezzard, P. and Balaban, R. (1995). Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine*, 34:65–73.
- Koch, K. M., Papademetris, X., Rothman, D. L., and de Graaf, R. a. (2006). Rapid calculations of susceptibility-induced magnetostatic field perturbations for in vivo magnetic resonance. *Physics in Medicine and Biology*, 51(24):6381–402.

-
- Kumar, A., Welte, D., and Ernst, R. R. (1975). NMR Fourier zeugmatography. *Journal of Magnetic Resonance*, 18:69–83.
- Lai, S., Reichenbach, J. R., and Haacke, E. M. (1996). Commutator filter: a novel technique for the identification of structures producing significant susceptibility inhomogeneities and its application to functional MRI. *Magnetic Resonance in Medicine*, 36(5):781–7.
- Langkammer, C., Krebs, N., Goessler, W., Scheurer, E., Yen, K., Fazekas, F., and Ropele, S. (2012). Susceptibility induced gray-white matter MRI contrast in the human brain. *NeuroImage*, 59(2):1413–9.
- Lauterbur, P. (1973). Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature*, 242:190–191.
- Lee, J., Shmueli, K., Fukunaga, M., van Gelderen, P., Merkle, H., Silva, A. C., and Duyn, J. H. (2010). Sensitivity of MRI resonance frequency to the orientation of brain tissue microstructure. *Proceedings of the National Academy of Sciences of the United States of America*, 107(11):5130–5.
- Li, S., Dardzinski, B. J., Collins, C. M., Yang, Q. X., and Smith, M. B. (1996). Three-dimensional mapping of the static magnetic field inside the human head. *Magnetic Resonance in Medicine*, 36(5):705–14.
- Lindemeyer, J., Oros-Peusquens, A.-M., Farrher, E., Grinberg, F., and Shah, N. J. (2011a). Orientation and Microstructure Effects on Susceptibility Reconstruction: a Diffusion Phantom Study. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 19, page 4516.
- Lindemeyer, J., Oros-Peusquens, A.-M., and Shah, N. J. (2011b). Stepwise filtering of background fields: optimising susceptibility reconstruction. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 24(1,supplement):507.
- Lindemeyer, J., Oros-Peusquens, A.-M., and Shah, N. J. (2012). Microstructure Phase Imaging at 9.4T of the Human Brain applying Stepwise Filtering of Background Fields. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 20, page 2329.
- Lindemeyer, J., Oros-Peusquens, A.-M., Vahedipour, K., and Shah, N. J. (2013). Quality-based UnwRap of SUBdivided Large Arrays (URSULA) at 9.4T. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 21, page 2494.
- Liu, C. (2010). Susceptibility tensor imaging. *Magnetic Resonance in Medicine*, 63(6):1471–7.
- Liu, J., Liu, T., de Rochefort, L., Ledoux, J., Khalidov, I., Chen, W., Tsiouris, A. J., Wisnieff, C., Spincemille, P., Prince, M. R., and Wang, Y. (2012). Morphology enabled dipole inversion for quantitative susceptibility mapping using structural consistency between the magnitude image and the susceptibility map. *NeuroImage*, 59(3):2560–8.
- Liu, T., Khalidov, I., de Rochefort, L., Spincemille, P., Liu, J., Tsiouris, A. J., and Wang, Y. (2011). A novel background field removal method for MRI using projection onto dipole fields (PDF). *NMR in Biomedicine*, 24(9):1129–1136.

- Liu, T., Spincemaille, P., de Rochefort, L., Kressler, B., and Wang, Y. (2009). Calculation of susceptibility through multiple orientation sampling (COSMOS): a method for conditioning the inverse problem from measured magnetic field map to susceptibility source image in MRI. *Magnetic Resonance in Medicine*, 61(1):196–204.
- Lodes, C., Felmlee, J., and Ehman, R. (1989). Proton MR chemical shift imaging using double and triple phase contrast acquisition methods. *Journal of Computer Assisted Tomography*, 13(5):855–61.
- Lorentz, H. A. (1916). *The Theory of Electrons*. B. G. Teubner, Leipzig.
- Marques, J. and Bowtell, R. (2005). Application of a Fourier-based method for rapid calculation of field inhomogeneity due to spatial variation of magnetic susceptibility. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 25B(1):65–78.
- Marques, J. P. R. F. (2004). *Effects of Dipolar Fields in NMR and MRI*. Phd thesis, University of Nottingham.
- Mittal, S., Wu, Z., Neelavalli, J., and Haacke, E. M. (2009). Susceptibility-weighted imaging: technical aspects and clinical applications, part 2. *American Journal of Neuroradiology*, 30(2):232–52.
- Ogawa, S., Lee, T. M., Kay, a. R., and Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, 87(24):9868–72.
- Rauscher, A., Sedlacik, J., Barth, M., Mentzel, H.-J., and Reichenbach, J. R. (2005). Magnetic susceptibility-weighted MR phase imaging of the human brain. *American Journal of Neuroradiology*, 26(4):736–42.
- Reichenbach, J. and Haacke, E. (2001). High-resolution BOLD venographic imaging: a window into brain function. *NMR in Biomedicine*, pages 453–467.
- Reichenbach, J. R. (2012). The future of susceptibility contrast for assessment of anatomy and function. *NeuroImage*, 62(2):1–5.
- Reichenbach, J. R., Venkatesan, R., Yablonskiy, D. A., Thompson, S. L., and Haacke, E. M. (1997). Theory and application of static field inhomogeneity effects in gradient-echo imaging. *Journal of Magnetic Resonance Imaging*, 7(2):266–279.
- Robinson, S., Grabner, G., Witoszynskyj, S., and Trattnig, S. (2011). Combining phase images from multi-channel RF coils using 3D phase offset maps derived from a dual-echo scan. *Magnetic Resonance in Medicine*, 65(6):1638–48.
- Robinson, S. and Jovicich, J. (2011). B0 mapping with multi-channel RF coils at high field. *Magnetic Resonance in Medicine*, 66(4):976–88.
- Roemer, P. B., Edelstein, W. a., Hayes, C. E., Souza, S. P., and Mueller, O. M. (1990). The NMR phased array. *Magnetic Resonance in Medicine*, 16(2):192–225.

-
- Ros, C., Witoszynskyj, S., Herrmann, K., and Reichenbach, J. R. (2008). Reconstruction of phase images for GRAPPA accelerated Magnetic Resonance Imaging. In *IFMBE Proceedings* 22, pages 803–806.
- Salomir, R., de Senneville, B. D., and Moonen, C. T. (2003). A fast calculation method for magnetic field inhomogeneity due to an arbitrary distribution of bulk susceptibility. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 19B(1):26–34.
- Schäfer, A., Forstmann, B. U., Neumann, J., Wharton, S., Mietke, A., Bowtell, R., and Turner, R. (2012). Direct visualization of the subthalamic nucleus and its iron distribution using high-resolution susceptibility mapping. *Human Brain Mapping*, 33(12):2831–42.
- Schäfer, A., Wharton, S., Gowland, P., and Bowtell, R. (2009). Using magnetic field simulation to study susceptibility-related phase contrast in gradient echo MRI. *NeuroImage*, 48(1):126–37.
- Scheffler, K. (2004). Fast frequency mapping with balanced SSFP: theory and application to proton-resonance frequency shift thermometry. *Magnetic Resonance in Medicine*, 51(6):1205–11.
- Schenck, J. F. (1996). The role of magnetic susceptibility in magnetic resonance imaging: MRI magnetic compatibility of the first and second kinds. *Medical Physics*, 23(6):815–50.
- Schweser, F., Deistung, A., Güllmar, D., Atterbury, M., Lehr, B. W., Sommer, K., and Reichenbach, J. R. (2011a). Non-linear evolution of GRE phase as a means to investigate tissue microstructure. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 19, page 4527.
- Schweser, F., Deistung, A., Lehr, B. W., and Reichenbach, J. R. (2011b). Quantitative imaging of intrinsic magnetic tissue properties using MRI signal phase: an approach to in vivo brain iron metabolism? *NeuroImage*, 54(4):2789–807.
- Sehgal, V., Delproposto, Z., Haddar, D., Haacke, E. M., Sloan, A. E., Zamorano, L. J., Barger, G., Hu, J., Xu, Y., Prabhakaran, K. P., Elangovan, I. R., Neelavalli, J., and Reichenbach, J. R. (2006). Susceptibility-weighted imaging to visualize blood products and improve tumor contrast in the study of brain masses. *Journal of Magnetic Resonance Imaging*, 24(1):41–51.
- Shannon, C. (1949). Communication In The Presence Of Noise. *Proceedings of the IRE*, 37(1):10–21.
- Shewchuk, J. R. (1994). An Introduction to the Conjugate Gradient Method Without the Agonizing Pain.
- Shmueli, K., de Zwart, J. a., van Gelderen, P., Li, T.-Q., Dodd, S. J., and Duyn, J. H. (2009). Magnetic susceptibility mapping of brain tissue in vivo using MRI phase data. *Magnetic Resonance in Medicine*, 62(6):1510–22.
- Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, 17(3):143–55.

- Sommer, K., Schweser, F., Deistung, A., and Reichenbach, J. R. (2012). Investigating non-susceptibility contributions to GRE phase contrast. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 20, page 2331.
- Strand, J. and Taxt, T. (2002). Two-dimensional phase unwrapping using robust derivative estimation and adaptive integration. *IEEE Transactions on Image Processing*, 11(10):1192–1200.
- Strand, J., Taxt, T., and Jain, a. K. (1999). Two-dimensional phase unwrapping using a block least-squares method. *IEEE transactions on Image Processing*, 8(3):375–86.
- Ullisch, M. (2012). *A Navigator-Based Rigid Body Motion Correction for Magnetic Resonance Imaging*. PhD thesis, RWTH Aachen University.
- Weckesser, M., Langen, K. J., Rickert, C. H., Kloska, S., Straeter, R., Hamacher, K., Kurlemann, G., Wassmann, H., Coenen, H. H., and Schober, O. (2005). O-(2-[¹⁸F]fluorethyl)-L-tyrosine PET in the clinical evaluation of primary brain tumours. *European Journal of Nuclear Medicine and Molecular Imaging*, 32(4):422–9.
- Wharton, S. and Bowtell, R. (2010). Whole-brain susceptibility mapping at high field: a comparison of multiple- and single-orientation methods. *NeuroImage*, 53(2):515–25.
- Wharton, S. and Bowtell, R. (2012). Fiber orientation-dependent white matter contrast in gradient echo MRI. *Proceedings of the National Academy of Sciences of the United States of America*, 109(45):18559–64.
- Wharton, S., Schäfer, A., and Bowtell, R. (2010). Susceptibility mapping in the human brain using threshold-based k-space division. *Magnetic Resonance in Medicine*, 63(5):1292–304.
- Wiesinger, F., Van de Moortele, P.-F., Adriany, G., De Zanche, N., Ugurbil, K., and Pruessmann, K. P. (2004). Parallel imaging performance as a function of field strength—an experimental investigation using electrodynamic scaling. *Magnetic Resonance in Medicine*, 52(5):953–64.
- Windischberger, C., Robinson, S., Rauscher, A., Barth, M., and Moser, E. (2004). Robust field map generation using a triple-echo acquisition. *Journal of Magnetic Resonance Imaging*, 20(4):730–4.
- Witoszynskyj, S., Rauscher, A., Reichenbach, J. R., and Barth, M. (2009). Phase unwrapping of MR images using Phi UN – a fast and robust region growing algorithm. *Medical Image Analysis*, 13(2):257–68.
- Wu, Z., Mittal, S., Kish, K., Yu, Y., and Hu, J. (2010). Identification of Calcification with Magnetic Resonance Imaging Using Susceptibility-Weighted Imaging: A Case Study. *Journal of Magnetic Resonance Imaging*, 29(1):177–182.
- Yushkevich, P. A., Piven, J., Cody Hazlett, H., Gimpel Smith, R., Ho, S., Gee, J. C., and Gerig, G. (2006). User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *NeuroImage*, 31(3):1116–1128.
- Zeeman, P. (1882). Über einen Einfluss der Magnetisierung auf die Natur des von einer Substanz emittierten Lichtes. *Verhandlungen der Physikalischen Gesellschaft zu Berlin*, page 127.

Zhao, Y., Anderson, A. W., and Gore, J. C. (2005). Computer simulation studies of the effects of dynamic shimming on susceptibility artifacts in EPI at high field. *Journal of Magnetic Resonance*, 173(1):10–22.

Appendix



Formulae

A few definitions and relations that are implicitly used throughout this thesis are defined in the following:

Mean value: $\text{mean}([x_1, \dots, x_n]) = \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$

Standard deviation: $\text{std}([x_1, \dots, x_n]) = \sigma_x = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2}$

Variance: $\text{var}([x_1, \dots, x_n]) = \sigma_x^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2$

Median: $\text{median}([x_1, \dots, x_n])$: Determine the central value of the ascendingly ordered values $x_1 \leq x_2 \leq \dots \leq x_{n-1} \leq x_n$. For odd n , this is $x_{(n-1)/2}$, for even n , the mean of the two central values is calculated $((x_{n/2} + x_{(n/2)+1})/2)$.

Rounding: $\text{round}(x)$: rounds x to the nearest integer value (e.g. $\text{round}(3.6) = 4$, $\text{round}(10.3) = 10$).

Convolution: The convolution of two functions, f and g , is defined by:
 $(f * g)(\vec{r}) = \int_{-\infty}^{\infty} f(\vec{r}')g(\vec{r} - \vec{r}') d\vec{r}'$.

Convolution Theorem: $\mathcal{F}(f * g) = \mathcal{F}(f) \cdot \mathcal{F}(g)$

For more detailed explanations please consult e.g. Bronstein et al. [2005].

B

Phase Imaging – Images From a Sample Case

The data illustrated on the following pages were acquired in a patient study at a 3 T hybrid (Siemens AG, Erlangen, Germany) *PET-MR* scanner. An overview matrix on phase for different slice positions and echo times is shown in Fig. B.1. Fig. B.2 shows the same dataset after masking was applied. Fig. B.3 shows the sum-of-squares of the phase gradients. In Fig. B.4 the noise variation of the phase data is visualised. The noise is calculated using mean-value corrected variance estimation as in Eq. 3.24.




Figure B.1: Example of an *in vivo* measurement from a patient study. From left to right: slices from inferior to superior (every 6th slice shown); From top to bottom: 1st, 3rd, 5th, 7th, 9th and 11th echo with echo times: 3.73, 11.75, 19.77, 27.79, 35.81, and 43.83 ms. The study focussed on the T_2^* decay of the signal, thus many echoes were recorded. Gray scale colourmap reaches from $-\pi$ to π .




Figure B.2: *in vivo* example, as shown in Fig. B.1, but confined by a brain mask. Noisy phase is excluded from the valid phase region.




Figure B.3: Sum-of-squares of the phase gradient, $\sqrt{(\partial_x \varphi)^2 + (\partial_y \varphi)^2 + (\partial_z \varphi)^2}$ for different slices (left to right) and echo times (top to bottom) according to the phase data shown above (Fig. B.1). Value range is $[0, 2\pi]$ in grayscale.




Figure B.4: Phase variance for different slices (left to right) and echo times (top to bottom) according to the phase data shown above (Fig. B.1).



FSE – Function Overview

The FSE function set, written in MATLAB® (The MathWorks Inc., Natick, USA), splits into fundamental functions, direct inversion functions and the minimisation set. Fig. 7.16 outlines the FSE framework.




Figure C.1: Scheme of the FSE framework. The functions on the left are interface functions meant for the user to call. The functions on the right embody the core functionalities, called by the interface functions.

The use and properties of the functions will be outlined in short including only their most important parameters:

generateDipoleK

$x, y, z, \vartheta, \varphi, s_x, s_y, s_z, \text{shift}, \text{spat}$

This core function generates discretised dipole kernels of size $x \times y \times z$ in k -space. The generation can either be performed directly in k -space or by Fourier-transform of a spatial dipole representation (`spat`). The rotation of the dipole with respect to the z axis can be defined: θ tilts the kernel about x towards the positive y axis, while ϕ rotates about the z axis (azimuth). Further, the kernel can be calculated for non-isotropic axis scales, s_x, s_y, s_z . Fourier shifting (circular shift of the array, so zero resides at index [1,1,1]) can be included.

generateDipoleKInv

$x, y, z, \vartheta, \varphi, s_x, s_y, s_z, \text{shift}, \text{spat}, t_{\text{cone}}, \text{mode}$

This function generates the inverse dipole kernel of size $x \times y \times z$ for direct inversion. `generateDipoleK` creates the kernel template, that is inverted, corrected for divisions by zero and truncated at t_{cone} . The truncated areas are – according to the selected mode – filled with zeros or the constant value, t_{cone} .

convoluteDipole

$\chi, \vartheta, \varphi, s_x, s_y, s_z, \text{spat}$

A susceptibility map, χ , is convoluted with the dipole kernel in frequency domain. The dipole specifications follow the parametrisation which is passed to `generateDipoleK`.

deconvoluteDipole

$b, \vartheta, \varphi, s_x, s_y, s_z, \text{shift}, \text{spat}, t_{\text{cone}}, \text{mode}$

This is the user-side handler for direct inversion. The parametrisation is passed to `generateDipoleKInv`.

structSusReco

$b, W, m, \text{angle(s)}, \text{regmode}, \lambda, \mu, W_T, W_{Gx(y,z)}, s_x, s_y, s_z, \text{spat}, \text{cosmoscnt}, \text{pad}, \text{padsize}, \text{CGparameters}, B_0, \omega_{\text{off}} \Rightarrow \text{conf}$

All configuration data is kept in a structure, `conf`, generated by this function. The core element is the fieldmap, b . The minimisation problem is parametrised by the `regmode` (no regularisation, Tikhonov, gradient or both), the parameters, λ and μ , and the weighting matrices, W, W_T and $W_{Gx(y,z)}$ as well as an optional mask, m . The problem geometry is parametrised by the scaling parameters, $s_{x(y,z)}$, and the `angle` (containing θ and φ). If a multi-angle mode (`cosmoscnt`) is chosen, the `angle` parameter is expected to contain as many angle tuples as fields were recorded, while b is expected to contain the corresponding (coregistered!) fields stacked in the 4th dimension. A zero-padding parameter is also included (`pad, padsize`). The `CGparameters` consist of: `iterations, reset, reslim, chlim`. B_0 field strength can be defined to enable quantitative result scaling and ω_{off} is an optional off-resonance to be subtracted before minimisation.

runSusRecoL2

`conf`, χ_{guess}

Handling the minimisation procedure, this function expects the complete configuration structure (`conf`), as supplied by `structSusReco`. The handler will first check and probably correct the configuration using `susCheckPrepConf`, then generate the constant term of the derivative of the minimisation equation with `evaluateSusL2DeriveConst` and after collecting all information, call `computeConjugateGradients3Sus` to compute the solution. Here an optional initial guess, χ_{guess} can be passed to the minimiser. After minimisation has finished, scaling of the result to the preconfigured field strength, B_0 , is performed. If zero-padding was applied, the result is unpadded before returning.

susCheckPrepConf

`conf`

Configuration is checked for integrity in this function. Some private matrices for minimisation, such as W^2 , are pre-calculated. The cosmos count, if applicable, is checked for consistency regarding angle-tuple count and number of fields passed over. Eventual padding is applied, using the configuration parameter, `padsize`. Most importantly, the dipole kernel – or for cosmos mode several dipole kernels – is calculated according to the angle and scaling information.

evaluateL2DeriveConst

`conf(b,  $\hat{d}$ ,  $W$ )`

According to the selected minimisation mode, the constant term of the minimisation equation is calculated once.

computeConjugateGradients3Sus

`conf(CGparameters)`

The core minimisation function is the conjugate gradients solver. Besides the constant term of the minimisation equation and an optional initial guess, the function requires several methodical parameters (`CGparameters`) such as the iteration count and the residual limit for stopping. Further a solution change limit can be defined, instead of or in combination with the residual limit. A solution vector reset count can also be specified. Several debugging or view options are offered. Each iteration step calls `evaluateL2DeriveFunction` with the full parametrisation structure several times, to obtain the momentary solution and residual. The function returns the minimised solution and optionally a residual map.

evaluateL2DeriveFunction

conf

The objective function of the minimisation is usually only called by `computeConjugateGradients3Sus` and embodies calculations for all possible minimisation modes. The terms added to the result are chosen regarding regularisation type (none, Tikhonov, gradients or both), application of mask and weighting matrices. For gradient regularisation, the gradients are evaluated inside of this function. In case of COSMOS computing, the result vector is extended in 4th dimension to carry the computed results for different angles.

runSusRecoCFS

conf

This is the direct minimisation inversion pendant to `runSusRecoL2` and was added to the framework just recently. It computes the minimisation in the closed-form solutions suggested by Bilgic et al. [2013], Eq. 7.43.

Most of the framework is self-explaining. The FSE GUI was designed as a quick test interface for susceptibility reconstruction techniques. It expects at least a fieldmap, b , as input, optionally also a mask, m , and for enabling the weighting methods the corresponding matrices, W , W_T and $W_{Gx(.y,z)}$. A snapshot of the running application is shown in Fig. C.2




Figure C.2: Snapshot of the FSE GUI showing a Tikhonov-regularised minimisation performed on a post mortem measurement at 3 T.

MUBAFIRE – Function Overview

Fig. D.1 shows the organisation structure of the MUBAFIRE framework. The core functions will be shortly described in the following table.




Figure D.1: Schematic workflow of background field correction with the MUBAFIRE framework. The left hand side shows the interface functions, while the right hand side shows the separate background removal methods applied on the fieldmap, *b*, stored in the interface function. The green methods in the bottom manipulate the mask, *m*. The blocks possessing a yellow glow, are the ones used in the optimised MUBAFIRE workflow, presented before. Note that the implementations of SPHINX and GaussianIntelligentNormalisation are based on code from Sebastian Hirsch, [Hirsch, 2009]

structFieldHom

$b, m, \text{mode}, (\vartheta, \phi),$
 $\text{resample}, \text{filter}$
 parameters
 $\Rightarrow \text{conf}$

This function generates the filter structure, conf , containing the source fieldmap, b , the mask, m , and the orientation of the FOV. The crucial component is the mode vector, containing an ordered list of numbers indicating the filters to be applied in each step. Further, configuration parameters for all filters are included into the the structure. The resample parameter decides, whether the correction process is performed at full resolution or just a fraction of it: $0 \rightarrow 2^0 = 1$ stands for full resolution, $1 \rightarrow 2^1 = 2$ for half, $2 \rightarrow 2^2 = 4$ quarter and so on.

runFieldHom

conf

The executing interface function expects the configuration structure (generated by structFieldHom), containing all information for executing the background correction. First, resampling is applied, if desired (resample). This means that the entire array as well as the mask are downsampled to $1/2^{(\text{resample})}$ of the native resolution. A runtime version of the (eventually downsampled) field and mask is used for applying the filters: b_d, m_d , while a reference field b_{ref} is kept in original (downsampled) form. The mode vector (mode) is opened, containing the r algorithm identifiers for r filtering steps. The identifiers are integer values, described in the mode table at the end of this section. Each filter is applied on the runtime field, b_d , and mask, m_d , changing them in every step. After all filters of the requested chain have been applied, the background field is determined by:

$$b_{\text{BG}} = b_{\text{ref}} - b_d \quad . \quad (\text{D.1})$$

The field is finally resampled to native resolution, $b_{\text{BG}} \rightarrow B_{\text{BG}}$, and subtracted from the input field, b , to determine the background-corrected field map:

$$b_{\text{hom}} = b - B_{\text{BG}} \quad . \quad (\text{D.2})$$

The filter can further produce debug output in form of cross sections and normalised histograms.

rudimentary filters

$\text{MEDwidth}, \text{MEAwidth},$
 GSwidth

Gaussian, median and mean filters. According to their name, the functions are called with the enlisted parameters from the configuration structure (conf).

GaussianIntelligentNormalisation

IGSwidth

Intelligent Gaussian filter (with edge-compensation, GF) called with standard deviation, IGSwidth .

fieldHomPoly

POLorder

Polynomial filter (POLF) called with the intended polynomial fitting order (POLorder).

SPHINX SPorder	Calling the spherical harmonic filter, SPHINX, with a maximum harmonic order of SPorder
SPHINXapprox SPAorder, SPAmask	Calling the SPHINX modification using a reference mask, SPAmask, for orthonormalisation of the harmonics (or loading of pre-calculated functions).
homFieldDipMin DIPpadfac, DIPmode, angle	Calling the pseudo-susceptibility background correction (DIPF), with a padding factor, DIPpadfac, the mode of dipole generation (spatial or frequency domain dipole source), DIPmode, and most importantly the angle of the z-axis towards B_0 .

MUBAFIRE mode table:

Mode	Algorithm
-1	No mode – stop
0	External function (expects handle)
1	Median Filter
2	Mean Filter
3	Gaussian Filter
4	Intelligent Gaussian Filter (GF)
5	Polynomial 3D
7	Spherical harmonics (SPHINX)
8	Spherical harmonics with reference mask (SPHINXapprox)
9	Dipole Filter (DIPF)
31	Thresholding
32	BroadenMaskHole (selective erosion filter)
33	Mask erosion (6 neighbour)
34	Mask dilation (6 neighbour)
[5 7 9]	MUBAFIRE
[5 7 9 31 32 33 (34) 9]	MUBAFIRE Local

A screen shot of the optional debugging output of the MUBAFIRE base algorithm is shown in Fig. D.2. For each individual step three orthogonal central slice plots and a width-normalised histogram for the masked region are shown.




Figure D.2: Example of the visual debugging output generated by the MUBAFIRE base algorithm. From left to right the raw data and the steps, POLF, SPHINX and DIPF, are shown. The top three rows show orthogonal central slices of the corrected volume, while the lowest row shows a width-normalised histogram with x-axis units: Hz.