Harmonic Phase Processing in Magnetic Resonance Susceptibility Imaging

by

Ryan James Simpson Topfer

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

Department of Biomedical Engineering
University of Alberta

© Ryan James Simpson Topfer, 2014

Abstract

The popularity of magnetic resonance imaging (MRI) owes much to its flexibility. Sensitive to a host of different biophysical phenomena, parameters of the scan can be fine-tuned to highlight specific pathologies. One such mechanism for generating image contrast is magnetic susceptibility—the material property that defines how an object will distort an applied magnetic field such as that of the MR scanner. In functional MRI (fMRI), for instance, the unique magnetic signatures of oxygenated (diamagnetic) and deoxygenated (paramagnetic) blood are what permit the indirect measure of neuronal activity. However, rather than measuring the susceptibility itself, fMRI registers haemodynamic change as subtle gains and losses in the signal magnitude due to haem-iron induced dephasing of the proton spins. Susceptibility mapping, an emerging area in MRI, attempts to retrieve a direct measure of bulk susceptibility itself. Generally this is done, not by means of the standard magnitude image, but through the phase component of the signal which, in the idealized case, relates the magnetic field perturbation by a simple multiplicative constant.

Several issues interfere with the construction of accurate susceptibility maps. Foremost is the obfuscation of the small-scale "local" field (i.e., that pertaining to susceptibility variation within tissue) by the so-called "background" field, which owes predominantly to the comparatively substantial susceptibility shift between tissue and air. Whether the goal is to produce qualitatively useful images of the local field, or to map the susceptibility itself, the local field must first be isolated from the background. To isolate overlapping signals which are a priori unknown, the unique characteristics of the expected signals need to first be codified. One family of methods for isolating the local field begins by asserting that, away from air-tissue interfaces, the background field should satisfy the partial differential equation of Laplace, whereas the local field should satisfy that of Poisson (viz., the background field should be harmonic, the local field non-harmonic). This classification informs a filtering technique based on the spherical mean value (SMV) property of harmonic functions: the mean of a harmonic function calculated over a spherical surface equates to the specific value taken by the function at the centre of the sphere.

This thesis introduces another property of harmonic functions to the task of local field estimation: a harmonic function can be locally expressed by means of a convergent power series (viz., it is an analytic function). This property is first employed to characterize the SMV kernel as an estimator for the central field value when the field is discretized. Analysis reveals that when the field data is of a finite spatial resolution, the aim of accurate elimination of the background field via the SMV is fundamentally at odds with the aim of preserving the local field. Fortunately, given the rapid decay of the background field and its derivatives with distance from its field sources, the discrete SMV is nevertheless a robust estimator for field geometries such as those observed in the brain.

In addition to the problem of finite image resolution, MR imaging of the head has finite spatial support as signal is generally absent in the skull and ever-absent in the surrounding air. The SMV cannot be used to estimate the background field wherever the spherical kernel extends beyond the edges of this support. Hence, in conventional SMV-filtering, field points within a distance from the edges equal to the radius of the employed kernel are discarded outright. This is a considerable obstacle to a number of clinical applications as it precludes analysis of features such as subdural haematomas and cortical lesions in pathologies such as multiple sclerosis. To recover local field across the edges of the brain, this study presents an extension to conventional SMV-filtering by appealing to the analytic nature of the background field: by obtaining an initial SMV-estimate of the harmonic background field over a reduced inner portion of the brain, a three-dimensional Taylor expansion was performed to extend field coverage to the edges of the brain. The method is quantitatively assessed through a numerical experiment and qualitatively demonstrated on in vivo human brain data acquired at 4.7 T. Using a kernel radius typical of conventional methods, the extension recovered on average 26 % more in vivo brain volume.

Preface

The third chapter of this thesis has been adapted from a collaborative work: Topfer R, Schweser F, Deistung A, Reichenbach JR, Wilman AH. SHARP Edges: Recovering Cortical Phase Contrast Through Harmonic Extension. *Magnetic Resonance in Medicine*; article first published online March 3, 2014 DOI: 10.1002/mrm.25148. Contributions from co-authors were essentially limited to critical reviews of drafts of the manuscript. F. Schweser's assistance in this regard as well as his insightful discussions are of particular note. The original experiment was conceived by the author; the lion's share of the computer programming and writing of the manuscript was also done by the author.

The work makes use of in vivo MRI data for which ethics approval was granted by the University of Alberta Research Ethics Board as part of the project "Evaluation of MRI in normal volunteers and patients with suspected neurological disease", study ID Pro00000906, re-approved as of July 23, 2013.

Acknowledgements

Specific thanks are owed to my supervisor Alan Wilman, for the support and encouragement; to my friend and colleague Hongfu Sun, for sharing several of the image reconstruction programs used in this thesis; to Mr. and Mrs. Sandhu of Stir Krazy, for supplying the innumerable curries consumed during the construction of this thesis; and, most of all, to my mom, for always sharing her lunch.

Contents

List of Figures

Chapter 1

Introduction: Physics, physiology, and the machine

Wherein we proceed in an orderly, coherent fashion: Beginning from the top with an analogy—the moment arm of the spinning top—we then pause for a moment to consider what sits atop our own shoulders before proceeding to get to the bottom of things, the ground state—incoherent and disordered.

> In the beginning... the earth was without form, and void; and darkness was upon the face of the deep. And the Spirit of God moved upon the face of the waters. And God said, Let there be light: and there was light.
>
> -Genesis:1:1-2(KJV)

A flippant remark—the celestial switch was flipped, and delivered was the spectral blessing of electromagnetism. A phenomenal induction, and the world was thereby amenable to primordial experimentation in nuclear magnetic resonance (NMR).

1.1 The proton top's spin

In 1922, Stern and Gerlach performed the now famous experiment whereby a beam of electrically neutral silver atoms was directed through an orthogonally applied magnetic gradient field $\nabla\mathbf{B}$. Rather than pass straight through, or spread continuously through a range of angles as per the expected behaviour under classical mechanics, the particles were observed to deflect in but 2 directions. The interpretation to emerge later was that the silver atom possessed a net magnetic moment $\boldsymbol{\mu}$ due to the inherent spin angular momentum of its unpaired electron and, when placed in a spatially varying magnetic field, it would therefore experience a force $\mathbf{F}$ equal to the negative gradient of the system's potential energy U:

$$\mathbf{F} = -\nabla U = -\nabla(\boldsymbol{\mu} \cdot \mathbf{B}). \tag{1.1}$$

Like the spinning top which, rather than collapsing abruptly under the force of gravity, will precess around the surface normal in order to conserve angular momentum, the nuclear spin, with its intrinsic angular momentum, will precess about the direction of an applied magnetic field. However, rather unlike the gyroscopic top, the orientation of a nuclear magnetic moment in an applied field can assume one of a discrete set of values (i.e. quantum states).

The proton as would later be demonstrated is, like the electron, a spin 1/2 particle, for which a strict binary of quantum states are available: spin "up" or spin "down". That is, the spin angular momentum of the particle is $\pm\hbar/2$, where $\hbar$ is Planck's constant ($\sim 6.626 \times 10^{-34}\,\mathrm{m^2\,kg\,s^{-1}}$) divided by 2π. In turn, this determines the magnetic moment through multiplication with an experimentally determined proportionality term known as the gyromagnetic ratio γ ($\sim 42.576\ \mathrm{MHz\,T^{-1}}$ for protons):

$$\mu_z = \begin{cases} \mu_{(+)} = \gamma\hbar/2, & \text{spin up} \\ \mu_{(-)} = -\gamma\hbar/2, & \text{spin down.} \end{cases} \tag{1.2}$$

In the classical picture, the product of γ with the applied field strength yields the precessional frequency ω of the spins, also known as the Larmor frequency.

At thermal equilibrium, the probability of finding a particle in a given state is given by a Boltzmann distribution. The most probable configuration corresponds to that which possesses the lowest energy. In the absence of an applied field, the two states are effectively degenerate and therefore equally probable, with transitions between states governed by fluctuations in their local magnetic environment (fluctuations which, in aqueous environs, are essentially random due to Brownian motion). Upon the application of an external magnetic field, however, an energy differential known as the Zeeman effect arises. If the field $\mathbf{B}$ is applied along the z direction with magnitude B_0, labeling the absolute temperature as T, the Boltzmann constant as k_B, and denoting the number densities of $\mathbf{B}$-aligned (spin up) and the $\mathbf{B}$-antialigned (spin down) proton spins by $\rho_{(+)}$ and $\rho_{(-)}$ respectively, the ratio of state-occupancies is then

$$\frac{\rho_{(+)}}{\rho_{(-)}} = \frac{exp(-U_{(+)}/k_BT)}{exp(-U_{(-)}/k_BT)} = exp(\gamma\hbar B_0/k_BT) \approx 1 + \frac{\gamma\hbar B_0}{k_BT} + ... \tag{1.3}$$

where the rightmost expression has involved a Taylor expansion. Since the magnetic energies of the nuclear systems studied in MRI are much less than the typical thermal energies k_BT (e.g. $T \sim$ room temperature, $B_0 \sim 1$ Tesla), the linear approximation should suffice [1]. Thus, the spin ensemble has a slight preference for net alignment with the applied field, such that if $\rho_{(-)} \approx \rho_{(+)} \approx \rho/2$, where ρ is the total proton spin density, then $\rho_{(+)} - \rho_{(-)} \approx \gamma\hbar B_0/2k_BT$ which elicits a net magnetization $\mathbf{M}$ (dipole moment per unit volume). Since $\mathbf{B} = B_0\hat{\mathbf{z}}$, where $\hat{\mathbf{z}}$ is the unit vector in the z direction, the transverse components of $\boldsymbol{\mu}$ continue to be governed by random thermal motion; averaged over a volume, they produce no net magnetization. Hence, the magnetization is simply given by the product of the

overabundance of "up" spins with their associated dipole moment:

$$\mathbf{M} = M\hat{\mathbf{z}} = (\rho_{(+)} - \rho_{(-)})\mu_{(+)}\hat{\mathbf{z}} = \frac{\gamma^2\hbar^2 B_0 \rho_0}{4k_B T}\hat{\mathbf{z}}. \tag{1.4}$$

The material property that relates the resulting macroscopic magnetization to the applied induction field is known as the magnetic susceptibility χ. In general, the magnetic susceptibility of a given material, and ultimately its magnetization, will be dominated by effects owing to the interactions among its electron spins, with the protons playing a much smaller role. Denoting the permeability of free space as μ_0,

$$\mathbf{M} = \frac{\chi \mathbf{B}}{\mu_0(1+\chi)} + \mathbf{M}_0, \tag{1.5}$$

where the last term captures any remanent (a.k.a. permanent) magnetization that might be present in the absence of an applied field. So, apply a magnetic field to a medium of non-zero magnetic susceptibility and the result is a net magnetization. Apply a magnetic field to a collection of unpaired proton spins and the result is a net magnetization along the direction of the field, with the spins precessing at the Larmor frequency:

$$\omega_0 = \gamma B_0, \tag{1.6}$$

an important relation, to which we will return momentarily.

1.2 The Iron Age and aging with iron

In an alternative scenario, in the beginning, there was Alan Hills 84001.

Alan Hills 84001, affectionately known as AH84001, is a Martian meteorite collected in 1984 in Allan Hills Antarctica by a crew of American geologists. The rock largely owes its fame to a paper published twelve years later in the acclaimed journal Science: "Search for Past Life on Mars: Possible Relic Biogenic Activity in Martian Meteorite ALH84001" [2]. As the title suggests, a number of chemical and structural similarities were found between the features of the meteorite and established fossils. Publications followed in other journals of high-repute, wherein new analyses were performed, again leading the authors to favour the pro-Martian-life hypothesis. One of the main arguments advanced by these pro-life parties was the morphologic similarity between the linear chains of the mineral magnetite (Fe_3O_4) observed in the meteorite and those produced by earthly bacteria [3].

Though the scientific debate regarding the genesis of these structures may be (however surprisingly) ongoing [4] it is certainly true that multitudinous molecular forms and functions of iron exist in living creatures. At the unicellular level, magnetotactic bacteria contain linear chains of magnetite which serve as a "biological bar magnet" [5], passively directing them toward the geomagnetic poles (a phenomenon known as magnetoreception) [6]. Magnetite-based magnetoreception is believed to play a role in the navigational abilities of birds and fish during migration [7, 8]. To access nutritious endolithic algae (embedded in choral or

rock), the dental enamel of the chiton mollusc is also made of magnetite (harder than our own hydroxyapatite enamel by about 1 Mohs point). Indeed, based on the mineralogical properties (e.g. crystal perfection and chemical purity) generally observed in the biogenic form of the mineral, Kirschvink and Hagadorn [9] posit that the ability of the varied phyla to synthesize this biomineral suggests a common ancestor dating back several hundred million years.

Our understanding of geologic history owes much to the stability of the ordered magnetic domains of magnetite under normal atmospheric conditions. Magnetite was the first material discovered to manifest a bulk magnetism—a phenomenon which owes to the unpaired electron spins of its ferric iron. Like the nuclear paramagnetism of the aforementioned ensemble of proton spins, the magnetic moments of magnetite arrange with a net alignment in the direction of an applied field, however, below the Curie temperature (858 K for magnetite) the moments are essentially frozen in place and will remain as such should the external magnetizing field be removed (i.e. positive $\mathbf{M}_0$ in (1.5)). Under normal atmospheric conditions, magnetite is therefore classified as a ferrimagnet. Beyond the Curie temperature, however, this magnetic ordering breaks down as the randomizing influence of thermal interactions win out. Beyond the Curie temperature, which applies to all materials possessing remanent magnetism, magnetite becomes paramagnetic ($0 < \chi < 1$), such that a net alignment of the magnetic moments becomes possible only in the presence an external magnetic field.

Whether the mineral exists in human tissue in any significant quantity (or at all) remains an open question. One study reported finding it (or a mineral analogue) "unequivocally" within the cores of Alzheimer's plaques [10]. An electron microscopy study reported finding small magnetite particles, 50 nm in diameter, in post mortem brain tissue [5]. However, an NMR relaxometry-based study [11] examined samples of brain ferritin from Alzheimer's patients and healthy controls for the existence of magnetite within the protein core and found nothing to support the hypothesis that magnetite might be stored there—concentrations above 1% were effectively ruled out. Quoting Schenck [12], in reference to the aforementioned the electron microscopy study, "As with other such studies, additional confirmation and studies to rule out an exogenous source for these particles is desirable. Such small particles cannot produce MR imaging artifacts, at least using conventional pulse sequences, and if ferromagnetic particles much larger than this were present it is likely they could be detected in this way. Such artifacts are not observed."

Though iron is an essential nutrient, necessary for the synthesis of DNA, for the transport of oxygen in blood, and for basic cell metabolism, free iron is highly reactive and therefore toxic. Complex mechanisms have evolved to regulate iron reserves. Within the cell, this is achieved via the iron storage protein ferritin. Unlike magnetite, ferritin does not possess remanent magnetism and is merely paramagnetic. A territory of the midbrain known as the substantia nigra ("black substance", for its high concentration of the pigment neuromelanin) is responsible for producing much of the brain's dopamine—a neurotransmitter understood to be involved in so-called reward pathways, and in initiating and modulating voluntary

movement via the nigrostriatal pathway. Ferrous iron is a crucial cofactor for the synthesis of L-DOPA (the precursor to dopamine), hence, the substantia nigra tends naturally to be rich in ferritin iron reserves. In Parkinson's disease, however, the dopaminergic cells of the substantia nigra die off, with the reduced striatal innervation ultimately leading to the motor dysfunction that characterizes the disease. Colocated with the cell death tend to be increased levels of non-haem iron and substantially decreased levels of neuromelanin, which is known to have a high affinity for binding free iron. A popular theory of the degenerative mechanism behind Parkinson's disease has it that the nigral cell death is caused by iron-induced oxidative stress [13, 14].

The other principal contributor to regional susceptibility variation in the nervous system is myelin, which is diamagnetic ($\chi < 0$), thus tending to repel an applied magnetic field. Maintained by oligodendrocytes in the central nervous system and by Shwann cells in the peripheral, myelin is the dielectric sheath wrapped around nerve axons which lowers the effective capacitance, thereby increasing the propagation speed of neural (electrical) impulses, known as action potentials. When myelin is damaged in disorders such as multiple sclerosis (MS) it impinges upon the ability of neurons to propagate action potentials. Depending on the degree of damage and the parts of the nervous system affected, the results can be debilitating or fatal. As in Parkinson's, colocated with the lesions tends to be an accumulation of a form of non-haem iron. Though MS is commonly thought of as a white matter disease, iron accumulation in grey matter structures such as the thalamus may also be involved [15, 16].

There is evidence to suggest other neurodegenerative disorders such as Huntington's disease [17], amyotrophic lateral sclerosis [18], and vascular dementia [19, 20] also have associated changes in the magnetic susceptibilities of neural structures. Yet, there is currently no available technique for in situ quantification of biologic magnetic susceptibility.

1.3 Ludwig Boltzmann vs. the RF amplifier

> And God said, Let there be a firmament in the midst of the waters, and let it divide the waters from the waters.
>
> -Genesis:1:6(KJV)

The human body is over 50% water by mass. For the purpose of MRI, this is a particularly good thing: the biologic ubiquity of hydrogen (its nucleus consisting of a single proton), with its relatively large gyromagnetic ratio, offer systems which can be probed by NMR.

The bulk of the typical clinical MR scanner consists of a massive superconducting solenoid. The sole purpose of this powerful static field is to rend apart the energy levels of the proton, and thereby increase the bulk magnetization (1.4). Simply stated, the goal in conventional MRI is to map the sample (i.e., subject/object to be imaged) magnetization, itself largely dependent on the density of protons. The magnetization evokes a

secondary "response" field (to which a mathematical form will be given in Chapter 2). Any attempt to measure the response field outright, however, is doomed to fail: the longitudinal field applied to magnetize the sample being on the order of a billion times greater than the response field in the same direction (the proverbial "needle in a haystack"). Acquisition of the signal must therefore occur, away from $B_0\hat{\mathbf{z}}$, in the transverse plane.

To generate a coherent transverse signal the sample is irradiated for a brief interval by a radio frequency (RF) pulse with a central frequency held at the Larmor rate of the precessing spins (1.6). The effect of this *resonant* pulse is to synchronize the spins and, in the classical picture, to pull the magnetization vector into the transverse plane. Through Faraday's law, the rotating transverse magnetization varies the transverse magnetic flux density which can then be detected by means of a conducting coil tuned to the Larmor frequency. Left at that, the measured signal (an oscillating voltage induced in the receiver coil) is referred to as the "free induction decay" (FID). Ignoring the electronics of the receiver system, the FID envelope is governed by the initial magnitude of the tipped magnetization and by two relaxation times (T_2 and T_2') associated with the loss of phase coherence as the diffusing spins sample various magnetic environments. The phase accumulation of a given spin over the time course $t = TE$, during which it moves through the variable B_z field with trajectory $\mathbf{r}(t)$, is given by the Larmor equation (1.6):

$$\phi(TE) = -\int_0^{TE} \omega(\mathbf{r}(t))dt = -\gamma \int_0^{TE} B_z(\mathbf{r}(t))dt. \tag{1.7}$$

The "spin-spin" relaxation time T_2, as the name suggests, results from the magnetic interactions among the nuclear spins as they move about. It depends on the diffusion length of the protons—more mobile protons will experience a greater variety of field strengths as they diffuse, leading to random phase accumulation that cannot be refocused. T_2' pertains to the larger scale static field inhomogeneity which is essentially constant over the proton's diffusion length. A third relaxation time, the "spin-lattice" relaxation T_1 time describes the recovery of the longitudinal magnetization after the RF pulse as it returns to its equilibrium state. In practice, however, in the simple FID experiment a combination of the first two terms governs the observed decay: The effective decay time is $T_2^* = (1/T_2 + 1/T_2')^{-1}$.

In general, all four relaxation times depend on field strength (though T_2 less so) and on the local chemical environment of the proton spins. Hence, all four relaxation times are reduced by the local presence of magnetic materials [21], though the degree to which they are affected will depend on the material. That the signal does not depend purely on the density of spins, but also on the impurities of the local environments in which they are confined does much to explain the success of in vivo MR in general. For example, by way of intracellular vs. extracellular compartments, distinguished through numerous chemical properties, *the waters are divided*. Since variations between the spin densities of neighbouring tissue compartments may be minimal, it is fortuitous that there should be other salient parameters involved in

colouring the observed signal.

The guiding principle behind MR imaging, as distinct from other applications of NMR, is that by superposing linear magnetic field gradients in three dimensions onto the main polarizing magnetic field, a unique correspondence can be achieved between proton-spin locale and precessional frequency: $\omega(\mathbf{r}) = \omega_0 + \gamma \mathbf{r} \cdot \nabla B_z$.

Naturally, the effect of these phase-encoding gradients is to dephase the signal. To undo the deleterious effect of the original gradients and obtain an undistorted image, the phase accumulation must somehow be "rewound". One way of achieving this is simply to apply an equal and opposite gradient following the original. This protocol is referred to as gradient recalled echo (GRE) imaging. During the play-out of the refocusing gradient in the frequency-encode direction the oscillating time-domain signal is recorded and, after a time *TE* since the original RF excitation, the signal refocuses and an "echo" is produced. The time between successive RF excitations is referred to as the repetition time *TR*. To ensure consistency among successive excitations and avoid the formation of unintended echoes, the transverse magnetization, once sampled, is often deliberately annihilated by "spoiling" (dephasing) the spins with an additional RF pulse. As described by Haacke [22], assuming a perfectly homogeneous RF pulse which tips the longitudinal magnetization an angle θ into the transverse plane, the effective RF-spoiled GRE signal response at the echo is

$$\mathbf{\Psi}(\mathbf{r}, TE) \approx \rho_0^*(\mathbf{r}) \frac{sin\theta[1 - e^{-TR/T_1(\mathbf{r})}]e^{-TE/T_2^*(\mathbf{r})}}{1 - cos\theta e^{-TR/T_1(\mathbf{r})}} e^{i\phi(\mathbf{r},TE)}, \tag{1.8}$$

where ρ_0^* denotes an effective spin density which absorbs the constants of (1.4) relating to the magnetization along with any terms relating to the RF coil electronics. In general, we will assume that the phase at the echo time represents the average field experienced by the spins within a given voxel. This turns the integral of (1.7) into a simple linear relationship, and the phase difference at the echo time between two tissues can be written as

$$\Delta\phi(\mathbf{r}) = -\gamma \Delta B_z(\mathbf{r}) TE. \tag{1.9}$$

So the excitement wanes. Even the so-called permanent magnet isn't permanent, and entropy ultimately wins out: Our stimulated system decays back toward its dull equilibrium, leaving the receiver array with nothing to detect but frenzied thermal noise. Into the cold, dark abyss we plumet... Here in brief is the crux of the matter and the subject of this thesis:

> *GRE imaging does not refocus static field inhomogeneity. Field perturbations owing to spatial variation in magnetic susceptibility are recorded in the GRE signal via the decay envelope and via the phase.*

Figure 1.3 demonstrates a number of the susceptibility effects commonly observed in GRE imaging.

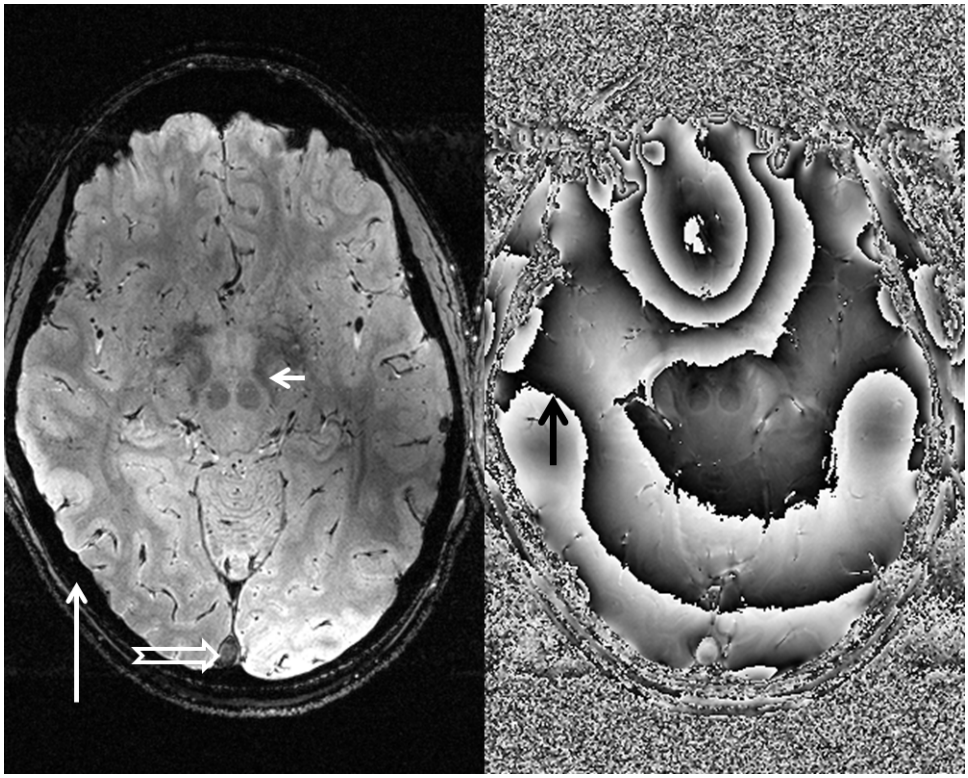

Figure 1.1: Complex reconstructed GRE magnitude (left) and phase (right). The small arrow points to the substantia nigra (SN), where the exceptional concentration of ferritin causes dephasing, manifested as reduced intensity in the magnitude. In the phase, the SN are here obscured by background contributions from macroscopic field inhomogeneity (e.g. air-tissue interfaces). However the red nuclei—also rich in paramagnetic iron—are visible in both phase and magnitude as the dark elliptic regions just medial to the SN. The proximity of the depicted slice to air-tissue interfaces (auditory canals and nasal passages) implies a large background field gradient, which gives rise to aliasing artifacts known as "phase wraps" (dark arrow). The long arrow points to the signal-starved region of the skull which, due to its short T_2 decay time, is all but absent at the echo time ($TE = 20$ ms). The signal-to-noise ratio (SNR) of the phase is proportional to the magnitude, hence, wherever there is a paucity for signal (air, bone) the phase is governed by random noise. The blood-oxygenation-level dependent (BOLD) effect, owing to the paramagnetism of de-oxyhaemoglobin, is visible around the veins and the sagittal sinus (notched hollow arrow).

1.4 Thesis overview

Though the aim of conventional MR imaging is to map a surrogate for the magnetization in the form of the signal magnitude, by mapping the field inhomogeneity via the phase, the hope is that one could solve the corresponding inverse problem to extract a measure of the underlying material susceptibility. As magnetic susceptibility is largely heterogeneous in tissue, it provides an endogenous source of image contrast. Beyond this qualitative benefit, there is significant clinical interest in developing a quantitative technique to map tissue susceptibility non-invasively for a number of reasons. Chiefly among them, as previously

mentioned, is that non-haem iron in the brain is understood to play a role in a host of neurodegenerative disorders for which the disease etiologies remain largely unknown.

A considerable obstacle to developing such a technique is that the sampled phase that relates to the tissue susceptibility structures of interest is, in effect, buried beneath a much larger "background" signal pertaining to susceptibility variation outside of tissue (e.g. Fig. 1.3). A number of techniques have been suggested to isolate this "local" phase of interest. Chapter 2 goes into greater detail introducing the problems at hand and reviews the basics of field and susceptibility mapping. Special attention is given to a popular filtering technique which employs a normalized spherical filter to isolate the local field. The technique, which takes several forms, is termed "spherical mean value" (SMV) filtering, as it adopts the SMV theorem of harmonic functions as its starting point. An additional property of harmonic functions is then introduced: namely, that a harmonic function is an analytic function, and can therefore be locally represented by its Taylor series expansion. By appealing to this analytic property, it is determined that the accuracy of the discrete SMV with regard to harmonic functions depends on image resolution and on the size of the adopted filter. In effect, not all digitized spheres are equal, and the properties of harmonic functions in continuous space need to be treated in light of the fact that MR phase (i.e. field) is sampled discretely. Nevertheless, the SMV-based technique is determined to be fairly robust under certain geometric conditions. In turn, this result does much to explain the success of the field extrapolation technique introduced in Chapter 3.

In addition to the problem of finite image resolution, MR imaging of the head has finite spatial support as signal is generally absent in the skull and ever-absent in the surrounding air. The SMV cannot be used to estimate the local field wherever the spherical kernel extends beyond the edges of this support. Hence, in conventional SMV-filtering, field points within a distance from the edges equal to the radius of the employed kernel are discarded outright. In brain imaging this is a considerable obstacle to a number of clinical applications as it precludes analysis of features such as subdural haematomas and cortical lesions in pathologies such as multiple sclerosis. To recover the local field across the edges of the brain, Chapter 3 presents an extension to conventional SMV-filtering by appealing to the analytic nature of the background field: By first obtaining an SMV-based estimate of the harmonic background field over a reduced inner portion of the brain, a three-dimensional Taylor expansion can then be performed to extend field coverage to the brain edges. The method is quantitatively assessed through a numerical experiment and qualitatively demonstrated on in vivo human brain data acquired at 4.7 T. Using a kernel radius typical of conventional methods, the extension recovered on average 26 % more in vivo brain volume.

Finally, Chapter 4 summarizes the research findings, discusses methodological limitations, and suggests a means of reformulating the technique introduced in Chapter 3.

Chapter 2

On the application of spherical averaging operators to a Dirichlet problem in MR phase imaging

Wherein the magnetic dipole field is dissected and the sphere sectioned into cuboids. An approach to an inverse problem, along with its inverse—problems with the approach—are reviewed.

2.1 Introduction

The essential feature to all forms of MRI is a strong, static, spatially invariant magnetic field. Immersing a sample of non-zero magnetic susceptibility $\chi(\mathbf{r})$ into such a field will cause the sample to magnetize which, in turn, induces a response field. Quoting Li [23], wherein the bulk magnetization is denoted $\mathbf{M}(\mathbf{r})$:

> By convention in MRI, the main field is applied in the z direction with a nominal field value B_0. We take a first-order approximation for [isotropic] non-ferromagnetic material ($|\chi| \ll 1$), and obtain $\mathbf{M}(\mathbf{r}) \approx \chi(\mathbf{r})/\mu_0 \cdot B_0\hat{\mathbf{z}}$, where $\hat{\mathbf{z}}$ is a unit vector in the z direction.

Denoting the angle between the position vector and the positive z-axis as θ, Li goes on to suggest a general form for the field perturbation along z owing to the magnetized susceptibility:

$$\Delta B(\mathbf{r}) = \frac{B_0}{4\pi} \int\limits_{\mathbf{r}' \neq \mathbf{r}} \chi(\mathbf{r}') \frac{3\cos^2\theta - 1}{|\mathbf{r}' - \mathbf{r}|^3} d^3\mathbf{r}'. \tag{2.1}$$

That is, the normalized perturbation along z, hereafter referred to as the relative field perturbation ($\mathbf{f}_\Delta := \Delta B(\mathbf{r})/B_0$), is given by the volume convolution (denoted $\otimes$) of the susceptibility with the z component of the unit dipole field d, scaled by the inverse of the

applied field strength [24]:[1]

$$\mathbf{f}_{\Delta} = \boldsymbol{\chi} \otimes \mathbf{d}, \qquad \text{where} \quad \mathbf{d} = \begin{cases} \dfrac{3\cos^2\theta - 1}{4\pi|\mathbf{r}|^3}, & \text{for } \mathbf{r} \neq \mathbf{0} \\ 0, & \text{for } \mathbf{r} = \mathbf{0}. \end{cases} \tag{2.2}$$

2.1.1 Mapping the magnetic susceptibility

The spatial relationships of (2.1) and (2.2) are nonlocal: variation in susceptibility distorts the magnetic induction at a distance. This gives rise to a complicated geometry dependence of field to susceptibility which dashes any hope of arriving at a simple closed form solution for arbitrary χ. To make matters worse, the field is sampled discretely over a spatially constrained volume (i.e., limited information about the continuous field distribution is available) and meaningful samples can be obtained strictly in regions where spin densities and T_2^* relaxation times are generous enough to yield signal above the level of noise at the echo time (i.e., signal from air is nonexistent, and that coming from bone is negligible in the most common imaging protocols). This is to say that the system relating the impoverished field data to underlying susceptibility is highly *underdetermined*.

The process of inverting the observed field data in attempt to recover the susceptibility has been termed *quantitative* susceptibility mapping (QSM). The unfortunate 'Q', however, belies the inherently relative nature of the process. Though tremendous effort goes into ensuring that the main field of the scanner is as uniform as possible, having a perfectly static and homogeneous effective field during image acquisition is a practical impossibility. However uniform the field of the empty scanner, introduction of magnetically interactive media (e.g. biologic tissue) into the bore will give rise to distortions. To minimize the deleterious effects of the resultant dephasing upon the MR signal, auxiliary applied fields are required. Active shimming is the iterative process of examining the free induction decay and compensating with additional applied fields in order to lengthen the decay envelope and boost the signal. The shimmed field nevertheless remains imperfect. Furthermore, ramping of the imaging gradients on and off during acquisition induces eddy currents in nearby conductors (e.g. cryostat) which then cause time-varying secondary fields via Lenz's law. Thus, although for the purpose of susceptibility mapping the field of interest is strictly the demagnetizing field of the object being imaged, external field inhomogeneities are nevertheless present. Since the effective $\mathbf{B}_0$ distribution is generally unknown, the GRE phase merely permits an estimation of the field perturbation rather than the magnetic induction itself.

Despite these limitations, interest in MR susceptibility mapping continues to grow, and has been doing so at an accelerated pace since a Fourier domain representation of the susceptibility-to-field operation was introduced [25, 26]. In Fourier domain, the nonlocal integration of susceptibility to field becomes a local multiplication: Denoting the Fourier

[1] Unless otherwise stated, the ubiquitous references to "fields" will be with respect to the field *projections* along $\mathbf{B}_0$ (i.e. the z component specifically).

transform operator as $\boldsymbol{\mathcal{F}}$ and the inverse operator as $\boldsymbol{\mathcal{F}}^H$

$$\boldsymbol{\chi} \otimes \mathbf{d} \approx \boldsymbol{\mathcal{F}}^H \mathbf{C} \boldsymbol{\mathcal{F}} \boldsymbol{\chi}, \qquad \text{where} \quad C(\mathbf{k}) = \begin{cases} \dfrac{1}{3} - \dfrac{k_z^2}{k_x^2 + k_y^2 + k_z^2}, & \text{for } \mathbf{k} \neq \mathbf{0} \\ 0, & \text{for } \mathbf{k} = \mathbf{0}. \end{cases} \tag{2.3}$$

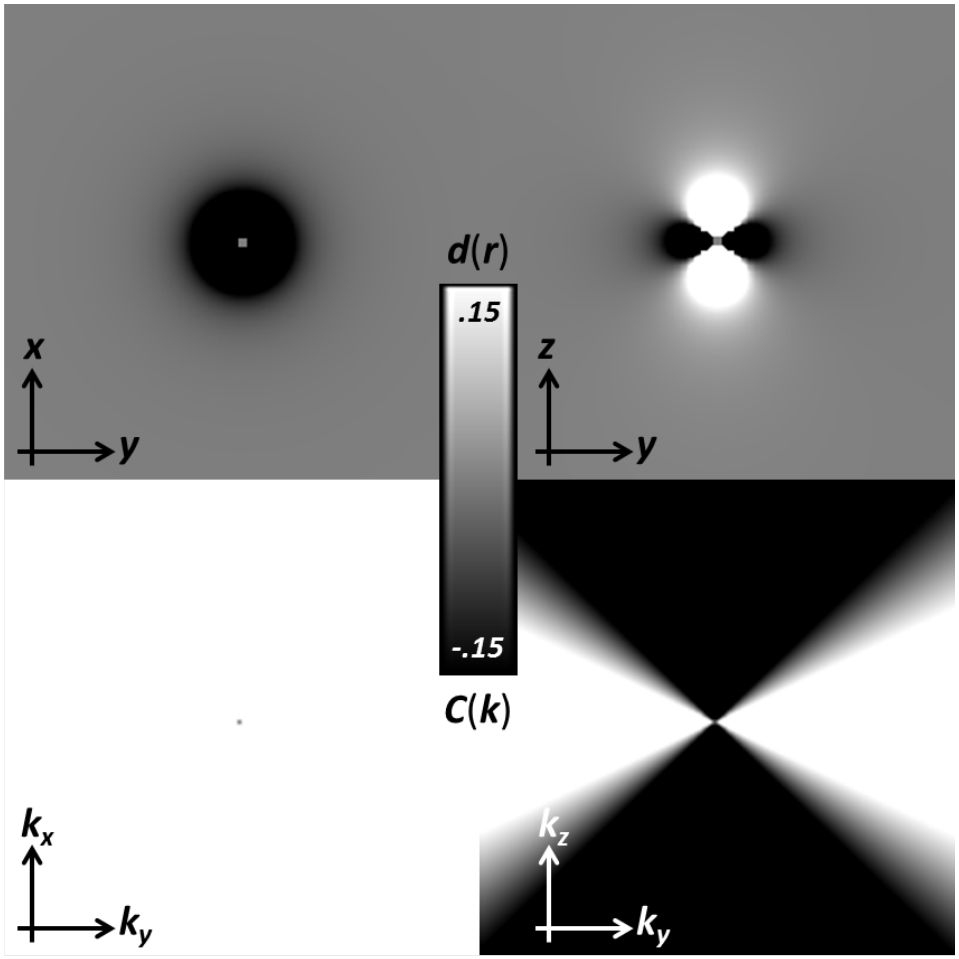

Figure 2.1: Magnetic dipole: Spatial and Fourier representations. Top: Unit dipole kernel $d(\mathbf{r})$ of (2.2) cross-sections in x–y and z–y. Bottom: Susceptibility-to-field transfer function $C(\mathbf{k})$ of (2.3) cross-sections in k_x-k_y and k_z–k_y. The two quantities are unitless and scaled to the same relative brightness.

The transfer function $C(\mathbf{k})$ tends toward zero on two conical surfaces in spatial frequency domain (k-space) wherever $k_z^2 = (k_x^2 + k_y^2)/2$. Since one is free to opt for an image resolution that avoids the zero-surfaces, it might appear that simple point-by-point division of the discrete Fourier transform (DFT) of $\mathbf{f}_\Delta$ by the transfer function ought to recover the susceptibility, provided one ignores (sets to zero) its DC component. Unfortunately, the zero-surfaces nevertheless interfere with reliable estimation of χ insofar as noise present in $\mathbf{f}_\Delta$ stands to be amplified in the division wherever $|C(\mathbf{k})|$ sinks below unity. In other words,

the problem is ill-posed, solutions are non-unique, and the inversion is unstable. Enter the panoply of regularization strategies [27–37]. However, more fundamental issues with the original problem formulation exist and should, sensibly, be of greater concern than the best choice of cost function.

The main issues with both the spatial (2.2) and k-space (2.3) representations of the relationship between field and susceptibility is that they assume in vivo susceptibility to be a scalar and, furthermore, that they incorporate a Lorentzian sphere correction to the magnetization.[2] As evinced by He and Yablonskiy, the field resulting from magnetic susceptibility inclusions depends on context rather than mere concentration [38]: "due to the anisotropic and inhomogeneous nature of the brain's cellular structure (especially considering elongated cells, such as axons and dendrites) [the Lorentzian sphere correction] should be modified since the susceptibility inclusions cannot be modeled as point magnetic dipoles anymore." In brief, by keeping the same volume susceptibility and varying only the arrangement of the sub-voxel susceptibility inclusions, the resulting field perturbation can be vastly different [39]. He and Yablonskiy thus proposed a generalized Lorentzian cylinder approach in order to account for the anisotropies of neural microstructure.

This introduces an angular dependence to the field perturbation based on the orientations of susceptility inclusions relative to the applied field. In myelinated white matter in particular the tensor nature of the susceptibility has been confirmed [40–42]. Since the relationship of field to susceptibility is nonlocal, the presence of magnetically anisotropic myelinated fibers in the brain likely means that even the predominantly isotropic magnetic structures will be inaccurately represented by the relationship given in (2.2) and (2.3).

2.1.2 Mapping the magnetic perturbation field

Beyond the aforementioned obstacles to transforming field to susceptibility, a number of issues complicate the process of arriving at a reliable field map based on the GRE phase. First, at the longer echo times adopted to optimise tissue contrast, the measured phase contains ambiguities—"wraps", whereby the phase exceeds $\pm\pi$ (e.g. Fig. 1.3). Prior to any other processing, these ambiguities need to be resolved by means of an unwrapping algorithm. In addition, the data are generally acquired by means of a receiver array, and due to the dielectric effects of the sample, and the wavelength effects owing to receiver geometry, each channel possesses its own phase offset. The offsets are generally spatially varying and difficult to determine, rendering channel combination for the phase somewhat complicated [43]. Furthermore, at a distance from the receiver elements, SNR tends to be low, defying proper unwrapping, and thus causing (along with other non-susceptibility effects, such as blood-flow [44] and partial-volume effects) spatially restricted outliers which confound unwrapping and channel combination. Even with the perfectly unwrapped, noiseless, and seamlessly channel-combined phase image, the assumptions behind (1.9)—that phase evolves *linearly*

[2]The Lorentz sphere correction makes the dipole kernel $\mathbf{d}$ zero at the origin in (2.2) and Fig. 2.1.

with time and that it represents the *average field* experienced by the spins within a given voxel—are only truly valid in the limit of high SNR and high resolution [34]. Given the host of preliminary issues, it is no wonder that the phase, until fairly recently, was often discarded before ever reaching the viewing console [22].

Despite the above complications, and despite the fact that the complete set of mechanisms at work in the generation of in vivo phase contrast may not yet be fully understood, the MR phase does depend on variation in *apparent* susceptibility. However, whereas intertissue variations in apparent susceptibility (inducing local field) deviate only slightly from the bulk susceptibility of water at body temperature, the difference between tissue and air is comparatively dramatic [45], inducing a background field which tends to obscure the small-scale field variation of interest. Beginning with the unwrapped phase and dividing by $\gamma B_0 TE$ (where γ is the proton gyromagnetic ratio, B_0 the strength of the applied field, and TE the echo time) yields an estimate of the relative field perturbation $f_\Delta(\mathbf{r})$ along the direction of the applied field:

$$f_\Delta(\mathbf{r}) = f_{\text{local}}(\mathbf{r}) + f_{\text{bkgr}}(\mathbf{r}) + \epsilon(\mathbf{r}). \tag{2.4}$$

For our purpose, ϵ includes anything present in the measurement which is unrelated to susceptibility (e.g., flow effects, noise [46], receiver offsets, chemical shift [47], contributions from macromolecule proton exchange [48–50] etc.). f_{local} and f_{bkgr} are, respectively, the fields due to magnetic susceptibility variation occurring inside and outside of the volume of interest (VOI, i.e., brain). Starting from Maxwell's equations, it can be shown [51] that, in the absence of noise and susceptibility variation, the longitudinal component of the magnetic induction expressed in the MR phase will be harmonic. In other words, f_{bkgr} is harmonic within the VOI, satisfying Laplace's equation:

$$0 = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2} + \frac{\partial^2 f}{\partial z^2}. \tag{2.5}$$

If the goal is to examine f_{local}, it must first be isolated from f_{bkgr}. This is particularly crucial for susceptibility mapping, as the problem of underdetermination must be addressed by constraining the local susceptibility solution set to reside within the VOI. Local field extraction techniques essentially fall into two camps: those that perform some form of high-pass filtering, and those that begin with an estimate of the background source geometry and proceed by fitting a susceptibility distribution via (2.3) to model the background field. An example of the latter is the Projection onto Dipole Fields technique (PDF) [52]. The simplest approach is to apply a generic low-pass filter (e.g. Gaussian or Hanning window) to the total phase, and then subtract the result of the low-pass from the total. However, the concern with this approach is that the frequency bands of local and background fields overlap, and thus, the removal of background generally comes at the expense of local field attenuation.

An essential property of harmonic fields in three-dimensional space is that they satisfy the spherical mean value (SMV) theorem [53]:

$$f(\mathbf{r}_0) = \frac{1}{4\pi R^2} \int\limits_{S(R)} f(\mathbf{r}) dS. \tag{2.6}$$

That is, the mean of a harmonic function calculated over the spherical surface $S(R)$ is equal to the specific value taken by the function at the centre of the sphere (positioned at $\mathbf{r}_0$, and where $|\Delta\mathbf{r}| = |\mathbf{r} - \mathbf{r}_0| = R$, the radius of the sphere).

The SMV property was introduced to MRI phase-imaging initially as a means to boost the SNR in regions of homogeneous susceptibility by averaging field measurements over a spherical shell [51]. More recently, the SMV property has been used to define a projection-style filtering scheme known as "Sophisticated Harmonic Artifact Reduction for Phase data" (SHARP) [54]. Notionally, it works as follows: Extraction of the SMV across the image is achieved via a convolution of the total field $\mathbf{f}_\Delta$ with a normalized sphere $\mathbf{S}$. Away from the air-tissue interfaces $\mathbf{f}_{\text{bkgr}}$ is harmonic, hence, the SMV should capture $\mathbf{f}_{\text{local}}$ at the central point, and subtraction of the SMV from $\mathbf{f}_\Delta$ should cancel $\mathbf{f}_{\text{bkgr}}$ entirely:

$$\begin{aligned}
\mathbf{f}_\Delta - \mathbf{S} \otimes \mathbf{f}_\Delta &= \mathbf{f}_{\text{bkgr}} - \mathbf{S} \otimes \mathbf{f}_{\text{bkgr}} + \mathbf{f}_{\text{local}} - \mathbf{S} \otimes \mathbf{f}_{\text{local}} + \boldsymbol{\epsilon} - \mathbf{S} \otimes \boldsymbol{\epsilon} \\
&= \mathbf{f}_{\text{local}} - \mathbf{S} \otimes \mathbf{f}_{\text{local}} + \boldsymbol{\epsilon} - \mathbf{S} \otimes \boldsymbol{\epsilon} \\
&= (\boldsymbol{\delta} - \mathbf{S}) \otimes (\mathbf{f}_{\text{local}} + \boldsymbol{\epsilon}),
\end{aligned} \tag{2.7}$$

where $\boldsymbol{\delta}$ is a spike of unit height at the centre of $\mathbf{S}$. What remains is a quantity which pertains strictly to the local field and to the noise, the sum of which can be estimated through a deconvolution. A variant of SHARP known as Regularization Enabled SHARP (RESHARP) was recently proposed [55]. Chapter 4 will discuss the distinction. Example local field images from the aforementioned techniques are shown in Fig. 2.2.

A concern with conventional SHARP[3], which adopts a radius for $\mathbf{S}$ on the order of a few millimetres, is that it implicitly conflates the field which is local relative to the boundary set by the sphere (i.e. owing to sources internal to the spherical surface) with that which is local relative to the air-tissue boundaries (i.e. owing to inter-tissue susceptibility variation). In fact, all fields owing to susceptibility sources outside the spherical surface are harmonic within it (be they due to air-tissue interfaces or tissue structures of interest) and thus, in the context of SHARP, whereby "harmonic field" is essentially equated with "background field", these fields are filtered out indiscriminantly. In a similar way, there is an apparent contradiction in SHARP in the decision to use a dense sphere over regions of spatially inhomogenous susceptibility: the outer shell includes susceptibility sources nearer the edge of the sphere which are external as far as the more central spherical shells are concerned.

[3]That is, apart from the bad acronym, which belies the fact all that is being done is *simple* averaging.

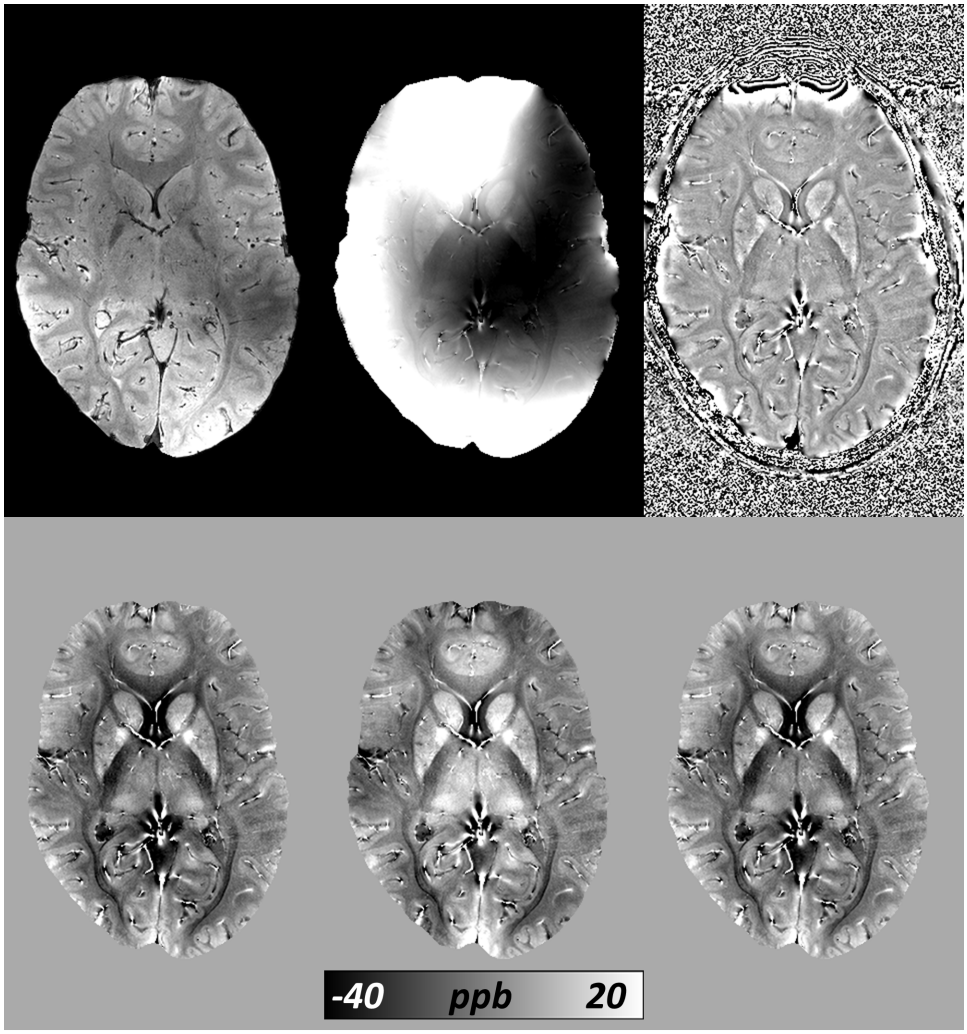

Figure 2.2: Current phase processing techniques in use. The advanced methods (second row, from left to right: PDF, SHARP, and RESHARP local field estimates) all require a binary mask derived from the magnitude (top left) to delineate the brain boundary and separate the local territory from the background and the noise. In addition, a necessary preliminary stage for the advanced techniques involves unwrapping the phase (top middle). Traditional homodyne high-pass filtering [44] (top right) requires neither mask, nor unwrapping, but generally attenuates the local field and often fails to completely overcome the issue of phase wrapping (e.g. toward the frontal sinuses).

Furthermore, it is unclear how the local susceptibility sources that overlap the outer edge of the sphere are dealt with in the SHARP scheme.

This notion of inclusion vs. exclusion based on the scale of the SMV operator forces one to consider the effect of spatial resolution, which ultimately determines what forms of discretized spheres are possible for the gridded data. In the discretized domain of the image, denoting the number of samples in the set defining the discretized spherical shell as N_S, and the positions of these samples by $\mathbf{r}_n$, the integral of (2.6) becomes a summation,

$$g(\mathbf{r}_0) = \frac{1}{N_s} \sum_{n=1}^{N_s} f(\mathbf{r}_n), \tag{2.8}$$

and the discrete spherical mean value (DSMV) conjecture underlying SHARP can be written as

$$g(\mathbf{r}_0) - f(\mathbf{r}_0) \approx 0. \tag{2.9}$$

To date, the size of the SMV operator (which, along with image resolution, determines N_S) has only been examined heuristically in the contexts of reducing measurement noise [51] and of extracting a measure of the local field [56]. These works take it for granted that given properties of continuous potential fields (i.e. (2.6)) translate without complication to the discretized space of the data.

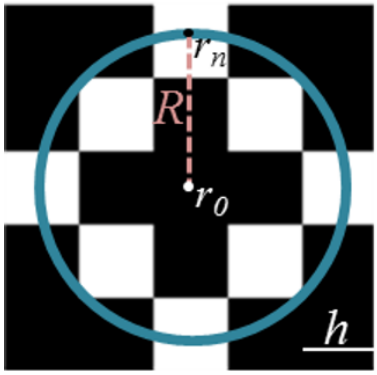

Figure 2.3: Mock-sphere in the finite grid.

The aim of this chapter, beyond giving a general introduction to the topic at hand, is to establish a framework for understanding how field estimation via the DSMV depends on the discretization and scale parameters: namely, spatial resolution and radius of the sphere.

2.2 Theory

A harmonic field is an analytic field, which means that at every point $\mathbf{r}$ in its domain, it can be locally represented by a convergent power series [57]. Appealing to the analytic nature of a harmonic field f, we represent the field at position $\mathbf{r}_n = (x_n, y_n, z_n)$ in terms of its Taylor expansion about $\mathbf{r}_0 = (x_0, y_0, z_0)$:

$$f(x_n, y_n, z_n) = \sum_{\alpha_x=0}^{\infty} \sum_{\alpha_y=0}^{\infty} \sum_{\alpha_z=0}^{\infty} \frac{(x_n - x_0)^{\alpha_x} (y_n - y_0)^{\alpha_y} (z_n - z_0)^{\alpha_z}}{\alpha_x! \alpha_y! \alpha_z!} \left[\frac{(\partial^{\alpha_x + \alpha_y + \alpha_z} f)}{\partial x^{\alpha_x} \partial y^{\alpha_y} \partial z^{\alpha_z}} \right]_{x_0, y_0, z_0}. \tag{2.10}$$

For notational ease, 1^{st} order partial derivatives of f in x, y, and z are denoted in the following as f_x, f_y, and f_z, with higher order partials denoted accordingly. Letting

$(\Delta x_n, \Delta y_n, \Delta z_n) = (x_n - x_0, y_n - y_0, z_n - z_0)$ and expanding out the second order terms explicitly,

$$\begin{aligned} f(\mathbf{r}_n) = \; & f(\mathbf{r}_0) + \Delta x_n f_x(\mathbf{r}_0) + \Delta y_n f_y(\mathbf{r}_0) + \Delta z_n f_z(\mathbf{r}_0) \\ & \ldots + \Delta x_n \Delta y_n f_{xy}(\mathbf{r}_0) + \Delta x_n \Delta z_n f_{xz}(\mathbf{r}_0) + \Delta y_n \Delta x_n f_{yx}(\mathbf{r}_0) \\ & \ldots + \Delta y_n \Delta z_n f_{yz}(\mathbf{r}_0) + \Delta z_n \Delta x_n f_{zx}(\mathbf{r}_0) + \Delta z_n \Delta x_n f_{zy}(\mathbf{r}_0) \\ & \ldots + \frac{\Delta x_n^2 f_{xx}(\mathbf{r}_0)}{2!} + \frac{\Delta y_n^2 f_{yy}(\mathbf{r}_0)}{2!} + \frac{\Delta z_n^2 f_{zz}(\mathbf{r}_0)}{2!} + \ldots \end{aligned} \tag{2.11}$$

Letting $\mathbf{r}_n$ be a point belonging to the set which defines the spherical shell of radius $R = \sqrt{\Delta x_n^2 + \Delta y_n^2 + \Delta z_n^2}$. Inserting (2.11) into (2.8), and noting that all terms involving a distance with an odd power will ultimately cancel (i.e., by spherical symmetry, for every Δx_n there exists a $-\Delta x_n$ to negate it), yields the expression

$$g(\mathbf{r}_0) = \frac{1}{N_s} \sum_{n=1}^{N_s} \left(f(\mathbf{r}_0) + \frac{\Delta x_n^2 f_{xx}(\mathbf{r}_0)}{2!} + \frac{\Delta y_n^2 f_{yy}(\mathbf{r}_0)}{2!} + \frac{\Delta z_n^2 f_{zz}(\mathbf{r}_0)}{2!} + \ldots \right) \tag{2.12}$$

Equivalently, by distributing the summation we arrive at

$$g(\mathbf{r}_0) = \frac{f(\mathbf{r}_0)}{N_s} \sum_{n=1}^{N_s} 1 \; + \; \frac{f_{xx}(\mathbf{r}_0)}{2! N_s} \sum_{n=1}^{N_s} \Delta x_n^2 \; + \; \frac{f_{yy}(\mathbf{r}_0)}{2! N_s} \sum_{n=1}^{N_s} \Delta y_n^2 \; + \; \frac{f_{zz}(\mathbf{r}_0)}{2! N_s} \sum_{n=1}^{N_s} \Delta z_n^2 \; + \; \frac{1}{N_s} \sum_{n=1}^{N_s} \ldots \tag{2.13}$$

Assuming an isotropic spacing h between grid points and noting, again by symmetry, that the squared displacements sum to the same positive constant, labelled $c_{2,0,0}$ below. The preceding expression, inserted into (2.9), leads to

$$g(\mathbf{r}_0) - f(\mathbf{r}_0) = \frac{c_{2,0,0}}{2! N_s} (f_{xx}(\mathbf{r}_0) + f_{yy}(\mathbf{r}_0) + f_{zz}(\mathbf{r}_0)) + \frac{1}{N_s} \sum_{n=1}^{N_s} \ldots \tag{2.14}$$

Invoking Laplace's equation (2.5), the term above in parentheses is zero. Though the 3^{rd} and all other higher *odd* order terms of the expansion necessarily cancel through spherical symmetry, all higher *even* order displacement terms—featuring even exponents exclusively—do not. Thus, expanding out the terms contained in the ellipsis up to 4^{th} order and absorbing redundant partial derivatives ($f_{xxyy} = f_{yyxx}$, etc.)

$$\begin{aligned} g(\mathbf{r}_0) - f(\mathbf{r}_0) = & \frac{1}{4} \sum_{n=1}^{N_s} (f_{xxyy}(\mathbf{r}_0) \Delta x_n^2 \Delta y_n^2 + f_{xxzz}(\mathbf{r}_0) \Delta x_n^2 \Delta z_n^2 + f_{yyzz}(\mathbf{r}_0) \Delta y_n^2 \Delta z_n^2) \\ & \ldots + \frac{1}{4!} \sum_{n=1}^{N_s} (f_{xxxx}(\mathbf{r}_0) \Delta x_n^4 + f_{yyyy}(\mathbf{r}_0) \Delta y_n^4 + f_{zzzz}(\mathbf{r}_0) \Delta z_n^4) \; + \; \sum_{n=1}^{N_s} \ldots \end{aligned} \tag{2.15}$$

$$g(\mathbf{r}_0) - f(\mathbf{r}_0) = \frac{c_{2,2,0}}{4}(f_{xxyy}(\mathbf{r}_0) \; + \; f_{xxzz}(\mathbf{r}_0) \; + \; f_{yyzz}(\mathbf{r}_0))$$
$$\ldots + \frac{c_{4,0,0}}{4!}(f_{xxxx}(\mathbf{r}_0) \; + \; f_{yyyy}(\mathbf{r}_0) \; + \; f_{zzzz}(\mathbf{r}_0)) \; + \; \sum_{n=1}^{N_s} \ldots \tag{2.16}$$

where constants $c_{2,2,0}$ and $c_{4,0,0}$ are positive constants for a given radius and resolution:

$$\begin{array}{rcccccc} c_{2,2,0} & = & \sum\limits_{n=1}^{N_s} \Delta x_n^2 \Delta y_n^2 & = & \sum\limits_{n=1}^{N_s} \Delta x_n^2 \Delta z_n^2 & = & \sum\limits_{n=1}^{N_s} \Delta y_n^2 \Delta z_n^2, \\ c_{4,0,0} & = & \sum\limits_{n=1}^{N_s} \Delta x_n^4 & = & \sum\limits_{n=1}^{N_s} \Delta y_n^4 & = & \sum\limits_{n=1}^{N_s} \Delta z_n^4. \end{array} \tag{2.17}$$

Considering the limit of $N_S \to \infty$, the summations become integrations over the spherical surface. Without loss of generality, centring the sphere at the coordinate origin,

$$\begin{aligned} c_{2,2,0|N_S\to\infty} &= \int_{S(R)} x^2 y^2 dS, \\ c_{4,0,0|N_S\to\infty} &= \int_{S(R)} x^4 dS. \end{aligned} \tag{2.18}$$

Transforming to spherical coordinates, where θ remains the polar angle, and the azimuth φ describes the angle between the projection of $\mathbf{r}$ onto the xy plane and the postive x axis,

$$x = r\,sin\theta\,cos\varphi, \qquad y = r\,sin\theta\,sin\varphi, \qquad z = r\,cos\theta,$$

$$\begin{aligned} c_{2,2,0|N_S\to\infty} &= \int_{-\pi}^{\pi}\int_{0}^{\pi} (r\,sin\theta\,cos\varphi)^2 (r\,sin\theta\,sin\varphi)^2 dS, \\ c_{4,0,0|N_S\to\infty} &= \int_{-\pi}^{\pi}\int_{0}^{\pi} (r\,sin\theta\,cos\varphi)^4 dS. \end{aligned}$$

Since we are examining a spherical shell of constant radius R, the surface element becomes $dS = R^2\,sin\varphi\,d\varphi\,d\theta$. Thus,

$$\begin{aligned} c_{2,2,0|N_S\to\infty} &= R^6 \int_{-\pi}^{\pi} sin^3\varphi\,cos^2\varphi\,d\varphi \int_{0}^{\pi} sin^4\theta\,d\theta, \\ c_{4,0,0|N_S\to\infty} &= R^6 \int_{-\pi}^{\pi} sin\varphi\,cos^4\varphi\,d\varphi \int_{0}^{\pi} sin^4\theta\,d\theta. \end{aligned} \tag{2.19}$$

In both cases, the integrand of the $d\varphi$ term consists of an odd function involving a sine multiplying an even function involving a cosine—the product of which being an odd function. Given the symmetric limits of integration, $c_{2,2,0|N_S\to\infty}$ and $c_{4,0,0|N_S\to\infty}$ necessarily vanish—the expected result for the *continuous* spherical mean value. In fact, all of the higher order even terms contained in the ellipsis of (2.16) can be organized in a like manner, with constant

coefficients $c_{\beta_1,\beta_2,\beta_3}$ pertaining to the displacements multiplying partial derivatives of f. It should be clear that *all* of these higher even-order $c_{\beta_1,\beta_2,\beta_3}$ coefficients will consist of *even* sine and/or cosine functions in the azimuth-specific integrand multiplying the *odd* sine from of the surface element dS. Therefore, in the continuous case, all higher order terms of the expansion vanish, and the mean value of f calculated over the surface of the sphere of constant radius R is indeed equal to value of f at the centre of the sphere. QED.

The discrete case is different: as $N_S \propto (R/h)^3$, clearly the only way to increase N_S without increasing the radius is to increase the resolution. Looking specifically to the dipole field of 2.2, the sum of the 4^{th} order partial derivatives corresponding to the coefficient $c_{2,2,0}$ is

$$f_{xxyy} + f_{xxzz} + f_{yyzz} = \frac{315}{4\pi} \frac{x^6 + y^6 - 2z^6 + 15z^2(x^2z^2 + y^2z^2 - x^4 - y^4)}{(x^2 + y^2 + z^2)^{\frac{13}{2}}}, \tag{2.20}$$

whereas the 4^{th} order non-mixed partials $(f_{xxxx} + f_{yyyy} + f_{zzzz})$ sum to -2 times the mixed term above. That is, the expression of (2.9) *does not*, in general, equate to zero. Retaining the spherical coordinates, in the discrete domain,

$$\begin{aligned} c_{2,2,0} &= \sum_{n=1}^{N_s} \Delta x_n^2 \Delta y_n^2 &&= R^4 \sum_{n=1}^{N_s} (sin^4\theta_n cos^2\varphi_n sin^2\varphi_n), \\ c_{4,0,0} &= \sum_{n=1}^{N_s} \Delta x_n^4 &&= R^4 \sum_{n=1}^{N_s} (sin^4\theta_n cos^4\varphi_n). \end{aligned} \tag{2.21}$$

Taking the absolute value of (2.16) to yield the approximation error of the discrete SMV $g(\mathbf{r_0})$, the R^4 of the two expressions above can be factored out to yield:

$$\frac{|g(\mathbf{r}_0) - f(\mathbf{r}_0)|}{R^4} \leq \tau, \tag{2.22}$$

where τ is some positive constant which absorbs all of the non-vanishing terms of the Taylor expansion. Correspondingly, the DSMV operator has an associated truncation error which, using "big O" notation, is $O(R^4)$.

2.3 Methods

To test the effects of discretization upon the DSMV, a simple numerical experiment was performed in C++. The field perturbation $f_\Delta(\mathbf{r})$ of the unit-susceptibility spike $\chi(\mathbf{r}) = \delta(\mathbf{0})$ (harmonic for all $\mathbf{r} \neq \mathbf{0}$, Fig. 2.1) was generated according to (2.2) in cubic array of 68^3 cells and normalized to a maximum value of unity. Isotropic grid spacings $h = 2^m$ mm were tested for $m = \{-1, \ldots, 3\}$. At each resolution, the field was generated and DSMV estimates were formed over the region of support for operators ranging in size from $R = h \rightarrow h_{max}$.

2.4 Results

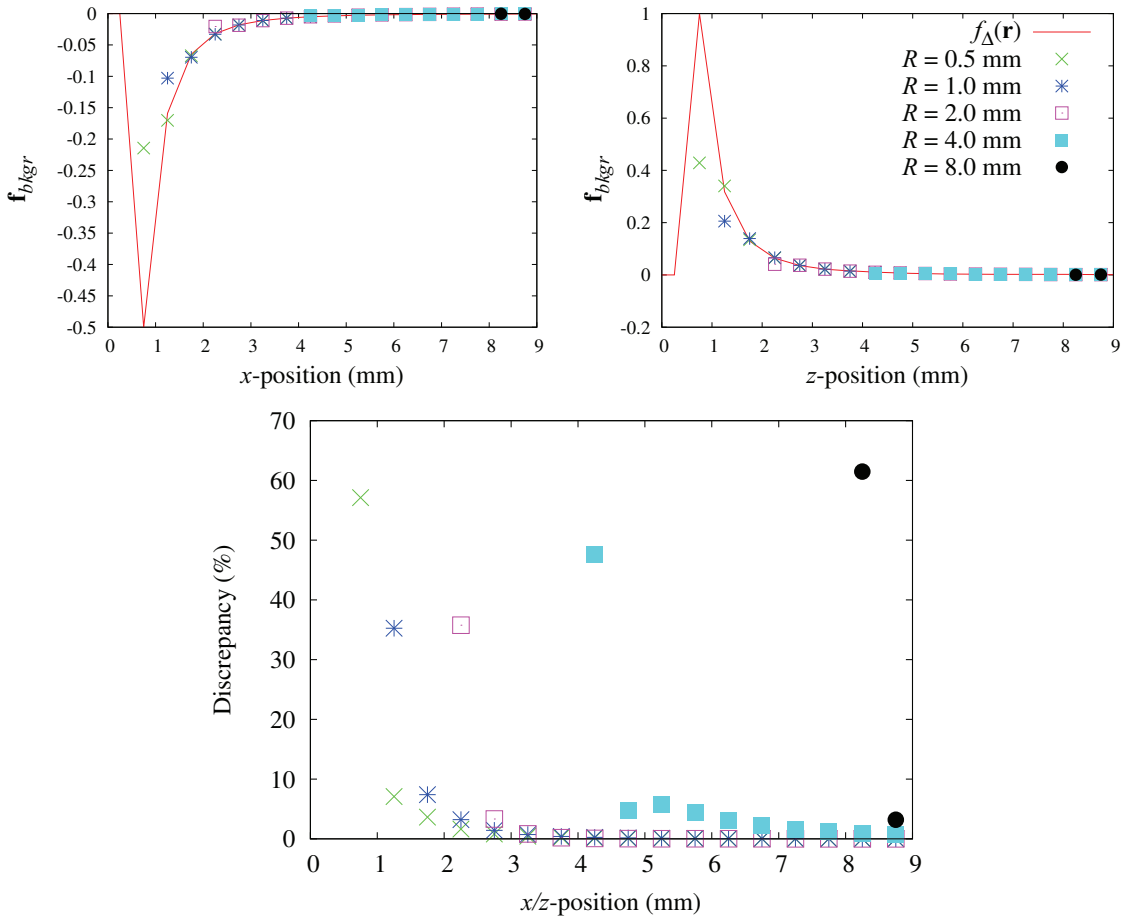

Figure 2.4: Top row: Profile of the field perturbation [line] and corresponding DSMV estimates [dots] of the harmonic dipole field along x and z beginning near the source at the origin. Bottom row: Percent discrepancy profile of the DSMV estimates from the expected relative field perturbation $f_\Delta(\mathbf{r})$ (the discrepancies for these particular x/z profiles are identical).

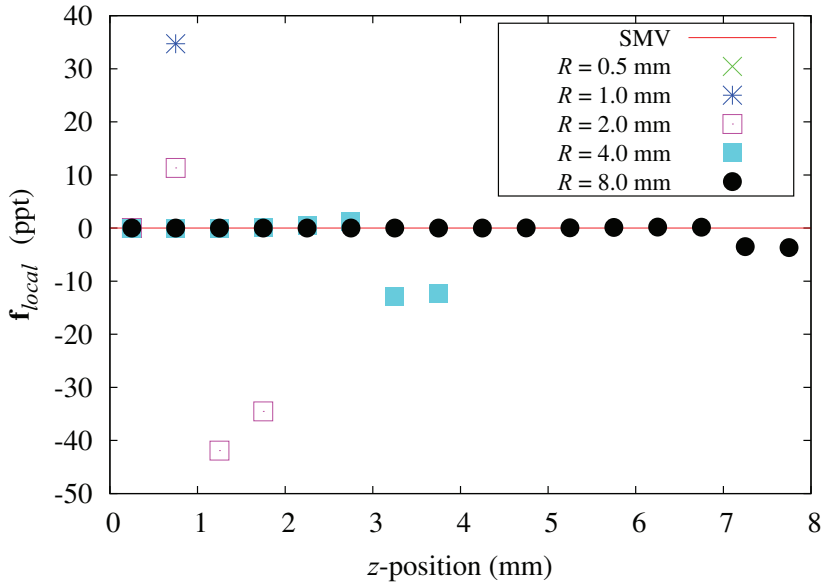

Figure 2.5: Profile of the expected solution for the continuous SMV [line] (uniformly zero for the enclosed source) and corresponding DSMV estimates [dots] in parts per thousand (ppt) along z for the case where the source is fully encapsulated by the DSMV operator.

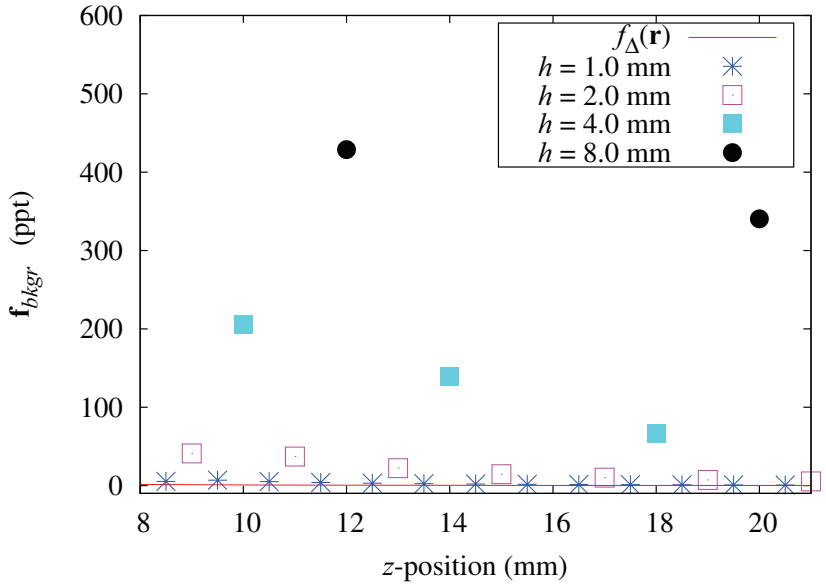

Figure 2.6: Profile of field perturbation along Z [line] and DSMV estimates [dots] varying grid spacing h for constant DSMV operator radius $R = 8$ mm.

2.5 Discussion

The degree to which the DSMV estimate of the field is accurate depends on the radius of the adopted DSMV. Figure 2.4 evinces this to be the case: For a fixed resolution, incorporating a greater number of points into the DSMV calculation does not improve the result as more sample points N_S necessarily connote a larger radius, and therefore a greater truncation error to the series expansion (2.11). Even in the simplest case—no noise, no local field—the presumption underlying SHARP (2.8) shows discrepancies over 60% when the DSMV calculation begins to incorporate samples in areas of rapid field change. Rather than stabilizing the process, additional samples—because they come at the expense of a larger R—necessarily amplify the error. The saving grace of the larger operators is that, in being further away from the source, the field and its derivates have largely decayed away, leaving the averaging to take place over regions that may be roughly constant over the spherical surface.

This is likely why this phenomenon has apparently not been noted in previous studies: In particular, since the brain does not interface with air but rather is separated from it by at least a few mm of connective tissue and skull, the background field within the brain is fairly flat. Hence, the DSMV may generally be fairly accurate insofar as $\mathbf{r}_0$ is distant from the magnetic background sources and that the $\mathbf{r}_n$ incorporated into the mean calculation are also outside of the regions of rapid field change. For example, beyond about 3 mm from the origin, discrepancies exhbited by the smaller operators in estimating $d(\mathbf{r})$ amounted to mere fractions of a percent (e.g. Fig. 2.4).

Figure 2.5 exhibits the case where the field source lies within the DSMV operator. Theory informs us that the SMV in the case of enclosed sources is zero. This is no longer quite true in the context of the discrete grid, though the observed deviations from 0 are indeed very small, on the order of parts per thousand. It is interesting to note, however, that the error in this case decreases with R, likely by virtue of the larger operators incorporating more samples in the slow-varying region of the field away from the origin, thereby curbing the influence of the perturbation at $\mathbf{0}$. DSMV estimation error increases with distance $\mathbf{r}_0$ as symmetry of the field is reduced (i.e. when the source is roughly at the centre of the sphere ($\mathbf{r}_0 \to \mathbf{0}$), every sample essentially has a complementary sample as a counterweight, thereby keeping the DSMV around zero.)

Figure 2.6 demonstrates the effect of grid spacing on the DSMV: as per expectations, for the fixed radius $R = 8$ mm, higher resolution (greater N_S) is better, but all DSMV operators generally perform well away from the source. Granted, the notion of spatial sampling in this trivial field model is abstracted from the practical case in MRI whereby the data are sampled in Fourier domain. Since acquisition time and SNR depend inversely on this sample rate, there are limits to the possible resolutions (e.g., for common clinical scanners, $\Delta h_{min} \sim 0.5$ mm). The necessary truncation of Fourier sampling in turn implies a $sinc(\mathbf{r})$ convolution in image space, thereby resulting in blurring and ringing artifacts. We can assume *a fortiori*

that the DSMV estimation in the more realistic sampling scenario will be somewhat worse.

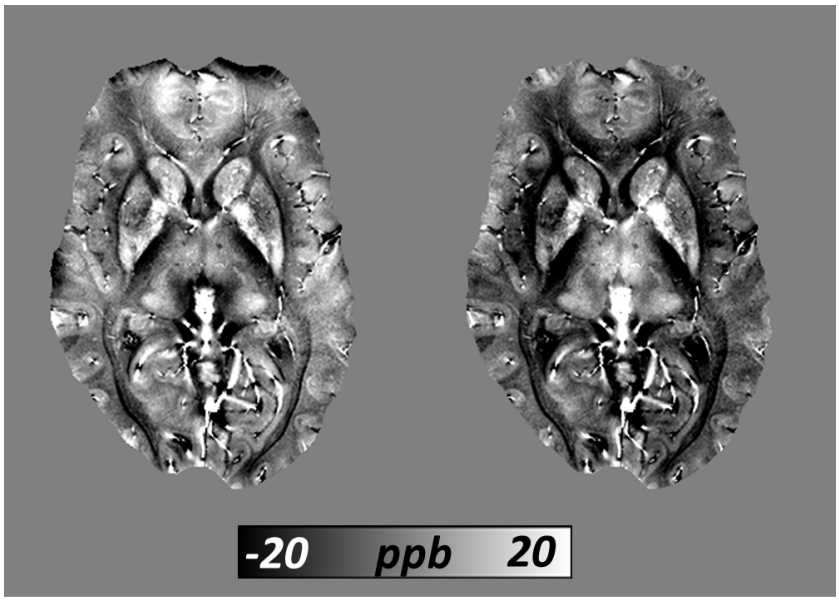

Figure 2.7: RESHARP local field estimates using DSMV kernels with $R = 1$ mm [left] and $R = 12$ mm [right]. These images (same data-set as Fig. 2.2, from a single-echo 2D SWI scan) have 0.5 mm^2 in-plane resolution with 1 mm slices. The employed kernels are dense, as opposed to spherical shells, with the coefficients of any given kernel invariant across its volume, with the exception of the central point, to which a value of 1 is added (corresponding to $\boldsymbol{\delta}$ of (2.7)).

2.6 Conclusion

MRI measures a form of the magnetic induction via the GRE phase, to which there at least two principal contributing sources: background susceptibility sources (e.g. air-tissue interfaces) and local susceptibility structures within tissue. The harmonic property of the magnetic background field has been proposed as a means to separate the two fields via SMV-filtering. Past works on the subject, however, have simply taken it for granted that the SMV property holds true, irrespective of filter size, when the field data and the spherical filter are discrete vector quantities rather than continuous functions. The fact that a harmonic field is also analytic offers a simple way of understanding the discrepancy between the continuous SMV and the discrete SMV of the mock-sphere. Contrary to what one might expect based on the idea that a greater number of samples should imply a better representation of quantity to be estimated, for a fixed resolution, this is not true of the DSMV estimate, which depreciates with filter size. Nevertheless, for the trivial model examined, provided R is relatively small and $\mathbf{r}_0$ is distant from magnetic background sources, the DSMV may be considered a good approximation.

Chapter 3

SHARP edges: Recovering cortical phase contrast through harmonic extension[1]

3.1 Introduction

Two approaches to eliminating the background field have recently been published: sophisticated harmonic artifact reduction for phase data (SHARP) [54] and projection onto dipole fields (PDF) [24, 52]. Unlike traditional high-pass filtering [44], which was based on the empirical observation that the measured phase is to some extent differentially composed of low and high spatial frequencies (contributed by the background and the local fields, respectively), SHARP and PDF look to the underlying physics for an approach which is less heuristic and less dependent on the particular form of the data. Both techniques begin with the recognition that the measured perturbation field consists of a superposition of local and background components. SHARP uses the spherical mean value theorem to extract the harmonic background field. On the other hand, PDF makes use of the approximate orthogonality between local and background fields, which, through the Hilbert projection theorem, designates a particular susceptibility distribution, in turn used to model the background field through a field-forward calculation [25, 26].

Each method has its own advantages and disadvantages. Of particular importance is that both methods necessitate a definition (i.e., binary input image) of the brain boundary, and both tend to produce invalid results in its vicinity, thereby limiting their application to an internal subsection of tissue. This shared pitfall is a considerable obstacle to several clinical applications as it precludes analysis of features such as pial vessels, cortical dysplasia, subdural haematoma, and cortical lesions in pathologies such as multiple sclerosis. To recover local field across the edges of the brain, this study presents an extension to conventional SHARP whereby another fundamental property of the background field is used:

[1]The contents of this chapter are adapted from Topfer R, Schweser F, Deistung A, Reichenbach JR, Wilman AH. SHARP Edges: Recovering Cortical Phase Contrast Through Harmonic Extension. Magnetic Resonance in Medicine 2014; Article first published online 3 March 2014 DOI: 10.1002/mrm.25148.

namely, its analyticity [57]. The method, Extended-SHARP (E-SHARP), is quantitatively assessed through a numerical field-forward experiment and qualitatively demonstrated on in vivo human brain data acquired at 4.7 T.

3.2 Theory

To eliminate the background field f_{bkgr}, SHARP invokes the spherical mean value property. However, the SMV[2] cannot be used to estimate the background field wherever the sphere overlaps the edge of the data support. Hence, in conventional SHARP, voxels in this region of overlap are wholly discarded (i.e., set to zero before the deconvolution stage). We refer to these discarded voxels as *edge points* (EP). The collection of voxels retained subsequent to the application of SHARP will be referred to as the *reduced VOI*.

Fortunately, the background field is analytic [57]. Therefore, within a neighborhood of location $\mathbf{r}_0$, for which $f_{\text{bkgr}}(\mathbf{r}_0 + \boldsymbol{\xi})$ is everywhere harmonic, the background field can be expressed as a convergent power series. Practically speaking, given an estimate of $f_{\text{bkgr}}(\mathbf{r}_0)$ (i.e., that provided by conventional SHARP), because $f_{\text{bkgr}}(\mathbf{r}_0 + \boldsymbol{\xi})$ is entirely harmonic within the region circumscribed by the SHARP kernel with radius $R = |\mathbf{r}_0 + \boldsymbol{\xi}|$ a three-dimensional (3D) Taylor expansion can be performed to extend the field coverage to the hitherto lost EP voxels:

$$f^{\text{EP}} = f^{\text{IP}} + \begin{bmatrix}\xi_x & \xi_y & \xi_z\end{bmatrix}\begin{bmatrix}f_x^{\text{IP}} \\ f_y^{\text{IP}} \\ f_z^{\text{IP}}\end{bmatrix} + \frac{1}{2!}\begin{bmatrix}\xi_x & \xi_y & \xi_z\end{bmatrix}\begin{bmatrix}f_{xx}^{\text{IP}} & f_{xy}^{\text{IP}} & f_{xz}^{\text{IP}} \\ f_{yx}^{\text{IP}} & f_{yy}^{\text{IP}} & f_{yz}^{\text{IP}} \\ f_{zx}^{\text{IP}} & f_{zy}^{\text{IP}} & f_{zz}^{\text{IP}}\end{bmatrix}\begin{bmatrix}\xi_x \\ \xi_y \\ \xi_z\end{bmatrix} + \dots \tag{3.1}$$

Here, f^{IP} denotes the background field at an *internal point* (IP) within the reduced VOI; f^{IP} denotes the background field of an EP voxel; f_i^{IP} denotes the first-order derivative in direction i ($i = x, y, z$) of the background field evaluated at an IP voxel and f_{ij}^{IP} denotes the second-order derivatives accordingly; ξ_i represents the EP-to-IP Euclidean distance in the i direction. Once the background field has been determined over the entire VOI, it can be subtracted from the measured field perturbation, ideally leaving only the local field due to tissue susceptibility. We refer to this process as *Extended*-SHARP.

3.3 Methods

3.3.1 E-SHARP Processing

Essential points of the processing scheme are illustrated in Fig. 3.1. For both in vivo and numerical data, the total field (masked, unwrapped, and scaled phase) was SMV-filtered (radius = 6 mm), eroded by the radius of the spherical kernel, and finally deconvolved, without regularization, to yield a local post-SMV field estimate. The local field estimate

[2]Technically this refers the *discrete* spherical mean but, having examined the distinction in detail in Chapter 2, for the sake of brevity, the 'D' is now simply implied.

was then subtracted from the total field to provide an estimate of the background field over the reduced brain volume. The hitherto lost EP voxels were paired with their nearest (in the Euclidean sense) IP neighbors for which first and second order spatial derivatives of the background field could be estimated by means of traditional 3-point central differences. For example, for the first derivative in the x-direction:

$$f_x^{\mathrm{IP}}(x_0, y_0, z_0) = \frac{f^{\mathrm{IP}}(x_0 + h_x, y_0, z_0) - f^{IP}(x_0 - h_x, y_0, z_0)}{2h_x} \tag{3.2}$$

where $f^{\mathrm{IP}}(x_0 + h_x, y_0, z_0)$ and $f^{\mathrm{IP}}(x_0 - h_x, y_0, z_0)$ are the background field measurements, respectively, one voxel ahead of, and one voxel behind (x_0, y_0, z_0) in the x direction; h_x is the voxel spacing in the x direction. Partial derivatives in y and z were calculated similarly, as were all second order derivatives—determined simply by applying the same 3-point difference scheme to each of the first order partials.

Using these derivatives and the IP-to-EP distances, the background field estimate was extended by means of a second order 3D Taylor expansion (3.1). The postexpansion background (extended background in Fig. 3.1) was then subtracted from the total field and the truncated singular value decomposition (TSVD)-like regularization of the original SHARP technique was applied to the result (threshold parameter $\lambda_{\mathrm{TSVD}} = 0.05$) [54, 56].[3] Specifically, regularization consisted of taking Fourier transforms of the spherical kernel and of the local field estimate; setting to zero the local field Fourier coefficients wherever the magnitude of the corresponding kernel coefficient was less than λ_{TSVD}; and finally taking the inverse Fourier transform of this thresholded form to yield the local field over the extended VOI. In the context of SHARP and E-SHARP, TSVD can be regarded as somewhat analogous to regularizing the deconvolution procedure (Eq. (9) in [54]) with a penalty on the L2 norm of the local field [55]. In this way, TSVD offers a convenient means of suppressing undesirable low-frequency components remaining in the local field estimate [59].

3.3.2 In Vivo Experiments[4]

Whole-head 3D multiple gradient echo data were acquired from five volunteers at 4.7 T (Varian, Palo Alto, CA) with the approval of the University of Alberta Research Ethics Board. Acquisition parameters were: field of view = $256 \times 160 \times 160$ mm^3; spatial resolution = $1.0 \times 1.0 \times 2.0$ mm^3; bandwidth = 90.1 kHz; repetition time (TR) = 40 ms; TE = 3/7/11/15/19 ms with unipolar readout; flip angle = 10. All datasets were processed in MATLAB (version 2012a, The MathWorks, Natick, MA) on a 16-processor computer (Quad-Core Xeon E5620, Intel, Santa Clara, CA) with 46 GB RAM. Images from individual receiver elements (four-channel array) were combined [43] and unwrapped [60], followed by a voxel-wise, magnitude-weighted, least-squares regression of phase to echo time to arrive at the final

[3] In fact, this is not true SVD. More on this next chapter.

[4] Credit is given to Hongfu Sun for reconstructing the original images of the total field and for sharing his susceptibility mapping program

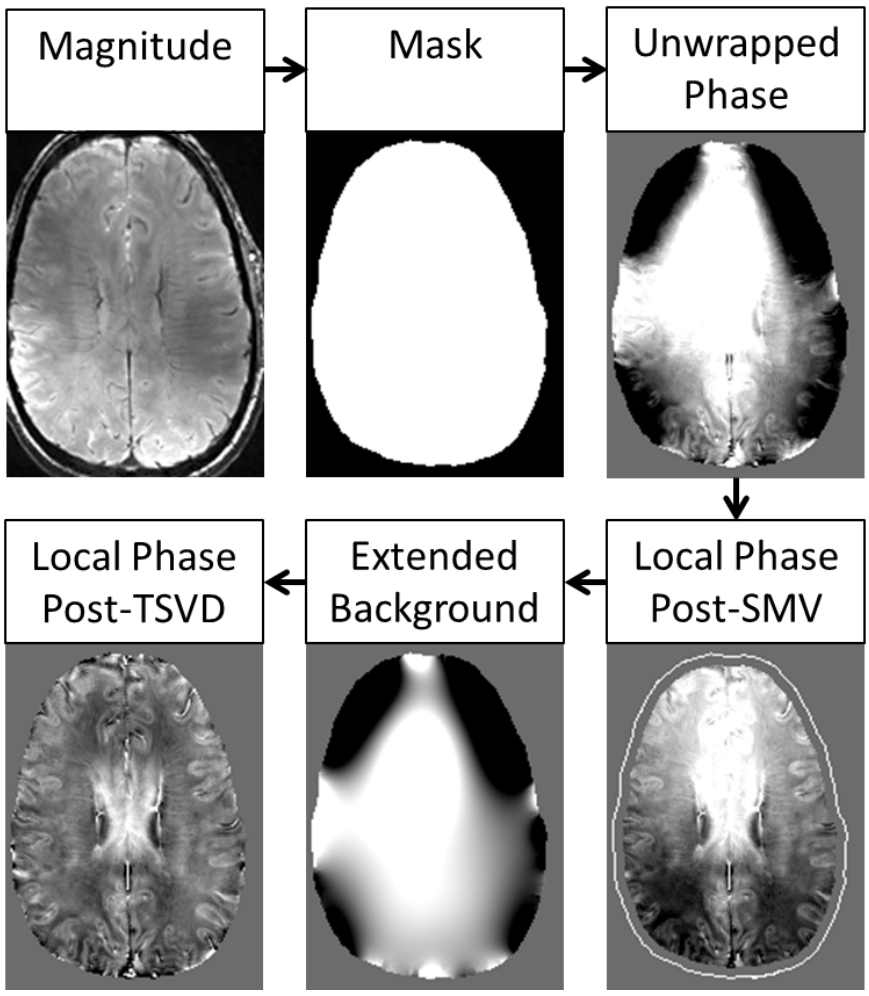

Figure 3.1: Simplified workflow. Magnitude data is passed to FSL's brain extraction tool [58], the output of which (*mask*) is used to constrain the unwrapping to the tissue of interest. The SMV kernel is applied to the unwrapped phase and deconvolved without regularization to produce the initial local phase estimate (local phase post-SMV). The outline indicates the edge of the mask (VOI) and the edge region lost in conventional SHARP. This local estimate is subtracted from the unwrapped total phase to produce an estimate of the background phase over the reduced VOI (not shown). The Taylor expansion (3.1) is used to recover the edge voxels (extended background). The result is subtracted out from the unwrapped total and SHARP-regularized to produce the final local phase post-TSVD.

field maps [61–63]. To obtain a binary brain mask for each dataset, magnitude data from the first echo were passed to FSL's brain extraction tool (BET) [58]. Phase discontinuities (jumps to high intensities) were often observed around the outer edges of the mask even after unwrapping. Whether due to shortcomings of Prelude, peculiarities inherent to the data, or the inclusion of nonbrain voxels by an overly generous BET mask, these values were, in any case, deemed unreliable. To define effective brain VOIs that excluded the edge outliers, FSL masks were eroded by between two to four voxels.

Finally, after phase processing, susceptibility maps were formed using the total variation inversion described in [35,64], with the regularization parameter 5×10^{-3} determined in [55].

3.3.3 Numerical Simulation

A numerical brain model was created by assigning a waterlike susceptibility of -9.4 ppm to an in vivo brain mask, while assigning an air-like value of 0 to the region outside the mask. This arrangement was to represent the susceptibility interfaces responsible for the harmonic background field. Although lacking susceptibility structures (e.g., veins and skull) typically associated with a more realistic field map, the simplistic model should suffice insofar as all fields owing to susceptibility sources outside of the brain are in fact harmonic within it, irrespective of the exact distribution of sources. Two internal susceptibility inclusions (susceptibility = -9.0 ppm; radii = 2 mm), simulating small spherical hemorrhages, were placed within the brain substrate: one at the center of the brain (IP-region), with the other at the edge (EP-region), such that it would be discarded post-SHARP. The susceptibility model was convolved with the unit dipole field to simulate the magnetic field [25, 26], and scaled to phase by the multiplicative factor $\gamma B_0 \mathrm{TE}$ (TE = 19 ms) to which zero-mean Gaussian noise was added (standard deviation = $\pi/4$ rad). Model phase quantities are shown in Figure 3.2a–c. To resemble the common in vivo case, the noisy total phase was masked by the brain VOI before beginning phase processing.

Accuracy of the processing scheme was assessed by means of the relative error between the noiseless model background phase ϕ_{model} and the background phase estimates ϕ_{est} given by SHARP and E-SHARP using zeroth, first, and second order expansions:

$$||\mathbf{M}(\boldsymbol{\phi}_{\mathrm{est}} - \boldsymbol{\phi}_{\mathrm{model}})||_2 / ||\mathbf{M}\boldsymbol{\phi}_{\mathrm{model}}||_2. \tag{3.3}$$

Error terms were ultimately calculated over specific VOIs: for SHARP, the masking operator $\mathbf{M}$ was strictly the reduced VOI (IP). For the E-SHARP error calculation, unless otherwise stated, $\mathbf{M}$ encompassed the brain VOI (IP $\cup$ EP), however, as the brain itself does not generally abut air in bulk, $\mathbf{M}$ was eroded by a single voxel so as not to contain the outermost voxels that defined the air-tissue interface (where the Laplacian of the background field would be nonzero).

3.4 Results

Simulation results are shown for a central slice in Figure 3.2d–i. The model background field (Fig. 3.2b) can be seen to be smooth and slowly varying away from transitions in susceptibility and is, therefore, well represented by the second order expansion. Error in the field estimation (Fig. 3.2i) is greatest in the immediate vicinity of the background source, where the field gradient is steepest. The effect of expansion order on the resulting background phase is illustrated in Figure 3.3. While relative error was 24 % for the zeroth order expansion background phase estimate across the brain VOI (IP $\cup$ EP), it was reduced

to 18 %, and further to 16 %, for first and second order expansions respectively. Across the reduced VOI alone (not including the EP voxels in $\mathbf{M}$ of (3.3)), the relative error of E-SHARP (second order expansion) was slightly higher than that of SHARP (15 % versus 13 %). The corresponding E-SHARP error over the EP voxels alone was 18 %.

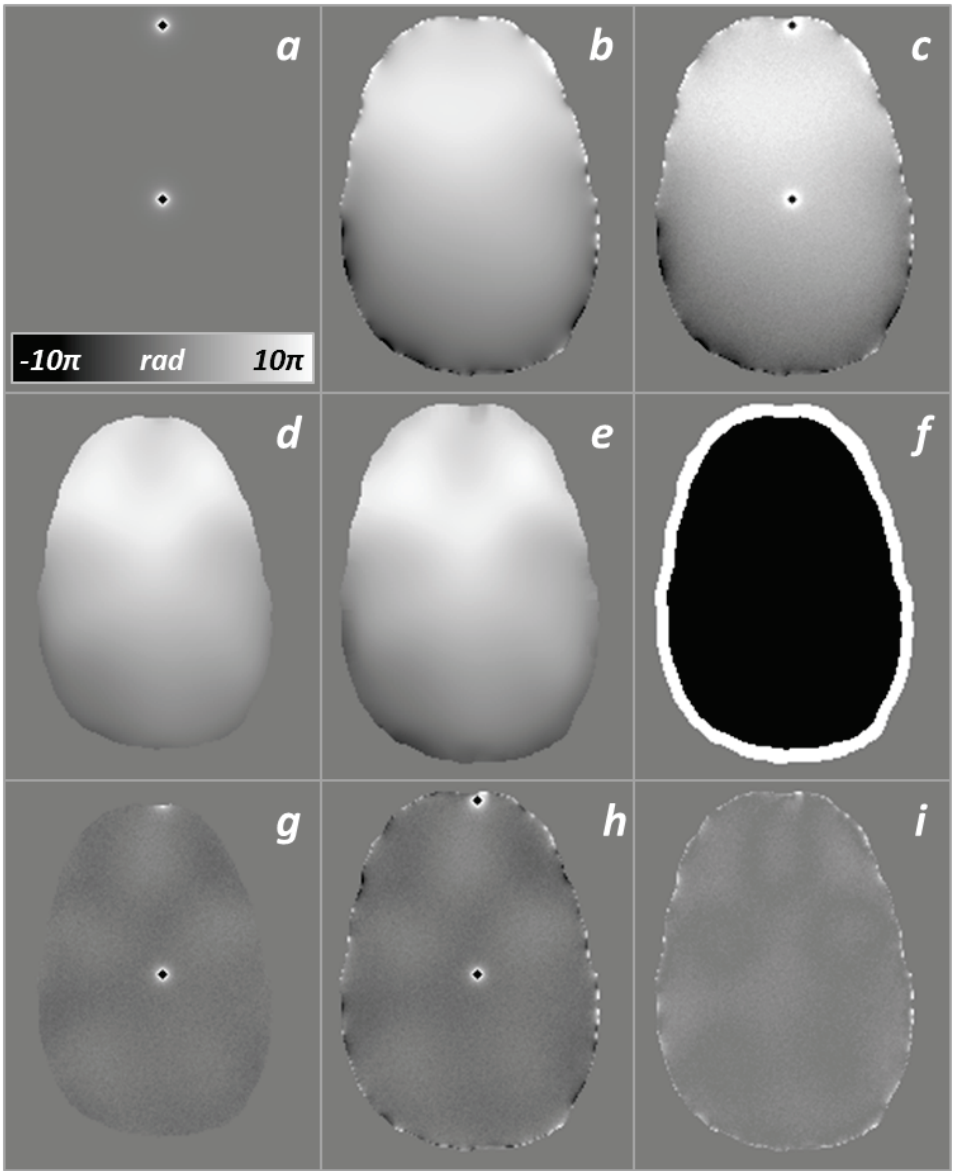

Figure 3.2: Numerical simulation. A central transverse slice is shown. The top row corresponds to model phase quantities: local (**a**), masked background (**b**), and noisy total (**c**). **d,e**: The background estimate over the reduced VOI and the full postexpansion form, from SHARP and E-SHARP, respectively. **g,h**: The local phase estimates courtesy of SHARP and E-SHARP, respectively. **i**: The error (absolute difference of (**h**) and (**a**)). All images are scaled to the same relative intensities, with the exception of (**f**), which depicts the internal (IP; black) and edge (EP; white) geometries involved.

Example results from a representative in vivo dataset are exhibited in 3.4. Compared

with conventional SHARP, E-SHARP recovered on average 26 % more brain volume (values for all five subjects: 27.6 %, 27.3 %, 27.1 %, 25.1 %, 25.4 %). This additional territory arises from the extension of the reduced VOI of SHARP by the radius of the kernel in all directions, thus revealing cortex and cortical veins [65,66] which were otherwise inaccessible. E-SHARP processing time was on average 3.4 min (range: 1.8–4.3 min).

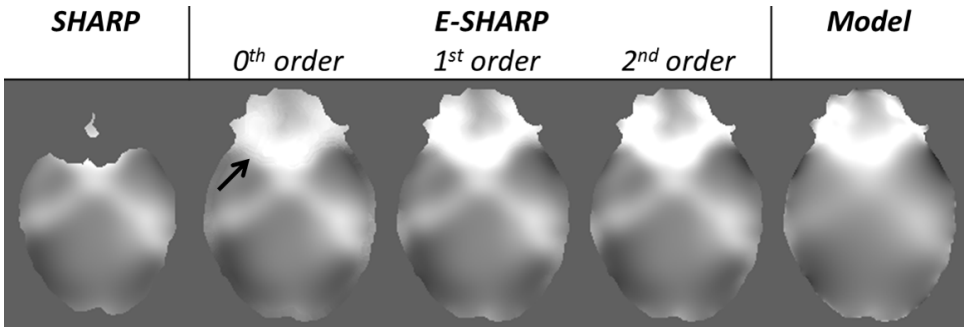

Figure 3.3: Effect of order of expansion on the estimated background phase (quantity shown) of the simulated data. Relative error over the reduced VOI was 0.13 for the SHARP background estimate. The corresponding errors were 0.24, 0.18, and 0.16 for zeroth, first, and second order expansions, respectively. As expected, by including higher order terms the discontinuities around the edge (arrow) are smoothed and at second order they are scarcely discernible.

Susceptibility maps derived from the two methods were similar across the reduced VOI. Representative examples are displayed in Figure 3.5, with differences shown in the rightmost column. While the SHARP susceptibility map (Fig. 3.5, left column) exhibited only a reduced portion of the sagittal sinus (middle row) and frontal white matter (top row), the region of susceptibility coverage in these regions was greatly enhanced by E-SHARP (middle column).

3.5 Discussion

This study demonstrated that by exploiting the analyticity of the harmonic background field, susceptibility and field maps of comparable quality to those made by conventional SHARP can be achieved with markedly greater spatial support (recovering, on average, 26 % more brain volume). Practical considerations concerning the technique are discussed in the following. First, for the Taylor series to actually equate to the dipole field, it would require, in theory, an infinite number of terms (i.e., derivatives). However, as the background field across the brain is characterized predominantly by the low-order terms, even the truncated form, as observed, can provide satisfactory results in most regions. Furthermore, because the SMV calculation is by definition indifferent to zero-mean Gaussian noise (phase noise in tissue may typically be of this form [46], the extracted background field is necessarily more or less noise free. This fact, combined with the slow-varying nature of the background field, generally make it a fairly safe quantity to subject to otherwise problematic finite-difference

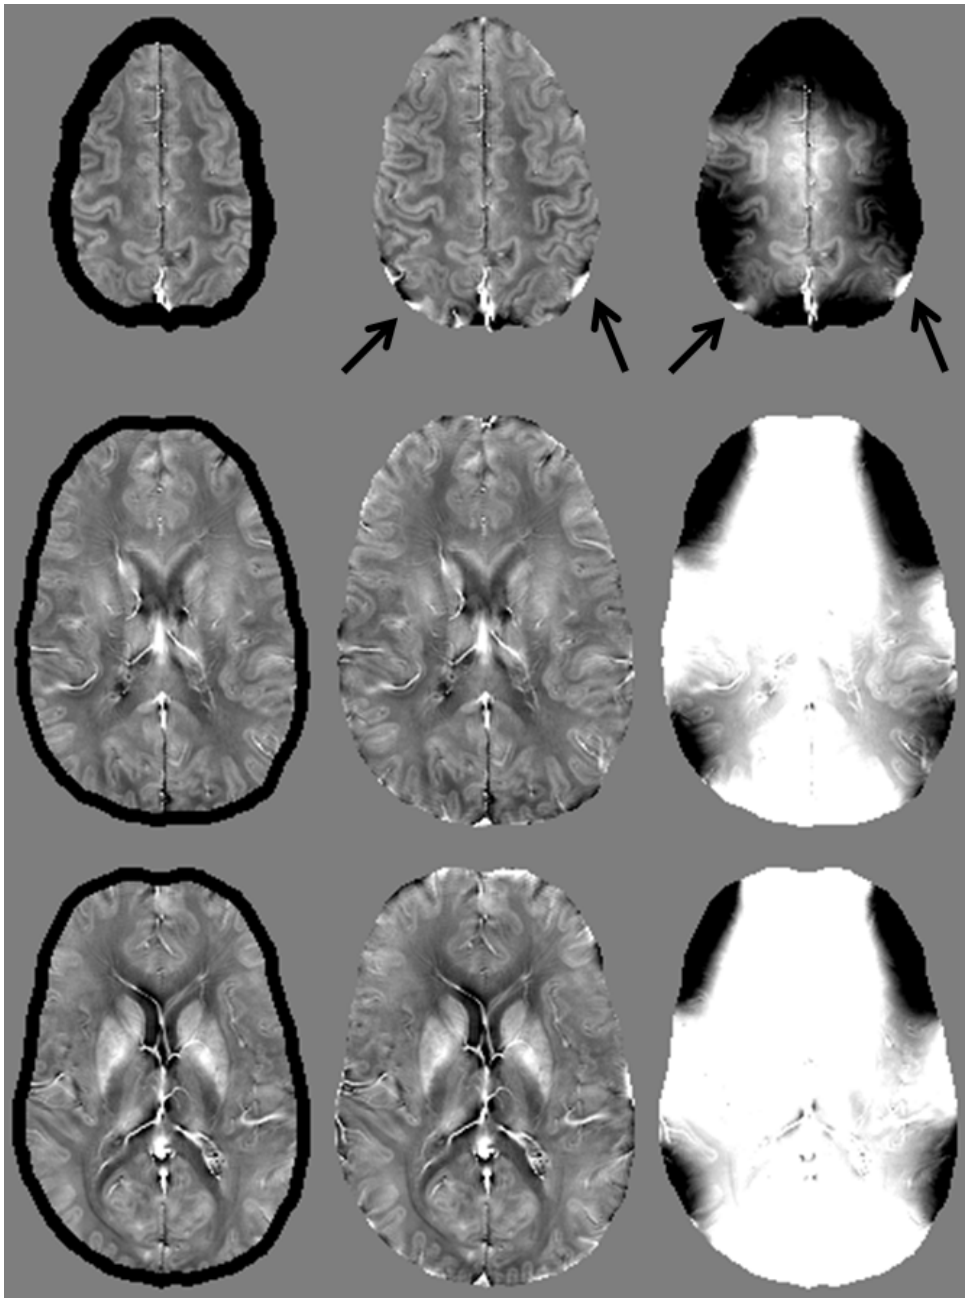

Figure 3.4: In vivo results: field processing. From left to right, the columns correspond to the conventional SHARP local field, E-SHARP local field (second order expansion), and total field, respectively. Arrows in the top row point to cortical veins, visible in the total field but lost in conventional SHARP (the black ribbon in the left column demarcates the EP region recovered by E-SHARP).

calculations.

In terms of series convergence, the second matter of note is the internal point (IP) about which the expansion is performed: in general, the greater the distance $|\boldsymbol{\xi}|$ between IP and EP, the more the field will vary between them and, accordingly, the more additional

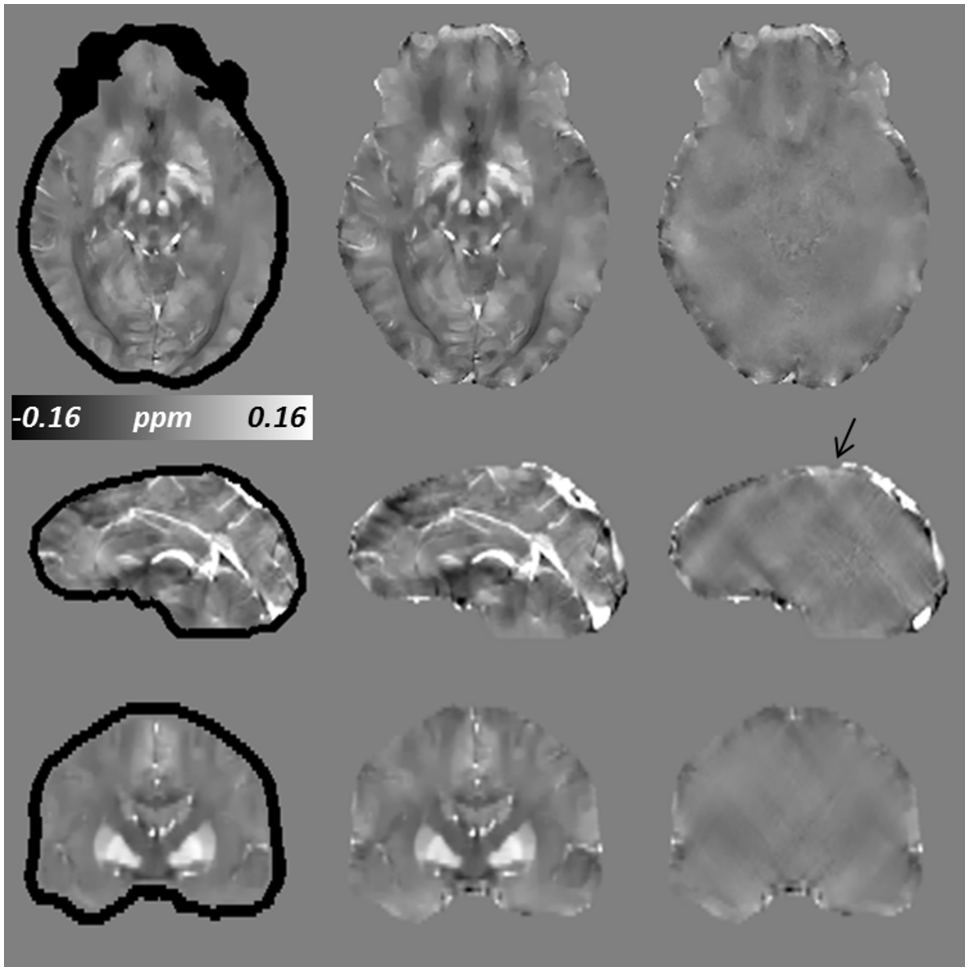

Figure 3.5: In vivo results: QSM. Columns correspond to the susceptibility maps formed using SHARP (left) for the phase processing and E-SHARP (centre). Again, the black ribbon in the left column demarcates the EP region recovered by E-SHARP. The difference is shown in the rightmost column, with the arrow in the middle row pointing to the artificial truncation of the sagittal sinus due to the naïve initial correction applied to the brain mask.

terms in the expansion will be required to compensate. In short, it is beneficial to assign IPs that are as near as possible to the EPs (truncation error being nil only in the limit of $|\boldsymbol{\xi}| \to 0$). However, the nature of finite-differences is such that a difference between points is required to estimate the derivative: using the 3-point central difference scheme, the first order derivative at x_0 requires field values at x_1 and x_{-1}. Given that the input to the expansion (the SMV-estimated background phase) has finite spatial support (the reduced VOI), we cannot calculate a first derivative directly at the edge of support, but rather one voxel removed from it; likewise, the nearest point to the edge of support at which we can calculate a second derivative is two voxels removed from it. In this way, by increasing the order of the expansion, so too must $|\boldsymbol{\xi}|$ be increased, which in some sense confounds convergence. Nevertheless, for the numerical simulation, the lowest relative error

was observed for the second order expansion (16 %). Thus, for the particular model field studied, the inclusion of higher order terms proved to be of greater benefit than the adoption of IPs that were slightly more proximal to the EPs.

Third, the quoted results depend intimately on the adopted parameters. For instance, as the error calculations incorporated the entire EP region, calculated errors were governed, not only by the order of expansion, but also by its particular geometry. A key factor in the determination of this geometry was the size of spherical kernel—a parameter that merits further study in its own right. So too does the regularization parameter [56]. Indeed, before regularization, the post-SMV local field estimate over the reduced VOI is identical between SHARP and E-SHARP and it is only subsequent to SHARP-style TSVD that there is a subtle but, nevertheless, calculable difference (2 % more error with E-SHARP) in this internal region. Ultimately, this may simply suggest that optimal SHARP parameters may not translate into optimal E-SHARP parameters. A more complete treatment of the interplay between these parameters, however, is beyond the scope of this work, for which the aim has been simply to demonstrate that by means of a slight modification to conventional SHARP one can arrive at a local field map with substantially expanded spatial support (e.g., Fig. 3.4).

Finally, although E-SHARP revealed extensive new territory, a portion of the edge remained missing due to the naïve correction applied to the initial brain mask. The sagittal view of the susceptibility maps (Fig. 3.5, middle row) evinces this issue. Because the removal of erroneous phase outliers was achieved at the expense of retaining the full brain VOI, parts of the edge region were erroneously truncated (e.g., the sagittal sinus; arrow in Fig. 3.5, right column). To solve this issue and retain the brain VOI to its fullest, a means of automatic detection and exclusion of problematic phase data will ultimately be required.

3.6 Conclusion

Extended-SHARP, an easily implemented adaptation to the postprocessing technique SHARP, can be used to determine the subset of missing field map values around the edges of the brain. Results suggest a new way of processing MR phase that may bring us one step closer to a reliable technique for whole-brain in vivo susceptibility assessment.

Chapter 4

Epilogue: To be (analytically) continued...

For roughly a decade, magnetic susceptibility-related MRI phase contrast has found clinical application in the form the susceptibility-weighted imaging (SWI) technique [44, 67–70]. Recently, efforts in this field have been increasingly directed toward quantitative susceptibility mapping (QSM), which seeks to collapse the blooming field distortions (a nonlocal and indirect effect) into the underlying material susceptibility itself [24, 27, 33, 36, 71, 72]. QSM may benefit our understanding of an array of neurodegenerative disorders for which paramagnetic iron is thought to play a role [13]. Consequently, considerable interest exists in improving the techniques for processing the requisite phase data.

A major obstacle to creating interpretable field and susceptibility maps that accurately depict tissue structures is the contaminating effect from the background field, such as that owing to air-tissue interfaces. This work has focused on spherical mean value (SMV) filtering as a means to isolate the local field of interest. The essential tool throughout, for both critically examining basic SMV-filtering in Chapter 2 as well as enhancing it in Chapter 3, has been the Taylor series and the analytic nature of the harmonic background field.

4.1 Conclusions

In Chapter 2, it was determined that the ability of the digitized mock-sphere to estimate the central value of a harmonic field depreciates with filter size. This is somewhat contrary to what one might expect based on the idea that a incorporating a larger number of samples into the estimation ought to improve the result. Rather, that assumption is strictly valid when the extra samples come courtesy of finer image resolution, but not when they derive from a larger filter.

In Chapter 3, an extension to the SMV-filtering technique SHARP was introduced and evaluated. Extended-SHARP was found to be of comparable accuracy to SHARP when applied to a simple numerical model, and qualitatively similar to SHARP when comparing in vivo field and susceptibility maps. It was also observed that although 2^{nd} order E-

SHARP worked best, the improvement from 1^{st} order was subtle (2 %), and even the 0^{th} order expansion worked reasonably well. The fact that the extension works as well as it does even at low-order—and moreover, the fact that SMV-filtering works at all—all owes to the rapid decay of the dipolar background field and its derivatives away from sources. That the brain is somewhat isolated from air in bulk means that the higher order terms to the background field have largely decayed away and are no longer significant contributing factors to the total field within the brain.

4.2 Limitations and future directions

In addition to the limitations of E-SHARP detailed in the Discussion of Chapter 3, there are at least 2 critical problems with the method.

Because the SMV cannot be calculated wherever the spherical kernel $\mathbf{S}$ overlaps with the edge of the data support, the support for the SHARP filtering operation is smaller yet and depends on the chosen size of S. Denoting this de facto support by masking operator $\mathbf{M}$ and the SHARP convolution operator by $\mathbf{L} = [(\boldsymbol{\delta} - \mathbf{S})\otimes]$, with the cost functional formalism, the process of solving for the noisy local field $\mathbf{f}_{\text{local},\epsilon}$ of (2.7) can be written succinctly as

$$argmin_{\mathbf{f}_{\text{local},\epsilon}} \|\mathbf{MLf}_{\text{local},\epsilon} - \mathbf{MLf}_{\Delta}\|_2^2. \tag{4.1}$$

The first problem is with SHARP itself, which, in order to perform some sort of deconvolution, simply ignores the masking term $\mathbf{M}$ being applied to $\mathbf{f}_{\text{local},\epsilon}$, and only inverts $\mathbf{L}$. Technically, $\mathbf{M}$ possesses an inverse only in the limit where the data support becomes the entire FOV and $\mathbf{M} = \mathbf{I}$, the identity matrix. Because in general $\mathbf{M} \neq \mathbf{I}$, the SHARP operator $\mathbf{ML}$ is no longer circulant and therefore, the deconvolution of SHARP is not true singular value decomposition. This effect of relaxing of the relevant boundary condition has been shown in [55] to be error-prone. Hence, the initial SHARP-style SMV estimation of the background field used by E-SHARP is marred with artifact. Regularization Enabled SHARP (RESHARP) corrects this issue by approaching the problem of local field determination by minimizing (4.1) iteratively via the conjugate gradient method.

The second issue is with the expansion itself. The edge extrapolation of E-SHARP may be considered a forward problem by which initial conditions of the "internal" background field are used to project the solution at a distance. Although the projected solution depends on the spatial variation of the internal field in all directions, lateral variation in the projected edge field is essentially incidental, depending mainly on the edge geometry and the adopted IP-to-EP pairing. An additional problem with E-SHARP is that the adopted one-to-one mapping of IP-to-EP is extrinsic and was adopted largely out of convenience. In general, the Taylor series expansion of $f(\mathbf{r_0})$ will converge about a neighbourhood of $\mathbf{r_0}$ provided that $f(\mathbf{r}_0 + \boldsymbol{\xi})$ remains harmonic.

Work is now underway to reformulate the analytic edge extension into a linear matrix operation such that it incorporates multiple field points into the extrapolation for a single

edge-point and also works to promote lateral smoothness of the edge field. This approach needs to be united with that of RESHARP to avoid the deconvolution artifact of SHARP, to which much of the quoted errors of Chapter 3 should be attributed. Ideally, the optimization associated with (4.1) can be reformulated to accomodate the extension such that the local field can be mapped across whole brain accurately and simply, in one fell swoop.

Bibliography

[1] E Mark Haacke, W Robert Brown, R Michael Thompson, and Ramesh Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. John Wiley & Sons, Inc., 1999.

[2] David S Mckay, Everett K Gibson Jr, Kathie L Thomas-keprta, Hojatollah Vali, Christopher S Romanek, Simon J Clemett, Xavier D F Chillier, Claude R Maechling, and Richard N Zare. Search for Past Life on Mars: Possible Relic Biogenic Activity in Martian Meteorite ALH84001. *Science*, 273(August):924–930, 1996.

[3] E. Imre Friedmann, Jacek Wierzchos, Carmen Ascaso, and Michael Winklhofer. Chains of magnetite crystals in the meteorite ALH84001: Evidence of biological origin. *P Natl Acad Sci USA*, 98(25), 2001.

[4] Jan Martel, David Young, Hsin-Hsin Peng, Cheng-Yeu Wu, and John D. Young. Biomimetic Properties of Minerals and the Search for Life in the Martian Meteorite ALH84001. *Annu Rev Earth Pl Sc*, 40(1):167–193, May 2012.

[5] J L Kirschvink, A Kobayashi-Kirschvink, and B J Woodford. Magnetite biomineralization in the human brain. *P Natl Acad Sci USA*, 89(16):7683–7, August 1992.

[6] Richard Blakemore. Magnetotactic Bacteria Magnetotactic Bacteria. *Science*, 190(4212):377–379, 1975.

[7] Hervé Cadiou and Peter A Mcnaughton. Avian magnetite-based magnetoreception: a physiologist's perspective. *J R Soc Interface*, 7 Suppl 2:S193–205, April 2010.

[8] Michael M Walker, Todd E Dennis, and Joseph L Kirschvink. The magnetic sense and its use in long-distance navigation by animals. *Curr Opin Neurobiol*, 12(6):735–44, December 2002.

[9] Joseph L Kirschvink, James W Hagadorn, and Division Geological. A Grand Unified Theory of Biomineralization. In *The Biomineralisation of Nano- and Micro-Structures*, pages 139–149. Wiley, 2000.

[10] Joanna F Collingwood, Ryan K K Chong, Takeshi Kasama, Lionel Cervera-gontard, Edward T Simpson, Mark A Smith, and Jon Dobson. Three-Dimensional Tomographic Imaging and Characterization of Iron Compounds within Alzheimer's Plaque Core Material. *J Alzheimers Dis*, 14:235–245, 2008.

[11] Yves Gossuin, Dimitri Hautot, Robert N Muller, Quentin Pankhurst, Jon Dobson, Chris Morris, Pierre Gillis, and Joanna Collingwood. Looking for biogenic magnetite in brain ferritin using NMR relaxometry. *NMR Biomed*, 18(7):469–72, November 2005.

[12] J F Schenck. Safety of strong, static magnetic fields. *J Magn Reson Im*, 12(1):2–19, July 2000.

[13] S.M. Hossein Sadrzadeh and Yasi Saffari. Iron and brain disorders. *(Pathology Patterns Reviews S64) Am J Clin Pathol*, 121(1):65–70, January 2004.

[14] Marco T Núñez, Pamela Urrutia, Natalia Mena, Pabla Aguirre, Victoria Tapia, and Julio Salazar. Iron toxicity in neurodegeneration. *Biometals*, 25(4):761–76, August 2012.

[15] Andrew J Walsh, Gregg Blevins, R Marc Lebel, Peter Seres, Derek J Emery, and Alan H Wilman. Longitudinal MR Imaging of Iron in Multiple Sclerosis : An Imaging Marker of Disease. *Radiology*, 270(1), 2014.

[16] Robert Zivadinov, Mari Heininen-Brown, Claudiu V Schirda, Guy U Poloni, Niels Bergsland, Christopher R Magnano, Jacqueline Durfee, Cheryl Kennedy, Ellen Carl, Jesper Hagemeier, Ralph H B Benedict, Bianca Weinstock-Guttman, and Michael G Dwyer. Abnormal subcortical deep-gray matter susceptibility-weighted imaging filtered phase measurements in patients with multiple sclerosis: a case-control study. *Neuroimage*, 59(1):331–9, January 2012.

[17] Eve M Dumas, Maarten J Versluis, Simon J a van den Bogaard, Matthias J P van Osch, Ellen P Hart, Willeke M C van Roon-Mom, Mark a van Buchem, Andrew G Webb, Jeroen van der Grond, and Raymund a C Roos. Elevated brain iron is independent from atrophy in Huntington's Disease. *Neuroimage*, 61(3):558–64, July 2012.

[18] Justin Y Kwan, Suh Young Jeong, Peter Van Gelderen, Han-Xiang Deng, Martha M Quezado, Laura E Danielian, John a Butman, Lingye Chen, Elham Bayat, James Russell, Teepu Siddique, Jeff H Duyn, Tracey a Rouault, and Mary Kay Floeter. Iron accumulation in deep cortical layers accounts for MRI signal abnormalities in ALS: correlating 7 Tesla MRI and pathology. *PLOS One*, 7(4):e35241, January 2012.

[19] Jens M Theysohn, Oliver Kraff, Stefan Maderwald, Markus Barth, Susanne C Ladd, Michael Forsting, Mark E Ladd, and Elke R Gizewski. 7 tesla MRI of microbleeds and white matter lesions as seen in vascular dementia. *J Magn Reson Im*, 33(4):782–91, April 2011.

[20] Grant McAuley, Matthew Schrag, Samuel Barnes, Andre Obenaus, April Dickson, Barbara Holshouser, and Wolff Kirsch. Iron quantification of microbleeds in postmortem brain. *Magn Reson Med*, 65(6):1592–601, June 2011.

[21] John F Schenck. Magnetic resonance imaging of brain iron. *J Neurol Sci*, 207(1-2):99–102, March 2003.

[22] E M Haacke, S Mittal, Z Wu, J Neelavalli, and Y-C N Cheng. Susceptibility-weighted imaging: technical aspects and clinical applications, part 1. *Am J Neuroradiol*, 30(1):19–30, January 2009.

[23] Lin Li and John S Leigh. Quantifying arbitrary magnetic susceptibility distributions with MR. *Magn Reson Med*, 51(5):1077–82, May 2004.

[24] Ludovic de Rochefort, Tian Liu, Bryan Kressler, Jing Liu, Pascal Spincemaille, Vincent Lebon, Jianlin Wu, and Yi Wang. Quantitative susceptibility map reconstruction from MR phase data using Bayesian regularization: validation and application to brain imaging. *Magn Reson Med*, 63(1):194–206, January 2010.

[25] Rares Salomir, Baudouin Denis de Senneville, and Chrit Tw Moonen. A fast calculation method for magnetic field inhomogeneity due to an arbitrary distribution of bulk susceptibility. *Concepts Magn Reson*, 19B(1):26–34, October 2003.

[26] J.P. Marques and R. Bowtell. Application of a Fourier-based method for rapid calculation of field inhomogeneity due to spatial variation of magnetic susceptibility. *Concepts Magn Reson Part B Magn Reson Eng*, 25B(1):65–78, April 2005.

[27] Karin Shmueli, Jacco a de Zwart, Peter van Gelderen, Tie-Qiang Li, Stephen J Dodd, and Jeff H Duyn. Magnetic susceptibility mapping of brain tissue in vivo using MRI phase data. *Magn Reson Med*, 62(6):1510–22, December 2009.

[28] Tian Liu, Pascal Spincemaille, Ludovic de Rochefort, Bryan Kressler, and Yi Wang. 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. *Magn Reson Med*, 61(1):196–204, January 2009.

[29] Bryan Kressler, Ludovic de Rochefort, Tian Liu, Pascal Spincemaille, Quan Jiang, and Yi Wang. Nonlinear regularization for per voxel estimation of magnetic susceptibility distributions from MRI field maps. *IEEE T Med Imaging*, 29(2):273–81, February 2010.

[30] Ferdinand Schweser, Andreas Deistung, Berengar W. Lehr, and Juergen R. Reichenbach. Differentiation between diamagnetic and paramagnetic cerebral lesions based on magnetic susceptibility mapping. *Med Phys*, 37(10):5165, 2010.

[31] Sam Wharton, Andreas Schäfer, and Richard Bowtell. Susceptibility mapping in the human brain using threshold-based k-space division. *Magn Reson Med*, 63(5):1292–304, May 2010.

[32] Sam Wharton and Richard Bowtell. Whole-brain susceptibility mapping at high field: a comparison of multiple- and single-orientation methods. *Neuroimage*, 53(2):515–25, November 2010.

[33] Wei Li, Bing Wu, and Chunlei Liu. Quantitative susceptibility mapping of human brain reflects spatial variation in tissue composition. *Neuroimage*, 55(4):1645–56, April 2011.

[34] Tian Liu, Cynthia Wisnieff, Min Lou, Weiwei Chen, Pascal Spincemaille, and Yi Wang. Nonlinear formulation of the magnetic field to source relationship for robust quantitative susceptibility mapping. *Magn Reson Med*, 69(2):467–76, February 2013.

[35] Bing Wu, Wei Li, Arnaud Guidon, and Chunlei Liu. Whole brain susceptibility mapping using compressed sensing. *Magn Reson Med*, 67(1):137–47, January 2012.

[36] Ferdinand Schweser, Karsten Sommer, Andreas Deistung, and Jürgen Rainer Reichenbach. Quantitative susceptibility mapping for investigating subtle susceptibility variations in the human brain. *Neuroimage*, 62(3):2083–2100, June 2012.

[37] Berkin Bilgic, Audrey P. Fan, Jonathan R. Polimeni, Stephen F. Cauley, Marta Bianciardi, Elfar Adalsteinsson, Lawrence L. Wald, and Kawin Setsompop. Fast quantitative susceptibility mapping with L1-regularization and automatic parameter selection. *Magn Reson Med*, Online, November 2013.

[38] Xiang He and Dmitriy a Yablonskiy. Biophysical mechanisms of phase contrast in gradient echo MRI. *P Natl Acad Sci USA*, 106(32):13558–13563, August 2009.

[39] Dmitriy a Yablonskiy, Jie Luo, Alexander L Sukstanskii, Aditi Iyer, and Anne H Cross. Biophysical mechanisms of MRI signal frequency contrast in multiple sclerosis. *P Natl Acad Sci USA*, 109(35):14212–14217, August 2012.

[40] Chunlei Liu. Susceptibility tensor imaging. *Magn Reson Med*, 63(6):1471–7, June 2010.

[41] Chunlei Liu, Wei Li, Bing Wu, Yi Jiang, and G Allan Johnson. 3D fiber tractography with susceptibility tensor imaging. *Neuroimage*, 59(2):1290–8, January 2012.

[42] Chunlei Liu and Wei Li. Imaging neural architecture of the brain based on its multipole magnetic response. *Neuroimage*, 67:193–202, February 2013.

[43] Simon Robinson, Günther Grabner, Stephan Witoszynskyj, and Siegfried Trattnig. Combining phase images from multi-channel RF coils using 3D phase offset maps derived from a dual-echo scan. *Magn Reson Med*, 65(6):1638–48, June 2011.

[44] Y Wang, Y Yu, D Li, K T Bae, J J Brown, W Lin, and E M Haacke. Artery and vein separation using susceptibility-dependent phase in contrast-enhanced MRA. *J Magn Reson Imaging*, 12(5):661–70, November 2000.

[45] John F. Schenck. The role of magnetic susceptibility in magnetic resonance imaging: MRI magnetic compatibility of the first and second kinds. *Med Phys*, 23(6):815–850, 1996.

[46] Hbkon Gudbjartsson and Samuel Patz. The Rician distribution of noisy MRI data. *Magn Reson Med*, 34:910–914, 1995.

[47] S C Chu, Y Xu, J a Balschi, and C S Springer. Bulk magnetic susceptibility shifts in NMR studies of compartmentalized samples: use of paramagnetic reagents. *Magn Reson Med*, 13(2):239–62, March 1990.

[48] Jie Luo, Xiang He, D Andre' D'Avignon, Joseph J H Ackerman, and Dmitriy A Yablonskiy. Protein-induced water 1H MR frequency shifts: contributions from magnetic susceptibility and exchange effects. *J Magn Reson*, 202(1):102–8, January 2010.

[49] Karin Shmueli, Stephen J Dodd, Tie-Qiang Li, and Jeff H Duyn. The contribution of chemical exchange to MRI frequency shifts in brain tissue. *Magn Reson Med*, 65(1):35–43, January 2011.

[50] Tobias Leutritz, Liane Hilfert, Karl-Heinz Smalla, Oliver Speck, and Kai Zhong. Accurate quantification of water-macromolecule exchange induced frequency shift: effects of reference substance. *Magn Reson Med*, 69(1):263–8, January 2013.

[51] L Li and J S Leigh. High-precision mapping of the magnetic field utilizing the harmonic function mean value property. *J Magn Reson*, 148(2):442–8, March 2001.

[52] Tian Liu, Ildar Khalidov, Ludovic de Rochefort, Pascal Spincemaille, Jing Liu, a John Tsiouris, and Yi Wang. A novel background field removal method for MRI using projection onto dipole fields (PDF). *NMR Biomed*, 24(9):1129–36, November 2011.

[53] Lin Li. Averaging of harmonic physical fields over an annular region enclosing field sources field sources. *Am J Phys*, 70(10):1029–1033, 2002.

[54] Ferdinand Schweser, Andreas Deistung, Berengar Wendel Lehr, and Jürgen Rainer Reichenbach. Quantitative imaging of intrinsic magnetic tissue properties using MRI signal phase: an approach to in vivo brain iron metabolism? *Neuroimage*, 54(4):2789–807, March 2011.

[55] Hongfu Sun and Alan H Wilman. Background field removal using spherical mean value filtering and Tikhonov regularization. *Magn Reson Med*, 2013.

[56] F Schweser, K Sommer, M Atterbury, A Deistung, B W Lehr, and J R Reichenbach. On the impact of regularization and kernel type on SHARP-corrected GRE phase images. In *Proceedings of the 19th Annual Meeting of ISMRM, Montreal, Quebec*, page 5165, 2011.

[57] Paul I. Richards. *Manual of Mathematical Physics*. Pergamon Press Ltd., Belfast, 1959.

[58] Stephen M Smith. Fast robust automated brain extraction. *Hum Brain Mapp*, 17(3):143–55, November 2002.

[59] F Schweser, M Atterbury, A Deistung, B W Lehr, K Sommer, and J R Reichenbach. Harmonic phase subtraction methods are prone to B1 background components. In *Proceedings of the 19th Annual Meeting of ISMRM, Montreal, Quebec*, page 2657, 2011.

[60] Mark Jenkinson. Fast, automated, N-dimensional phase-unwrapping algorithm. *Magn Reson Med*, 49(1):193–7, January 2003.

[61] Guillaume Gilbert, Geneviève Savard, Céline Bard, and Gilles Beaudoin. Quantitative comparison between a multiecho sequence and a single-echo sequence for susceptibility-weighted phase imaging. *Magn Reson Imaging*, 30(5):722–30, June 2012.

[62] Bing Wu, Wei Li, Alexandru Vlad Avram, Sung-Min Gho, and Chunlei Liu. Fast and tissue-optimized mapping of magnetic susceptibility and T2* with multi-echo and multi-shot spirals. *Neuroimage*, 59(1):297–305, January 2012.

[63] Christian Denk and Alexander Rauscher. Susceptibility weighted imaging with multiple echoes. *J Magn Reson Im*, 31(1):185–91, January 2010.

[64] Michael Lustig, David Donoho, and John M Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med*, 58(6):1182–95, December 2007.

[65] Ronit Agid, Ilan Shelef, James N Scott, Richard I Farb, and Venography Ctv. Imaging of the intracranial venous system. *Neurologist*, 14(1):12–22, 2008.

[66] Audrey P. Fan, Berkin Bilgic, Louis Gagnon, Thomas Witzel, Himanshu Bhat, Bruce R. Rosen, and Elfar Adalsteinsson. Quantitative oxygenation venography from MRI phase. *Magn Reson Med*, February 2013.

[67] J R Reichenbach, R Venkatesan, D J Schillinger, D K Kido, and E M Haacke. Small vessels in the human brain: MR venography with deoxyhemoglobin as an intrinsic contrast agent. *Radiology*, 204(1):272–277, July 1997.

[68] J R Reichenbach and E M Haacke. High-resolution BOLD venographic imaging: a window into brain function. *NMR Biomed*, 14(7-8):453–467, 2001.

[69] E Mark Haacke, Yingbiao Xu, Yu-Chung N Cheng, and Jürgen R Reichenbach. Susceptibility weighted imaging (SWI). *Magn Reson Med*, 52(3):612–8, September 2004.

[70] Andreas Deistung, H.J. Mentzel, Alexander Rauscher, S Witoszynskyj, W.A. Kaiser, and Jürgen R Reichenbach. Demonstration of paramagnetic and diamagnetic cerebral lesions by using susceptibility weighted phase imaging (SWI). *Z Med Phys*, 16(4):261–7, 2006.

[71] Andreas Deistung, Andreas Schäfer, Ferdinand Schweser, Uta Biedermann, Robert Turner, and Jürgen R Reichenbach. Toward in vivo histology: a comparison of quantitative susceptibility mapping (QSM) with magnitude-, phase-, and R2*-imaging at ultra-high magnetic field strength. *Neuroimage*, 65:299–314, January 2013.

[72] Jürgen R Reichenbach. The future of susceptibility contrast for assessment of anatomy and function. *Neuroimage*, 62:1311–1315, January 2012.