

Improved Magnetic Induction Tomography

Submitted December 2022, in partial fulfillment of
the conditions for the award of the degree **Doctor of Philosophy In Physics.**

Logan Turton
4342649

Supervised by T. M. Fromhold.

School of Physics and Astronomy
University of Nottingham

I hereby declare that this dissertation is all my own work, except as indicated in the text:

Signature _____

Date _____ / _____ / _____

I hereby declare that I have all necessary rights and consents to publicly distribute this dissertation via the University of Nottingham's e-dissertation archive.

Public access to this dissertation is restricted until: DD/MM/YYYY

Abstract

Imaging through magnetic induction has inherent advantages over rival technologies, including being safe, contactless, and capable of both the location and characterisation of underlying electrical properties. Applications are wide-ranging, varying from geophysical investigations of the sub-surface, through security screening of baggage, to medical imaging. This research aims to improve the quality of image reconstructions from inversion.

A derived set of magnetic fields use orthogonality to maximise the available information. Necessary generalisations of existing forward models then include accounting for magnetometer measurements and inhomogeneous magnetic fields. Adapted sensitivity formulations move to a rectangular matrix with scaling computational advantages as image resolution increases. Well suited to augmentation, this also allows the combination of multiple magnetic profiles into a single inversion. An examination of relevant inverse theory covers appropriate analysis tools and a library of algorithms. Individually algorithms have varying convergence properties, computational requirements, and regularising effects, with the specific choice subsequently problem dependent.

An approach to produce a set of the aforementioned magnetic fields with a single modular system by varying the driving currents is evaluated, alongside a framework to quantitatively assess whether a particular profile is a unique component of that set. Various designs demonstrate the robust nature of the design procedure and the influence of parameters such as component composition, driving current precision, and the regularisation term.

Finally, an evaluation considers the culmination of the preceding work applied to three-dimensional inversion in MIT. Increasing the number of spatially unique magnetic fields leads to a demonstrable improvement in image quality, with significant enhancement possible using improved orthogonality. Orthogonality also appears to reduce the computational effort required to reach semi-convergence for multiple algorithms, reducing inversion time. Expanding the number of fields explicitly enables overdetermined problems to be constructed, which show reduced noise sensitivity as the number of orthogonal fields increases. This behaviour holds for various, less ideal conductivity distributions, as well as for low conductivities and higher resolutions. A demonstration also covers the applicability of numerous novel inversion algorithms and the potential associated with optimisation.

Acknowledgements

This endeavour would not have been possible without my parents, my dad, who inspired my interest in science, and my mum, whose memory I hope I am making proud. I would also like to thank the rest of my family, particularly my Gran & Nanna, for their unwavering support, encouragement and pride, which has allowed me to make some tiny contribution to the collective knowledge of the human race.

I'm extremely grateful to my supervisor Professor Mark Fromhold. His support and guidance through turbulent times have been invaluable, while his immense knowledge and faith in me have given me the confidence to tackle the insane challenge of doctoral study. Many thanks are owed to Paul Wilkinson and Jonathan Chambers at the British Geological Survey, who provided insightful suggestions and a profound belief in my work. I would also be remiss in not mentioning my undergraduate supervisor, Professor Matthew Brookes, from whom I learned invaluable skills. He was instrumental in building the resilience required to begin my academic journey.

I am deeply indebted to the invaluable contributions, teaching, and guidance of Michael Packer, without which I would have ended my study in the earliest stages. His ingenious suggestion directly enabled my idea to vary field landscapes from a modular induction system. I am also grateful to everyone in C19 who has suffered and supported me throughout the years, Chris Morley, Jonathan Gosling, Peter Hobson, Monty Clark, Alister Davis, and everyone else. You all have continuously provided guidance, even with simple questions and without you, this would never have been written.

I would also like to recognise Dave Lazenby & Alex Berry, the physics teachers who sparked the journey to reach this point, Chris Tompkins for the extensive support, and Stu Poole for his friendship.

Last, but far from least, my deepest gratitude goes to my partner Lou your unending belief, continuous encouragement, and tremendous understanding have truly made this work possible.

Glossary

ART Algebraic Reconstruction Techniques.

CART Column-Action Reconstruction Technique.

CAV Component Averaging.

CG Conjugate Gradient.

CGLS Conjugate Gradient Least Squares.

dGSVD Damped Generalised Singular Value Decomposition.

DROP Diagonally Relaxed Orthogonal Projections.

dSVD Damped Singular Value Decomposition.

EIT Electrical Impedance Tomography.

EM Electromagnetic.

ERT Electrical Resistivity Tomography .

FISTA Fast Iterative Shrinkage-Thresholding Algorithms.

GCV Generalised Cross Validation.

GMRES Generalised Minimal Residual.

GSVD Generalised Singular Value Decomposition.

IRN Iteratively Reweighted Norm.

LR Low-Rank.

MIT Magnetic Induction Tomography.

MRNSD Modified Residual Norm Steepst Descent.

MULTISSIM3 Structural Similarity Index.

NNFCGLS Non-Negative Flexible Conjugate Gradient Least Squares.

OPM Optically Pumped Magnetometer.

RMS Root Mean Squared.

RMSLE Root Mean Squared Logarithmic Error.

SIRT Simultaneous Iterative Reconstruction Techniques.

SSIM Structural Similarity Index.

SVD Singular Value Decomposition.

TGSVD Truncated Generalised Singular Value Decomposition.

TSVD Truncated Singular Value Decomposition.

Contents

Abstract i

Acknowledgements iii

Glossary iv

1 Introduction 1

1.1 Tomography and Electromagnetic Imaging 1

1.2 Aims, Objectives & Description of the Work 5

2 History and Developments in Magnetic Induction Tomography 7

2.1 Frequency Domain MIT - The 'Slingram' Method for Geophysics 9

2.1.1 A Classic Device: EM34-3 9

2.1.2 Technological Development 12

3 Forward Modelling 16

3.1 Magnetoquasistatics 18

3.2 Target Magnetic Fields for Magnetic Induction Tomography 19

3.3 Generating Spherical Harmonic Magnetic Fields 22

3.4 Basis Set Expansion 24

3.5 Three Dimensional Current Sources & The Finite Element Method 25

3.6 Traditional MIT 29

3.7 Quasi-static Eddy-Currents 31

3.8 Approximate Formulations for Prototyping - Vector Potential 35

3.9 Approximate Formulations for Prototyping - Magnetic Field 39

4 Inverse Theory & Image Reconstruction 42

4.1 Singular Value Decomposition . . . 44

4.2 An Introduction to Regularised Inverse Methods . . . 47

4.3 Truncated Inverse Methods . . . 49

4.4 Regularisation Parameter Selection and The L-Curve . . . 51

4.5 Iterative Regularisation and Image Reconstruction Algorithms . . . 54

4.5.1 Conjugate Gradient Algorithms . . . 55

4.5.2 Algebraic, Column-Action & Simultaneous Iterative Algorithms . . . 56

4.5.3 Further Inversion Algorithms & Combination Approaches . . . 59

4.6 Resolution Matrix . . . 60

4.7 Magnetic Field & Image Quality Metrics . . . 61

5 Induction Systems For Orthogonal Magnetic Field Generation 64

5.1 Induction Systems - Localised Applications . . . 65

5.1.1 Experimental System: Four Coils . . . 65

5.1.2 Theoretical System 1: Cylindrical Coil Array . . . 73

5.1.3 Theoretical System 2: Closed-Cylindrical Coil Array . . . 85

5.1.4 Theoretical System 3: Hexagonal Array . . . 93

5.2 Induction Systems - Scanning Applications & Three-Dimensional Current Sources . . . 103

5.2.1 Array Composition . . . 104

5.2.2 Three-Dimensional Current Sources . . . 106

5.2.3 Filament Current Sources . . . 110

5.2.4 Comparison & Analysis . . . 111

5.3 Summary . . . 116

6 Eddy-Currents & Inversion Problems In MIT 118

6.1 Application of Quasi-Orthogonal Magnetic Fields . . . 118

6.1.1 Varying The Conductivity Distribution . . . 143

6.2 Optimising Sensor Positions . . . 149

6.3 Other Problems . . . 156
6.3.1 Increasing The Number of Magnetic Fields . . . 156
6.3.2 Increasing The Number of Measurements . . . 160
6.3.3 High Resolution/Low Conductivity with Novel Algorithms . . . 162
6.4 Future Work . . . 174
6.5 Conclusion . . . 176

7 Summary **178**

Bibliography **183**

Appendices **207**

A Spherical Harmonic Look-Up Tables **207**

B Biot-Savart Law & Numerical Validation Problems **215**
B.0.1 Validation 1: Square Loop . . . 216
B.0.2 Validation 2: Curved Geometries . . . 218
B.0.3 Validation 3: Least-Squares Optimised Currents . . . 220

C Numerical Solutions to Integral Expressions **222**

D Frequency Domains and the Skin Depth **225**
D.1 Helmholtz Equation . . . 225
D.2 A Conductive Sphere . . . 227
D.3 A Two Object System . . . 229

E Regularisation, Preconditioning, and Prior-conditioning Matrices **234**

Chapter 1

Introduction

1.1 Tomography and Electromagnetic Imaging

Tomography concerns creating an image of some physical property through any penetrative wave, often but not exclusively achieved using 2D slices. The most popular is probably X-Ray Tomography - a "CT scan". Electromagnetic (EM) tomography methods aim to create images of the passive EM properties present in a region of investigation, where the dominant response is usually conductivity. Numerous approaches exist, with often physically identical techniques having varying application-dependent naming conventions.

Prevalent examples would be; Electrical Resistivity Tomography (ERT), a widely used geophysical technique utilising surface measurements of electrical resistivity responses; and the mathematically identical Electrical Impedance Tomography (EIT), used for non-invasive imaging of biological resistivity. The link between geophysical techniques and their smaller-scale counterparts is a particular source of inspiration in this thesis. Simply put, this is because geophysical problems are typically harder to solve. They have a large-scale variable nature, complicated by increased complexity from significant, often overwhelming[1, 2], and varied noise contributions[3, 4]. Frequently they are massively underdetermined, meaning there are significantly fewer measurements than resolved parameters. Robust techniques and meticulous survey design appropriate for a geophysical problem are therefore highly likely to benefit to smaller-scale problems. A variety of rele-

vant geophysical techniques exist which the reader may investigate using [5, 6, 7, 8], with combinations often finding practical application in improving image quality[9].

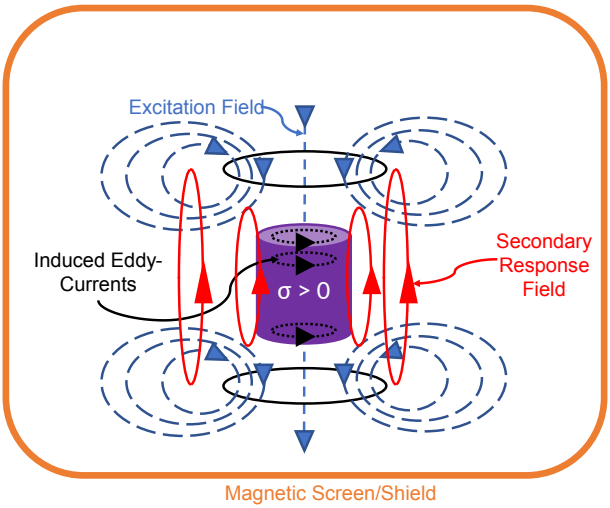


Figure 1.1: The basic procedure of a MIT investigation.

Magnetic Induction Tomography (MIT) is a contactless electromagnetic imaging process capable of imaging a complete set of a materials' passive EM properties: conductivity (σ), permeability (μ), and permittivity (ϵ)[10]. MIT attempts to image electrical properties using some variation on a simple physical relationship. An induction system applies a time-varying primary magnetic field to a region. This interacts with material, inducing eddy currents and a subsequent secondary magnetic field, dependent on the aforementioned properties, Figure 1.1. For inversion, a sensitivity formulation is necessary to relate measurements of the secondary magnetic field to the underlying electrical properties[11, 12]. Numerous naming conventions exist for this technique, including but not limited to: Mutual Inductance Tomography (MIT again), Electromagnetic Topography (EMT), Electromagnetic Induction Tomography (EIT) and Eddy-Current Tomography[10, 13, 14].

The simplest conceptually relevant geophysical EM technique is DC electrical resistivity surveying, discussed in [15, 16, 17], and [18]. A set of electrodes with galvanic contact induce currents with the resulting potential measured at other electrodes[15]. Knowledge of the injected current and geometric distribution can be related to apparent resistivity - a measure of the material resistivity where injection passes through a uniform half-space.

This value is some average over the region of interrogation, an integrated value dependent on all passive-electrical properties[19], and is a measure relevant to any contactless EM study. That averaging also increases with measurement distance, limiting feature resolution[16].

Classical applications rely on four electrodes in an extensive list of application dependent configurations[15, 20] which have been revolutionised by modern technological developments[21]. This inspires this MIT investigation, with the use of orthogonal spatial excitation[22] and optimised detection being an established approach to improve image quality.

MIT-based geophysical exploration techniques have two rough categories, frequency domain methods which make measurements of the induced magnetic field and the associated response, and time domain (or transient) methods which make measurements of the decay of the associated response[23]. Frequency Domain systems have physical advantages over their transient counterparts[24], the most fundamental being greater dynamic range in the presence of noise[25]. Response contamination in the presence of the primary field generally restricts investigation depth, whereas transient measurements occur while the primary field is 'off'. This, theoretically, allows for the application of stronger transmitters and the detection of weaker responses, which may explain the dominance of transient systems in modern commercial applications[26].

Both MIT approaches have several fundamental advantages over rival EM techniques. They are contactless, allowing surveying to occur in motion enabling increased survey size and speed. Typically scalable, with imaging volume implicitly linked to the coil-sensor array dimensions. They are also non-ionising, meaning they can be applied safely in delicate ecosystems without shielding. This thesis focuses on frequency domain techniques, motivated by the aforementioned dynamic range, which historically suffer from low resolution, with analysis frequently performed through direct interpretation of apparent conductivity rather than quantitative inversion. See Equation 2.3 in the next chapter

for more details on the conventional approach.

Geophysical applications of techniques suited to measurements of conductivity are incredibly broad. Natural geologic studies include hydro-geological investigations of aquifers/ groundwater stores[27], mineral exploration[21] and petro-physical investigations [28]. Man-made investigation can then incorporate near-surface waste dumps[29], submerged conductors, building foundations, archaeological investigations[30] and identification of unexploded ordinance (UXO)[16, 21]. Combined, these cover a conductivity range of up to 25 orders of magnitude (diamond - copper)[31], though this will be site-specific.

Investigations may also concern marine environments[32], often at very low, near-static, frequencies[33]($\approx 0.02 - 10$Hz). Practicalities in this context typically restrict the applicability of magnetic methodologies[34, 35] especially when attempting to translate those systems to land[36, 37, 38, 39]. Airborne approaches also exist, which can rapidly image much larger areas than their counterparts. Like land-surveying, modern systems are dominated by transient methodologies[26]. For more information, see great review articles in [26, 40, 41, 42, 43]. Semi-airborne methods have also been considered to tackle anthropic noise, survey precision, and increase investigation depth[44, 45].

The design and modelling of a theoretical MIT system capable of producing varied spatial excitation is non-trivial. The complexity further increases for an experimental device which may include: transmission coil(s), signal generators, power supplies, amplifiers, sensors or detector coils, data-acquisition tools, compensation circuits (to filter out the primary field) and translation stages to control measurement positions precisely. Coupling these together is the associated control software, signal processing techniques, alongside underlying physical models with sufficient computational resources to enable high-accuracy modelling. An example is in Figure 1.2. A geophysical approach also necessitates a portable system, with high-precision location mapping, and optimisation to minimise the time involved in large-scale surveying.

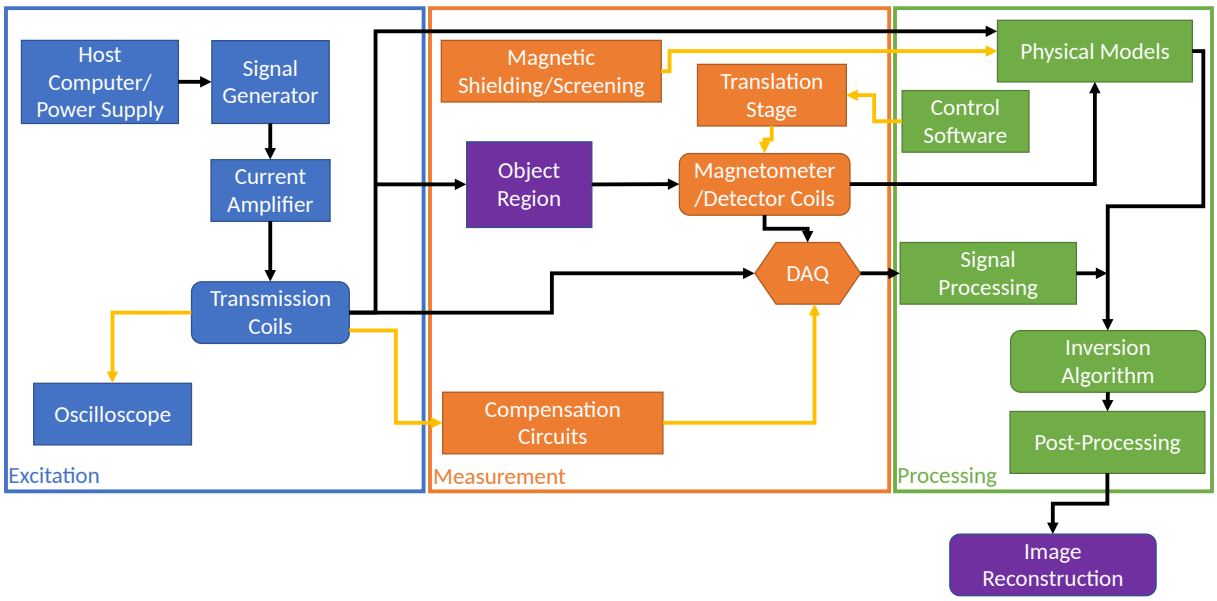


Figure 1.2: A block diagram of the components of a simple MIT system which performs image reconstruction through an inversion procedure. Assuming a static set of transmission coils, individual amplitude modulation of time-varying source currents is also necessary to produce variable spatial excitations. This modulation must have high phase stability to ensure field generation performs as intended, and, as evidenced in Section 5.1.2, this modulation must have significant precision.

1.2 Aims, Objectives & Description of the Work

This thesis aims to perform a theoretical investigation into the MIT problem, attempting to improve the quality of image reconstructions resulting from an ill-posed inverse problem. A heuristic approach has been taken, exploring a variety of free parameters in the formulation with a focus on the application of spatially varying magnetic field landscapes. Beginning with theoretical developments, identifying an theoretically infinite set of orthogonal target fields which should maximise the information content in measured responses, and a design procedure which aims to practically produce multiple of those fields from a singular modular system by varying the driving currents.

A framework is developed, providing a novel Green's function approach which enables full-phasor modelling of MIT with inhomogeneous magnetic fields. It also enables the necessary three-axial magnetometer measurements and provides an associated sensitivity formulation to allow volumetric inversion for conductivity. The design of the sensitivity formulation permits multiple profiles to be combined into a single inverse problem by aug-

mentation whilst minimising the computational requirements by moving to a rectangular sensitivity matrix. A further approximate scheme is introduced, ideal for prototyping because: it is independent of conductivity and requires substantially fewer computations than a comparative full-phasor model. Contrasted with a conventional FEM approach, these provide substantial versatility to explore various free parameters in the system.

In chapter 4 this is advanced to consider linear inversion, including detailing various concepts which are often a stumbling block for new researchers working with ill-posed physical inverse problems such as this one. The associated theory applied to MIT is frequently application specific, so various analysis tools are presented, alongside a range of appropriate algorithms, up to some state-of-the-art choices. The resolution matrix is introduced, which is another standard tool utilised in the quantitative analysis of geophysical inversion problems, followed by a brief discussion of image-quality metrics.

Subsequently, the text considers the proposed modular magnetic field design methodology, attempting to highlight some key pitfalls. A series of examples examine the influence of a variety of parameters, including: the number of modular elements, the assortment of generated fields, determining relative field quality and uniqueness, the necessary current control precision, the influence of regularisation, as well as the adaptability of the excitation patterns to different volumes.

The final chapter then combines the theoretical framework with the fields generated through the design procedure to solve the linear inversion problem in MIT. It details the impact of applying the realistically generated quasi-orthogonal fields from the previous chapter to the linear inversion problem in MIT, identifying benefits associated with varied spacial interrogation and subsequently increasing the orthogonality. This is combined with modern inversion techniques and considerations of computational optimisation to identify the areas of improvement investigated through this research.

Chapter 2

History and Developments in Magnetic Induction Tomography

The oldest recognised application of MIT involved measuring the response signal directly and can be used to detect faults within conductive objects. A region of constant conductivity should produce a constant response. A localised change in the response then identifies a region of differing conductivity. An early application of this methodology, sometimes known as eddy-current detection, was in aviation. Any aircraft in flight undergoes considerable stress, with cyclical stress fatigue cracks are able to form and propagate outwards from the regions of highest stress[46, 47]. These pose potential safety risks, and their propagation increases maintenance costs if undetected. Eddy-current detection was a well established methodology to detect such cracks by 1980, estimated to save millions of dollars a year[46]. Eddy-current imaging methods hold clear advantages for imaging metallic objects as their high-conductivity leads to a large, easily measurable response.

Recent research in metallic imaging has attempted to produce an automated setup[48]. A planar Helmholtz induction system, coupled with four hundred sensor coils, and a dual-phase lock-in amplifier, allows measurements of both in-phase and out-of-phase response components. Images of the variation in the detected phase were plotted automatically with LabVIEW control software, avoiding necessary inverse procedures required to explic-

itly characterise the conductivity. Related approaches are ubiquitous, with most authors performing experimental work attempting to avoid noise-magnifying inverse procedures by producing plots of measured data rather than passive-electrical properties. Examples undergoing a regularised inversion process, such as [49], are likely to omit the magnitude values because these are significantly harder to reproduce in practice.

Another recent proposal in metallic imaging is Rotational MIT. An equal distribution of sensors are placed in an annulus. Responses from an induction coil are measured, the system is rotated, and further measurements are taken. The result is a simple method to increase the number of independent measurements without additional excitation elements. This demonstrates increased resolution where signal-to-noise ratios are greater than 20dB, likely restricting its use in low-conductivity environments[50].

Medical applications of MIT have seen significant research recently, with a number of malignant tissue properties detectable utilising MIT at different frequencies[51]. Modern works have considered miniaturising the induction system, including just a single coil[52, 53], and accelerating imaging time. This can be achieved by relocating a single coil at a fixed RF frequency and collecting associated coil loss data. Coil loss data considers the resistive loss due to inductive coupling between the coil and sample region. The variations in coil loss data can then be related analytically to a conductivity distribution[52]. Optical tracking of this single coil was later introduced to accelerate imaging time[53].

In many MIT systems, the out-of-phase component holds all information with respect to the conductivity response. Highly-phase stable devices are then integral to any MIT approach which makes such measurements, typically associated with regions of low conductivity or experiments at low excitation frequencies. A highly phase-stable detection amplifier was developed in [54] for use in the medical sphere at ambient temperatures from 35 - 50°C. This is equally applicable in high-frequency measurements of metallic objects where small drifts in temperature with conventional equipment will lead to large drifts

in the phase measurement[14]. It is noteworthy that the amplifier and phase-detection method required is highly dependent on the exact magnetic fields induced and subsequently measured[14].

The introduction of Optically Pumped Magnetometers (OPMs) holds the potential to revolutionise MIT, with highly accurate three-axis detection potentially possible in unshielded environments without requiring considerations of background noise[55, 56]. An OPM is essentially a group of atoms optically pumped into a known quantum state, sensitive to the magnetic field which needs to be measured. These are then probed by a resonant laser which may allow for sensitivities up to $10^{-15}\frac{T}{\sqrt{Hz}}$ in low frequency magnetic fields[57] - particularly suitable to MIT. These devices are highly sensitive to small magnetic fields and utilise weak lasers in a room temperature setup to perform continuous measurements with low power draw. One existing OPM applied to MIT operates at $2-6\times10^{-12}\frac{T}{\sqrt{Hz}}$[56].

2.1 Frequency Domain MIT - The 'Slingram' Method for Geophysics

2.1.1 A Classic Device: EM34-3

The best-known geophysical application of MIT is likely the frequency-domain 'Slingram' method[58]. The first documented use of this method appears to be from 1948, in a company founded by German physics professor Freidrich Forster, citing technologies developed during the second world war. Investigations are confined to the quasi-stationary limit at low induction numbers, recognised by a World Trade Organisation international standard[59]. For geophysical uses, this typically limits the frequency of the induced field to below a few kilohertz. Measurements of the in-phase ($H_p(r'')$) and quadrature ($H_s(r'')$) components of the vertical magnetic field are made, usually at several frequencies[60]. A ratio of the two leads to apparent conductivity. Contour maps of the apparent conductiv-

ity may be interpreted directly to infer subsurface parameters, bypassing a quantitative connection to the generating conductivity through the use of inversion[16]. This includes differentiating between conductive fault and shear zones.

The Geonics EM34-3 represents a simple example of such a system, with two (conventionally horizontal co-planar) magnetic dipoles as the receiver and transmitter, shown in Figure 2.1. Three inter-coil separations, at 10,20 and 40m, with a fixed frequency, generate responses from shallow depths between 6-30m[61]. Contour mapping of the resulting apparent conductivity measurements and careful analysis of anomalies in the data can detect underground features such as cavities and groundwater aquifers[62]. For this system, a response parameter, Q, can be defined as[63],

$$Q = r\sqrt{\frac{\sigma\mu\omega}{2}}, \tag{2.1}$$

where r is the inter-coil separation, μ is the permeability of the region (assumed to be constant), and ω is the angular frequency of the source magnetic field. If $Q << 1$ is true for all interrogated conductivities, a fixed coil separation and frequency can be used to image at a fixed depth. This approximation will only hold for ground conductivities under approximately 1000 mS/m and requires systems where the coil separation is significantly less than the skin depth[61] - the aforementioned limit of low induction numbers. The 'skin depth' (δ) is defined as the depth over which the primary magnetic field strength decays to 1/e of its max value[60];

$$\delta = \sqrt{\frac{2}{\omega\mu\sigma}} = \frac{r}{Q}. \tag{2.2}$$

The skin depth is also a measure of the characteristic length-scale of the investigation. In this limit the ratio of the two magnetic fields, is directly proportional to the apparent subsurface conductivity[61],

$$\sigma_a = \frac{4}{\mu_0\omega r^2}\left(\frac{H_s(r'')}{H_p(r'')}\right). \tag{2.3}$$

A commonly encountered issue in this setup is related to the relative strengths of the



Figure 2.1: The Geonics EM34-3 is a frequency-domain geophysical MIT system. The transmission and receiver coils (orange) have a 63cm radius and width of 5cm, with the total system weighting 20.5kg. A lightweight shielded cable (black) is used as a reference to remove the primary field from the measurements. The system can image conductivities in the range of 10-1000mS/m with a measurement accuracy of $+/-5\%$ at 20mS/m. An 'XL' system which can induce a stronger primary magnetic field exists for areas with significant noise components. Image from www.geomatrix.co.uk.

primary and quadrature magnetic fields[16]. The quadrature component is often several orders of magnitude smaller than the primary field, necessitating some way of filtering out primary contributions. In simple setups, geometric nulling by ensuring transmission and receiver coils are mutually orthogonal can be sufficient to overcome this issue[25, 64, 65]. The quadrature component is also easily masked due to environmental noise, systematic[66] & random calibration errors, temperature instability or even the minute variations in the magnitude of the applied currents due to an unstable power source[67, 68, 69].

Additional electromagnetic information can be acquired using multiple dipole modes[61] - the typical application of the EM34-3, known as a horizontal dipole-mode (HDM), has the transmission coil placed with a horizontal alignment to the subsurface. The vertical dipole-mode (VDM), requiring a ninety-degree rotation of the transmission coil, typically

has twice the investigation depth at the cost of near-surface resolution[16, 70]. With further depth information also generated by varying excitation frequency[16], with a good reference on multi-configuration devices available is in[15]. Investigations have also considered non-standard coil orientations [71], identifying additional information for three-dimensional inversions can be extracted via a forty-five-degree coil inclination, but practically there is little gain to such non-standard coil orientations.

Despite being an early example of slingram surveying, the Geonics EM34-3 still has its place today. A recent survey investigating seawater intrusion in the Martil-Alila plain, Morocco[72] found the EM34-3 in both horizontal and vertical dipole modes presents a powerful tool. Several advantages over traditional gravity and electrical surveys are noted, including speed, cost-effectiveness, flexibility and sensitivity to small changes in conductivity. D. Gomez-Ortiz et al.[73] applied the system with both dipole modes to image thermal anomalies in dry rock systems, noting that highly-resistive areas are associated with high-temperature zones. They found the slingram survey to be a reliable tool for providing both location and the lateral/vertical extent of thermal anomalies.

A recent study by Ame Thato Selepeng et al.[74] considered interpretations of the 'unphysical quantity' of negative conductivity. The negative values can be useful in the sub-surface characterisation of highly conductive zones. As an anisotropic effect, it can be removed from data by data acquisition in multiple dipole modes. Despite the advances in inversion codes[75] others using this system (or similar) simply use it with a direct interpretation to identify areas of interest for subsequent investigation[76, 77].

2.1.2 Technological Development

The development of a geophysical MIT system is non-trivial. Investigations depths may vary from tens of centimetres to hundreds of kilometres, including numerous environments, making singular system design near-impossible. One commercial focus has been tackling the difficulty of measuring quadrature components. 'Fixed-boom' systems involve a fixed

transmitter-receiver separation and mounting apparatus. They allow automated primary field cancellation but limit investigation depth[60], and may be complemented by filters and compensation circuits. An example would be the EM31-MK2 system, Figure 2.2.



(a) The EM31-MK2 fixed-boom system.



(b) The EM31-MK2 and its GPS tracker.

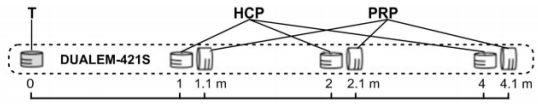
Figure 2.2: The EM31-MK2 is a standard fixed-boom system. Transmission and detection coils have a fixed separation of 3.66m at 9.8kHz. It images a similar conductivity range with comparable accuracy to the EM34-3, limited to depths up to 6m whilst being roughly half the weight - suitable for surveying by a single operator. Images from www.geomatrix.co.uk.

Approximate compensation is common, but error-prone in practice[1]. Source-receiver separation errors impact the in-phase measurement significantly and are often magnified during large surveys[78, 79], whereas the out-of-phase component is typically stable in their presence[80]. Electrical instabilities exist in much commercial instrumentation[81], inspiring the development of the DUALEM-1S instrument which also demonstrated self-calibration and internal data storage[82, 83]. Later developments towards three-dimensional subsurface characterisation[60] include multi-frequency generators[84], multi-coil induction systems[85, 86], continuous data & position acquisition tools (GPS Tracking) and airborne investigations[70, 43].

Arguably the DualEM-421S, Figure 2.3, is the state-of-the-art[87, 88, 89, 90]. It is multi-configuration with six receivers of varying orientations, operating at different frequencies below 10Hz. It demonstrates improved data acquisition and measurement precision, including allowing for magnetic susceptibility measurements[91, 92, 93], but struggles with small conductivity contrasts[94]. The CEM-2[95] in an alternative using three co-planar coils, with pulse width modulation allowing the combination of multiple frequencies into a single waveform with arbitrary intervals. The third 'bucking' coil measures and cancels the primary field. Similar multi-receiver approaches are seen in marine environments[1, 96, 97], suggesting inherent advantages to multiple simultaneous measurement positions[75]. This consistency forms part of the inspiration to formulate an MIT problem which allows for an arbitrary set of tri-axial measurements.



(a) Typical use of the device.



(b) Coil configurations which allow six simultaneous measurements.

Figure 2.3: The DualEM-421S is a high sensitivity device and contains six horizontally (HCP)/ vertically (PRP) aligned detection coils. It has a maximum penetration depth of roughly 10m, and data is acquired at up to 10km/h with high lateral resolution. Images from DUALEM Marketing and ramboll.com.

Theoretical developments include Sternberg et al.[25] presenting Alternating Target Antenna Coupling (ATAC). Controlled dipole rotations allow multiple response measurements. During rotations, the source-target coupling varies, whilst the primary field contribution does not. This allows the removal of that contribution from the data and results in a thirty-percent improvement in depth resolution. Differential Target Antenna Cou-

pling (DTAC) generalises this with independent transmitter/receiver rotations and larger magnetic moments. Other works have considered multi-frequency resonant networks, which identify resonances via a Fourier transform of a low-frequency square pulse[98] to maximise information. Inductive Source Resistivity (ESR) combines both electric and magnetic field measurements from a single source to provide complementary information, though it has seen little practical use[99, 100, 101].

Modern investigations often apply multiple techniques either as calibration steps[102, 103] or into joint inversions[16, 104, 105, 106, 107] to overcome intrinsic methodological limitations. But MIT is still used alone [30, 108, 109] including research to extend the range of problems it can be applied to[110, 111].

Developments have also concerned themselves with the ubiquitous issue of noise. A major source in any MIT study is the electronics involved[112, 113]. This could be approached using physical shielding[114], passive interference filtering[115, 116, 117, 118, 119, 120], active interference filtering[81, 121, 122, 112], as well as improved, highly stable, induction circuitry[123, 124, 125, 126]. Modelling noise is a key issue which can include excluding features which influence EM transmission[1, 127, 128] or effects such as frequency-dependent induced polarisation[23].

Chapter 3

Forward Modelling

This chapter begins with the basic framework for modelling in Electromagnetism and introduces the hypothesis that orthogonal fields should be beneficial for the MIT inverse problem. Given this, a standard set of well-known orthogonal magnetic fields are detailed. To produce these an algorithm is derived, capable of taking various arbitrary inputs covering coil geometries, target field, and generation volume, then returning an appropriate set of currents through individual coils to produce that field. The ill-posed nature of this calculation is overcome by utilising a minimisation coupled with the current norm. A further development considers generalising this to utilise three-dimensional distributions of current density. Finally, a mathematical framework to model the entire MIT problem is introduced, aiming to overcome the limitations of a traditional approach when applying numerous inhomogeneous magnetic fields. The first development is a full-phasor model extended to a Greens' function solution with various computational advantages. A secondary development then covers an approximate phase-separated regime, which links closely to existing geophysical models for the MIT problem. This has additional computational advantages because the sensitivity is conductivity-independent.

For any researcher attempting to get involved in MIT-related research, the confusing literature, terminology, and theory are often significant stumbling blocks to making progress with an idea. Even with analytic expressions, the computational requirements of a complete model are almost guaranteed to be extreme. The following theory section attempts

to explain all the related components of a specific model in a coherent framework. Individual components are not all that complex - combining them into a coherent system is where the difficulty lies.

For this model in Magnetic Induction Tomography, a complete mathematical description can be broken into two sections. The forward model - considers most of the physics, including the generated magnetic fields, and their interaction with material to produce a measurable response. The Inverse Model - relating the sensitivity calculation and measurements to some unknown image of the passive electric properties by applying inversion algorithms.

Using Maxwells' equations, given in differential form:

$$\text{Gauss' Law} \qquad \vec{\nabla} \cdot \vec{E}(\vec{r}) = \frac{\rho(\vec{r})}{\epsilon(\vec{r})}, \tag{3.1a}$$

$$\text{Gauss' Law of Magnetism} \qquad \vec{\nabla} \cdot \vec{B}(\vec{r}) = 0, \tag{3.1b}$$

$$\text{Faraday's Law} \qquad \vec{\nabla} \times \vec{E}(\vec{r}) = -\frac{\partial \vec{B}(\vec{r})}{\partial t}, \tag{3.1c}$$

$$\text{Ampere's Law} \qquad \vec{\nabla} \times \vec{B}(\vec{r}) = \mu(\vec{r})\vec{J}(\vec{r'}) + \mu(\vec{r})\epsilon(\vec{r})\frac{\partial \vec{E}(\vec{r})}{\partial t}. \tag{3.1d}$$

Here, $\vec{B}(\vec{r})$ and $\vec{E}(\vec{r})$ are the spatially varying multi-dimensional vector electric and magnetic fields in region $\vec{r}$, $\rho(\vec{r})$ is the volume charge density, $\epsilon(\vec{r})$ is the permittivity of the region, $\vec{J}(\vec{r'})$ is the vector current density providing the source field from some induction geometry, $\vec{r'}$, and $\mu(\vec{r})$ is the spatially varying permeability of the region. From Gauss' Law of Magnetism, one can then immediately conclude from the grad curl vector identity the existence of a magnetic vector potential, $\vec{A}(\vec{r})$, defined by;

$$\vec{B}(\vec{r}) = \vec{\nabla} \times \vec{A}(\vec{r}). \tag{3.2}$$

To be uniquely defined, the divergence of this quantity must also be defined, with various choices in use depending on the specific application. This choice of divergence is known

as a gauge choice and will be explicitly specified where relevant.

3.1 Magnetoquasistatics

Magneto-quasistatics applies to any situation where sources are static, varying slowly, or repetitively in time. The natural choice of gauge for this problem is the Coulomb Gauge,

$$\vec{\nabla} \cdot \vec{A}(\vec{r}) = 0. \tag{3.3}$$

Under these conditions, all time-derivatives approach zero, leading Equation (3.1d) to simplify to,

$$\vec{\nabla} \times \vec{B}(\vec{r}) = \mu(\vec{r})\vec{J}(\vec{r'}). \tag{3.4}$$

Even in conductive media it is often true to say that the permeability varies little, if at all, spatially so $\mu(\vec{r}) = \mu_r = \mu_0\mu_r$ is true, where μ_0 is the permeability of free space. Substituting Equation (3.2) in and combining the definition of the curl of the curl with the gauge choice leads to,

$$\nabla^2\vec{A}(\vec{r}) = -\mu_r\vec{J}(\vec{r'}), \tag{3.5}$$

an integral solution for which can be found through the method of Greens' functions,

$$\vec{A}(\vec{r}) = \frac{\mu_r}{4\pi}\int\int\int \frac{\vec{J}(\vec{r'})}{|\vec{r}-\vec{r'}|}d^3\vec{r'}. \tag{3.6}$$

Reusing the definition in Equation (3.2) leads to the well known Biot-Savart law,

$$\vec{B}(\vec{r}) = \frac{\mu_r}{4\pi}\int\int\int \frac{\vec{J}(\vec{r'}) \times (\vec{r}-\vec{r'})}{|\vec{r}-\vec{r'}|^3}d^3\vec{r'}. \tag{3.7}$$

This can then be used in combination with the superposition principle, allowing the quasi-static magnetic field generated by an arbitrary number of current sources to be calculated. Assuming that the current sources - or coils - are constructed of filament (i.e infinitesimally

thin) wires, such $\vec{J}(\vec{r'}) = Id\vec{r'}$,

$$\vec{B}(\vec{r}) = \frac{\mu_r}{4\pi} \int\int\int \frac{Id\vec{r'} \times |\vec{r} - \vec{r'}|}{|\vec{r} - \vec{r'}|^3} d^3\vec{r'}, \tag{3.8}$$

where I is some constant current through wire element of direction $d\vec{r'}$.

The primary focus of this thesis considers a theoretical, simulation-based study with custom codes. Appendix B details an analytic solution to Equation (3.8) for a general wire segment, confirming the numerical implementation holds for a simple square current loop and extends to curved geometries.

3.2 Target Magnetic Fields for Magnetic Induction Tomography

In MIT an ideal approach would be the ability to create a small, localised, oscillating 'hot-spot' magnetic field scanning across a region of interest; this focused field forms eddy-currents only in a small volume. Responses could then be directly attributed to a specific location in space, allowing easy three-dimensional mapping of the passive electric properties. Fundamental physical limitations in electromagnetism mean this is practically impossible to achieve.

In the thought experiment, the applied magnetic field never self-interacts, leading to an unambiguous relation between the response and underlying passive electrical properties. A proposal is that a set of magnetic fields with zero overlap - which are orthogonal - would therefore be an ideal way to investigate a region of interest. Each successive magnetic field provides unique information, with increased information from magnetic field profiles of greater complexity.

One possible source of appropriate magnetic fields is inspired by traditional 'shimming'

approaches to produce a uniform magnetic field. If the only source currents in the system are inside the coil(s), the region of investigation is current free and so[129];

$$\vec{\nabla} \times \vec{B} = \vec{\nabla} \cdot \vec{B} = 0. \tag{3.9}$$

Under this assumption a magnetic scalar potential (Ψ) can also be introduced;

$$\vec{B} = -\vec{\nabla}\Psi = \vec{\nabla} \times \vec{A}. \tag{3.10}$$

Using the Coulomb gauge, the following relationships hold;

$$\vec{\nabla} \cdot \vec{\nabla}\Psi = \vec{\nabla} \times \vec{\nabla} \times \vec{A} = 0, \tag{3.11}$$

and,

$$\vec{\nabla} \times \vec{\nabla} \times \vec{A} = -\vec{\nabla}^2\vec{A}. \tag{3.12}$$

There are then three Laplace equations defining the magnetic fields inside the current free volume,

$$\vec{\nabla}^2\Psi = 0,\ \vec{\nabla}^2\vec{A} = 0,\ \vec{\nabla}^2\vec{H} = 0. \tag{3.13}$$

This means, given a solution to the Laplace equation for Ψ, there is also a solution for $\vec{A}$ and $\vec{H}$.

An ideal solution to the Laplace equation, well known to any undergraduate studying the hydrogen atom, comes in the form of spherical harmonics - an infinite, mutually orthogonal set of mathematical objects. In magnetic shimming, a particular generated magnetic field is decomposed into a set of spherical harmonic coefficients. A common approach for uniform fields is to then design a shim system which cancels the leading-degree error harmonic, improving the homogeneity of the generated profile.

This tells us that there exists a set of infinite, mutually orthogonal functions, which

are a natural solution to the governing equations for the production of magnetic fields. Choosing a particular system wide target harmonic for Ψ then leads to related orthogonal functions in $\vec{A}$ and $\vec{H}$, and established theory exists to produce specific spherical harmonic magnetic fields.

For real, measurable, magnetic fields which avoid complex phase relationships the set of 'real' spherical harmonics are a pertinent choice. Other texts may refer to these as 'tesseral' spherical harmonics or spherical harmonics of sine/cosine type which are defined as,

$$T_n^m(r,\theta,\phi) = \binom{r^n}{r^{-n-1}} Y_{n,m}(\theta,\phi) = \begin{cases} \binom{r^n}{r^{-n-1}} X P_{n,|m|}(\cos\ \theta)\sin(|m|\ \phi), & \text{if } m<0, \\ \binom{r^n}{r^{-n-1}} X P_{n,m}(\cos\ \theta), & \text{if } m=0, \\ \binom{r^n}{r^{-n-1}} X P_{n,m}(\cos\ \theta)\cos(m\ \phi), & \text{if } m>0. \end{cases} \tag{3.14}$$

Here X is some normalisation constant, $Y_n^m(\theta,\phi))$ is the normalised angular portion of a spherical harmonic, and $P_{n,m}$ is the associated Legendre polynomial of degree n, order m. There are many possible choices of normalisation, here it is chosen so the overlap integral between the angular portion of two harmonics is,

$$\int\int T_{n,m}(r=1,\theta,\phi){T^*}_{n',m'}(r=1,\theta,\phi) = \delta_{n,n'}\delta_{m,m'}, \tag{3.15}$$

where $\delta_{n,n'}$ is the Kronecker delta function and the normalisation is,

$$X_i = (-1)^m\sqrt{2}\,\sqrt{\frac{(2n+1)}{(4\pi)}\frac{(n-|m|)!}{(n+|m|)!}}. \tag{3.16}$$

Various contrasting naming conventions and normalisation terms are in use across the literature, including [129], which instead refer to a subset of the real spherical harmonics, $m \neq 0$, as tesseral harmonics with $m = 0$ known as zonal harmonics. The functions utilised in this thesis are tabulated for reference in Appendix A.

3.3 Generating Spherical Harmonic Magnetic Fields

There is a volume of well-established theory to design induction systems capable of producing specific spherical harmonic magnetic fields with high fidelity. Production of multiple magnetic fields of this type presents a design problem in a practical setting - one cannot realistically couple numerous unique coil designs together into a device with a reasonable footprint. This body of work considers the potential to use a single modular induction system composed of simple building blocks to produce multiple magnetic fields by varying the amplitude of current driven to individual components.

This can be achieved by utilising an inverse method to solve the governing system of equations defined in Equation (3.8) for current, provided knowledge of both a geometry, and a generation volume. Differentiating Equation (3.8) with respect to current leads to a related equation for the magnetic flux sensitivity to current,

$$\frac{\partial \vec{B}(\vec{r})}{\partial I} = \vec{M}^I(\vec{r}) = \frac{\mu_r}{4\pi} \int\int\int \frac{d\vec{r'} \times (\vec{r} - \vec{r'})}{|\vec{r} - \vec{r'}|^3} d^3\vec{r'}. \tag{3.17}$$

In terms of a discrete approximation of the integral expression, this can then be re-expressed as a linear matrix equation,

$$\vec{B}_j = \sum_{k=1}^{K} I_k \frac{\partial \vec{B}_j}{\partial I_k} = \sum_{k=1}^{K} I_k \ \vec{M}^I_{jk}, \tag{3.18}$$

where $\vec{M}^I_{jk}$ represents the sensitivity to the overall magnetic field from wire element, k, at measurement position(s) j, and I_k is the associated current element which generates a magnetic field $\vec{B}_j$.

The inverse procedure can then be formulated as a least squares minimisation, with a cost function for a specific target field $B^{desired}$ being,

$$\Phi = \sum_j (\vec{B}^{desired}_j - \sum_{k=1}^{K} I_k \ \vec{M}^I_{jk})^2. \tag{3.19}$$

Generally, this is an ill-posed problem due to the existence of infinite, unique, current distributions which can equally generate an arbitrary magnetic field and so requires regularisation - a subject tackled in more detail in the inverse theory portion of this thesis. A simple approach is the addition of any penalty term which varies quadratically with current, commonly known as Tikhonov regularisation,

$$\Phi = \sum_{j}(\vec{B}_{j}^{desired} \; - \; \sum_{k=1}^{K} I_k \; \vec{M}_{jk}^{I})^2 - \lambda I_k^2, \tag{3.20}$$

where λ is some adjustable weighting parameter. This weighting is commonly known as a minimisation with respect to the current norm. It determines the balance between a smooth distribution of current across all sources K, and a distribution which reproduces the target field. Minimising this cost function leads to,

$$I_k = (\vec{M^I}_{jk}^{\top} \; \vec{M^I}_{jk} \; - \lambda \; \mathbf{I}^{\top} \; \mathbf{I})^{-1} \; (\vec{M^I}_{jk} \; \vec{B}_{j}^{desired}), \tag{3.21}$$

where $\mathbf{I}$ is the identity matrix. λ physically represents the relative weighting of minimisation of the current-norm across the system compared to the target-field minimisation. Large λ values result in a well-conditioned inverse problem with a smooth distribution of source currents to the detriment of accurate generation of $\vec{B}_{j}^{desired}$. Conversely, the product of a smaller value is a less well-conditioned problem, with a nuanced distribution which may vary abruptly to produce $\vec{B}_{j}^{desired}$ more accurately. This is valid for any physically attainable choice of $\vec{B}_{j}^{desired}$, giving a best approximation where a specific system cannot accurately produce that field.

Choosing a specific spherical harmonic as the target systems' magnetic scalar potential,

$$I_k = (\vec{M^I}_{jk}^{\top} \; \vec{M^I}_{jk} \; - \lambda \; \mathbf{I}^{\top} \; \mathbf{I})^{-1} \; (\vec{M^I}_{jk} \; \vec{\nabla}\Psi) \tag{3.22}$$

$$= (\vec{M^I}_{jk}^{\top} \; \vec{M^I}_{jk} \; - \lambda \; \mathbf{I}^{\top} \; \mathbf{I})^{-1} \; (\vec{M^I}_{jk} \; \vec{\nabla} T_n^m(r, \theta, \phi)). \tag{3.23}$$

In [129], Equations [13], [14] and [15] give expressions for the target magnetic fields with $m \geq 0$, which are extended in an Appendix of [130] to cover the complete set $m = -n : n$. A subset of the fields used in this research can be found tabulated in Appendix A.

The result in Equation (3.23) is an expression which can be applied to utilise a single induction geometry, composed of an array of simple building blocks, to produce a range of approximately orthogonal magnetic fields. Section 5.1 considers a variety of possible systems which may adequately solve this design problem. Appendix B.0.3 details testing undertaken to ensure this reproduces known results.

3.4 Basis Set Expansion

It is often useful to express the fields, $[B_x(\vec{r}),\ B_y(\vec{r}),\ B_z(\vec{r})]$, which result from the current distributions produced with Equation (3.23) in terms of a series of basis functions. The result is a set of coefficients describing a particular profile, providing a memory-efficient manner to store descriptions of numerous functions. The most appropriate basis set is the spherical harmonics already introduced in Equation 3.14,

$$\vec{F}(\vec{r}) = \sum_{n=1}^{\infty} \sum_{m=-n}^{n} C_{n,m} \begin{pmatrix} r^n \\ r^{-n-1} \end{pmatrix} Y_{n,m}(\theta, \phi), \tag{3.24}$$

where $\vec{F}(\vec{r})$ is any arbitrary vector function and $C_{n,m}$ are weighting coefficients, - calculated by taking an overlap integral (Equation (3.15)) with a series of harmonics, repeated over each vector component of the field. Alternatively, an inversion method, which minimises a cost-function of the form Equation (3.19), can be used to derive these coefficients. A specific field is then defined by $3 \times \tilde{N}$ coefficients where $\tilde{N}$ is the number of basis functions.

Ideally, this uses an infinite number of harmonics, but in practice truncating at the $n+2^{th}$ degree, where n relates to the degree of the target harmonic, is sufficient. Numerical experiments undertaken during this research suggest this is within 0.01% accuracy. Beyond

efficient storage, this can also minimise the computational strain of calculations. Combining this expansion for $\vec{A}$ with a Laplace expansion of the greens function reduces the integration problem to a series of overlap calculations weighted by coefficients $C_{n,m}$.

A useful generalisation of this concept is to replace the scalar basis functions with their gradients which also form an orthonormal basis,

$$\vec{F}(\vec{r}) = \sum_{n=1}^{\infty} \sum_{m=-n}^{n} C_{n,m} \vec{\nabla} T_n^m(r, \theta, \phi). \tag{3.25}$$

The result is $\tilde{N}$ coefficients which are ideal for the analysis of the $\vec{B}$ results from Equation (3.23). In the perfect case, there is a single non-zero coefficient corresponding to the target harmonic, otherwise, the coefficient distribution provides a simple comparison between individual fields.

3.5 Three Dimensional Current Sources & The Finite Element Method

The eddy-current response in MIT is proportional to the strength of the primary magnetic field. Producing a stronger primary field will therefore enable a greater dynamic range in the detectable responses. In this thesis it is proposed that solid three-dimensional current sources may allow the production of a stronger magnetic field than a comparable wound source with the same footprint.

Assuming an individual coil contains no free charges, Equation (3.1a) reduces to the Laplace equation,

$$\vec{\nabla}^2 V(\vec{r}) = 0, \tag{3.26}$$

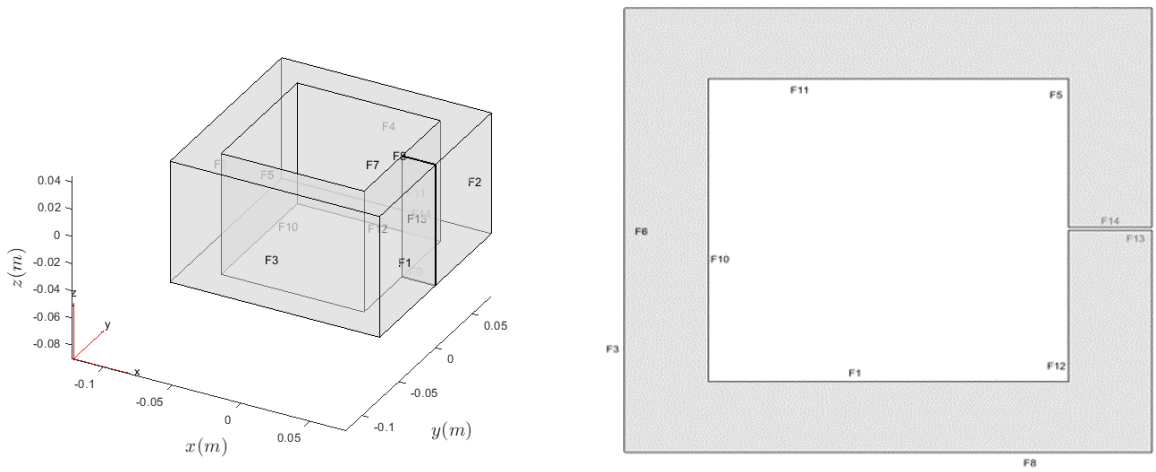
where $V(\vec{r})$ is some electric scalar potential abiding by the Coulomb gauge. The electric field is a standard gradient operation on the potential,

$$\vec{E}(\vec{r}) = -\vec{\nabla}V(\vec{r}). \tag{3.27}$$

Subsequently, the current density can be found,

$$\vec{J}(\vec{r}) = \vec{\sigma}(\vec{r})\vec{E}(\vec{r}), \tag{3.28}$$

where $\vec{\sigma}(\vec{r})$ is the conductivity of the source material. This could be spatially varying and anisotropic, but it is assumed sources are a block of uniform conductivity for simplicity - $\vec{J}(\vec{r}) = \sigma\vec{E}(\vec{r})$. To drive a current through a coil, an infinitesimal cut is made in the geometry to form two driving faces, demonstrated in Figure 3.1 for a square coil.



(a) (b)

Figure 3.1: An example rendering of the *cut coil* geometry.

Any typical source then has two key boundary conditions. Along the two driving faces a known potential difference is applied, resulting in a Dirichlet boundary condition. Elsewhere, the electric field should be parallel to the source-air interface, corresponding to a

zero Neumann boundary condition. Mathematically this states,

$$\text{Dirichlet:} \qquad V(\vec{r_b}) = \Gamma, \qquad \Gamma \in \mathbb{N}, \tag{3.29a}$$

$$\text{Neumann:} \qquad \frac{\partial V(\vec{r_b})}{\partial \vec{n}} = 0, \tag{3.29b}$$

where $(\vec{r_b})$ represents points to which the boundary condition is applied, Γ is some constant voltage difference between faces, and $\partial\vec{n}$ is the normal derivative.

Deriving an expression relating Equation (3.23) and Γ is necessary to evaluate whether a modular array of solid sources are capable of producing the target fields. An approximate way to achieve this is via Ohms law,

$$\Gamma \approx I\Omega_R \approx \frac{Id}{\sigma A_{sec}}, \tag{3.30}$$

where Ω_R is the electrical resistance of a solid source, d is approximated by the path length of an electron flowing due to Γ, and A_{sec} is the area of a driving plate.

A flexible approach to solving Equation (3.26) is the Finite Element Method (FEM) which, aided by CAD design and open-source optimised meshing tools such as GMSH, enables the calculation of the current density inside a generic three-dimensional geometry. The FEM is a well studied numerical method so only a brief introduction will be provided here, with introductory papers found useful by the author in [131, 132, 133, 134] as well as a more detailed textbook in [135]. The implementation here was hand built in MATLAB utilising the aforementioned third-party meshing software and some features of the inbuilt Partial Differential Equation Toolbox[136].

There are four stages: Dividing the geometry into elements, deriving the governing equations for a typical element, assembling all elements in the solution region, and solving the resultant system of matrix equations. Over a particular element, the functional approxi-

mation is,

$$f_\zeta = \sum_i f_\zeta^i N^i, \tag{3.31}$$

where N^i is a set of basis functions at individual nodes, edges or cells depending on the specific formulation.

Having assembled a model suitable for the current density through arbitrary geometries, the governing PDE (Equation (3.5)) is then solved to find the related magnetic field. The FEM can again be applied, which requires a further generalisation to allow two additional degrees of freedom in each calculation, noting the magnetic vector potential is simply separable in Cartesian coordinates to $[A_x A_y A_z]$. Compared to an integrated approach such as in Equation (3.8), FEM also guarantees implicit field continuity, variable boundary conditions, and knowledge of the field in the entire simulation domain.

Boundary conditions must be carefully implemented when formulating any FEM model. Assuming the Coulomb gauge there are four internal boundary conditions: the normal component of $\vec{B}$ should be continuous, the tangential component of $\vec{B}$ will be discontinuous by some surface current density $\vec{K}$, both components of the associated vector potential should be continuous, and the normal derivative of the potential will be discontinuous by $\vec{K}$[137]. Often $\vec{K} = 0$ is true and will be assumed here for simplicity. The edges of the simulation domain are usually treated differently, however, they may be treated as far-field boundaries with a Dirichlet condition using $\Gamma = 0$. This is a physical assumption that the source magnetic field is fully enclosed by the simulation domain - valid as long as the volume is appropriately large for the sources in a specific application. If this assumption is invalid there is likely to be some subtle relationship which will be very application specific, particularly where magnetically shielded environments are involved.

3.6 Traditional MIT

Any magnetic induction tomography based investigation must also consider the related responses from the region of interest due to the time-varying magnetic fields applied by the induction system. This response relates to the interaction between that field and the passive electrical material(s) with which it interacts. A good starting point is a discussion of a traditional approach and its apparent limitations in the context of our system of multiple orthogonal magnetic fields.

As developed so far in this thesis, this begins with a coil model for the excitation field[138, 139]. Using a topological approach individual coils can be modelled as the surface of an open-ended cylinder which carries some tangential current $\vec{J}_s$ whilst acting as a perfect insulator along the axial direction. Multiple coils of this type, often in an annular array, are then contained, either in a perfect electrically conducting shield that zeros the tangential component of the electric field, or one which is sufficiently large to allow far-field boundary conditions.

In phasor form, where all fields and sources are of the form $\vec{E}(\vec{r},t) = \vec{E}(\vec{r})e^{i\omega t}$, Maxwells' equations can be expressed as:

$$\nabla \times \big(\frac{1}{\mu}\nabla \times \vec{E}(\vec{r},t)\big) + i\omega(\sigma(\vec{r}) + i\omega\epsilon(\vec{r}))\vec{E}(\vec{r},t) = -i\omega\vec{J}_s(\vec{r'},t). \tag{3.32}$$

The displacement current can be neglected when $\sigma(\vec{r})$ is sufficiently large compared to $\omega\epsilon(\vec{r})$, combining this with a vector potential formulation where $\vec{E}(\vec{r},t) = -i\omega\vec{A}(\vec{r},t)$ then leads to a governing equation for eddy-responses,

$$\nabla \times \big(\frac{1}{\mu}\nabla \times \vec{A}(\vec{r},t)\big) + i\omega\sigma(\vec{r})\vec{A}(\vec{r},t) = -i\omega\vec{J}_s(\vec{r'},t). \tag{3.33}$$

The remaining component of a forward model in MIT is the sensitivity[140] - the relationship between measured responses and the underlying material which produces them.

Conventionally this is the linearised change in the voltage measured at a detector coil due to a perturbation in conductivity which occurs whilst $\vec{J}_s$ from an excitation coil is held constant.

Solving the governing equation with a FEM, the discrete model of a vector potential inside a specific element, $\vec{A}_\zeta$, can be expanded in terms of a set of basis functions N^i containing an imaginary component[141],

$$\vec{A}_\zeta = \sum_i \vec{A}_\zeta^i N^i. \tag{3.34}$$

Taking an excitation-detection coil pair (t, q), the relationship between the measured value at q and the conductivity of a specific element is[138, 139],

$$\frac{\partial V_{t,q}}{\partial \sigma_\zeta} = -\frac{\omega^2}{I_t I_q} A_\zeta^t \left(\int_{\Omega_\zeta} \sum_i N_\zeta^i \cdot N_\zeta^{i\top} \right) A_\zeta^q, \tag{3.35}$$

where Ω_ζ is the volume of an element. Conveniently, $\left(\int_{\Omega_\zeta} \sum_i N_\zeta^i \cdot N_\zeta^{i\top} \right)$ relates directly to the stiffness matrix commonly generated as part of a FEM formulation.

This formulation does not fit with the proposed use of spherical harmonic magnetic fields produced by a modular system. The explicit coupling of two coils does not provide a clear way to account for a field produced by an amalgamation of sources per Equation (3.23). Pickup coils measure only the perpendicular component of magnetic flux, compared to responses which are inherently anisotropic and therefore require three-axis measurements. Even triaxial workarounds[142] survey the average field over their area, making them poorly suited to inhomogeneous fields[143].

Implicit use of the FEM also introduces further issues. In three dimensions, conforming to complex geometries requires numerous elements. The stiffness matrix is a square matrix of dimensions $[\zeta \times \zeta]$, adding significant computational requirements in an inverse problem, even when limited to matrix methods in a well-suited programming language.

This then scales drastically if generalised to multiple magnetic fields - each composed of a unique $[\zeta \times \zeta]$ matrix. There is no clear approach to combine individual matrices into a coherent inversion problem - augmentation leads to incredibly large matrices, separated inversions necessitate further post-processing. For geophysics, where ζ becomes substantially larger, this quickly becomes unworkable at any realistic resolution.

Formally when considering the entire domain, zero-Dirichlet conditions exist at the boundaries, producing a singular stiffness matrix. Eliminating these adds additional computations consisting of either: a stiff-spring approximation, or a nullspace linear algebra approach[136] to produce an invertible sensitivity map. Formulations which use a separated region of interest[139] still require measurement positions and that region inhabit the same volume, often necessitating the implicit inclusion of conductive coils in that formulation[144].

3.7 Quasi-static Eddy-Currents

To overcome the limitations associated with a traditional approach, this proposal first re-derives the governing equation. Using the phasor convention to solve the time-derivative, Equation (3.1c) becomes,

$$\vec{\nabla} \times \vec{E}(\vec{r}, t) = i\omega \vec{B}(\vec{r}, t). \tag{3.36}$$

Taking the curl of Equation (3.1d) leads to,

$$\vec{\nabla} \times \vec{\nabla} \times \vec{B}(\vec{r}, t) = \mu_r \vec{\nabla} \times \vec{J}(\vec{r'}, t) + \mu_r \vec{\nabla} \times \epsilon(\vec{r}) \frac{\partial \vec{E}(\vec{r}, t)}{\partial t}. \tag{3.37}$$

Assuming a scalar permittivity, further substituting in Equation (3.28) and utilising the trivial solution to the time integral in phasor form leads to,

$$\vec{\nabla} \times \vec{\nabla} \times \vec{B}(\vec{r}, t) = \mu_r(\sigma(\vec{r}) + i\omega\epsilon)\nabla \times \vec{E}. \tag{3.38}$$

Combining Equations (3.36) and (3.1b) and a well-known vector identity this then becomes,

$$-\nabla^2 \vec{B}(\vec{r},t) = i\omega\mu_r \vec{B}(\vec{r},t)(\sigma(\vec{r}) + i\omega\epsilon). \tag{3.39}$$

Following the work introduced in [145] then generalised in [146], assume that $\vec{B}(\vec{r},t)$ is a linear superposition of two components. $\vec{B}_i$ is the impressed field produced by current sources external to the region of interest, $\vec{r}$, which is presumed to have zero curl. $\vec{B}_e$ is the eddy or induced field, arising in response to the time-varying impressed field. This leads to,

$$-\nabla^2 \vec{B}_e(\vec{r},t) = i\omega\mu_r(\sigma(\vec{r}) + i\omega\epsilon)(\vec{B}_i(\vec{r},t) + \vec{B}_e(\vec{r},t)). \tag{3.40}$$

Under a quasi-static assumption, where ω is small, and therefore the displacement current can be neglected, this expression simplifies to the vector-Helmholtz expression presented in [146],

$$-\nabla^2 \vec{B}_e(\vec{r}) = i\omega\mu_r\sigma(\vec{r})(\vec{B}_i(\vec{r}) + \vec{B}_e(\vec{r})). \tag{3.41}$$

Once the induced field is known, the eddy-current density $\vec{J}_e$ is easily calculated,

$$\vec{\nabla} \times \vec{B}_e(\vec{r}) = \mu_0 \vec{J}_e(\vec{r}). \tag{3.42}$$

The quasi-static approximation imposes strict limits on the highest operating frequency of the impressed magnetic field, requiring that the system is significantly smaller than a single wavelength in free space. The proportionality between wavelength and skin depth leads to a simple approximation - the skin depth should be larger than the interrogated conductive objects, allowing edge effects to be neglected, and ensuring $\vec{B}_i$ is not dissipated internally.

FEM formulations are often computationally expensive, especially with multiple impressed fields. Mesh quality is also important, which may require a different mesh for each individual impressed field, ultimately leading to a cumbersome modelling procedure. For rapid prototyping, an integral formulation is often advantageous. Re-expressing the

governing partial differential equation,

$$-\nabla^2 \vec{B}_e(\vec{r}) = k(\vec{r})^2(\vec{B}_i(\vec{r}) + \vec{B}_e(\vec{r})), \tag{3.43}$$

with $k(\vec{r}) = \sqrt{-i\omega\mu_r\sigma(\vec{r})}$. This vector-Helmholtz equation may be solved using the method of Greens' functions in an approach analogous to formulating Equation (3.8). The most appropriate Greens' function will vary with application, with three relevant candidates being[147, 148],

$$\text{Stationary Wave:} \quad \frac{-\cos(k(\vec{r})\ |\vec{r} - \vec{r''}|)}{4\pi|\vec{r} - \vec{r''}|}, \tag{3.44a}$$

$$\text{Incoming Wave:} \quad \frac{-e^{(+ik(\vec{r})|\vec{r}-\vec{r''}|)}}{4\pi|\vec{r} - \vec{r''}|}, \tag{3.44b}$$

$$\text{Outgoing Wave:} \quad \frac{-e^{(-ik(\vec{r})|\vec{r}-\vec{r''}|)}}{4\pi|\vec{r} - \vec{r''}|}, \tag{3.44c}$$

where $\vec{r''}$ are measurement positions in n-dimensional space. For the induced field in a quasi-static limit, the outgoing wave is a natural choice, which leads to the Sommerfield radiation boundary condition - ensuring no energy can come from infinity to the source, which is described mathematically as,

$$\lim_{|\vec{r}\ |\to\infty} |\vec{r}\ |^{\frac{n-1}{2}} \left(\frac{\partial}{\partial|\vec{r}\ |} - k(\vec{r})^2\right)\vec{B}(\vec{r}) = 0, \tag{3.45}$$

with n = 1,2,3 being the spatial dimensions. This assumes that generated fields are finite at their origin and decay to zero at an infinite distance. The resulting expression for the induced field at some measurement positions is,

$$\vec{B}_e(\vec{r''}) = -\frac{1}{4\pi}\int\int\int \frac{\vec{B}_i(\vec{r})k(\vec{r})^2 e^{-ik(\vec{r})|\vec{r}-\vec{r''}|}}{|\vec{r} - \vec{r''}|} d\vec{r}^{\,3}, \tag{3.46}$$

which, to the authors' knowledge, has not been applied to MIT. This full-phasor model allows any impressed magnetic field to be related to the subsequent vector response generated due to an arbitrary structure of material(s) with passive-electric properties.

A useful approximation may be to assume the spatially varying passive electrical properties, conductivity in this case, behave as a piece-wise linear constant:

$$k(r) \rightarrow \sum_{\zeta} k_{\zeta}, \tag{3.47}$$

which is valid under the assumption that a region is divided into small enough discrete segments. This results in,

$$\vec{B_e}(\vec{r''}) = -\sum_{\zeta} \frac{k_{\zeta}^2}{4\pi} \int_{\Omega_{\zeta}} \frac{\vec{B_i}(\vec{r}) e^{-ik_{\zeta}|\vec{r}-\vec{r''}|}}{|\vec{r}-\vec{r''}|} d\vec{r}^{\,3}. \tag{3.48}$$

Numerous numerical approaches exist to solve this integral equation, discussed in more detail in Appendix C, with Gaussian quadrature providing an obvious high-accuracy approach. The sensitivity, $\frac{\partial B_e(\vec{r''})}{\partial k_{\zeta}}$, is then most easily found through a perturbation approach. Experimentally this involves sweeping a small test object over the imaging region, then measuring $\vec{B_e}$ at each location[149].Numerically, this translates to successive small perturbations in each element, k_{ζ}, the size of which determines the model precision. However, for any practical number of elements, the computational requirements for calculations become extreme[149].

The key result here is a rectangular sensitivity matrix with dimensions of $[\zeta \times 3\ n_{r''}]$, where $3\ n_{r''}$ is the number of measurements taken with a three-axis magnetometer. Typically $\zeta \gg 3\ n_{r''}$, making this matrix significantly less computationally expensive to invert than a conventional square matrix. When applying multiple magnetic fields, augmentation can be utilised, combining all information into a single coherent matrix, requiring one inversion and producing one image from all measurements.

This formulation, in principle, allows an arbitrary set of impressed magnetic fields from an indiscriminate induction system to be related to the subsequent response measured at a point-like sensor due to any particular structure of material(s) with some conductive

properties. The resulting sensitivity can then be calculated coherently, representing all impressed fields in a compact system of equations ready for inversion.

3.8 Approximate Formulations for Prototyping - Vector Potential

Both of the forward modelling and sensitivity formulations discussed are computationally expensive to apply. The latter has significant computational costs to find the sensitivity matrix, whilst the former results in a significantly larger square matrix which is costly to invert. These issues are magnified when including multiple magnetic fields, particularly during system prototyping.

An alternative formulation takes a similar approach to traditional MIT, beginning with Amperes' law, Equation (3.1d). Under the quasi-static approximation with displacement current neglected and a homogeneous permittivity,

$$\vec{\nabla} \times \vec{B}(\vec{r}) = \mu_r \vec{J}_{Total}. \tag{3.49}$$

This contains a linear combination of current density terms,

$$\vec{J} = \vec{J}_s + \vec{J}_e + \vec{J}_v, \tag{3.50}$$

where $\vec{J}_v$ are current densities related to the motion of a conductive object in the impressed field. For a stationary system applying Equation (3.28) leads to,

$$\vec{\nabla} \times \vec{B}(\vec{r}) = \mu_r \vec{J}(\vec{r'}) + \mu_r \sigma(\vec{r}) \vec{E}(\vec{r}). \tag{3.51}$$

Substituting in the vector potential from Equation (3.2) combined with the Coulomb gauge and a vector identity,

$$-\nabla^2 \vec{A}(\vec{r}) = \mu_r \vec{J}(\vec{r'}) + \mu_r \sigma(\vec{r}) \frac{\partial \vec{A}(\vec{r})}{\partial t}. \tag{3.52}$$

The gauge choice here is only valid where the induced magnetic field is not dissipated inside an object due to the skin effect, else edge effects begin to dominate the system, and the electric potential can no longer be neglected. Applying the phasor convention to solve for the time-dependence,

$$-\nabla^2 \vec{A}(\vec{r}) = \mu_r \vec{J}(\vec{r'}) + i\omega\mu_r \sigma(\vec{r}) \vec{A}(\vec{r}). \tag{3.53}$$

Using a series expansion the magnetic vector potential can be split into an in-phase and out-of-phase component, vectors $\vec{a}$ and $i\vec{b}$ respectively. Leading to two coupled governing equations,

$$\text{Real:} \quad -\nabla^2 \vec{a}(\vec{r}) = \mu_r \vec{J}(\vec{r'}) - \omega\mu_r \sigma(\vec{r}) \vec{b}(\vec{r}), \tag{3.54a}$$

$$\text{Imaginary:} \quad -\nabla^2 \vec{b}(\vec{r}) = \omega\mu_r \sigma(\vec{r}) \vec{a}(\vec{r}). \tag{3.54b}$$

In the limit $\vec{a} \gg \vec{b}$ Equation (3.54a) becomes identical to Equation (3.5) but in terms of the in-phase component,

$$\nabla^2 \vec{a}(\vec{r}) = -\mu_r \vec{J}(\vec{r'}), \tag{3.55}$$

which can be solved using a standard Greens' function solution, $\vec{a}(\vec{r}) = \frac{\mu_r}{4\pi} \int \frac{\vec{J}(\vec{r'})}{|\vec{r}-\vec{r'}|} d^3r'$. This regime specifies then that the in-phase component is related directly to the applied magnetic field, whilst the out-of-phase component is the sole component dependent on the passive electric properties of the region of interest, in this case, conductivity.

Taking a Greens function solution to the out-of-phase component,

$$\vec{b}(\vec{r''}) = \frac{\mu_r \omega}{4\pi} \int\int\int \frac{\sigma(\vec{r})\ \vec{a}(\vec{r})}{|\ \vec{r''} - \vec{r}\ |} d^3r, \tag{3.56}$$

where $\vec{b}(\vec{r''})$ is then the out-of-phase response measured at some detection points $\vec{r''}$.

Explicitly, this derivation relies on $\mu_r\omega\sigma(\vec{r})$ being suitably small. In simple geometries this depends on the skin depth being small enough that the magnetic field fully penetrates the object. Under the influence of multiple quasi-orthogonal magnetic fields combined with complex conductive structures, the skin depth becomes poorly defined, necessitating a more rigorous definition of the applicability of this model. Equation (3.55) requires explicitly that,

$$|\vec{\nabla}^2\vec{a}| \;\geq\; |\mu_r\omega\sigma\vec{b}|, \tag{3.57}$$

from which a sufficiently small $\mu_r\omega\sigma(\vec{r})$ can be estimated for a specific magnetic field. Appendix D provides some information on the limitations of this model by investigating two simple systems with Equation 3.41.

To formulate a sensitivity relationship in terms of conductivity, again assume it behaves as a piece-wise linear constant, hence,

$$\vec{b}(\vec{r''}) = \frac{\mu_r\omega}{4\pi}\sum_{\zeta}\sigma_\zeta\int_{\Omega_\zeta}\frac{\vec{a}(\vec{r})}{|\,\vec{r''}-\vec{r}\,|}d^3r, \tag{3.58}$$

where Ω_ζ is again the volume of a specific element within a discrete model. Differentiating this expression with respect to conductivity then describes the sensitivity of the out-of-phase response, at some measurement positions $\vec{r''}$, due to some external impressed magnetic field $\vec{a}(\vec{r})$.

$$\frac{\partial\vec{b}(\vec{r''})}{\partial\sigma_\zeta} = \frac{\mu_r\omega}{4\pi}\int_{\Omega_\zeta}\frac{\vec{a}(\vec{r})}{|\,\vec{r''}-\vec{r}\,|}d^3r. \tag{3.59}$$

The result is another sensitivity formulation which implicitly leads to a sensitivity matrix of dimensions $[\zeta \times 3\ n_{r''}]$ for three-axis measurements. Compared to the perturbation approach this fully-integration dependent model also requires significantly less computations to construct, whilst depending only on a specified discrete volume, the input field(s) and a choice of measurement positions. This means a single calculation can be used for a range of conductive geometries and magnitudes, provided the required approximations

are met - perfect for system prototyping.

In this framework the forward model is a simple linear matrix equation:

$$\vec{b}_\kappa = \sum_\zeta \sigma_\zeta \vec{M}^A_{\zeta\kappa}, \tag{3.60}$$

where $\vec{M}^A_{\zeta\kappa}$ is the sensitivity, $\frac{\partial \vec{b}(\vec{r''})}{\partial \sigma_\zeta}$, for elements ζ and measurements κ.

Equation (3.58) and the subsequent sensitivity formulation can also make use of the previously introduced series expansion. This leads to,

$$\vec{b}(\vec{r''}) = \frac{\mu_r \omega}{4\pi} \sum_{n=0}^{\infty} \sum_{m=-n}^{n} C_{nm} \sum_\zeta \sigma_\zeta \int_{\Omega_\zeta} \frac{T_n^m(\vec{r})}{\mid \vec{r''} - \vec{r} \mid} d^3r. \tag{3.61}$$

Taking a harmonic expansion of the Greens' function then simplifies the computation of the integral,

$$\vec{b}(\vec{r''}) = \mu_r \omega \sum_{n'=0}^{\infty} \sum_{m'=-n'}^{n'} \sum_{n=0}^{\infty} \sum_{m=-n}^{n} C_{nm} \sum_\zeta \sigma_\zeta \frac{T_{n'}^{m'}(\vec{r''})}{2n'+1} \int_{\Omega_\zeta} T_n^m(\vec{r}) T_{n'}^{m'}(\vec{r}) d^3r. \tag{3.62}$$

In a real-world experiment, the applicability of this and the upcoming model is restricted by numerous factors, including: source current density, source frequency, detector sensitivity, and the dynamic range of detection. At a small scale, safety concerns limit the strength of source currents. There may be a trade-off between a magnetic field which completely penetrates the object and one which generates a detectable response. Modern optically pumped magnetometers will significantly increase the applicable range of this model but may require appropriate magnetic shielding to utilise.

3.9 Approximate Formulations for Prototyping - Magnetic Field

For a real-world implementation of MIT, there is an obvious flaw in the preceding formulation because measurements are made of the magnetic field, not a vector potential. In a geophysical problem, there are further issues associated with reconciling the predicted data and a conventional approach utilising the apparent conductivity, Equation (2.3). Two separate approaches can derive a related model in terms of the magnetic field. Using the definition of a magnetic vector potential,

$$\vec{H}_s(\vec{r''}) = \vec{\nabla} \times \vec{b}(\vec{r''}), \tag{3.63}$$

recalling $\vec{H}_s(\vec{r''})$ is the out-of-phase or quadrature component of magnetic flux as measured at a detector. Substituting in Equation (3.58), applying the vector identity

$$\vec{\nabla} \times \left(f(r)\vec{G}(r)\right) = f(r)\vec{\nabla} \times \vec{G}(r) - \vec{G}(r) \times \vec{\nabla} f(r),$$

and noting the coordinate dependence of the ∇ operator leads to,

$$\vec{H}_s(\vec{r''}) = \frac{\omega}{4\pi} \sum_{\zeta} \sigma_{\zeta} \int_{\Omega_{\zeta}} \frac{\vec{a}(\vec{r}) \times (\vec{r''} - \vec{r})}{\mid \vec{r''} - \vec{r} \mid^3} d^3r. \tag{3.64}$$

The sensitivity formulation in terms of magnetic flux is then simply,

$$\frac{\partial \vec{H}_s(\vec{r''})}{\partial \sigma_{\zeta}} = \vec{M}_{\zeta\kappa} = \frac{\omega}{4\pi} \int_{\Omega_{\zeta}} \frac{\vec{a}(\vec{r}) \times (\vec{r''} - \vec{r})}{\mid \vec{r''} - \vec{r} \mid^3} d^3r. \tag{3.65}$$

This is the formulation used for investigatory MIT inversion problems later in this thesis. The same expression, but with an explicit i term denoting the out-of-phase nature of this response, can be found in a geophysical framework by following the derivation in [150], generalised to a generic source magnetic field and an unspecified set of detector positions. The primary field can be found through Equation (3.8) where $\vec{r} \rightarrow \vec{r''}$, which leads naturally to calculations of the apparent conductivity.

The layered nature of this model, coupled with the highly ill-posed related inversion problems, necessitates a need to minimise the number of computations whilst maximising the related numerical accuracy, particularly for synthetic studies. To understand why, consider the components of a complete model for spherical harmonic fields in magnetic induction tomography.

An induction system is constructed of several arbitrarily shaped current sources, each requiring many elemental segments to calculate the resulting magnetic field. These generate numerous independent, approximately orthogonal magnetic fields, which interact with a discrete region composed of many three-dimensional voxels to produce some response, measured at an array of positions. A secondary calculation constructs an inversion map - usually the sensitivity calculation from the forward model with fewer voxels to enable image reconstruction. Extensions to geophysics then require repeated modelling at numerous source locations while considering dramatic increases in the modelling volume, multiplying the number of computations.

The approach used in Equation (3.62) could be applied here to solve Equation (3.65). The result is: a superposition of analytic composites to find $\vec{a}(\vec{r})$ (Appendix B), followed by a solution to Equation (3.24), then a nested summation combined with a solution to the radial integration up to the $n/n'^{\text{ th}}$ degree. In practice taking the expressions for $\vec{a}(\vec{r})$, then solving the integral directly with Gaussian Quadrature (Appendix C) is often more efficient. Multi-dimensional arrays and matrix mathematics can be further manipulated for computational efficiency.

In summary, this chapter has covered the key equations imperative to modelling an MIT problem. It provides a novel algorithm designed to produce an arbitrary magnetic field given a system of coils and a generation volume by varying the current through individually controlled elements. It extends the existing approaches to model the responses,

and sensitivity to inversion, to allow the utilisation of a system of inhomogeneous fields. Greens function solutions to the full-phasor system, and a simplified approximate regime are given, intending to minimise the computational expense of both the forward and inverse models.

Chapter 4

Inverse Theory & Image Reconstruction

The overarching aim of this chapter is to present the complexity of ill-posed inversion problems intuitively in the context of MIT, including numerous tools to analyse the system, select an appropriate algorithm, find a solution, and quantitatively assess that solution. It also intends to highlight a frequent disparity between theory and practical reality, where apparently invalid algorithms have proven to be incredibly effective in subsequent work.

Considerations include a library of appropriate algorithms ranging from those suited to simple smooth problems, increasing the complexity, to those with sharp discontinuities. Each imposes a particular subspace to imply some desired property of the solution, and as a result, there is no 'best' choice. Each also obtains a balance between speed, quality, and stability of the solving process, with this work identifying optimisable parameters able to be adjusted to achieve a desired outcome. The discussion covers methods to select the regularisation parameter and, subsequently, an appropriate quantitative solution without manual intervention given the infinite number of permutations and semi-convergent nature of many of these approaches. Individual algorithms come from other research fields, up to some state-of-the-art solutions, with appropriate modifications where necessary. Most are wholly novel to MIT and typically exhibit significant computational advantages

to the widely used Tikhonov inversion, removing computational bounds concerning the maximum resolution attainable and the number of measurements involved.

The inverse problem in MIT works to find the causal passive electric responses which produce a set of eddy-magnetic observations. Formulations like Equation (3.60) lead to a linear inverse problem,

$$\sigma_\zeta = M_{\zeta\kappa}^{-1} b_\kappa, \tag{4.1}$$

where the lack of vector notation in $M_{\zeta\kappa}$ identifies this may refer to a problem formulated in terms of the magnitude of the magnetic field. A trivial solution to this problem does not exist, for there are infinite conductivity profiles which could generate any given b_κ. This is often referred to as an ill-posed problem in the Hadamard sense[151], which is true for most inverse problems in Electromagnetism.

Regardless of formulation, the sensitivity matrix $M_{\zeta\kappa}$ is also ill-conditioned, suggesting individual measurements (columns) are nearly-linearly independent[152]. Any tiny perturbation in the observations leads to a colossal change in the estimated causal properties. Deriving a related matrix which is well conditioned also fails to guarantee an adequate solution[152]. Further imposing additional constraints like confining the solution norm by some parameter gives rise to a non-linear dependence on the parameter, so the parameter choice is obscure even for an ideal problem[152].

From a purely mathematical perspective, the inverse of any matrix is also only well defined for a square, or well-determined set of equations ($\kappa = \zeta$). Where this isn't true a standard approach is to minimise a least squares cost function, like Equation (3.20),

$$\min_{\sigma_\zeta} ||M_{\zeta\kappa}\sigma_\zeta - b_\kappa||_2. \tag{4.2}$$

The solution norm is then, the value of this function for a particular solution, σ_ζ^{sol}.

4.1 Singular Value Decomposition

The Singular Value Decomposition (SVD) is a standard tool which provides some useful insights into the consequences of ill-conditioning and consists of a decomposition of $M_{\zeta\kappa}$[153],

$$M_{\zeta\kappa} = \mathbf{U}\ \eta\ \mathbf{V}^\top = \sum_i \mathbf{u}_i\ \eta_i\ \mathbf{v}_i^\top \qquad i = 1, ..., \min(\zeta, \kappa), \tag{4.3}$$

where $\mathbf{U} \in \mathbb{R}^{\zeta\times\zeta}$ and $\mathbf{V} \in \mathbb{R}^{\kappa\times\kappa}$ are matrices with orthonormal columns such, $\mathbf{U}\mathbf{U}^\top = \mathbf{V}\mathbf{V}^\top = \mathbf{I}$, with the columns known as the left-singular and right-singular vectors of the sensitivity matrix respectively, while $\eta \in \mathbb{R}^{\zeta\times\kappa}$ is a diagonal matrix with non-negative decreasing valued entries known as the singular values. For a system of κ singular values, the ratio of η_1/η_κ is known as the condition number of the sensitivity matrix[154, 153]. From the SVD, consider the following quantities[152, 155],

$$\left.\begin{aligned} M\mathbf{v}_i &= \eta_i u_i \\ ||M\mathbf{v}_i||_2 &= \eta_i \end{aligned}\right\} \qquad i = 1, ..., \min(\zeta, \kappa) \tag{4.4}$$

A relatively small η_i specifies that there exists a linear combination of the columns of $M_{\zeta\kappa}$ where $||M\mathbf{v}_i||_2 = \eta_i$ is also small. As $\eta_i \to 0$, the associated singular vectors therefore become increasingly oscillatory[152]. The forward problem becomes,

$$b_\kappa = M_{\zeta\kappa}\sigma_\zeta^{sol} = \sum_i \eta_\kappa (V_{\zeta\kappa}^\top\ \sigma_\zeta^{sol})\ U_{\zeta\kappa}. \tag{4.5}$$

The multiplication by η leads to a damping effect on those oscillatory singular vectors. This is reversed in the inverse problem, amplifying any small errors in b_κ so that they dominate the final solution[152]. A visual representation of this amplification is in Figure 4.1 for a linear inversion of a specific MIT problem. The solution (right) is noise corrupt with highly oscillatory variations in magnitude despite being an accurate least-squares solution.

An additional insight gained from applying the SVD is that it demonstrates that Equa-

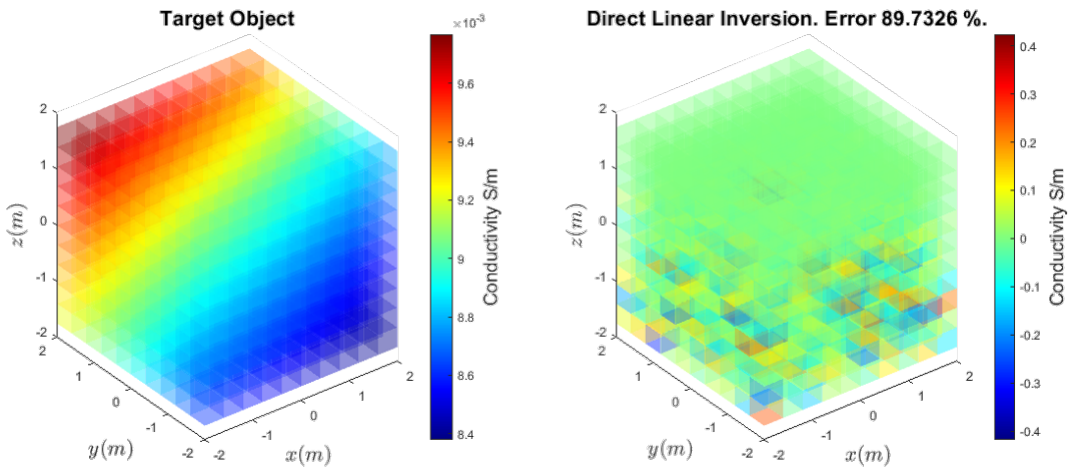


Figure 4.1: A perfect reconstruction (left) compared to a direct linear inverse solution (right) which suffers from noise amplification. A subsequently defined induction system, Figure 5.35, producing nineteen magnetic fields, is coupled with twenty-five point-like sensors detecting the three axial components of the magnetic response from Equation 3.48 giving $\kappa = 1425$. Combined with Gaussian white noise at a level equivalent to the noise floor of a Lakeshore Telsameter, the target image is a $[10 \times 10 \times 10]$ grid of cubic voxels, such $\zeta = 1000$ in the inversion.

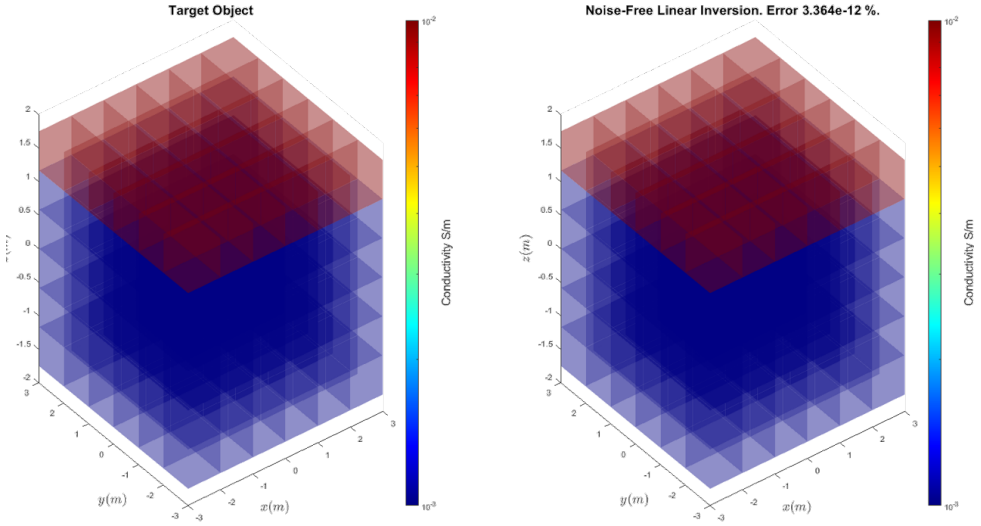
tion (4.1) and the "generalised" least-squares cost function result in identical solutions[152],

$$\sigma^{LSQ} = \sum_{\kappa} \frac{U_{\zeta\kappa}^{\top} b_{\kappa}}{\eta_{\kappa}} V_{\zeta\kappa}. \tag{4.6}$$

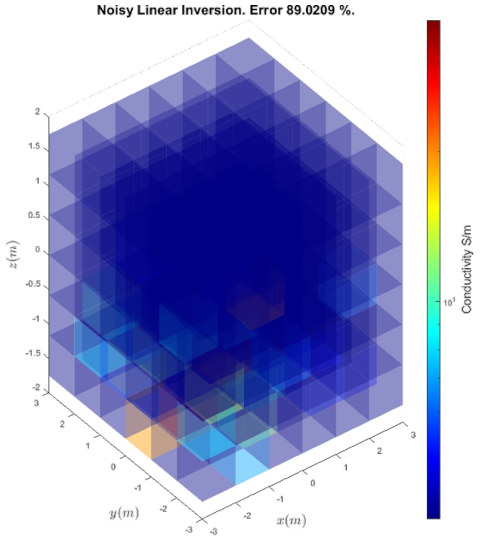
For synthetic studies like those considered here, another key concept is the inverse crime[156]. Within the confines of a particular model, one can produce a set of perfect, noise-free observations. When inverted these may lead to flawless solutions which are in fact physically unattainable. This links directly to the ill-conditioning of the problem.

Some examples for three-dimensional MIT inversions are provided in Figures 4.2 and 4.3, showing this is relatively ubiquitous and can allow logically impossible inverse solutions to be found. Both of these demonstrations use the induction system in Figure 5.35 to produce signals.

The inverse crime is easily overcome, but it is important to do so. In many cases simply varying the discretisation between the forward and inverse models can add a large enough



(a) A flawless image can be produced by identical forward/inverse models.



(b) With Gaussian noise at 0.1% of the signal magnitude.

Figure 4.2: A demonstration of the inverse crime in a simple, synthetic, three-dimensional MIT problem which varies smoothly from 10^{-2} to 10^{-3}S/m. Though easily overcome, there is clear potential to get carried away with seemingly perfect results which can never be practically reproduced. The quoted error values use the MULTISSIM3 metric[157].

perturbation. Best practice is to also include some additive noise with the observations which are to be inverted,

$$b_{\kappa}^{invert} = b_{\kappa}^{synthetic}\mathbf{e}_{\kappa}, \tag{4.7}$$

where $\mathbf{e}$ is some measure of noise - typically a white or Gaussian distribution.

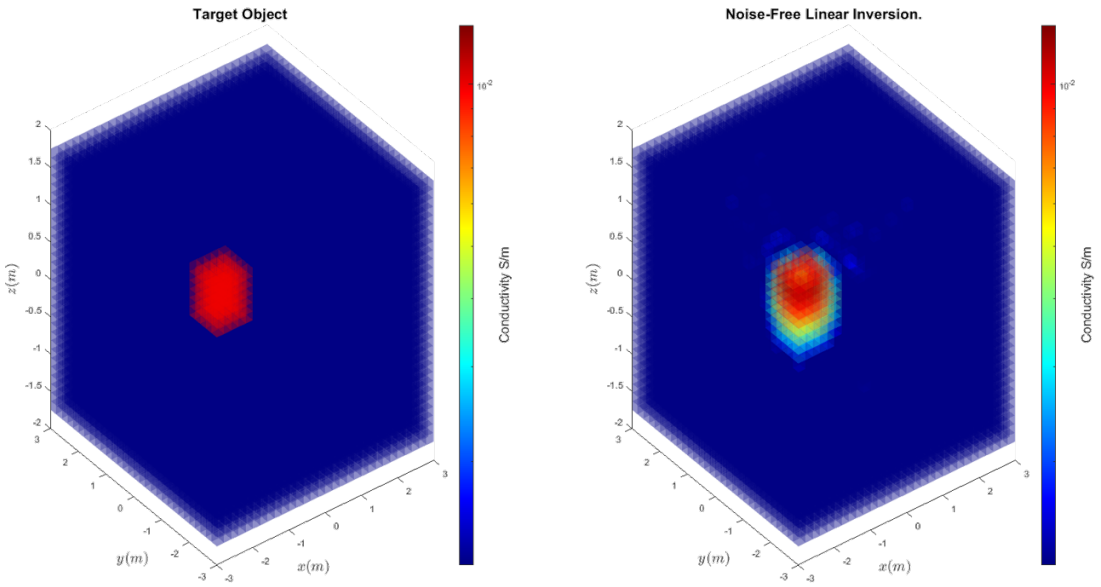


Figure 4.3: The inverse crime is ubiquitous and enables logically impossible problems, here $\zeta = 1,000,000$ and $\kappa = 4$. Though there is some blurring, the object and magnitudes are reproduced definitively, as is an overwhelming majority of the background.

4.2 An Introduction to Regularised Inverse Methods

The discussion demonstrates that a trivial solution using standard linear algebra is unsuitable. Regularisation and associated methods provide a more sophisticated approach allowing information about the desired solution to be incorporated, which aims to extract a stable, favourable choice. This also means the initial set of linear equations is no longer being solved, instead, a balance between matching the set of linear equations and any additional constraints is sought[152].

The best known approach is Tikhonov regularisation[158] which was briefly introduced for Equations (3.21) and (3.23). This selects a solution by considering the addition of a side constraint,

$$\sigma_{\zeta}^{\lambda} = argmin\left\{ \left|\left| M_{\zeta\kappa}\sigma_{\zeta} - b_{\kappa} \right|\right|_2^2 + \lambda^2 \left|\left| \mathbf{L}(\sigma_{\zeta} - \sigma_{\zeta}^{*}) \right|\right|_2^2 \right\}, \tag{4.8}$$

where λ is known as the regularisation parameter, $\mathbf{L}$ is the regularisation or Tikhonov matrix, and σ_{ζ}^{*} is some initial guess of the solution. λ controls the relative importance of the side constraint relative to the least-squares solution to the system of equations. A

more familiar expression arises by solving this with matrix mathematics,

$$\sigma_{\zeta}^{\lambda} = \left(M_{\zeta\kappa}^{\top} M_{\zeta\kappa} + \lambda \mathbf{L}^{\top}\mathbf{L}\right)^{-1}(M_{\zeta\kappa} b_{\kappa} + \lambda \mathbf{L}^{\top}\mathbf{L}\sigma_{\zeta}^{*}), \tag{4.9}$$

which simplifies to the form seen in Equation (3.21) when the initial guess is zero. The regularisation matrix in a Tikhonov approach specifies the form of additional constraints - the desired properties of a subsequent solution. The most common form has already been encountered in this thesis, $\mathbf{L} = \mathbf{I}$, known as L_2 regularisation, giving preference to solutions with smaller 2-norms[159]. For an image reconstruction problem such as this, this regularisation implies some smoothness to the final solution, which may lead to blurring or bleeding around regions of variable conductivity.

This matrix has additional uses beyond acting as a direct side constraint in the problem. In iterative and truncated inverse methods, additions such as preconditioning, prior-conditioning, and enrichment can benefit both convergence speed and the quality of the final solution. Some implementations may be valueless in a Tikhonov inversion, but highly successful in an iterative scheme or visa-versa. Regardless of the application, a regularisation matrix imparts additional information about the expected solution into the inverse procedure, restricting the available parameter space to select a more desirable solution.

Consequently the success of any regularised inverse procedure will always be very problem specific. Even in MIT a conventional problem considering fault detection in metals could be well suited to an edge-detection algorithm, whereas a large-scale conductivity profile encountered in a geophysical survey may vary in a manner better suited to a smoothness constraint. Appendix E provides a more comprehensive discussion of the variety of forms this matrix may take.

It can be shown using the SVD that the purpose of a regularisation method is to dampen, adjust, or remove the contributions from the smallest singular values[152]. This requires

some filter factors, f_i, be introduced to the least squares solution,

$$\sigma_{\zeta}^{Reg} = \sum_{\kappa} f_i \frac{U_{\zeta\kappa}^{\top} b_{\kappa}}{\eta_{\kappa}} V_{\zeta\kappa}, \qquad L = \mathbf{I} \tag{4.10}$$

assuming the regularisation matrix is the identity for simplicity. The choice of filter factors then should decay to zero so the contributions from smaller singular values of η_{κ}, which amplify noise, are removed from the final summation. The difference between effectively any regularisation or inverse approach is how the filter factors are defined and subsequently applied[155].

From Equation (4.8) with $\mathbf{L} = \mathbf{I}$ the filter factors take the form[152],

$$f_i = \frac{\eta_{\kappa}^2}{(\eta_{\kappa}^2 + \lambda^2)}, \tag{4.11}$$

which illustrates the regularisation has no influence for $\lambda > \eta_1$[155].

One major issue associated with Tikhonov inversions is the computational requirements, with just the inversion step of a $[\zeta \times \zeta]$ matrix scaling as $\mathcal{O}(\zeta^3)$[160].

4.3 Truncated Inverse Methods

The SVD of the system of equations and related filter factors leads directly to a more intuitive approach to regularisation. Instead of adding additional assumptions to the model, high-frequency singular values could be ignored completely. This is known as a truncated method and a number of variations exist. The simplest is perhaps the Truncated SVD (TSVD) using filter factors of the form[161],

$$f_i(\eta_{\kappa}) = \begin{cases} 1, & \text{if } \kappa < \lambda, \\ 0, & \text{if } \kappa > \lambda, \end{cases} \tag{4.12}$$

and so λ acts as a threshold[162]. This is also the only regularised solution which has no null-space component which overlaps with the sensitivity matrix, and thus is the only method which does not impart some a priori information through the regularisation matrix restricting the available parameter space[161].

The General Singular Value Decomposition (GSVD) is an extension of the SVD and considers a matrix decomposition of the matrix pair $[\mathbf{M}\ \mathbf{L}]$[163, 164, 165]. For a matrix $\mathbf{L} \in \mathbb{R}^{p\times\kappa}$ in a general ill-posed problem where $\eta \geq \kappa \geq p$ is true, this takes the form

$$M_{\zeta\kappa} = \widetilde{U} \begin{pmatrix} \widetilde{\Sigma} & 0 \\ 0 & I_{\kappa-p} \end{pmatrix} X^{-1}, \quad \mathbf{L} = \widetilde{V}(\widetilde{A}\ ,\ 0)X^{-1}, \tag{4.13}$$

where $\widetilde{U} \in \mathbb{R}^{\eta\times\kappa}$ and $\widetilde{V} \in \mathbb{R}^{p\times p}$ are orthonormal matrices similar to those in the SVD, $X \in \mathbb{R}^{\kappa\times\kappa}$ is a non-singular matrix and $\widetilde{\Sigma}$ and $\widetilde{A}$ are $p \times p$ diagonal matrices, with ordered non-negative entries. The ratio of the values in $\widetilde{\Sigma}$ and $\widetilde{A}$ are then used to define the generalised singular values of the decomposed matrix pair[163, 164, 165].

The final $\kappa - p$ columns in X capture the basis vectors contained in the null-space of a chosen regularisation matrix $\mathbf{L}$. An important characteristic for any $\mathbf{L}$ matrix is that known features of the desired solution can be represented by vectors in the matrix null-space and doing so ensures they are not damped in the regularised reconstruction[166]. Hence X provides insight into the regularising properties of a specific matrix and the type of solution it gives preference to.

As with the TSVD a threshold can then be applied to produce a regularised reconstruction known as the Truncated General Singular Value Decomposition (TGSVD)[165]. This approach allows for a hybrid between direct truncation of noisy components of the model, and the addition of a priori information from $\mathbf{L}$.

A pair of related approaches are the damped Singular Value Decomposition (dSVD)[167]

and damped General Singular Value Decomposition(dGSVD)[152]. These are more nuanced with the filter factors of 1 and 0 replaced with smoothly decaying filter factors,

$$f_i(\eta_\kappa) = \frac{\eta_\kappa}{\eta_\kappa + \lambda}. \tag{4.14}$$

The key feature is that these decay at a significantly slower rate than those introduced by the related Tikhonov methods. In theory this results in less filtering per truncation step, allowing a more optimal solution to be found by sampling a more finely-tuned parameter space.

A final notable method is Truncated Total Least Squares (TTLS), designed for the situation where neither the sensitivity matrix or data are known precisely - this can often be relevant when modelling and working with data from a real-world experiment[168, 152].

4.4 Regularisation Parameter Selection and The L-Curve

In Tikhonov and truncated approaches to finding an inverse solution, there is a regularisation parameter, λ. The influence of that parameter is highly non-linear, and the choice determines the trade-off between selecting an accurate solution to a given system of equations or one that meets given a priori assumptions.

A convenient and intuitive graphical approach is known as the L-curve, with an early application in [169]. This characteristic curve is a log-log scale plot of the residual norm, $||M_{\zeta\kappa}\sigma_\zeta^{reg} - b_\kappa||_2$, to the regularised (semi-)norm, $||\mathbf{L}\sigma_\zeta^{reg}||_2$, for a range of valid values of λ, Figure 4.4.

When λ is small, most filter factors are close to unity and perturbation errors will dominate the final solution. Conversely, as λ becomes large, the filter factors become

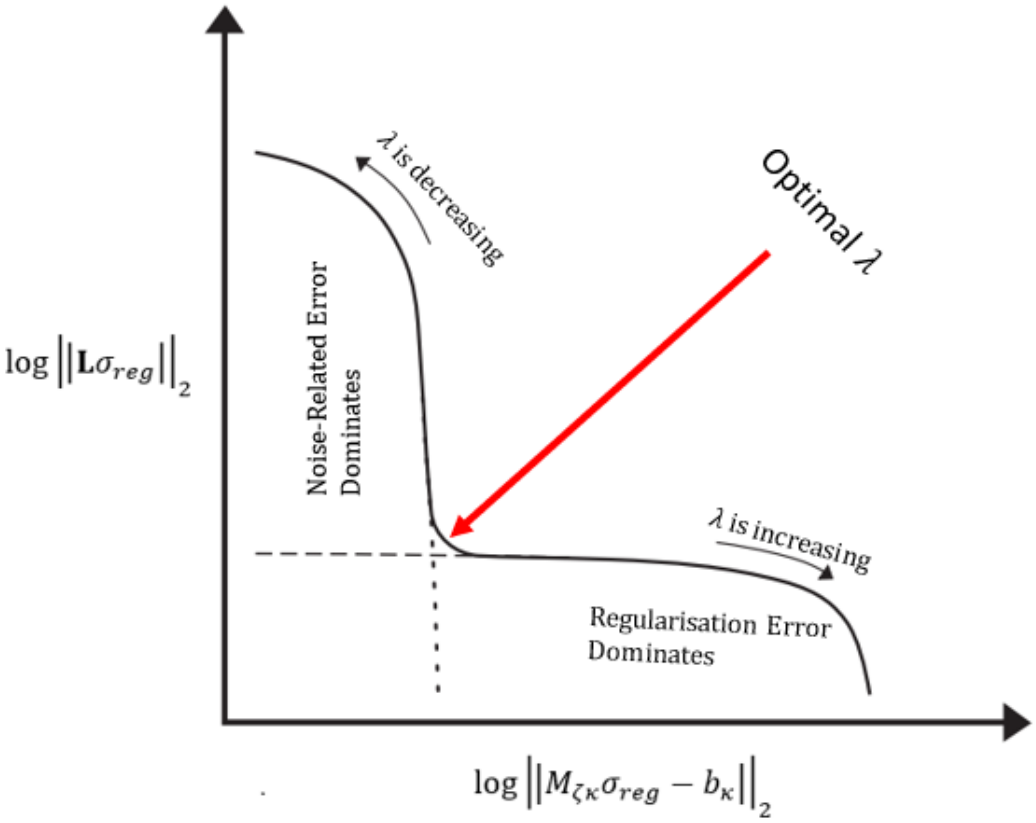


Figure 4.4: The generic shape of an L-curve is given in terms of the MIT modelling parameters used in this thesis. This curve is idealistic and therefore is unlikely to be found in a practical application. Two regions are apparent; one where the perturbation error, or error associated with noise quantities in the model, dominates the final solution; one where the regularisation error, or the smoothness implied by the side constraint imposed on the model dominates, and varying the regularisation parameter has little influence on the solution.

very small, and the regularisation error then dominates the result. An optimal selection, in terms of a balance between the solution to the linear system of equations and the imposed side constraint, is then located at the point of maximal curvature on the log-log plot, as shown for an ideal inverse problem in Figure 4.4[170]. This represents one method of selecting the optimal amount of regularisation for an unknown problem, though the exact form varies between Tikhonov and Truncated methods. Part of the convenience is that no estimate of the measurement error is required to select a value.

A robust approach, it provides a visual reference which can impart additional insight into a problem. For example, if the shape deviates significantly, this may identify that

Tikhonov methods are unable to solve a specific problem for any λ value.

Depending on the problem, other approaches may be more practical. When both ζ and κ are large, computing a solution at a near continuum of regularisation parameters is computationally expensive. If an estimate of the measurement error can be made, a quick and simple estimation of the parameter comes from the discrepancy principle amounting to choosing λ such,

$$||M_{\zeta\kappa}\sigma_{\zeta}^{Reg} - b_{\kappa}||_2^2 = ||\mathbf{e}||_2, \tag{4.15}$$

is true[152, 162] when $||\mathbf{e}||_2$ is the estimate of the error. The quality of the resulting solution is then implicitly linked to the accuracy of the estimated error. In MIT this may apply when the dominant source of error is associated with a commercial magnetometer which has been locally calibrated to account for any environment variables. A closely related approach is known as the Miller Criterion which includes a statistical estimate of the distribution of the solution[162].

A statistical approach leads to Generalised Cross Validation (GCV). The premise is that a well-chosen solution should be able to accurately predict any single excluded measurement from the vector b_{κ}. A regularisation parameter should then minimise the GCV function,

$$G \equiv \frac{||M_{\zeta\kappa}\sigma_{\zeta}^{Reg} - b_{\kappa}||_2^2}{\left(trace[\mathbf{I}_{\zeta} - M_{\zeta\kappa}M_{\zeta\kappa}^{I}]\right)}, \tag{4.16}$$

where $M_{\zeta\kappa}^{I}$ is a matrix for which $\sigma_{\zeta}^{Reg} = M_{\zeta\kappa}^{I}b_{\kappa}$ is true[152, 162]. This is widely used and, using the SVD and filter factors, the computation of the denominator can be significantly simplified[155]. It is also often used as a stopping rule for iterative inversion algorithms because of its statistical rigour.

All approaches to parameter selection are imperfect. In a hypothetical test case where the sought solution is already known, and a perfect continuum of λ values can be trialled with infinitely precise mathematics, a better value can almost certainly be selected.

4.5 Iterative Regularisation and Image Reconstruction Algorithms

Tikhonov methods have found wide application in historical MIT problems, with the application of TSVD also seen in texts such as [171]. These schemes are advantageous because they are: robust, easily applied, involve minimal algorithmic complexity, and for the limited data available in conventional MIT, often require few computations to produce an image. But those approaches scale cubically with the number of variables, often restricting MIT to low-resolution, two-dimensional images, with as few as 32 measurements[49, 48]. In part, this explains the rarity of geophysical MIT inversions.

Modern image reconstruction algorithms designed for large-scale ill-posed problems form an entire active research field. The algorithmic complexity of these algorithms historically means their application to novel inversion problems like this is unusual. Thankfully a variety of open-source modular packages exist to overcome this, such as [155, 172, 173, 174, 175]. A brief introduction to a subset of the available algorithms follows each included because they have found some application to MIT during this body of research. Numerous permutations exist when considering the various constraints and coupling to regularising matrices like those in Appendix E.

A simple iterative least-squares method is the addition of a non-negative constraint,

$$min\left|\left|M_{\zeta\kappa}\sigma_{\zeta} - b_{\kappa}\right|\right|_2 \;\; subject\ to \;\; \sigma_{\zeta} \geqslant 0. \tag{4.17}$$

This functions by computing a possible solution, adjusting it through increasingly smooth steps, and then stopping at a solution which meets both the constraint and solves the least-squares problem[176].

4.5.1 Conjugate Gradient Algorithms

The conjugate gradient (CG) or conjugate gradient least squares (CGLS) approach can be applied directly to an ill-posed system of equations, with a regularising effect dependent on the number of iterations[177, 152, 155]. λ is then the number of iterations. At the λ^{th} step,

$$min\left|\left|M_{\zeta\kappa}\sigma_{\zeta}^{\lambda} - b_{\kappa}\right|\right|_2 \quad subject\ to \quad \sigma_{\zeta}^{\lambda} \in \mathcal{K}_{\lambda}(M_{\zeta\kappa}^{\top}M_{\zeta\kappa}, M_{\zeta\kappa}^{\top}b_{\kappa}), \tag{4.18}$$

where $\mathcal{K}_{\lambda}$ is the Krylov subspace associated with the λ^{th} step of the algorithm applied to $M_{\zeta\kappa}^{\top}M_{\zeta\kappa}\sigma_{\zeta}^{\lambda} = M_{\zeta\kappa}^{\top}b_{\kappa}$. This functions as a Tikhonov variant with a side constraint.

A Krylov subspace is defined explicitly as,

$$\mathcal{K}_{\lambda}(M_{\zeta\kappa}^{\top}M_{\zeta\kappa}, M_{\zeta\kappa}^{\top}b_{\kappa}) = span\left\{M_{\zeta\kappa}^{\top}b_{\kappa}, M_{\zeta\kappa}^{\top}M_{\zeta\kappa}M_{\zeta\kappa}^{\top}b_{\kappa}, \ldots, (M_{\zeta\kappa}^{\top}M_{\zeta\kappa})^{\lambda-1}M_{\zeta\kappa}^{\top}b_{\kappa}\right\}, \tag{4.19}$$

where span means every vector in the subspace is a linear combination of the constructed basis[155, 178]. The subspace involves a series of vector multiplications, with calculations $\mathcal{O}(\zeta)$ opposed to $\mathcal{O}(\zeta^3)$ in a Tikhonov approach[178, 179]. This calculation efficiency has various consequences, such as, in a limiting case, identical hardware may solve a higher-resolution problem.

The filter factors relate to the Ritz polynomial at a certain iteration,

$$P_{\lambda}(\eta) = \prod_{j=1}^{\lambda} \frac{\left(\theta_j^{(\lambda)}\right)^2 - \eta^2}{\left(\theta_j^{(\lambda)}\right)^2}. \tag{4.20}$$

where $\left(\theta_j^{(\lambda)}\right)^2$ are known as the Ritz values and η are the non-zero singular values[152]. Ritz values specify the λ eigenvalues of $M_{\zeta\kappa}^{\top}M_{\zeta\kappa}$ restricted to the specific Krylov subspace, $\mathcal{K}_{\lambda}(M_{\zeta\kappa}^{\top}M_{\zeta\kappa}, M_{\zeta\kappa}^{\top}b_{\kappa})$. The filter factors are then[152],

$$f_j^{(\lambda)} = 1 - P_{\lambda}(\eta_j), \qquad j = 1, \ldots, \lambda. \tag{4.21}$$

A useful variant is Brakhage's ν-method[180]. Analogous to dSVD, it is less sensitive to the iteration number and can theoretically select a more optimal solution. It considers two problem-independent update vectors which rely on λ and some specified value $0 < \nu < 1$. The filter factors can be expressed as Jacobi polynomials, and the integration of preconditioning matrices is possible[155].

Theoretically, CG methods are implicitly derived to solve systems of equations with a symmetric positive definite sensitivity matrix only. This suggests that the dense, rectangular, augmented matrices proposed as part of this investigation into multiple orthogonal source fields are not well suited to CG solutions. Despite this, some data demonstrating their practical superiority over Tikhonov methods is in Chapter 6.

From the theoretical note, various generalisations principally better suited to the problem exist. The bi-CG algorithm is a simple extension where both the original linear system and the conjugate system, $b^*_\kappa = \sum_\zeta \sigma^*_\zeta M^\top_{\zeta\kappa}$, are solved for at each step, with two residual vectors per iteration[181]. A further generalisation is the bi-CG stabilised method which introduces a secondary ill-posed inversion algorithm to perform repair steps. Repair steps attempt to improve solution quality, and the natural form considers the use of first-order polynomials[182], with variants such as bi-conjugate gradients stabilised (l) extending this to higher orders.

4.5.2 Algebraic, Column-Action & Simultaneous Iterative Algorithms

Algebraic Reconstruction Techniques (ART) typically proceed as row-action methods, acting upon each measurement sequentially during an individual iteration. Column-Action Reconstruction Technique (CART) methods are similar, acting instead on each voxel in sequence. Simultaneous Iterative Reconstruction Techniques (SIRT) meanwhile solve all equations at the same time during a particular iteration, analogous to Tikhonov methods. MATLAB implementations of all the algorithms come from [174].

Any ART method is differentiated only by the sequential process order of individual measurements. Each step updates the entire solution for the $\kappa = n^{th}$ row of the transpose of the sensitivity matrix as[173],

$$\sigma \leftarrow \sigma + \nu_\lambda \frac{b_n - M_n^\top \sigma}{||M_n^\top||_2^2 + \alpha_\lambda} M_n^\top, \tag{4.22}$$

where ν_λ is the relaxation parameter and α_λ is an additional damping parameter. ν_λ is the primary convergence parameter, taking a value in the range [0,2] which may be fixed or vary with λ. The best known is the Kaczmarz method, where ν_λ is fixed and each iteration consists of a ordered pass through all measurements, $n = 1, 2, \ldots, \kappa$[173, 183, 184].

The Randomised Kaczmarz method[173, 185] attempts to speed convergence by randomising the order of n, something well suited to large-sparse over-determined matrices. The symmetric Kaczmarz method consists of a double sweep[173, 186], with one forward and the other inverse, $n = 1, 2, \ldots, \kappa - 1, \kappa, \kappa - 1, \ldots, 2, 1$. An iteration-dependent relaxation parameter ($\nu_\lambda \to 0$ as $\lambda \to \infty$) forces the algorithm to converge to the weighted-least squares solution[187] as $\lambda \approx \infty$.

Side constraints, such as upper and lower inequality bounds on σ, can also be included. Non-negativity is of particular relevance given $\sigma \geq 0$ is true in nearly all settings[74]. An upper bound can also accelerate convergence and improve the final image quality, with appropriate choices set by either the problem at hand or physical limits. Permutations including different constraints then formulate at least twelve algorithms, all of which can be further tuned by adjusting ν_λ and α_λ.

A straightforward alternative is CART methods which consider an update vector at the $\zeta \equiv n^{th}$ column, or an individual voxel,

$$\sigma_n \leftarrow \sigma_n + \nu_\lambda \frac{M_n^\top (b - M\sigma)}{||M_n||_2^2 + \alpha_\lambda}. \tag{4.23}$$

The sequence is adjustable, and the various constraints may be integrated, leading to algorithms with dissimilar inversion properties. Theoretically, this algorithm is well suited to our MIT problem. Convergence analysis[188] suggests it is best suited to those where individual measurements are independent - guaranteed with perfectly orthogonal magnetic fields. It also approaches the optimal solution faster during early computations, before update size slows, allowing early stopping without majorly affecting image quality compared to other algebraic techniques[189, 188, 190].

SIRT algorithms, in general form, at a particular iteration are given by,

$$\sigma^{\lambda+1} = \sigma^{\lambda} + \nu_{\lambda} T M_{\zeta\kappa}^{\top} D(b_{\kappa} - M_{\zeta\kappa}\sigma_{\zeta}^{\lambda}), \tag{4.24}$$

where $\sigma^{\lambda+1}$ denotes the new iteration vector, σ^{λ} denotes the current iteration vector or initial guess at $\lambda = 0$. Once again, various constrained permutations exist for each algorithm. Individual methods can be differentiated by the choice of $[T, D]$ the simplest of which is the classical Landweber method, $T = D = \mathbf{I}$,

$$\sigma^{\lambda+1} = \sigma^{\lambda} + \nu_{\lambda} M_{\zeta\kappa}^{\top}(b_{\kappa} - M_{\zeta\kappa}\sigma_{\zeta}^{\lambda}). \tag{4.25}$$

The Cimmino method[173, 191, 192] is a little more nuanced. Each iterate is the average of the projections of the previous iterate on all affine hyperplanes[193, 194, 195] defined by the series of measurements, κ. Two extensions consider that the equal weighting of each projections' contribution to the solution is less than ideal in large-sparse systems (such as geophysical MIT). Component Averaging (CAV)[196] and Diagonally Relaxed Orthogonal Projections (DROP)[197] both take heuristic approaches to include any sparsity-related information. In dense systems, these reduce to Cimminos' method with additional computational overheads[173].

All of the iterative algorithms introduced are semi-convergent. As the number of iterations increases, the solution quality will improve, before decaying away to noise-corrupted

results comparable to an un-regularised choice. Various 'stopping rules' exist to estimate an ideal iteration number, including the previously introduced discrepancy principle and GCV. More information is found in [198] or [155, 172, 173, 174, 175].

4.5.3 Further Inversion Algorithms & Combination Approaches

The entire IR Tools [175] package is implemented as part of this exploration into algorithms suitable for our MIT problem, the mathematical details of which will be omitted for brevity. Practically hybrid least-squares and related algorithms are particularly interesting, designed to avoid the characteristic semi-convergent behaviour of typical inversion algorithms. They are effectively convergent, bypassing the need for a stopping rule. Functionally they apply a specific penalisation to the least squares problem after restriction to a specific Krylov subspace (e.g Equation 4.19 for CGLS), instead of building that subspace from a chosen penalisation.

To ensure a point of reference, both where algorithms may arise in subsequent inverse problems, and to those which this researcher found useful but the limited format of a thesis did not allow to be formally quantified, a list of IR tools derived algorithms are:

- Non-Negative Flexible CGLS (NNFCGLS), which functions with a preconditioner which changes with each iteration, and includes the integration of various regularisation matrices.
- Enriched Least Squares, a modified CG algorithm using a low-dimensional Krylov subspace to solve the penalised least-squares problem whilst incorporating additional information about the desired solutions. This can be coupled with various regularisation matrices and use a custom choice of enriched subspace.
- Constrained Least Squares, a projected-restarted method which includes box/energy constraints and the various regularisation matrices.
- Iteratively Reweighted Norm (IRN) approach, constructing a 1-norm solution through inner/outer iterations and any other algorithm as an initial guess.
- Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) which solves the Tikhonov

problem in a computationally efficient manner, well suited to large-scale problems even with dense matrix data, and may include box and/or energy constraints.

- Modified Residual Norm Steepest Descent method (MRNSD), a non-negatively constrained steepest-descent algorithm.
- Various generalised minimal residual methods (GMRES) apply Arnoldi iterations, including a converging hybrid version.

Further developments of state-of-the-art algorithms which function as extensions of IR-tools include the Hybrid JBD approach[199] designed for edge-enhancement and the Krylov low-rank (LR) solvers[200] designed for computational efficiency. There is also a package of 'lasso' methods, which focus on applying a 1-norm in iterative solutions, found in [201].

Given various inverse solutions enabled by this library of algorithms, it is plausible to generate multiple good solutions. In principle, the major difference between algorithmic solutions should be perturbations in elements related to magnified noise contributions in the problem, particularly for those which make similar assumptions. Further improvements in image quality could come from combining solutions in a post-processing step, such as taking standard averages. Care is required as each final image is no longer an explicit solution to the system of equations, likely compelling a manually observed process.

Given that noise corruption is a significant source of error in any solution, standard noise filters designed for image reconstruction may also be relevant. Three-dimensional image reconstructions limit the available choices, but: Gaussian noise filters, box filters, median filters, and band-pass filters have all been considered.

4.6 Resolution Matrix

A resolution matrix is a geophysical tool, which uses a square kernel function to provide insight into the influence various components of the forward model have on the ability to resolve parameters in the inverse problem[202].

There are two component matrices required, the forward map

$$\mathbf{F}_{\zeta\zeta} = M_{\zeta\kappa} M_{\zeta\kappa}^{\top}, \tag{4.26}$$

and the generalised inverse which inverts that matrix using some explicit scheme,

$$\mathbf{F}_{\zeta\zeta}^{\dagger} = (M_{\zeta\kappa} M_{\zeta\kappa}^{\top} + \lambda \mathbf{L}^{\top}\mathbf{L})^{-1}, \tag{4.27}$$

for a Tikhonov regularised inverse. The resolution matrix is then,

$$R_{\zeta\zeta}^{M} = \mathbf{F}^{\dagger}\mathbf{F}. \tag{4.28}$$

For the perfect resolution of all parameters, the result is the identity matrix. Columns give insight into the imaging of anomalies present in a specific element through the combined process of measurement and then inversion. Diagonal values, often known as the Model Resolution, represent the resolution of individual voxels[203, 204]. A value of unity represents perfect resolution whilst zero denotes a cell which is unresolved[203, 204]. The mean of these is known as the Information Efficiency, or average information content[202], describing the average resolution of the entire model.

The choice of regularisation in the inverse map imparts some bias to the results. The use of the identity matrix involves minimal assumptions and enables a focus on varied forward modelling. λ is again a variable with a non-linear dependence, selected with the L-curve. An open avenue for investigation is exploring the impact of various regularisation schemes, for example, by comparing this result to those from other strategies.

4.7 Magnetic Field & Image Quality Metrics

Both image reconstruction and assessment of magnetic field generation necessitate qualitative methodologies to assess result quality. Investigations have considered several interchangeable metrics. None of these is capable of quantifying the variation in image quality

as well as physical observation of the result.

The simplest error metric considered is the Root Mean Squared (RMS), which for a specific solution σ_{sol} is,

$$RMS = \sqrt{\frac{\sum(\sigma_{sol} - \sigma_{true})^2}{\zeta}}, \tag{4.29}$$

where ζ in this context is the number of elements (pixels, voxels etc.) in the final image. For an vector magnetic field this becomes,

$$RMS = \frac{\sum \sqrt{(B_{x,sol} - B_{x,targ})^2 + (B_{y,sol} - B_{y,targ})^2 + (B_{z,sol} - B_{z,targ})^2}}{\zeta}, \tag{4.30}$$

where ζ in this context is the number of points sampled in the discrete representation of the field. Another metric which is often useful for non-negative results which vary over multiple orders of magnitude is the Root Mean Squared Logarithmic Error given by,

$$RMSLE = \sqrt{\frac{\sum(log(\sigma_{sol} + 1) - log(\sigma_{true} + 1))^2}{\zeta}}. \tag{4.31}$$

Another applicable metric is a normalised overlap integral, posed as,

$$OVERLAP = \frac{\int(\sigma_{sol}\ \sigma_{true})\ dV}{\int(\sigma_{true}\ \sigma_{true})\ dV}, \tag{4.32}$$

leading to a value between 0-1, with 1 representing a perfect image. For consistency and ease of use, it is often beneficial to convert this into the range 1-0 with zero representing a perfect image. In a discrete problem, this becomes a sum over a pair of dot products between individual elements.

For two-dimensional images, a metric defined explicitly in terms of image quality is the Structural Similarity Index (SSIM)[205]. Combining local image structure (i.e patterns of pixels), luminance, and contrast into a single metric intends to represent a subjective

human quality score. Quantitatively this becomes,

$$SSIM(x,y) = [s(x,y)]^{\alpha} \cdot [l(x,y)]^{\beta} \cdot [c(x,y)]^{\gamma} \begin{cases} \text{Structure}, & s(x,y) = \frac{\Sigma_{xy}+C_1}{\Sigma_x \Sigma_y + C_1}, \\ \text{Luminance}, & l(x,y) = \frac{2\bar{\mu_x}\bar{\mu_y}+C_2}{\bar{\mu_x}^2+\bar{\mu_y}^2+C_2}, \\ \text{Contrast}, & c(x,y) = \frac{2\Sigma_x\Sigma_y+C_3}{\Sigma_x^2+\Sigma_y^2+C_3}, \end{cases} \quad (4.33)$$

where Σ_{xy} is the covarience of images x,y, Σ_x, Σ_y are the local standard deviations, $\bar{\mu_x}$, $\bar{\mu_y}$ are the local means and $\alpha, \beta, \gamma, C_1, C_2, C_3$ are constants. The resulting value ranges between 0-1, where 1 represents a perfect match.

The image reconstructions in this thesis consider volumetric images, leading to the Multi-SSIM3 metric[157]. Assessment concerns a combination of luminance, contrast, and structure terms upon several versions of the volume, at multiple scales. Designed to closely match a human perception this has a value between 0-1, which is re-scaled with zero representing a perfect image for consistency.

Overall, this chapter has presented the complexity of ill-posed inverse problems in the context of MIT. It has covered several tools to analyse the system, select an appropriate algorithm, find a solution, and quantitatively assess that solution. It provides a library of algorithms and identifies links between them and the properties of a desired solution. Individual algorithms are of varying complexity, novel to MIT, and include some state-of-the-art approaches. All exhibit computational advantages to traditional, widely used, approaches in MIT and should remove computational bounds on both the maximum attainable resolution and the number of measurements made.

Chapter 5

Induction Systems For Orthogonal Magnetic Field Generation

Equation (3.23),

$$I_k = (\vec{M^I}_{jk}^{\top}\ \vec{M^I}_{jk}\ - \lambda\ \mathbf{I}^{\top}\ \mathbf{I})^{-1}\ (\vec{M^I}_{jk}\ \vec{\nabla} T_n^m(r,\theta,\phi)),$$

is derived in an effort to utilise a single modular system of coils to produce a range of orthogonal spherical harmonic magnetic fields in a compact manner. Even confining the investigation to simple building blocks like square, circular, and saddle coils produces a complex design problem. The goal of this chapter is to develop a deeper understanding of this novel field generation philosophy and establish some intuition for pertinent questions such as:

- What influence does the coil composition have on orthogonality properties?
- How many magnetic fields can be produced?
- What approach can be used to identify unique magnetic fields?
- How precise does a current control system need to be to produce these fields?
- What limitations/issues are associated with the generated current distributions?
- How should regularisation be applied?

There are two broad categories for applications; localised where the system may enclose the volume of interest, similar to X-ray screening or MRI scanners; and portable applica-

tions, for investigation of larger regions such as geophysical surveys or shipping containers. Regardless of the application, these systems are scalable, allowing analysis on normalised scales to be adjusted depending on end-user requirements within physical limitations.

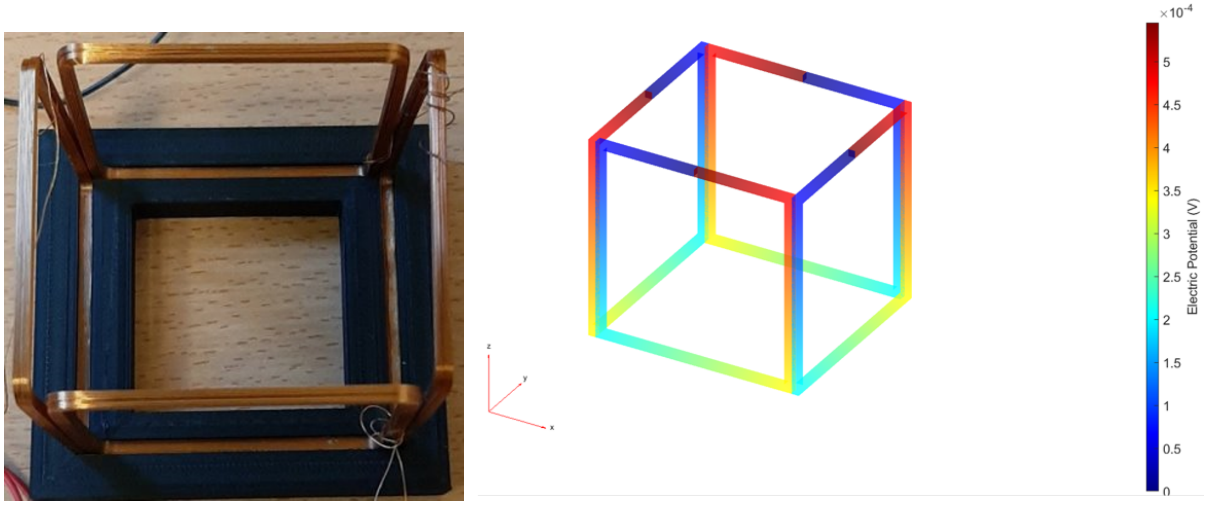
5.1 Induction Systems - Localised Applications

5.1.1 Experimental System: Four Coils

One approach to produce multiple, independently orthogonal, magnetic fields is using four square coils, orientated along the X and Y axes of a volume like Helmholtz pairs. This is predominantly a theoretical analysis used to inform a linked miniature experimental setup - designed to verify that AC spherical harmonic fields, coupled with magnetometer measurements, are beneficial to the MIT imaging problem.

This first section presents a robust approach to identify whether a particular magnetic field can be considered mutually unique within a set of generated profiles - whether each is 'quasi-orthogonal'. It proceeds to highlight that calculating currents for various spherical harmonic targets reproduces intuitive results and provides added value by producing additional unexpected field profiles. The resulting fields are spatially dissimilar but have minimal orthogonality properties, providing a baseline set of fields. Later, a comparison is made between these and higher quality results more representative of the targets to highlight the influence of improving the orthogonality of the applied fields in the MIT inverse problem.

Construction of individual coils consists of a ten-by-ten layered structure of wire gauge no. 31 with a single set of terminals, wrapped into a square coil with an inner side length of 50mm and an outer side length of 54mm. Four of these are mounted in an insulating base as shown in Figure 5.1a. A signal generator produces a sinusoidal signal at some fixed frequency, which passes through a bipolar amplifier, and then an amplitude control circuit, powering the coils. At 50KHz using a parallel configuration of the maximum current



(a) The experimental induction system.

(b) The electric potential for a parallel input current of 130mA.

Figure 5.1: The simple bi-axial Helmholtz setup utilised as a proof of principle of MIT with a modular induction system, and the associated simulated model of the electric potential for coils powered in parallel.

of $130mA$, there is a mutual inductance of $2.1 \pm 0.1mH$ and a resistance of $11.9 \pm 0.1\Omega$, for which an electric potential model produced in MATLAB is in Figure 5.1b. Field measurements use a Lakeshore F71 Teslameter with a triaxial mounted probe, which features a sensitivity of 0.1μT at frequencies up to 70kHz.

To apply Equation (3.23) to this system, consider targeting the set of magnetic fields described by spherical harmonics from $n = 1, m = -1$ to $n = 6, m = 6$ - a total of 48. In subsequent work these are numbered linearly in order of appearance, for example, $n = 1, m = -1$ would be field 1, $n = 2, m = 2$ would be field 8 and so on. A subset of these are tabulated in the Appendix A for clarity. A generation volume must also be specified, which defines the imaging region or scanner bore in the MIT problem. This must be carefully selected, balancing a region small enough that all component coils can interact to control the spatial distribution of the magnetic field, but large enough that the subsequent volume is of practical use. The algorithmic results depend implicitly on this volume and the geometric properties of sources, so any variations will alter the resulting current distributions in general. Here a centralised cube with a side length one-quarter of an individual current source is used. The resulting sensitivity matrix applies Equation B.1

with 27,000 points, resulting in a 81000×4 augmented M^I_{jk} term. The condition number is 5.8 which is strongly linked to the simplicity of the design, and makes the regularisation term relatively unimportant.

A result of the simplicity of this induction system is a limited ability to produce magnetic fields which are mutually unique and well-representative of the targeted harmonic functions. One approach to analyse this is through the RMS, comparing the normalised spatial profile of the three-dimensional field targeted in Equation (3.23) to the equivalent produced by taking the resulting current distribution and solving Equation (3.18) using Equation 4.30. The results for the target set are in Figure 5.2.

The specific normalisation used to calculate these values here and throughout the thesis is: taking the target function divided by its magnitude, and the actual field divided by its magnitude before calculating the component-wise residuals. We note in passing that the target fields are simply mathematical constructs, and their magnitudes are arbitrary, but have been scaled to provide physically realistic values which may not be entirely consistent with the generated fields throughout this chapter. In reality reproducing the spatial distribution and variation in magnitude specifies whether the target orthogonal magnetic field has been generated.

This indicates how well each target field is reproduced. Identifying the unique, quasi-orthogonal, results concerns simply selecting fields below some arbitrary cutoff denoting a 'good' approximation. Four obviously meet this criterion: the best being two of the gradient fields $n = 2, m = 0$ at 0.35% error, then $n = 2, m = 2$ at 0.53% error; followed by the two axially symmetric linear fields at $n = 1, m = \pm 1$ with 5.5% error. These are intuitive given what is effectively a pair of Helmholtz coils producing X-Y symmetry, and utilise unitary currents which could be selected manually, values which do not change with double the generation volume. That volume adjustment does impact the RMS values, which increase to 2.7%, 3.5% and 18% respectively as edge effects are incorporated.

Beyond an indication, the RMS has limited value in assessing these magnetic fields. Higher-degree harmonics have spatially varying magnitudes, and areas of weaker field strength have smaller residual values with a suppressed contribution to the RMS. There is no information pertaining to the uniqueness of individual fields; large values may relate to poor harmonic approximations, but the spatial distribution of magnetic field still provides additional information; in certain cases, Equation (3.23) suggests identical current distributions for different targets, which have unrelated RMS despite being physically identical.

Convention would suggest mapping the current distributions and the resulting field simulations to perform this evaluation[129]. Whilst valid, this approach quickly becomes impractical; 48 volumetric images of the three-axial vector fields exist, along with corresponding current distributions; these require permutative comparison and manual separation. Magnification of complexity occurs with an increasing number of coils as more complex fields can be produced. A tolerance-based comparison of the current distributions can simplify things, for instance, normalising distributions to $\pm$1A, with a tolerance of 0.5A reduces this to a set of 16 possibly unique magnetic fields. The physics of field generation necessitates further analysis though, for example, two dissimilar distributions can produce an effectively identical magnetic profile.

Applying Equation (3.25), for example, using a least-squares minimisation of the simulated field against a set of targets is a more robust approach to performing this analysis. As a limited system, basis functions up to $n = 6$ are sufficient to characterise the resulting magnetic fields to 0.01% accuracy, resulting in a matrix of 48x48 coefficients for all targeted magnetic fields. Diagonal elements relate to the target field, which ideally is the dominant coefficient providing a quick reference to the success of Equation (3.23). In general, a row-wise comparison of the harmonic composition is necessary. Similarities in characterising coefficients can then highlight similar profiles and those which potentially

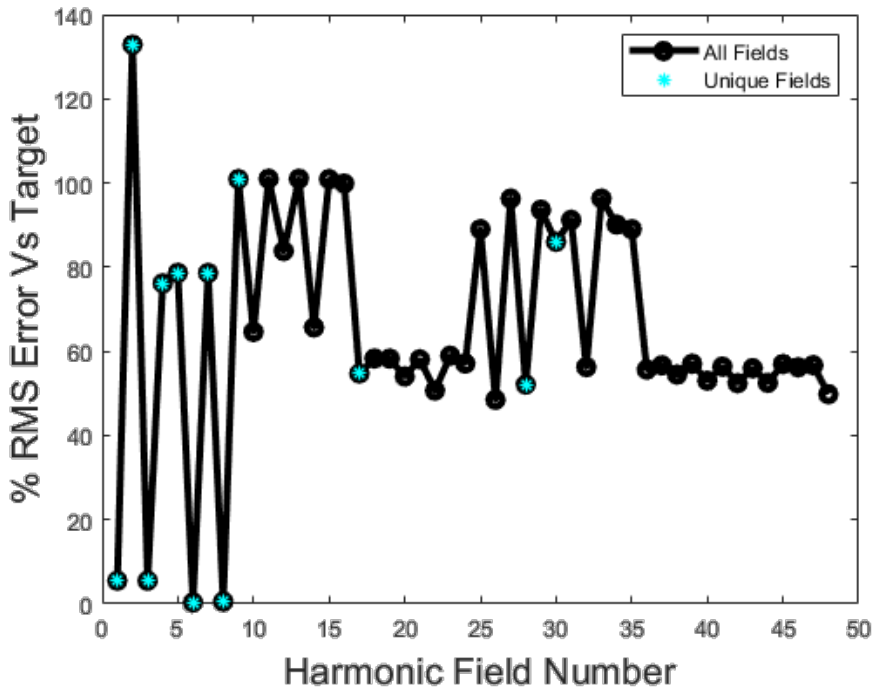


Figure 5.2: The various RMS values for each magnetic field produced by the four coil induction system compared to the targeted spherical harmonic. Production of four intuitive harmonics is easily achieved, but no conclusive information pertains to the usefulness of the remaining forty-four beyond they are poor representations of their target. The mean of these values over all 48 fields is 66%, indicating a limited ability to reproduce the target profiles. A row-wise comparison of coefficients resulting from Equation (3.23) leads to twelve unique fields, highlighted with blue stars, with an improved mean RMS of 47%.

have significant overlaps. Using a tolerance results in twelve unique magnetic fields, some of which are degenerate - arising from reversing the polarity of current at all four coils. Conventional MIT uses degenerate fields to extract additional information, and multiple activated coils are likely to increase the signal-to-noise ratio vindicating their usefulness. An overlap integral of each profile with the remaining eleven provides further quantitative confirmation of relative uniqueness, Figure 5.3. Despite some significant overlaps between some fields, each is generally suitably dissimilar. Even where the overlaps are large, nearing 0.7, each field will provide some variation in interrogation compared to a traditional MIT approach. In turn, this will provide unique information to the inverse problem. Negative overlaps are of less concern because, as previously discussed, degenerate fields are known to be valuable in MIT.

These can be labelled with respect to the target harmonic, though the simplicity of this de-

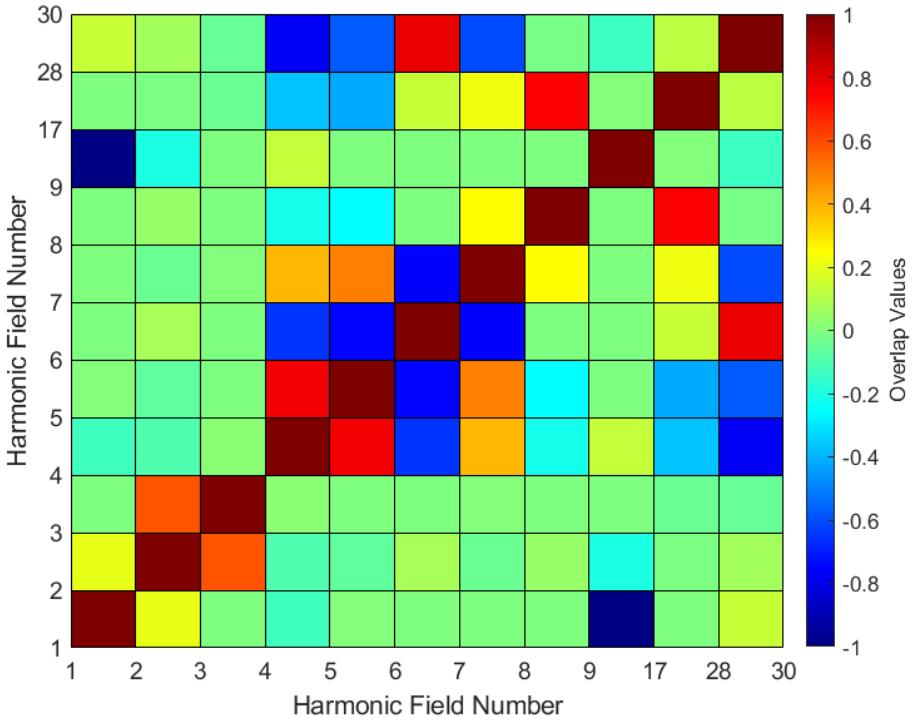


Figure 5.3: The values resulting from a permutative calculation of the overlap integrals between all twelve selected fields. The diagonal represents the overlap of a particular magnetic field with itself and has the expected value of unity. The symmetry between each half (above and below the diagonal) further mirrors expectation - for example, the overlap between field one with field two, compared to field two with field one, should have the same result.

sign means the majority are not accurate depictions of their namesake: $n = 1,\ m = -1 : 1$, $n = 2,\ m = -2 : 2$, $n = 3,\ m = -3$, $n = 4,\ m = -3$, $n = 5,\ m = -2, 0$, highlighted with blue stars in Figure 5.2. A less intuitive unique excitation pattern relates to the $n = 5,\ m = 0$ field given in Figure 5.4, highlighting how the algorithm can select non-integer current ratios for multiple sources to produce unique magnetic field data. Note that the $n = 1,\ m = 0$ target field is a linear-z profile which intuitively cannot be generated by this induction system, yet in terms of this analysis it is a spatially unique field.

The experimental setup utilised a single manually controlled power supply to control the sinusoidal current supplied to all coils. Amplitude adjustments involve introducing resistors or polarity changes between the amplifier and a specific coil, limiting the available ratios to a subset of the twelve unique fields - the four accurate spherical harmonics ($n = 1, m = \pm 1$, $n = 2, m = 0$ and $n = 2, m = 2$).

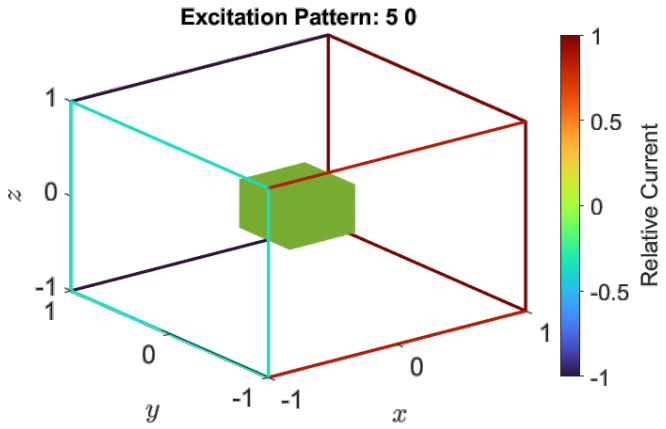
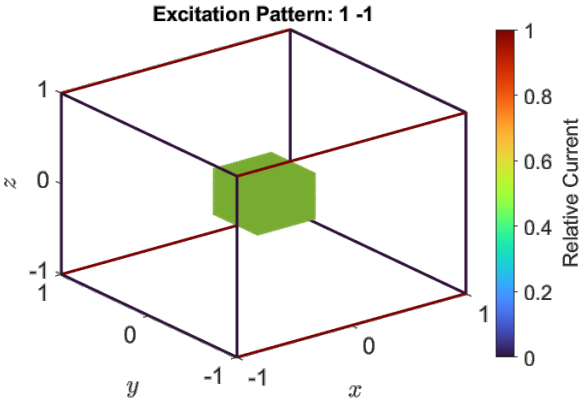
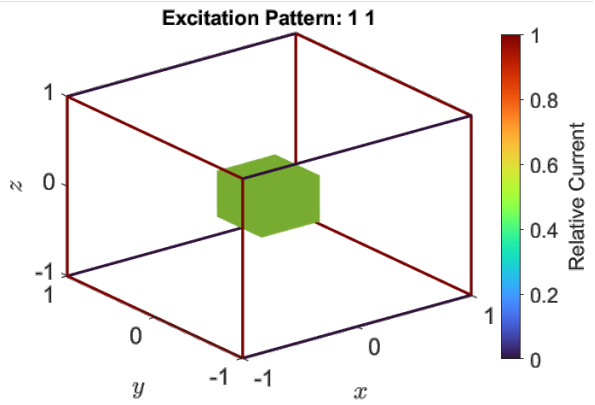


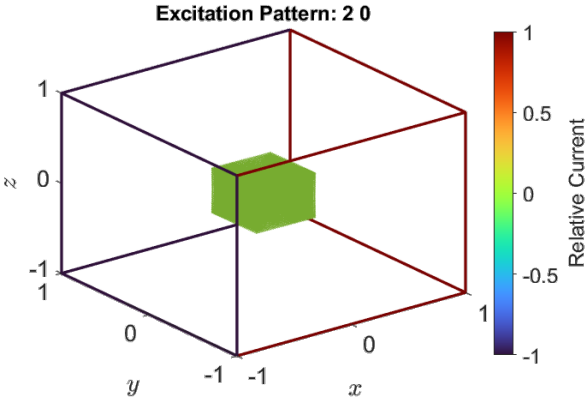
Figure 5.4: The coil excitation pattern for the $n = 5$, $m = 0$ target harmonic.



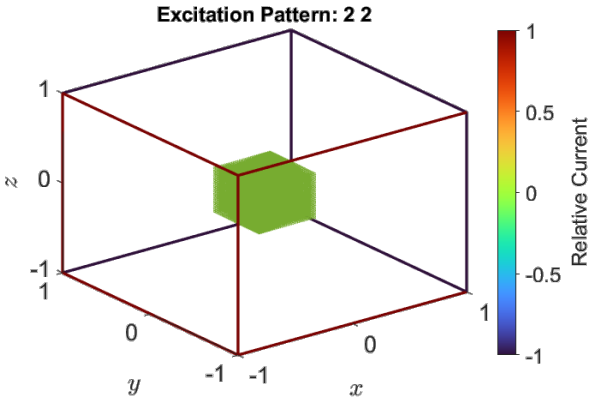
(a) The coil excitation pattern for $n = 1$, $m = -1$.



(b) The coil excitation pattern for $n = 1$, $m = 1$.



(c) The coil excitation pattern for $n = 2$, $m = 0$.



(d) The coil excitation pattern for $n = 2$, $m = 2$.

Figure 5.5: The set of normalised current ratios implemented in the experimental MIT system, with the green block identifying the field generation region. This region is chosen as an ideal internal volume for the production of spherical harmonic magnetic fields.

The excitation patterns are found in Figure 5.5 and are intuitive given the induction

geometry and rotational symmetries of the target magnetic profiles. The lack of a translation stage impeded precise mapping, and as Helmholtz/anti-Helmholtz magnetic fields, the experimental field mapping results are disinteresting beyond qualitative verification. The omission of the field maps is justified for this reason and in the interest of brevity.

In summary, this section presents a robust approach to analyse magnetic fields generated by the optimisation of currents for a simple fixed coil geometry. As a limiting case, one expects the potential to produce only a fractional subset of possible spherical harmonic magnetic fields. Those provide an intuitive verification, with the analysis identifying several unexpected, mutually unique profiles. Despite being poor representations of the target spherical harmonics, their relative independence should add value in providing varied spatial interrogation in an MIT problem. They also form a set of independent, but not orthogonal, fields, providing a baseline to compare to more orthogonal fields in the next chapter.

5.1.2 Theoretical System 1: Cylindrical Coil Array

Consider an open-ended cylindrical topology with analogies in MRI or millimetre wave scanners. The most appropriate conforming modular components are circular and saddle coils, where the number of elements relates directly to the complexity of the generatable set of fields. The complexity increases with the harmonic degree, requiring more components and/or a more nuanced distribution of elements.

This section aims to highlight the versatility of the current distribution calculations and the framework to quantify the resulting magnetic fields with respect to the orthogonal target functions. It demonstrates that it is feasible to produce more complex spherical harmonic fields with reasonable accuracy whilst establishing some bounds on the current control required for generation. Further work highlights the influence of λ (the relative weighting of the current-norm) on both fidelity and the field-per-unit current.

The symmetries in the highest-degree target fields, $n = 6$, and the number of spatial variations can inform the distribution of modular components. A structure divided into a six-by-six grid of saddle coils along the azimuth, with two circular end-coils to provide additional transverse field control is proposed, seen in Figure 5.6. The circular end coils have a radius of 2 normalised units, separated by 4 units, with individual saddle coils having an arc length by z-extent of approximately 2.1 x 0.47 units.

Equation (3.23) is then applied using a centred cubic volume with dimensions of one normalised unit. Individual inductors are assumed to contain one-hundred wire turns, and the results are normalised to ± 1A. This covers a volume with approximately half the radius of the enclosing cylinder, and double the z-extent of the dominant component saddle coil. Dimensions relative to the system are equivalent to the four coil analysis. Using the same calculation process leads to an 81000×38 augmented M^I_{jk} term with a condition number of 2.69×10^3. The choice of an appropriate λ value varies with the target field, so an L-curve estimates a value, which is then scaled uniformly to give preference to high

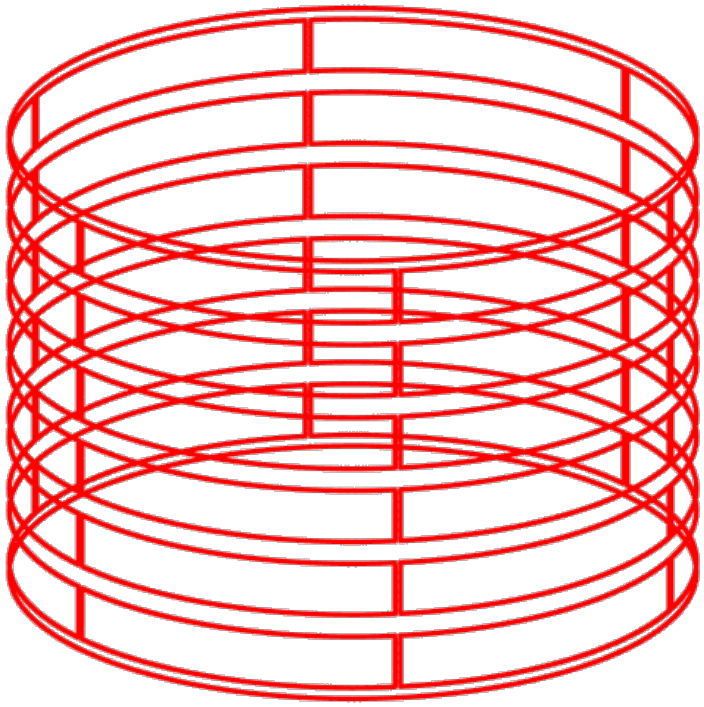


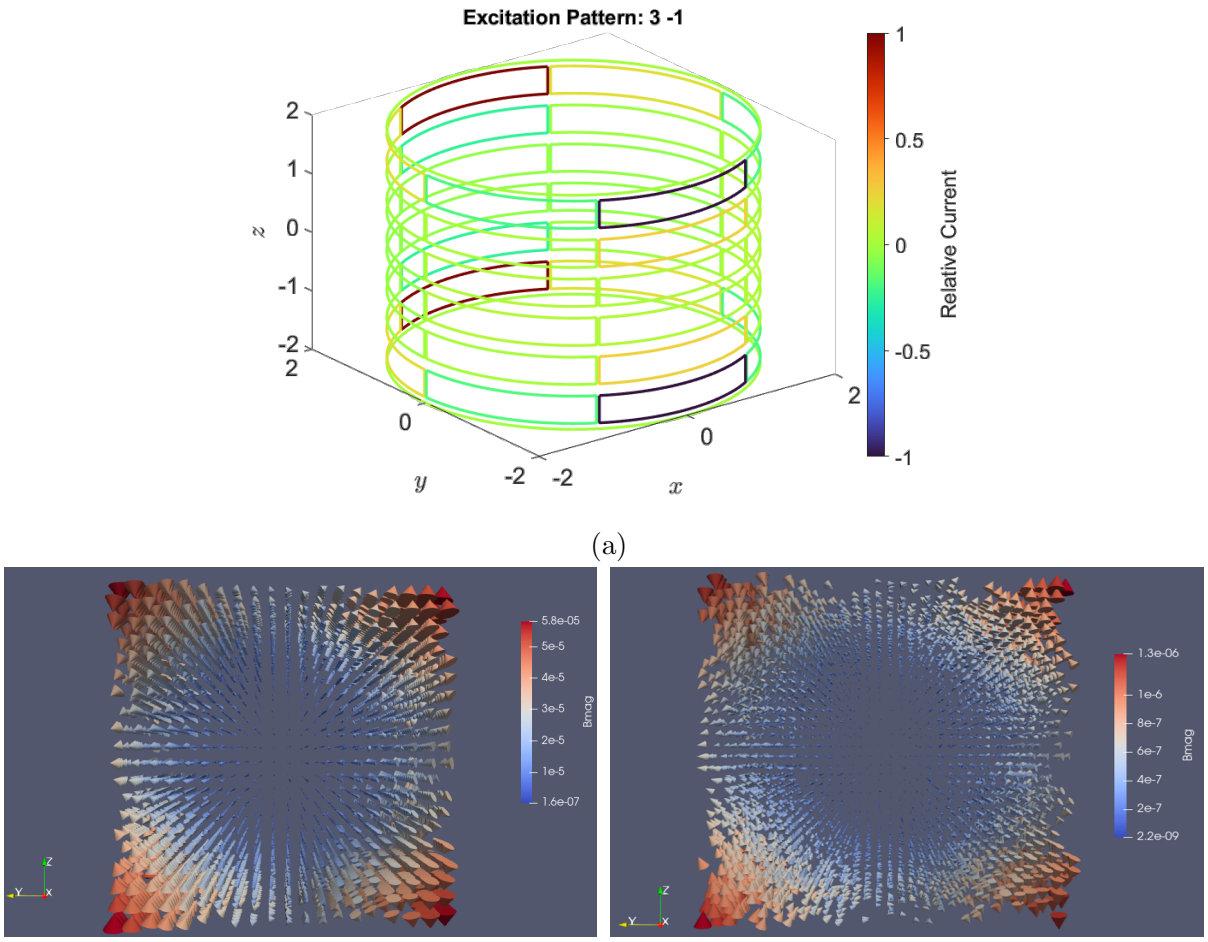
Figure 5.6: A coil structure composed of a six-by-six grid of saddle coils distributed around the azimuth coupled with two circular end coils. Each has individual current control so, by activating them in different sequences, multiple magnetic fields can be produced from a single geometry.

fidelity fields - see the discussion around Figure 5.16 for more details - with values in the range $2.6 \times 10^{-16} \leq \lambda \leq 4.3 \times 10^{-13}$.

A subset of the excitation patterns corresponding to the harmonics at $n = 3, 4, 5, 6$, $m = -1, -2, -4, 2$ are provided in Figures 5.7, 5.8, 5.9 and 5.10. Each subplot presents the excitation pattern, then quiver plots of the target, and simulated magnetic fields from that current distribution.

The related RMS data is provided in Figure 5.11, calculated is discussed in Section 5.1.1, demonstrating an obvious improvement relative to the previous system. For quick reference, the average RMS, dubbed Mean Field Fidelity henceforth, is an indicative value, 8.5% here.

The advantage of determining current distributions via Equation 3.23 is immediately



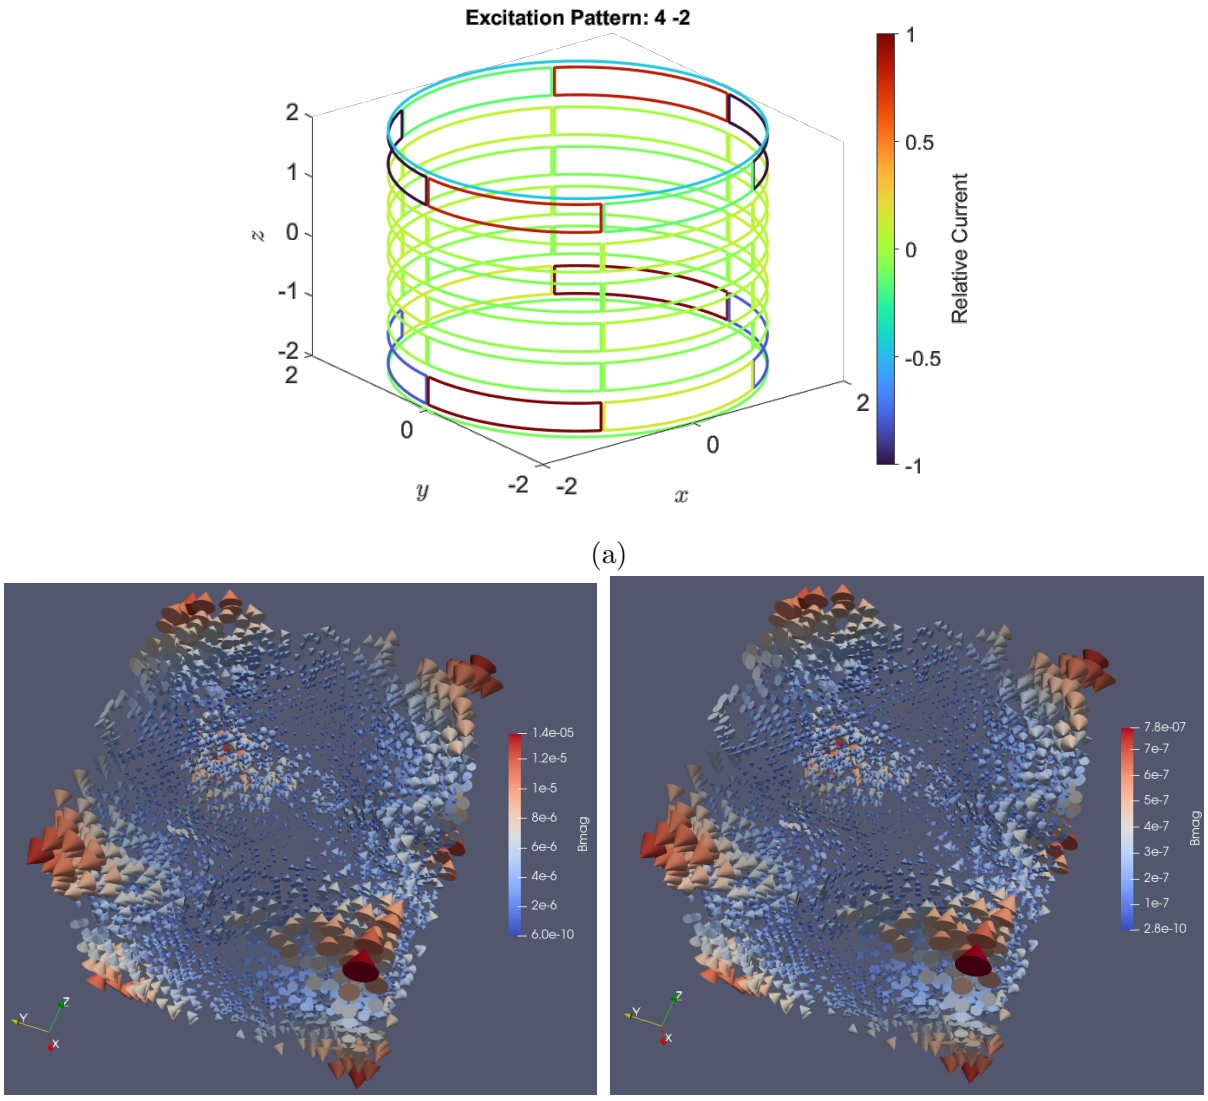
(a)

(b) (c)

Figure 5.7: 5.7a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.7b) The target function related to the gradient of the real spherical harmonic n = 3 m = -1. 5.7c) The reconstruction from the excitation pattern with 0.7% RMS error, colour bar gives values in Tesla.

apparent for a system with thirty-eight sources. Even with an innate understanding of the links between the geometries and generated fields, selecting an appropriate distribution for the simplest linear fields is not a trivial task. With forty-eight fields of increasing complexity, this quickly becomes impractical.

As this system can produce some $n = 6$ spherical harmonic fields with good RMS accuracy, it becomes relevant to consider another thought experiment. Assume a system exists and is capable of producing *all* spherical harmonics with $m = 0$ perfectly over a unit cube centred at the origin. As the degree increases, so do both complexities in the spatial magnetic field landscape variation and the associated gradient in field strength



(a)

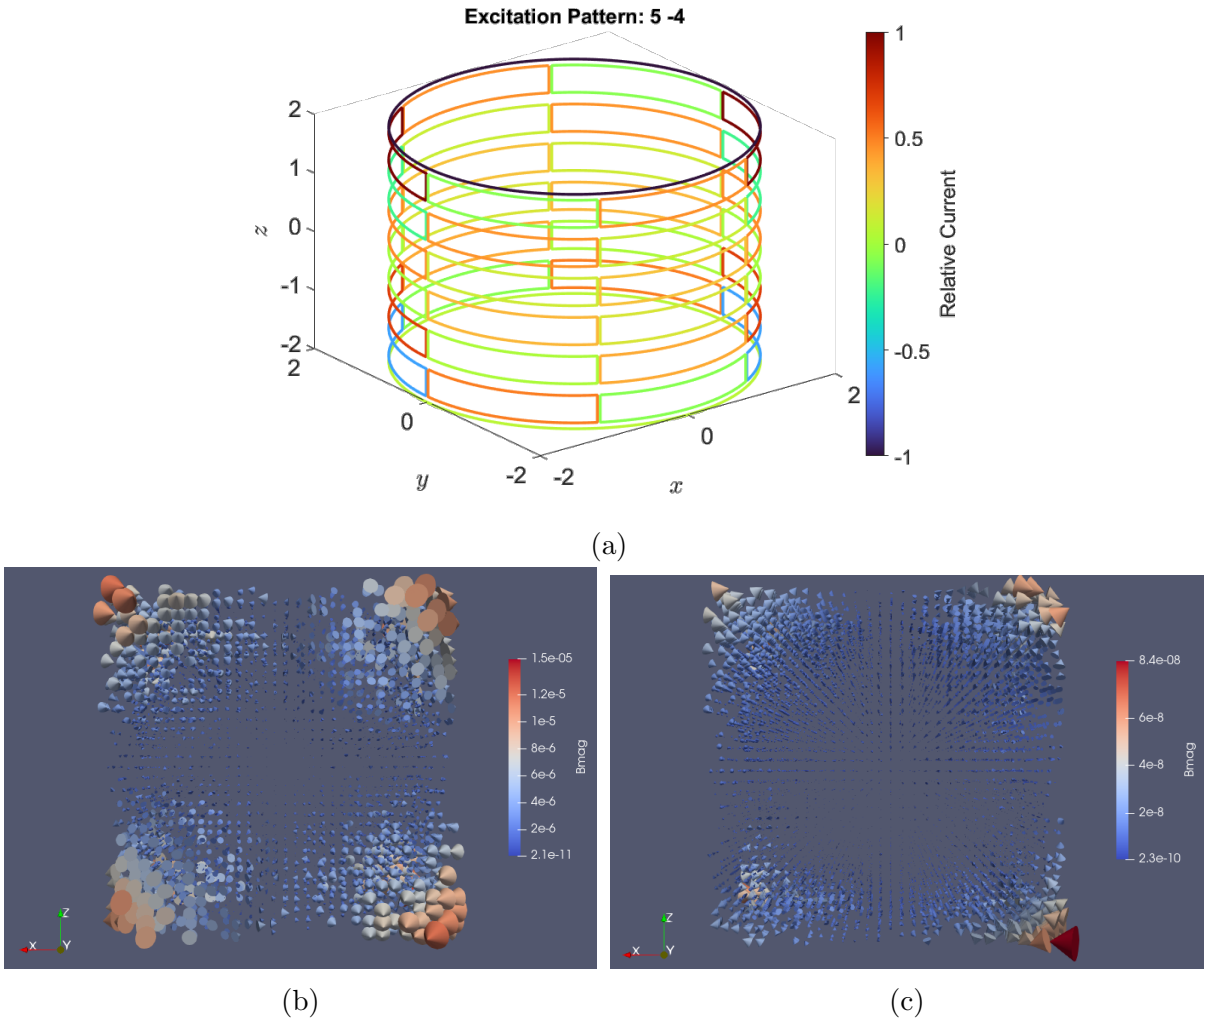
(b)

(c)

Figure 5.8: 5.8a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.8b) The target function related to the gradient of the real spherical harmonic n = 4 m = -2. 5.8c) The reconstruction from the excitation pattern with 1.6% RMS error, colour bar gives values in Tesla.

across any specific volume of interest. Examples for $n = 3$ and $n = 9$ are in Figure 5.12, with a physical scaling as discussed in Section 5.1.1. Both are plotted with an identical colour map to ease visual comparison.

The low degree harmonic varies in field amplitude by roughly three orders of magnitude over the volume - from 7.9×10^{-6}T to 9.4×10^{-9}T. The comparable high degree harmonic ranges by twelve orders of magnitude - from 1.0×10^{-5}T to 2.0×10^{-17}T - over the same volume. Contrasting the two the bulk response is smaller in the $n = 9$ field but



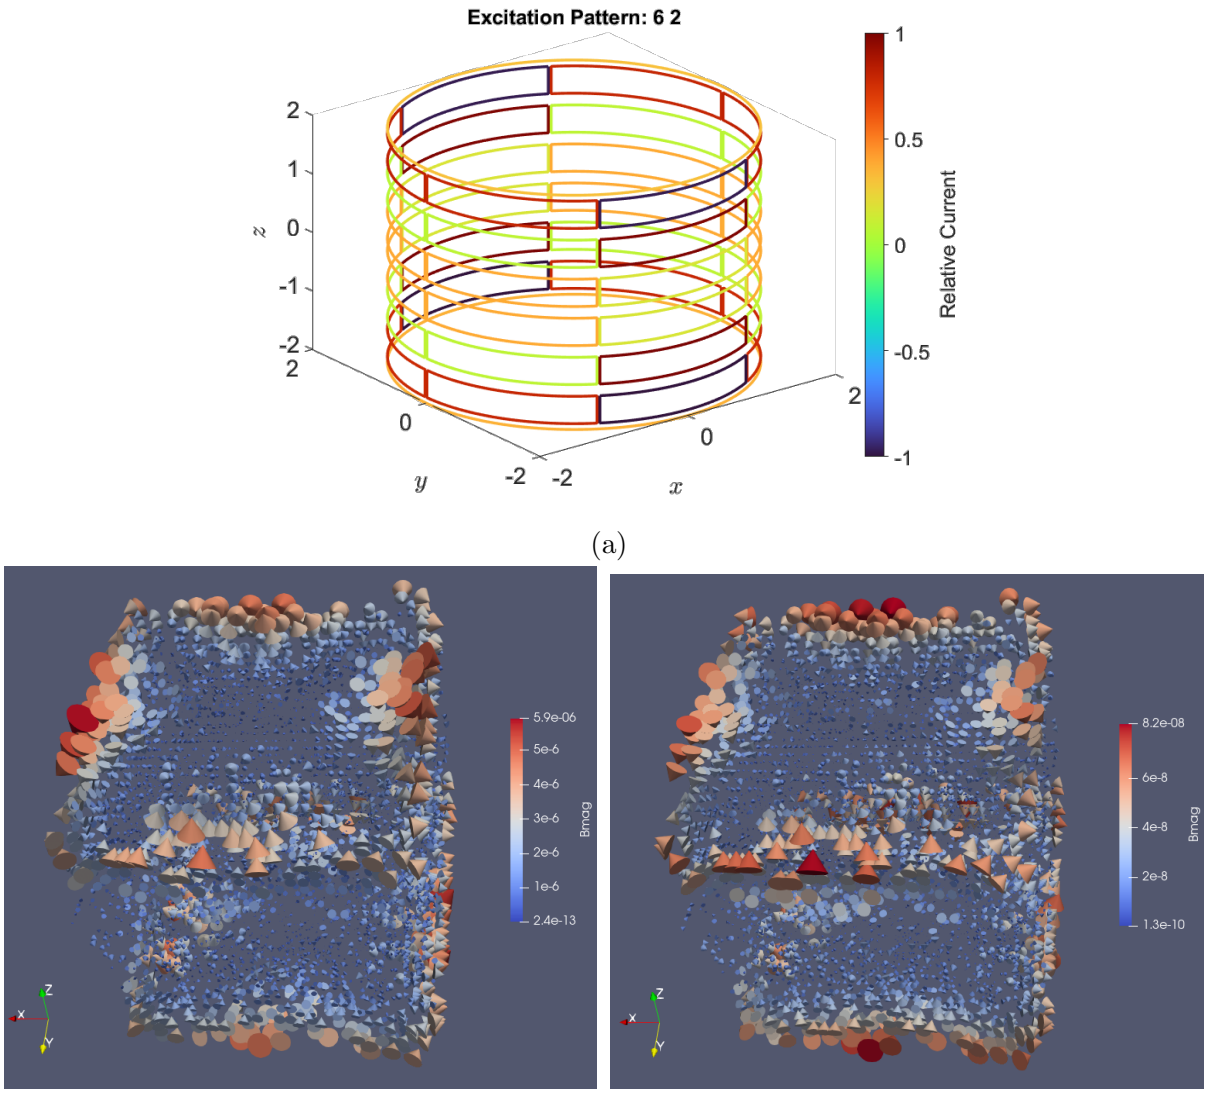
(a)

(b) (c)

Figure 5.9: 5.9a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.9b) The target function related to the gradient of the real spherical harmonic n = 5 m = -4. 5.9c) The reconstruction from the excitation pattern with 14% RMS error, colour bar gives values in Tesla.

contains an increase in local information at the volumes' extrema.

From a pure mapping perspective $n = 9$ hits the dynamic limit of fields which can be fully characterised, even by the most idealised superconducting quantum interference device[206] with an appropriate uniform increase in amplitude. For volumetric MIT allowances must be made for: contactless measurements, the maximum generation region relative to inductors, the weak integrated out-of-phase response, and the dynamic range in the detectable conductivity. Higher degree harmonics then become impractical for anything but the characterisation of minute edge-based variations, coupled with computational complexity, this sets a natural truncation point at $n = 6$. Higher degree fields



(a)

(b) (c)

Figure 5.10: 5.10a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.10b) The target function related to the gradient of the real spherical harmonic n = 6 m = -1. 5.10c) The reconstruction from the excitation pattern with 6.4% RMS error, colour bar gives values in Tesla.

may find use in other applications, such as non-destructive fault analysis, where the effect is a decrease in the effective 'spot size' of interrogation.

The RMS values in Figure 5.11 identify three outliers which may require analysis to ensure they are useful members of an orthogonal set. Each has symmetries which are not adequately reproduced in the angular components of this design. Uniqueness is more accurately determined by Equation (3.25). Successful generation of some $n = 6$ fields requires a basis of $n_{max} \geq 8$. Taking $n = 9$ there is a matrix of 48x99 coefficients charac-

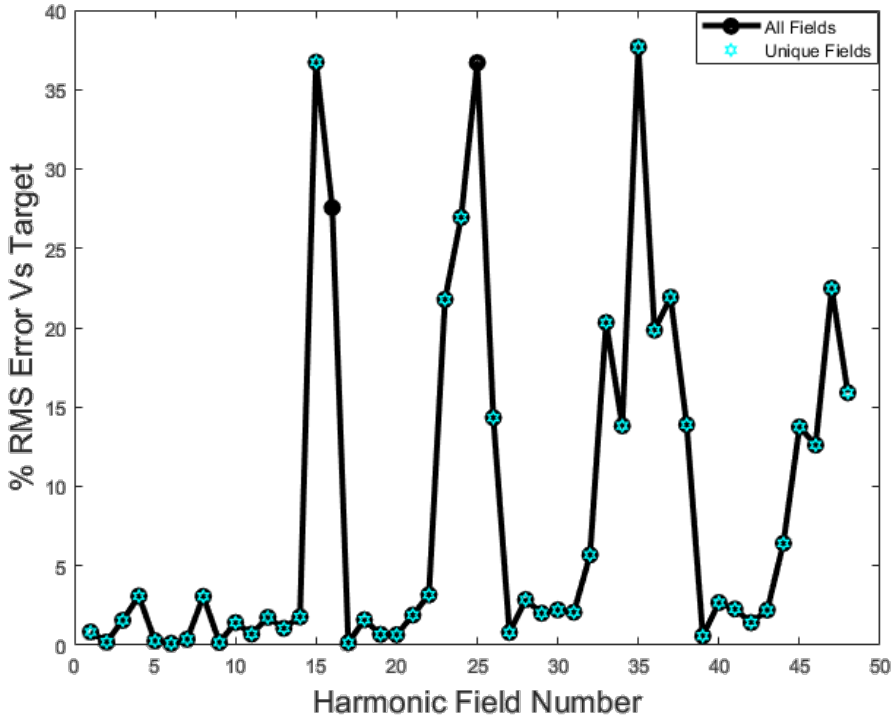
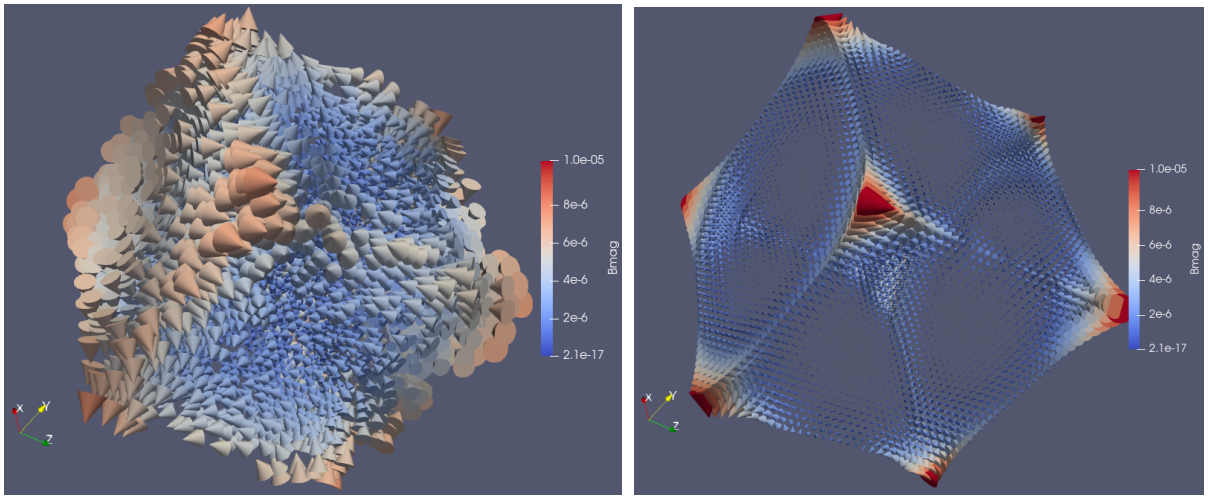


Figure 5.11: A plot of the various RMS values for each magnetic field produced by the induction system compared to the targeted spherical harmonic. The Mean Field Fidelity over all 48 fields is 8.5%. There are three distinct RMS outliers which warrant additional investigation at field numbers 15, 24 and 35 ($n = [3, 4, 5]$, $m = [3, 4, 5]$ respectively). Equation 3.25 coupled with overlap calculations suggest two non-unique fields, 16 and 26 ($n = 4, m = -4$ & $n = 5, m = -5$ respectively), leaving 46 quasi-orthogonal magnetic fields with a Mean Field Fidelity of 7.5%.



(a) The target field related to the $n = 3, m = 0$ spherical harmonic.

(b) The target field related to the $n = 9, m = 0$ spherical harmonic.

Figure 5.12: A comparison of a perfect 'low' degree spherical harmonic magnetic field and a perfect 'high' degree spherical harmonic magnetic field, plotted in Tesla for a setup which assumes one-hundred wire turns in the current sources.

terising the entire system, which a tolerance-based comparison reduces to ten potentially similar fields - in five pairs.

A normalised overlap integral between a pair gives an additional measure of similarity, where orthogonal results have values $\approx 1 \times 10^{-16}$ for finite-precision numerical mathematics, and identical results approach unity. The systems' complete results are in Figure 5.13. From this, two field-pairs have significant overlap (exceeding 0.5): $n = 4, m = -4$ and $n = 6, m = -4$, and $n = 5, m = -5$ and $n = 5, m = -1$. The field with the highest RMS value from these pairs is removed, leading to the 46 unique magnetic fields identified in Figure 5.11. These results again highlight the limits of an RMS evaluation, where the two highest RMS fields are unique, but one 10% *better* is not. Most of these are high-quality spatial approximations of the target harmonics, where most of the error is associated with producing variations in magnitude over the volume.

Field configurations in Figures 5.7, 5.8, 5.9 and 5.10 typically have a good spatial match in the magnetic distribution, though Figure 5.9 has fewer areas of peak field than intended. Most of the error is associated with the variation in field strength over the generation volume. For Figures 5.7 and 5.8, the peak value is around twenty times smaller than expected, increasing to three/four orders of magnitude in Figures 5.9 and 5.10 respectively. This is relatively ubiquitous throughout this work and limits the absolute orthogonality of the generation procedure.

The lower-quality magnetic fields have some functional overlap with the remaining generated set which must be considered to characterise uniqueness. Through Equation (3.25), and overlap calculations, we find a maximal similarity of 12% and 28% for Figures 5.9 and 5.10 respectively, identifying each provides some unique information.

The generation volume is also relatively small, which links explicitly to the subsequent interrogation region. As scaleable physics, theoretically, a larger induction system is

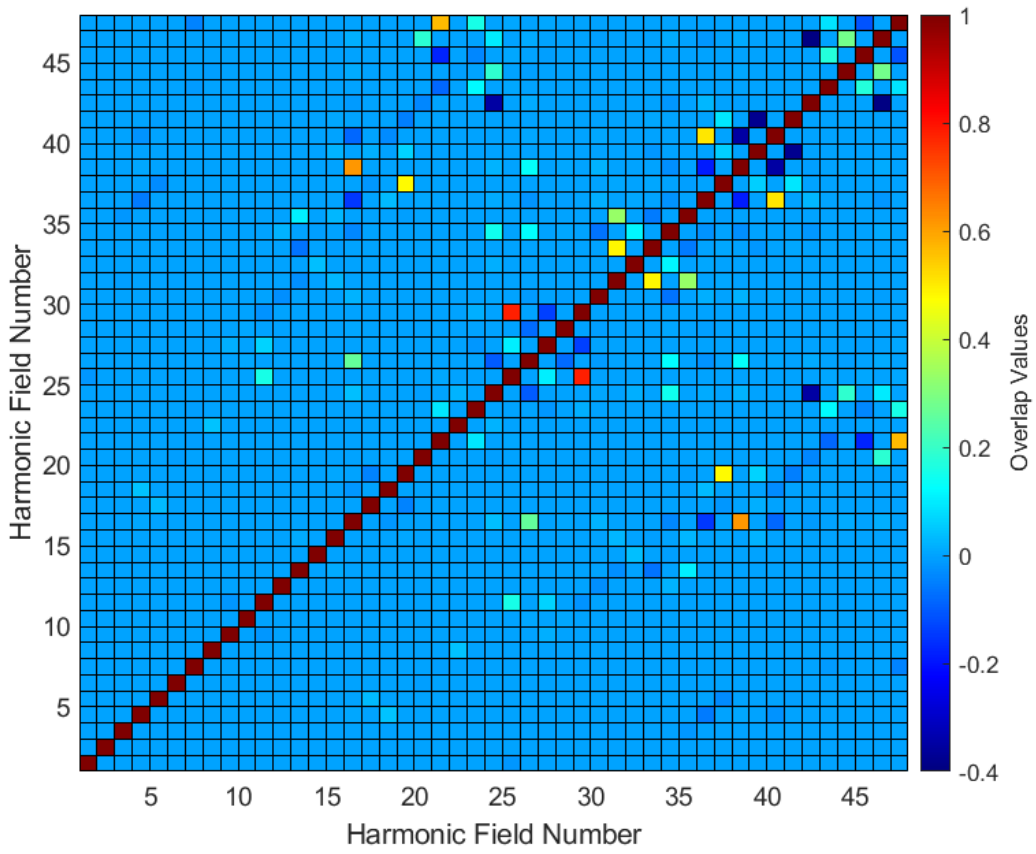


Figure 5.13: The values resulting from a permutative calculation of the overlap integrals between all forty-eight evaluated fields. Compared to the previous four-coil system, Figure 5.3, there is significantly less overlap between individual magnetic fields suggesting this system provides results which are significantly closer to true orthogonality. Two magnetic fields have an overlap value greater than 0.5 and therefore have relatively poor orthogonality.

sufficient to target any specific volume. These results demonstrate a clear, intuitive improvement in field quality, variability, and orthogonality by introducing additional source elements.

The focus on idealised, high-fidelity magnetic fields does mean these results are relatively unstable in the presence of errors in the current supply. The variation in field fidelity with the precision in current control is in Figure 5.14. Fluctuations beyond ± 0.0001 adversely impact field quality, though at three decimal places this difference is small, reducing the Mean Field Fidelity to 9.1%. Further losses in precision will have a significant impact, with only the harmonics up to $n = 2$ maintaining a reasonably consistent fidelity value

with control up to $\pm 0.01A$.

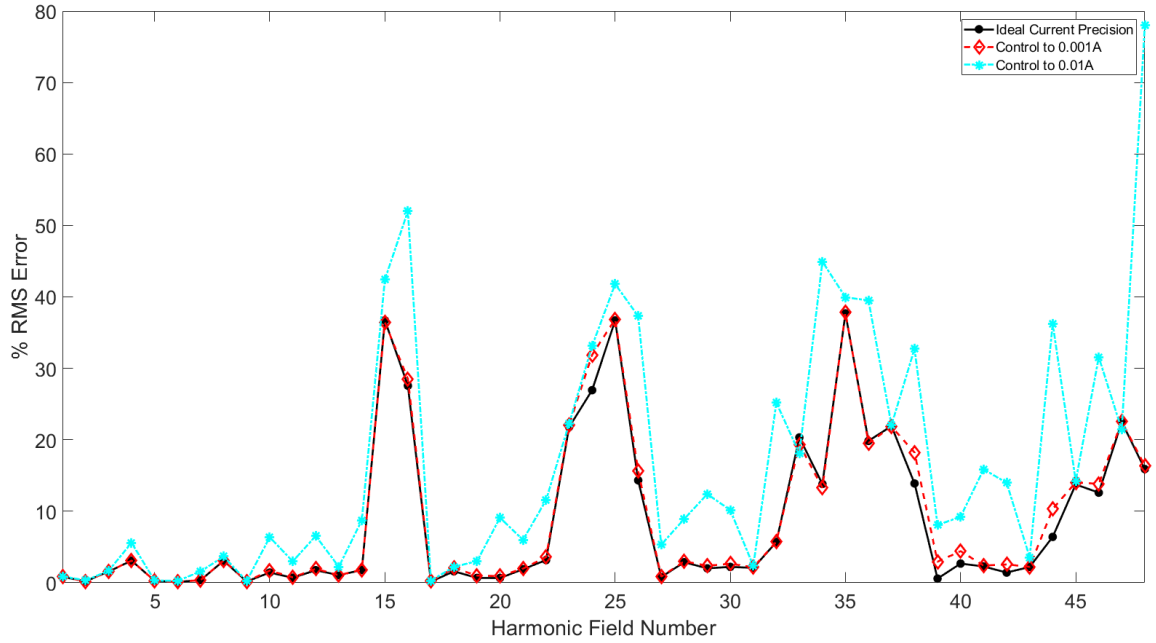


Figure 5.14: The variation in RMS fidelty across the range of individual magnetic fields with fluctuations in the precision of the supplied current. There is little distinguishable difference between an idealised system and one precise up to 4 decimal places relative to a maximal current of $\pm 1A$ (max deviation of $\pm 0.1\%$), so the data is excluded for ease of interpretation.

The high condition number for the inversion problem means that the final results are highly-nonlinear with respect to λ. One consequence of the subsequent ill-posed nature of the problem is that the repeatability of results is notoriously problematic, with dependencies on the available machine precision and specific hardware contributing to the final set of current ratios. The extreme variability is easily highlighted by taking the λ values and varying them by a scale factor, $10^{\pm 1}$. For 10^{-1}, the range becomes $2.6 \times 10^{-17} \leq \lambda \leq 4.3 \times 10^{-14}$ and the Mean Field Fidelity value falls to $39.50 \pm 23.14\%$. Adding further smoothing to the current distribution, shifting regularisation to $2.6 \times 10^{-15} \leq \lambda \leq 4.3 \times 10^{-12}$ changes the Mean Field Fidelity value to $38.22 \pm 26.77\%$. An overlay of the resulting fidelity plots for each magnetic field with this variation in regularisation is in Figure 5.15. From this data, it is apparent that a quasi-optimal value results from a uniformly scaled L-curve value for each magnetic field - the same scaling was used for the previous four coil setup and will be used for any upcoming systems. Some improvement occurs for the

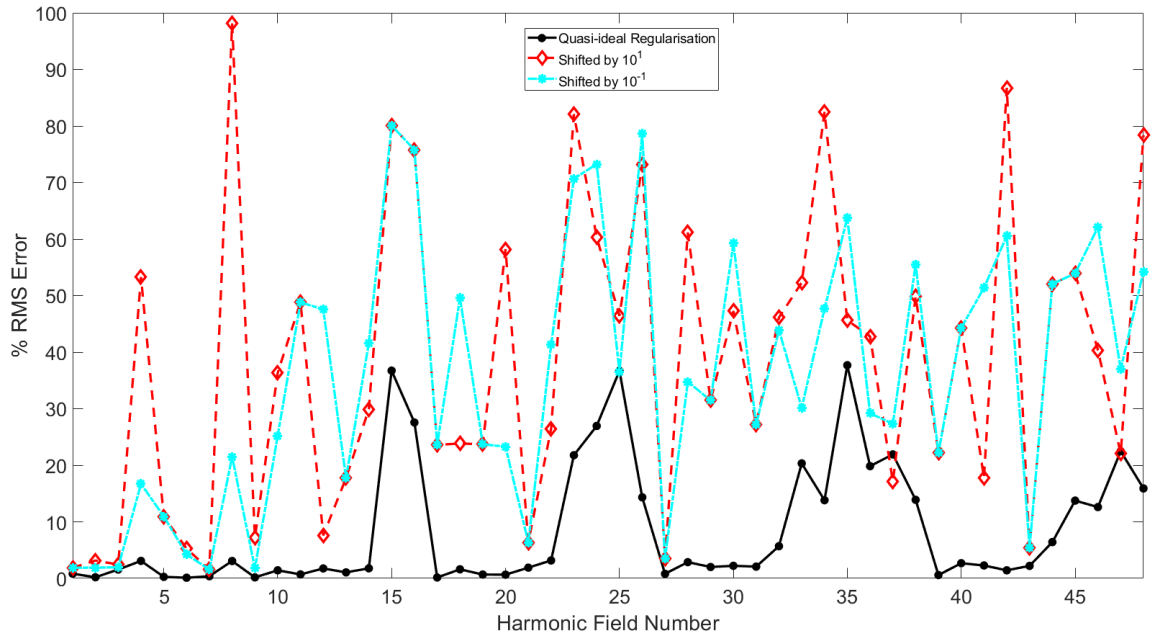


Figure 5.15: A plot concerning the variation in RMS fidelity for each of the individual magnetic fields when the the amount of regularisation is varied.

$n = 6, m = \pm 5$ profiles (fields 37 & 47) due to an increase in the regularisation suggesting some benefits are associated with manual refinement of λ. That said a *best* choice of value in this approach is dependent on: the number of coils, their shapes & spatial distribution (e.g inter-coil separation(s)), the target generation volume & the number of points sampled in that volume, as well as the previously mentioned uncontrollable variables such as specific hardware & available machine precision. Combined this means practically, especially for the system prototyping considered in this thesis, manual selection of the ideal λ value is incredibly inefficient, if not entirely impractical.

Beyond the impact on the quality of each magnetic field, another approach to the analysis of the influence of the regularisation parameter is the consider the physical impact in terms of the magnetic field strength produced per unit current - $\frac{B_0}{N \sum_k |I_k|}$ as a function of λ. Choosing a specific field - the linear $n = 1, m = 0$ for simplicity - B_0 is then the characteristic field strength in the region of uniformity, Figure 5.16.

The general trend is much as expected - increasing regularisation improves the power efficiency of the current distribution due to an increase in smoothing. In the central re-

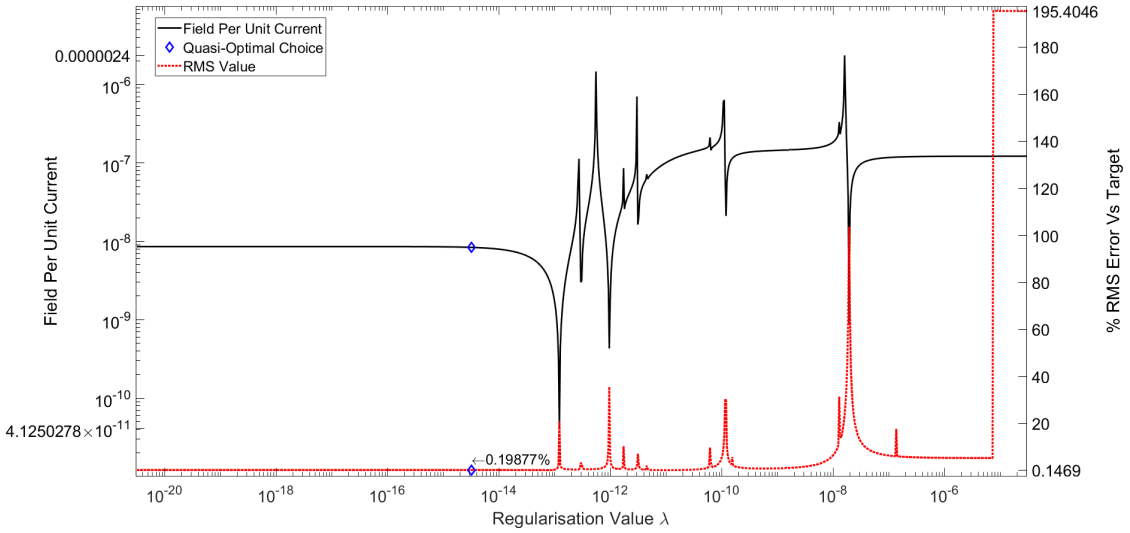


Figure 5.16: A plot concerning the variation in the field per unit current for the $n = 1, m = 0$ magnetic profile with different levels of regularisation. The left-hand axis gives the field per unit current data (Black Line), whilst the right-hand axis provides the RMS value relating to the target magnetic field (Red Dotted Line). The blue diamonds highlight the data analysed in the preceding sections of this chapter. The relative RMS value is also included with an arrow because it deviates from the limiting values on the y-axis.

gion, the oscillations identify modes which are more/less power efficient due to inter-coil interactions. The increased power efficiency typically comes at the cost of reduced fidelity, as a smooth distribution is unable to reproduce the target magnetic field profile accurately. This plot also identifies that the RMS value can be improved slightly by further varying the amount of regularisation used, though this improvement is just 0.05% and so is not likely to significantly impact the orthogonality properties of the specific magnetic field.

In conclusion, the presented system consists of thirty-eight modular coils, primarily of the saddle type. It can produce forty-six unique spherical-harmonic type magnetic field profiles, with the majority accurate depictions of the target. This is achieved by simply varying the currents through individual coils, assuming an ability to control current to a precision of at least $\pm 0.001A$ relative to a normalised max value of $1A$. λ has an expected highly non-linear dependence, and the choice is incredibly influential on the orthogonality properties of the final generated profiles, so selection requires care.

5.1.3 Theoretical System 2: Closed-Cylindrical Coil Array

Next, the induction system explored here considers a slightly more restrictive topology, with several coils enclosing each end of the cylinder. This small sub-array intends to provide greater transverse control of the magnetic field profile, leading to a higher fidelity on some magnetic fields. Potentially, the application of this design is similar to the previous cylindrical one, where it is easy to imagine sliding ends allowing objects of interest to be placed inside the system, before they are replaced, and interrogation occurs.

This section further highlights the versatility of the current distribution calculations when considering variable excitation geometries in a single system. It also demonstrates that the framework established in the previous two sections is robust and presents a design capable of producing forty-eight near-orthogonal magnetic fields, each being a good approximation to a particular spherical harmonic function.

To settle on a final design an iterative procedure explored a series of potential coil distributions and their subsequent fidelity for the entire set of magnetic fields. The overall volume is identical to that in Figure 5.6, with the same field generation space such that the two problems are directly comparable. Recalling the Mean Field Fidelity value already introduced, a series of coil distributions can be compared by varying the number of coils across/around the three distinct axes, and then performing an analysis with that value. The result is saddle coils of various arc and side lengths and, for simplicity, symmetric distributions of either: a single coil, a three-by-three, or a three-by-three grid of circular elements at the end-caps.

A total of 192 arrays were examined, with some designs containing an equal number of coils. For example, a system with fourteen induction coils may use two circular end coils, a discretisation of four around theta, and three along the axis. Denoting this structure as [2*1,4,3], different systems of fourteen could also be constructed as [2*1,3,4] or [2*4,2,3]. In seeking a coil design which produces the highest-accuracy spherical harmonic

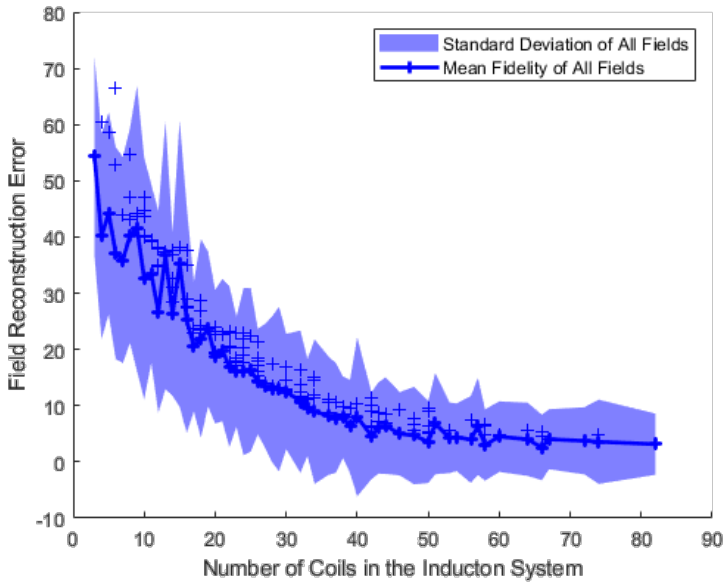


Figure 5.17: An iterative procedure exploring the ability of a discrete cylindrical coil structure, composed primarily of saddle coils, to reproduce a target set consisting of the first forty-eight spherical harmonic magnetic fields, using Equation (3.23).

field generation across the entire set, typically overlap analysis (Equation (3.15)) and inversion calculations are then confined to the design with the lowest RMS value. Figure 5.17 holds the complete data set, covering all 192, with the associated Mean Field Fidelity and standard deviation across the entire set of forty-eight magnetic fields.

In general, the trend here follows expectations. With an increasing number of coils, there is an increase in magnetic field reconstruction accuracy, coupled with diminishing returns as the size, and therefore region of influence, of individual saddle coils decreases. Various designs are of interest depending on the specific application considered. For example, seeking the simplest capable of producing a high-quality approximation to the set of forty-eight magnetic fields, one could consider the design with the least amount of coils with a Mean Field Fidelity below 10%. In that case, two, structured as [2*4,5,5] and [2*9,5,3], both achieve this aim with thirty-three coils. The former achieves this with a slight increase in field accuracy, $9.67 \pm 8.60\%$, whilst the latter uses fewer saddle coils with slightly worse accuracy, $9.71 \pm 10.07\%$.

The theoretical aim of this thesis is to produce a set of perfectly orthogonal magnetic

fields from a modular induction system, for which the design with the lowest possible Mean Field Fidelity is sought. From Figure 5.17 this is one using sixty-six induction coils, structured [2*9,6,8], with a Mean Field Fidelity of 2.2% which is a significant improvement on the previous design. For this design, Equation (3.23) is again ill-posed with a condition number of 2.3×10^4. The fidelity plot for all forty-eight targeted magnetic fields

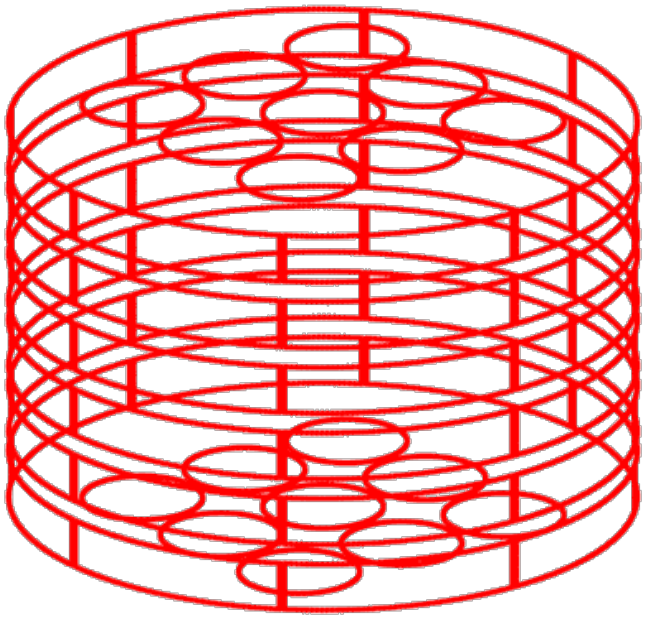


Figure 5.18: An induction system designed for a security application of MIT, and derived from the iterative procedure in Figure 5.17, consisting of sixty-sixty induction coils - forty-eight saddle coils and eighteen circular coils structured [2*9,6,8]. This system again assumes individual current control for each coil, allowing a set of excitation patterns to be produced relating to multiple quasi-orthogonal magnetic fields.

is in Figure 5.19, with only six having an RMS greater than 5%, compared to eighteen in the equivalent thirty-eight coil system. The regularisation applied uses the same scaled L-curve approach discussed previously, leading to a range of $4.0 \times 10^{-17} \leq \lambda \leq 6.0 \times 10^{-15}$. The upper value of 6.0×10^{-15} relates to the $n = 6, m = 6$ field profile and was manually adjusted from the standard L-curve scaled value used for the other magnetic fields, which suggested an overly-smooth value of 1.2×10^{-13}. This overly-smoothed value resulted in a magnetic field which had significant overlap with a variety of the other generated profiles, most prominently with the $n = 6, m = 2$ field. The adjusted result still overlaps most significantly with that field but with a nominal overlap of at least two orders of magnitude

less (0.98 vs 0.007 for an overlap integral normalised to unity for identical profiles). The overlap values in Figure **??** represent a vivid improvement over the thirty-eight coil system.

Two magnetic fields allow for qualitative analysis & discussion in this section - one

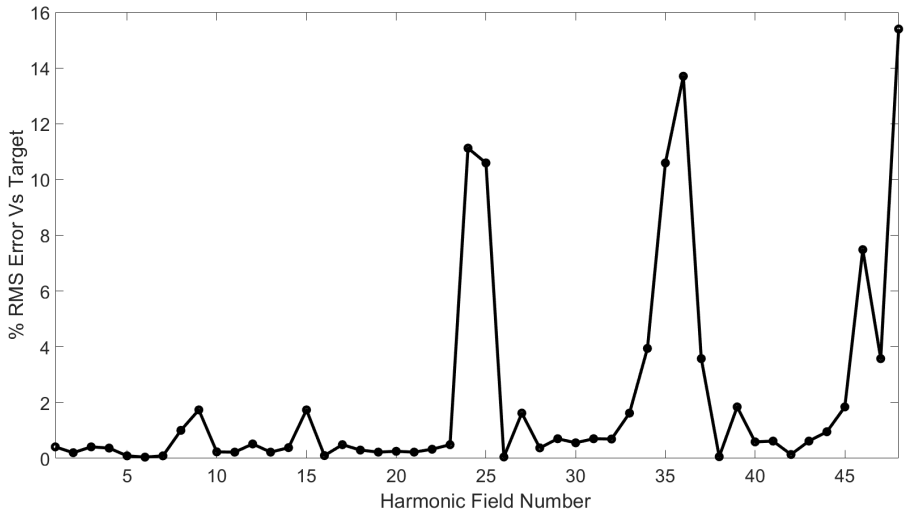


Figure 5.19: A plot of the various RMS values for each magnetic field produced by the closed-cylindrical induction system compared to the targeted spherical harmonic field. The Mean Field Fidelity for this data is 2.2%, and twenty-three are produced with an RMS less than 0.5%.

high-quality and one low-quality field. The low-quality field is easily identified in Figure 5.19, as the $n = 6, m = 6$ target spherical harmonic. It has a 15% error as shown in Figure 5.21, and is the only apparent outlier from an RMS evaluation of all forty-eight magnetic fields. The $n = 6, m = -4$ field is the best accuracy sixth degree magnetic field, with an associated RMS of 0.06% as shown in Figure 5.22.

As comparable systems, the fidelity values in Figure 5.11 for the thirty-eight coil system, Figure 5.6, can be compared directly to Figure 5.19. There is a substantial improvement in fidelity, with only the fields already accurately generated failing to improve significantly in the presence of additional current sources. For direct comparison of RMS values, the $n = 3, m = -1$ profile has improved from 0.70% to 0.22%, the $n = 4, m = -2$ profile from 1.6% to 0.3%, the $n = 5, m = -4$ profile from 14% to 0.05% and the $n = 6, m = 2$ profile from 6.4% to 0.95%. This identifies a clear benefit to the introduction of addi-

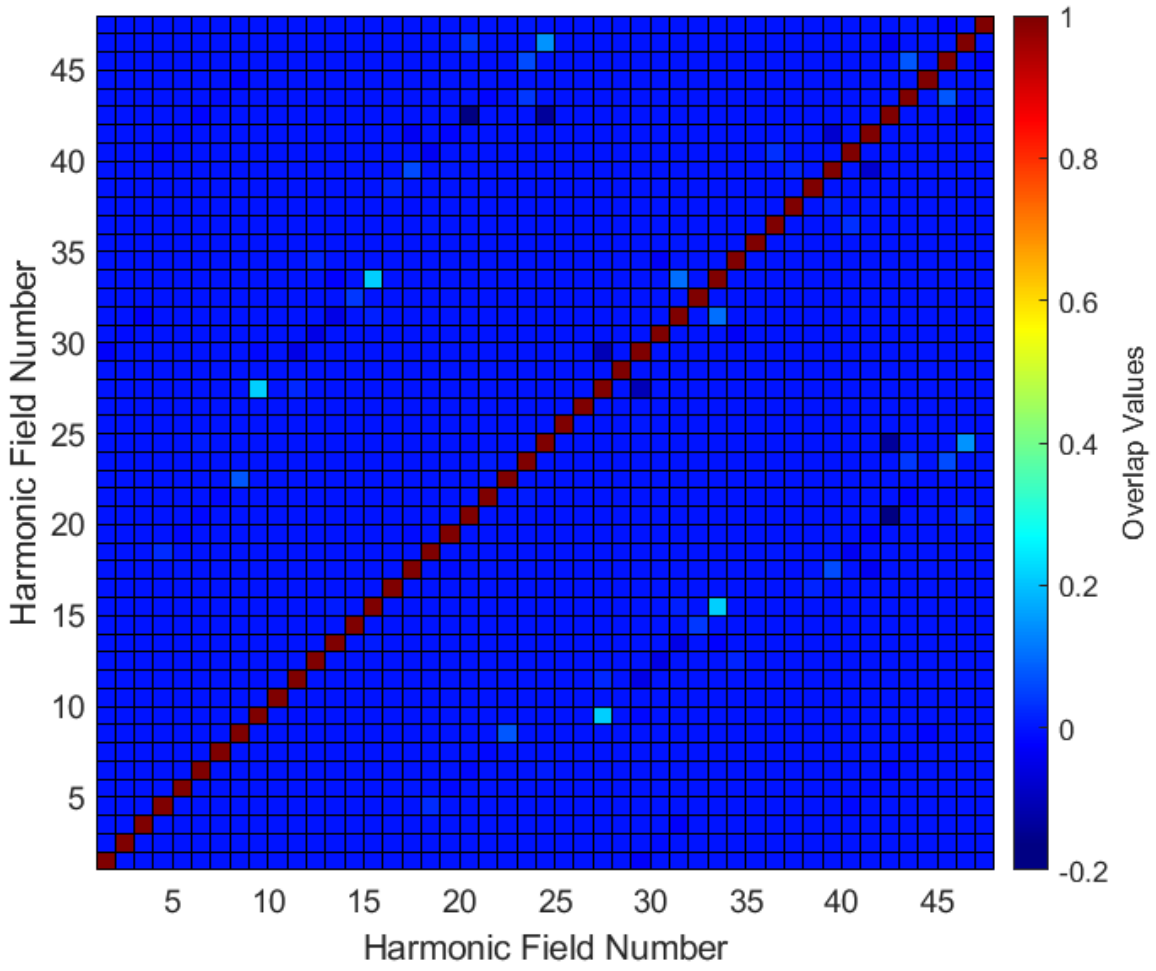
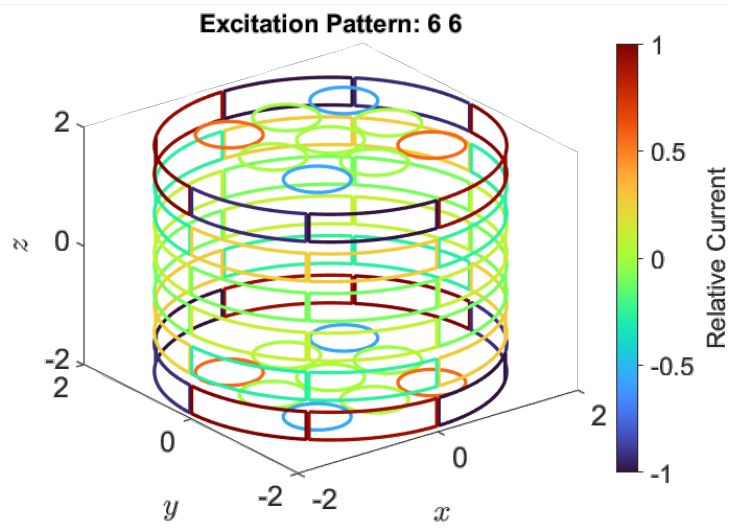


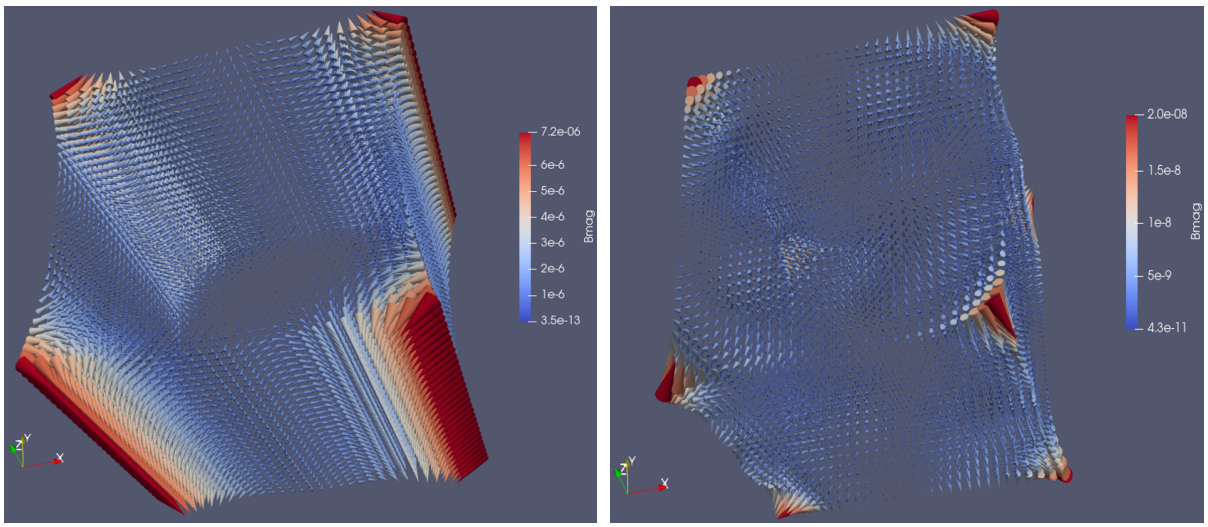
Figure 5.20: The values resulting from a permutative calculation of the overlap integrals between all forty-eight produced fields. Compared to the thirty-eight coil system, Figure 5.13, there is a vivid improvement in the orthogonality properties. With a most significant overlap of around 0.2, this set of magnetic fields should be ideal for investigating the influence of quasi-orthogonal interrogation in MIT with a system which should be relatively easy to construct.

tional current sources, balanced against practical requirements such as the development of a control system capable of handling precision control of roughly sixty percent more individual time-varying source channels.

The condition number scales proportionally, with a sixty percent increase in the number of sources corresponding to a sixty percent worsening of the conditioning of the problem. The size of this value has wider implications on the stability solutions from Equation (3.23) with respect to the value of λ as highlighted in Figure 5.15, suggesting a simple L-curve based Tikhonov method may be less appropriate for complex systems with many



(a)

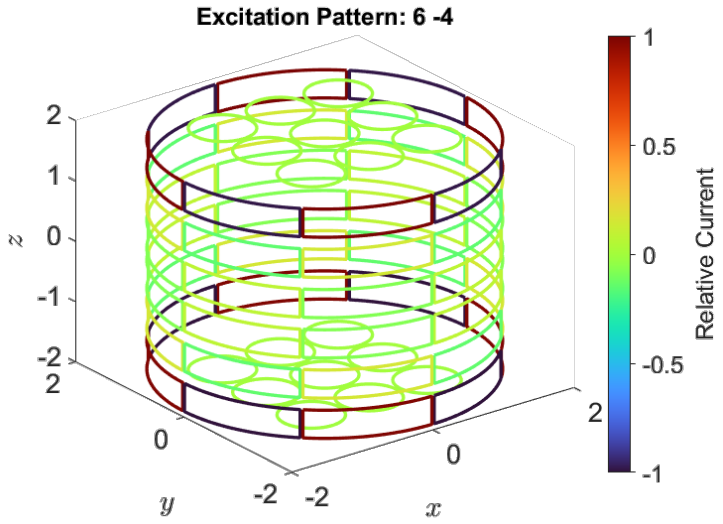


(b) (c)

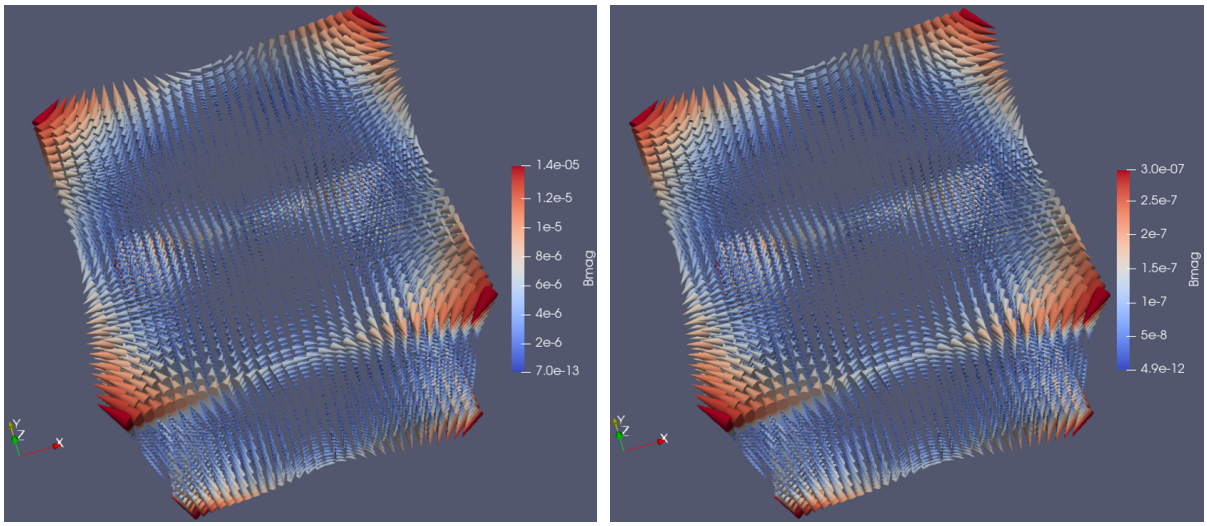
Figure 5.21: 5.21a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.21b) The target function related to the gradient of the real spherical harmonic n = 6 m = 6. 5.21c) The reconstruction from the excitation pattern with 15% RMS error, colour bar gives values in Tesla.

current sources. The obvious solution to this may be to apply a more nuanced inversion algorithm to minimise the initial least-squares problem rather than the simple smoothed Tikhonov identity approach used in Equation (3.23).

This is practically possible by utilising the existing codes associated with this thesis, and would likely represent an avenue of interesting future research, particularly as the investigation moves away from rapid design prototyping. Many of the algorithms in Chapter 4 will be applicable, but something which samples a finely-tuned parameter space like



(a)



(b) (c)

Figure 5.22: 5.22a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.22b) The target function related to the gradient of the real spherical harmonic n = 6 m = -4. 5.22c) The reconstruction from the excitation pattern with 0.06% RMS error, colour bar gives values in Tesla.

Brakhage's ν-method may be the most obvious candidate for initial investigation. The formulation of Equation (3.23), and the subsequent link to the Biot-Savart law by a simple matrix operation (Equation (3.18)) means that any specific solution can be related to the various ranking methodologies' in a computationally efficient manner. The problem is then particularly well suited to iterative inversion algorithms because the intended outcome of a *perfect* solution is known before the inverse algorithm is applied.

Knowledge of the desired result allows, assuming sufficient computational resources, any

iterative algorithm to be run well past the point of semi-convergence, with the solutions at each iterate stored. These can then be ranked in post-processing and the best possible solution selected heuristically from those rankings to find an ideal current distribution to produce a specific field profile. This avoids the application of sub-optimal stopping rules or heuristic estimation of the number of iterations entirely.

Moving on to analyse specific magnetic fields, the low-quality result, $n = 6, m = 6$, is perhaps of most interest. When applying the same selection criteria for λ as used elsewhere (leading to $\lambda = 1.2 \times 10^{-13}$), the resulting profile has an RMS value of 35%, making it a more significant outlier for this system & worse than the result from the preceding thirty-eight coil system. The resulting magnetic field using that criteria matched the spatial distribution of the target well in the corners but poorly over the remaining volume, with specific issues relating to significant unwanted bulk responses on the faces of the sampled volume. Through manual adjustment of λ, the selected profile in Figure 5.21 is the best ranking option within the model constraints, with significant improvement to the RMS value and the system-wide overlap. Visually this is perhaps surprising because the spatial distribution of the *improved* profile deviates from the target field at most of the sampled positions, again highlighting the limitations of an RMS analysis in this field generation problem. The result also has a weaker peak field and a smaller range in strength across the volume whilst providing a lower field per unit current. A reduced field per unit current relates, in part, to the reason this field wasn't suggested by the L-curve, recalling that the regularisation term favours power efficiency.

Figure 5.22 is easier to evaluate. Following the previous design, this field fails to produce the correct variation in magnitude. The result is some functional overlap with the other components of this set of fields, peaking at 1.9% with the generated approximation to the $n = 4, m = -4$ harmonic. Given the harmonic symmetries, this is sensible but is small enough that the field can be classified as quasi-orthogonal.

Overall this design successfully produces a larger set of higher-quality spherical harmonic magnetic fields compared to the equivalent system in Figure 5.6. The improvement is generally quite significant, though three are produced with a slightly lower accuracy when applying the same regularisation criteria. The result is a system which can accurately generate all forty-eight targeted quasi-orthogonal magnetic fields. These provide a set of realistically generated high-accuracy fields, which are applied in the next chapter as the ideal excitations whilst investigating the influence of orthogonality in the MIT inverse problem.

5.1.4 Theoretical System 3: Hexagonal Array

This theoretical system was designed during an exploration of iterative improvements to the experimental four coil setup. Relative to an imaging problem built around a set of quasi-orthogonal magnetic fields, that system has two obvious flaws. One relates to the limited number & quality of the magnetic fields produced by the system, whilst the other concerns the small generation volume where those fields have the desired properties. Costs constraints and control system challenges confined this to a theoretical analysis, with evaluation of the first forty-eight spherical harmonic fields.

This section focuses on the impact of the generation volume in Equation 3.23. It considers whether a larger relative internal volume is feasible for field generation and the stability of the excitation patterns concerning that volume. It establishes advantages to dynamically produced excitation patterns, suggesting results should always be tailored to end-user requirements.

The design consists of a sequence of rectangular coils encompassing an elongated cylindrical region for a security application. The geometry, Figure 5.23, is composed of a total of eighteen individually controlled coils arranged into a three-by-six grid. In normalised units, individual coils have a length of two along XY and three along Z. Attempting to quantify the intended improvements over the equivalent four coil design, the evaluation

of the field generation problem uses two stages. In stage one, the system targets an equivalent evaluation volume. In stage two, that volume is increased, and the stage one excitation patterns are contrasted with some dynamically produced results.

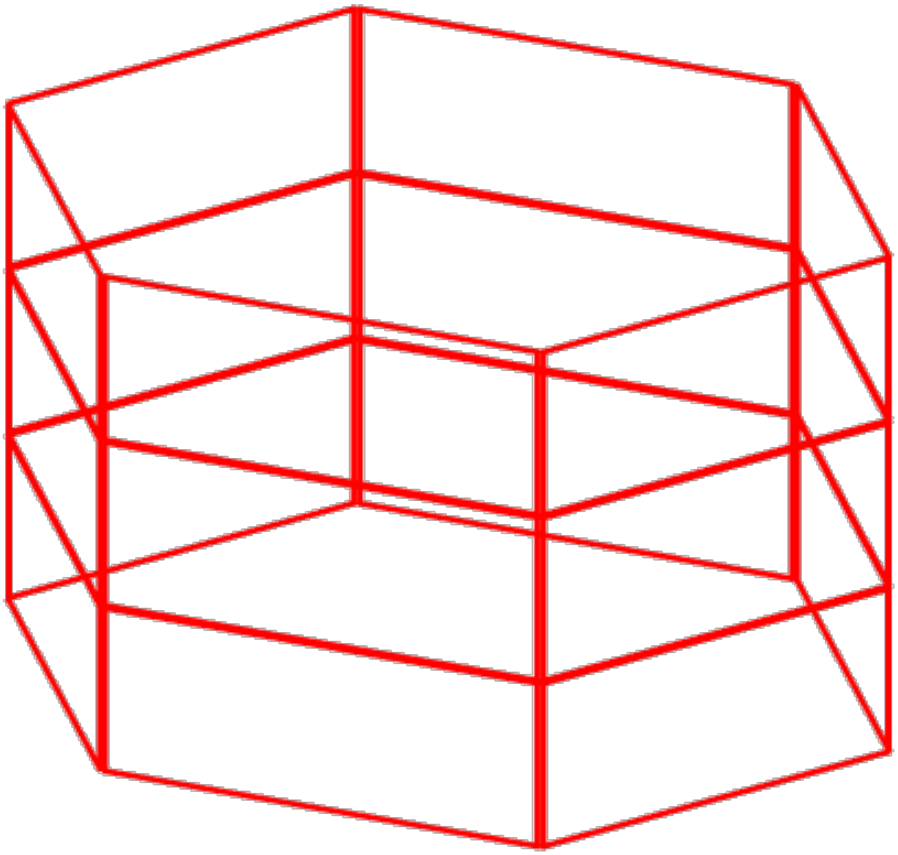


Figure 5.23: An induction system constructed from a three-by-six grid of rectangular coils.

For stage one, Equation (3.23) is then applied to a centred internal cuboid shaped volume with dimensions again set to be twenty-five percent those of a single inductor, $[0.5 \times 0.5 \times 0.75]$ in normalised units. The condition number for this system is 2.4×10^3, which using the same scaled L-curve approached discussed in the previous designs, leads to a range of regularisation values from $1.1 \times 10^{-16} \leq \lambda \leq 6.0 \times 10^{-14}$. The fidelity values are found in Figure 5.24.

In stage two, the volume is increased significantly to $[1.4 \times 1.4 \times 6]$ in normalised units. This is an 8-fold increase in the Z volume and a roughly 3-fold increase in the XY volume.

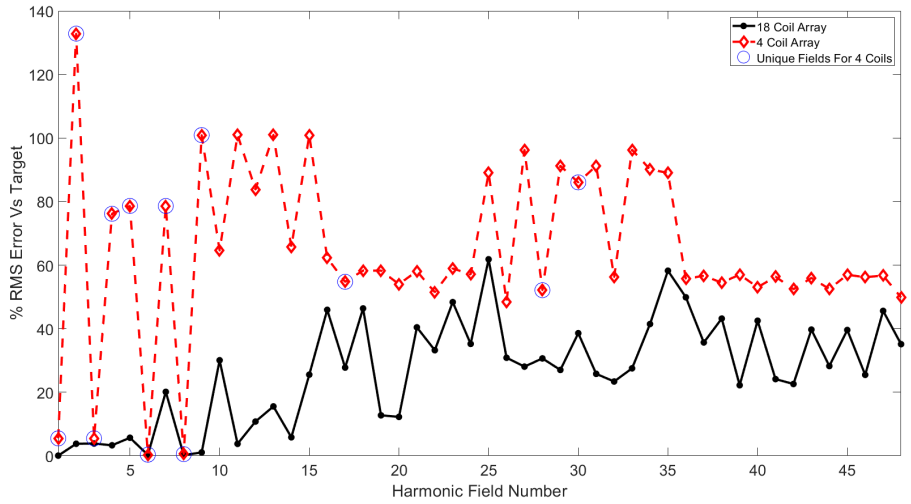


Figure 5.24: The distribution of RMS fidelity values, relative to the target harmonics, for this system (black) and the four coil setup (red), concerning the stage one analysis. Blue circles highlight the *useful*/unique fields generated by the four coil array for easier comparison. The Mean Field Fidelity for this system, including all forty-eight magnetic profiles, is 27%.

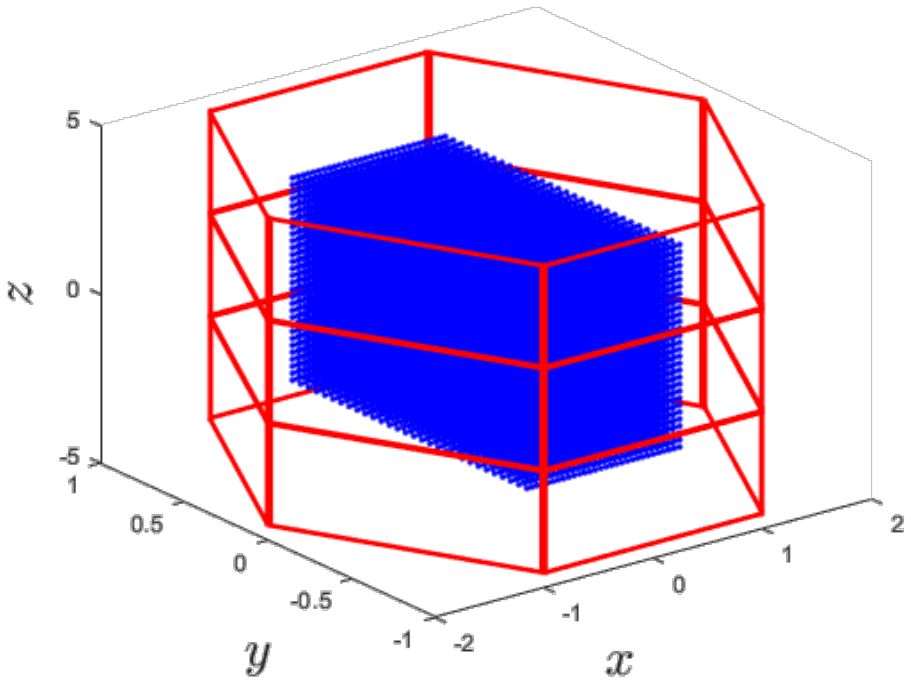


Figure 5.25: A view of the induction geometry including the region over which currents are calculated in stage two to generate target magnetic fields, and a set of normalised axes over which it is evaluated.

Relative to a single inductor, this is double the Z extent and seventy percent the XY extent, whilst compared to the volume of the scanner, this becomes $[35 \times 70 \times 60]\%$, with the squashed X-Y extent relating to the distribution of coils highlighted in Figure 5.25. The resulting contrast of fidelity values is in Figure 5.26.

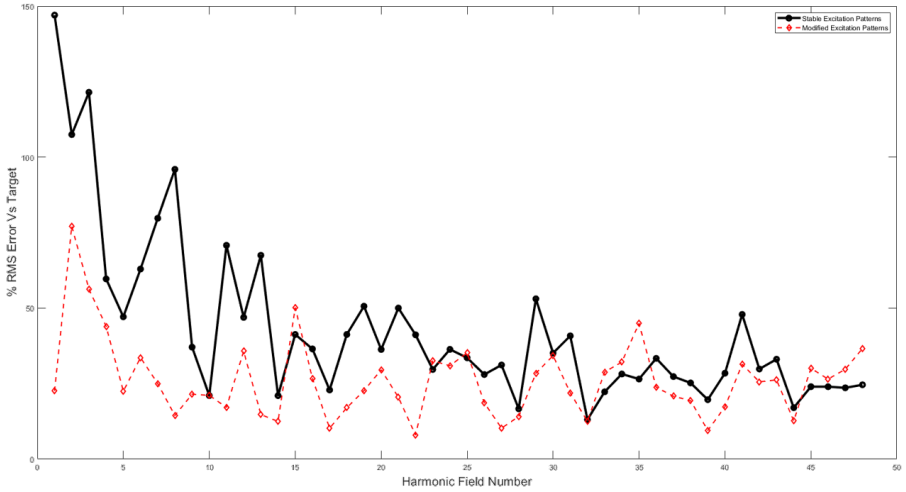


Figure 5.26: The distribution of RMS fidelity values for the hexagonal induction system in the stage two system analysis, considering excitation patterns for an increased interrogation volume. The static excitation patterns (black) lead to a Mean Field Fidelity of 43%, whereas the dynamically produced excitations' (red) reduce this to 26% for regularisation values in the range $2.4 \times 10^{-14} \leq \lambda \leq 7.4 \times 10^{-14}$.

The stage one excitation patterns are evaluated using the same process detailed for the thirty-eight coil cylindrical system, considering the RMS, the distribution of coefficients from Equation (3.25), and normalised overlap integrals between the generated results. This identified twenty-nine unique profiles, which have a Mean Field Fidelity value of 22% and are identified in Figure 5.27.

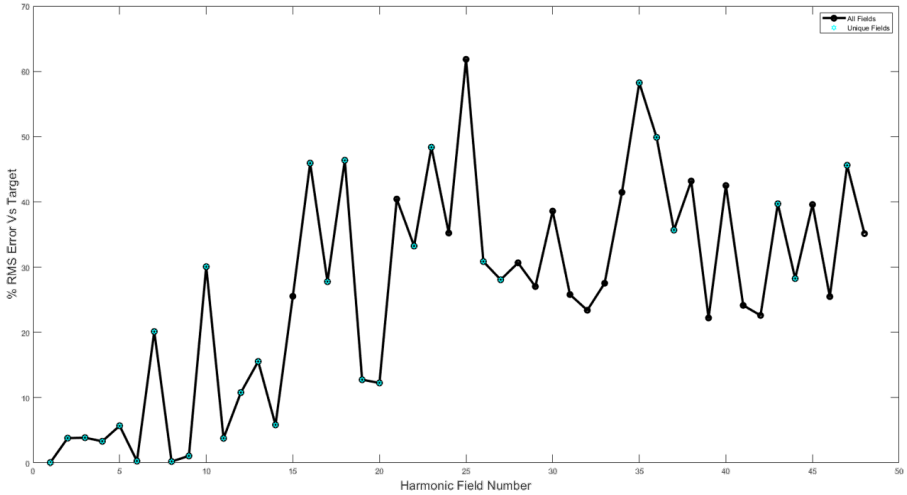
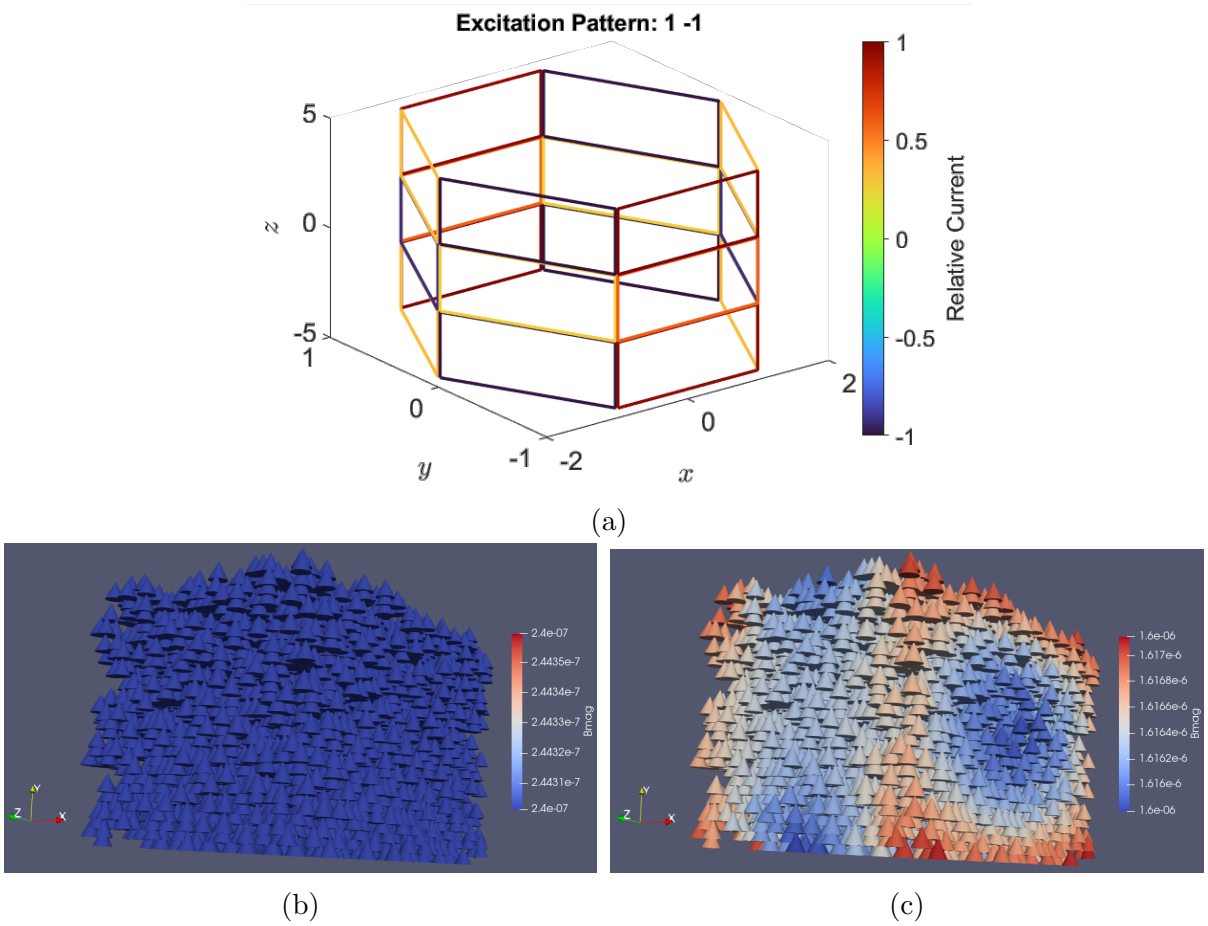


Figure 5.27: A plot of the various RMS values for each magnetic field produced by the hexagonal induction system in the stage one volume compared to the targeted spherical harmonic field, with the unique profiles identified by cyan stars. The Mean Field Fidelity for the unique data from the 18 coil array is 22%.

The highest quality magnetic field from the stage one results is the $n = 1, m = -1$ harmonic with an RMS of 0.06%, demonstrating an improvement of nearly a factor of ten in RMS compared to the equivalent field from the four coil model. The lowest quality unique field is (loosely) related to the $n = 5, m = 5$ harmonic with an RMS of 58%. The excitation pattern, target, and generated magnetic fields of the best & worst reproduced spherical harmonic fields from the filtered set can be seen in Figures 5.28 and 5.29 respectively.

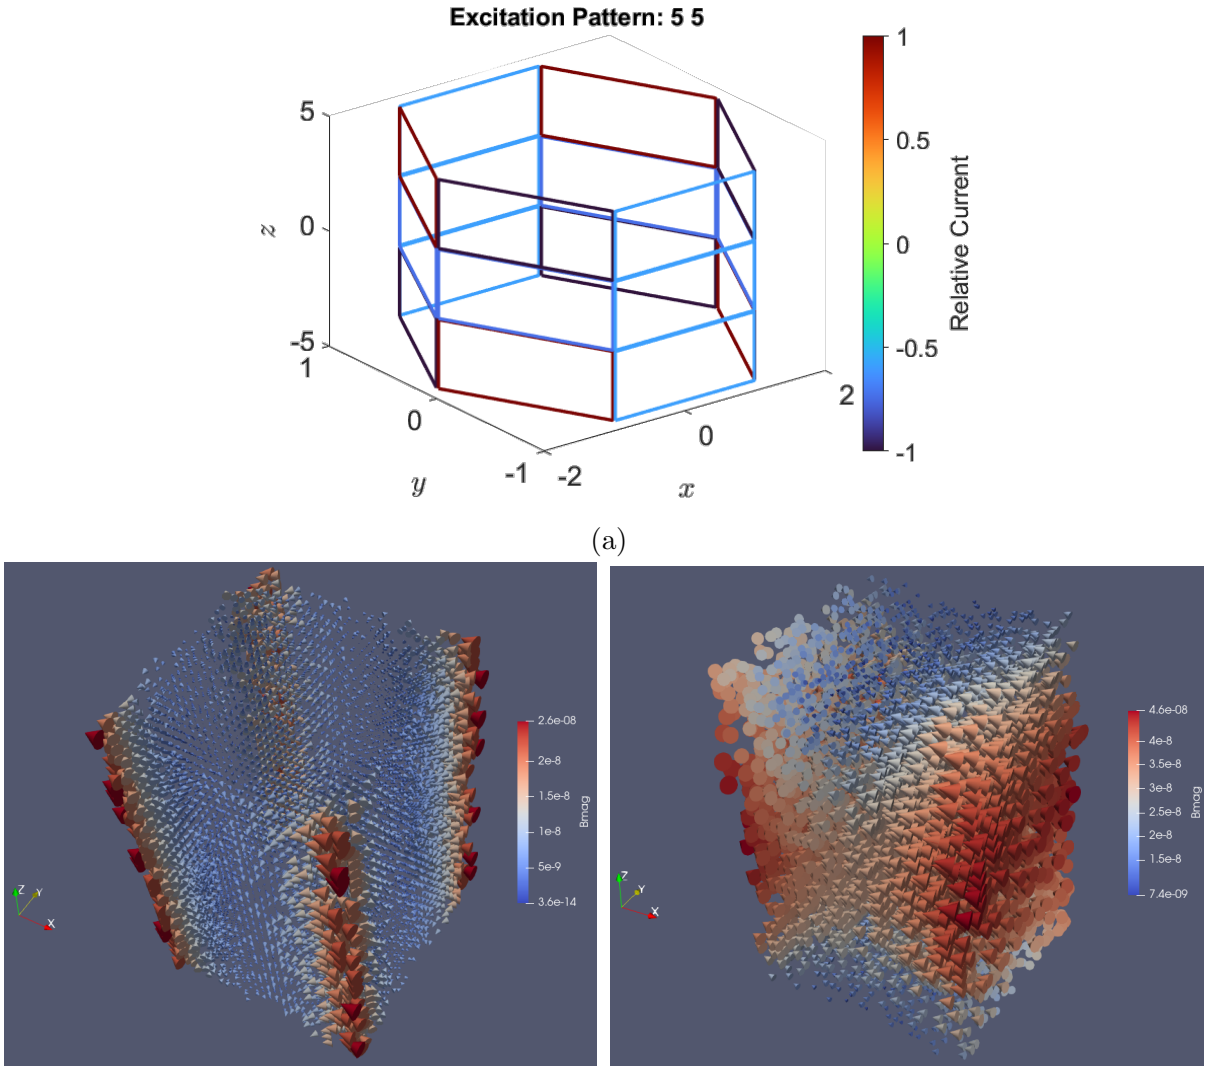


(a)

(b) (c)

Figure 5.28: 5.28a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.28b) The target function related to the gradient of the real spherical harmonic n = 1 m = -1. 5.28c) The reconstruction from the excitation pattern with 0.06% RMS error, colour bar gives values in Tesla.

For stage 2, both sets of excitation patterns, Figure 5.26, analysed through the tolerance-based coefficient analysis suggest the potential to produce thirty unique magnetic profiles



(a)

(b)

(c)

Figure 5.29: 5.29a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.29b) The target function related to the gradient of the real spherical harmonic n = 5 m = 5. 5.29c) The reconstruction from the excitation pattern with 58% RMS error, colour bar gives values in Tesla.

over a significantly larger volume - though the specific thirty targets are different in the two problems. Given the pre-existing analysis on the 'unchanged' excitation patterns, combined with the need for brevity and the clearly improved Mean Field Fidelity values, magnetic field results are presented only for the modified excitation patterns (red dotted line in Figure 5.26). The unique profiles in that data are identified in Figure 5.30.

The highest quality unique magnetic field from the stage two results is the $n = 4, m = 2$ with an RMS of 8.0%, Figure 5.31. The lowest quality unique field is the $n = 1, m = 0$

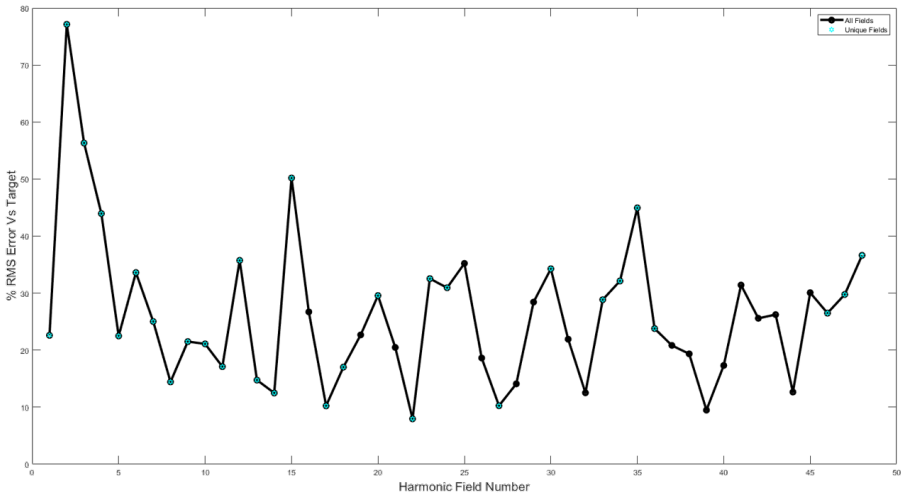
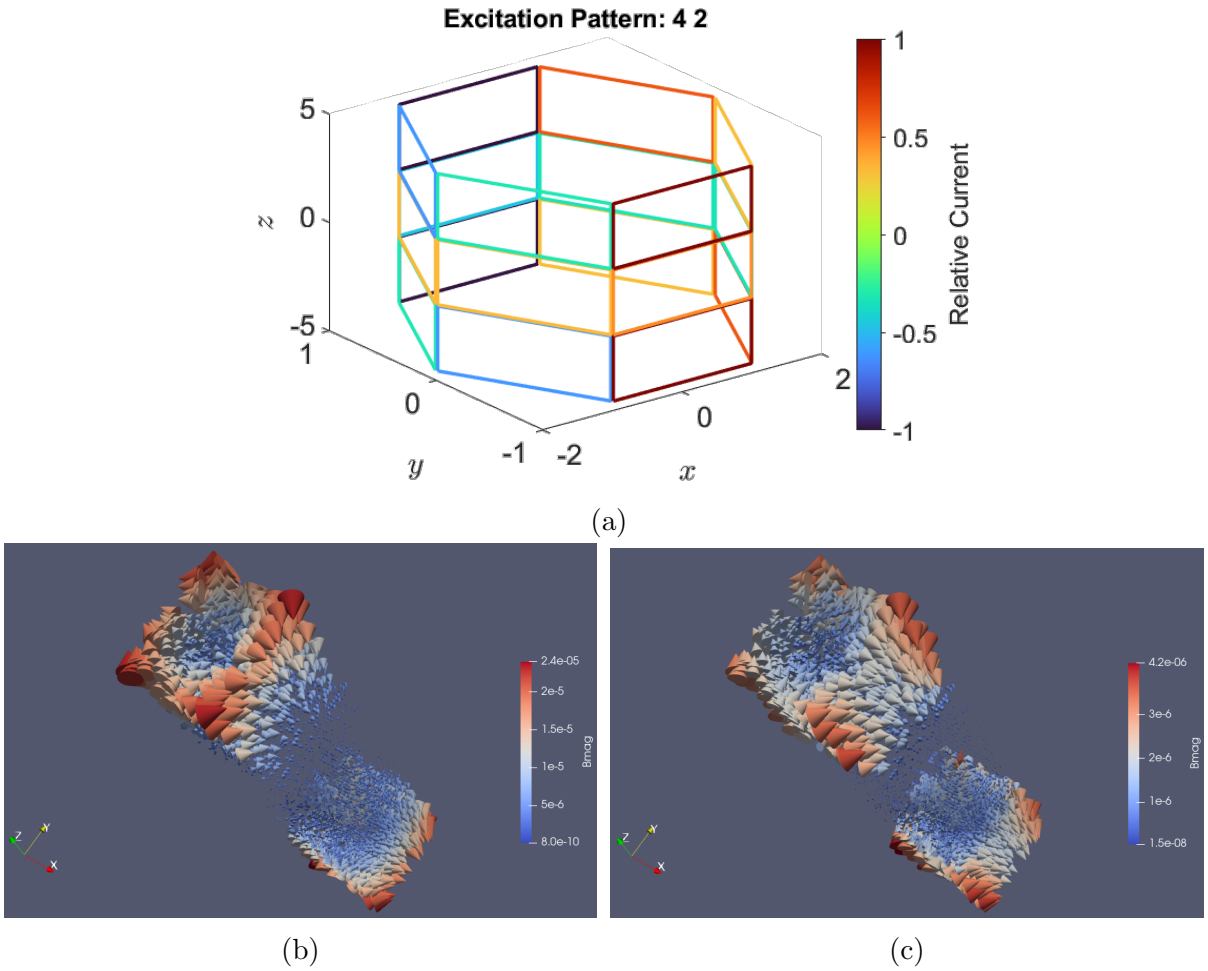


Figure 5.30: The RMS fidelity distribution for all target harmonics, as produced by the dynamically calculated excitation patterns with an increased volume, with unique profiles identified by cyan stars. The Mean Field Fidelity for the unique data from the eighteen coil array is 29%, which has increased compared to the analysis of the forty-eight fields because several low RMS profiles are not unique.

with an RMS of 77%, Figure 5.32.

Interpreting Figure 5.24 with considerations of the functional form of individual magnetic fields, this induction system is biased towards control of the magnetic field along the Y axis. Consider the $n = 1$ $m = \pm 1$ fields. As profiles with simple rotational symmetries, an unbiased system would produce appropriately symmetric results, such as those in Figures 5.5a and 5.5b. At higher degrees, the axial functional dependencies become non-linear, contributing significantly to the relative improvement in field reconstruction quality for certain profiles. For other harmonics, like $n = 2$, $m = 1$ (harmonic field number 7), the target is wholly independent of the Y axis, and this bias inhibits the field quality. It and the related diagonal coils also restrict the XY extent of well-controlled magnetic field.

The results for the stage one analysis identify some successes. The setup can produce over double the number of unique magnetic fields, though this requires many more individually controlled sources. Individual magnetic fields are generally of higher quality, with little RMS improvement only on the two second-degree profiles already well produced by four coils. The flaws of a sole-RMS analysis are clearly identified, with various relatively



(a)

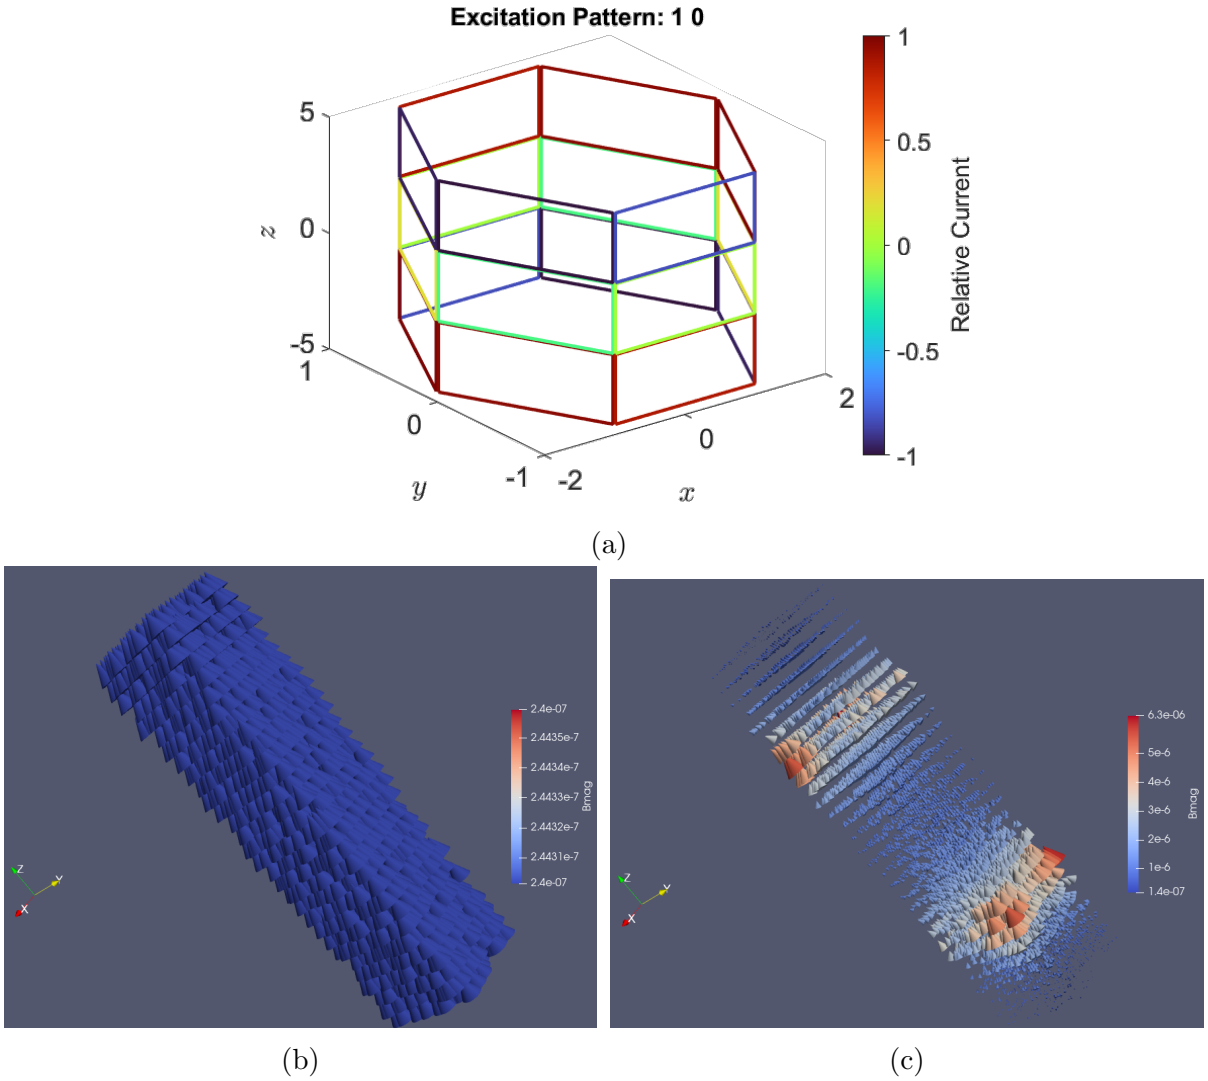
(b)

(c)

Figure 5.31: 5.31a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.31b) The target function related to the gradient of the real spherical harmonic n = 4 m = 2. 5.31c) The reconstruction from the excitation pattern with 8.0 % RMS error, colourbar gives values in Tesla.

low error fields failing to exhibit the expected quasi-orthogonal properties, whilst others with much higher RMS values are approximately unique. There is no obvious cut off, with the second lowest RMS magnetic field, harmonic 39, having significant overlap with at least one *simpler* profile.

The high-quality magnetic field, Figure 5.28, has few notable features and is highly uniform as is implied by the RMS value. The excitation pattern is of more interest. Axially opposing coil pairs behave in unison, and the excitation pattern has the expected symmetries predicted from extrapolating a Helmholtz system. The variation in current contributions is surprising, with *distant*, off-axis components clearly important, ensuring



(a)

(b) (c)

Figure 5.32: 5.32a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.32b) The target function related to the gradient of the real spherical harmonic n = 1 m = 0. 5.32c) The reconstruction from the excitation pattern with 77 % RMS error, colourbar gives values in Tesla.

the high uniformity relative to the four coil result. The result in Figure 5.29 is interesting because it is a low-quality spherical harmonic but exhibits quasi-orthogonality. Visual analysis suggests it is likely that there is degeneracy with the $n = 1, m = 1$ result, but these are vector quantities. This result acts in the opposite direction, leading to a normalised overlap value of approximately -0.9. As degenerate fields are already widely used in MIT, there is clearly some value associated with this result, vindicating its inclusion. The ten percent which doesn't overlap will also capture additional information, though quantifying the significance is for a future investigation.

Given the capability to produce more quasi-orthogonal fields, the next stage considers whether an expansion of the interrogation region is possible. The results, Figure 5.26, use the scaled L-curve values leaving space for future field refinement by varying λ. Contrasting the two data sets identifies a limited capability to produce accurate spherical harmonics inside a larger volume. The dynamic excitation patterns generally perform better, quantitatively highlighted by an improvement in the Mean Field Fidelity. In either case, thirty unique quasi-orthogonal fields are produced in a volume over sixty times larger, allowing the conclusion of success in the stage two analysis.

The high-quality stage two field, Figure 5.31, helps identify both the advantages and disadvantages of the greater relative z-volume. The central band of six coils can influence the centre of the volume somewhat independently, aiding the production of the null field region characteristic of higher degree harmonics. The outer regions also demonstrate a good spatial match, with the majority of the error fitting the theme of poorly controlled variations in magnitudes.

The low-quality result, Figure 5.32, behaves similarly. Individual layers act with some independence, and the spatial distribution has the expected uniformity in B_z. The massive RMS value relates to the significant variations in field strength. Interestingly this profile does lead to some magnetic field *hot-spots*, which perhaps makes the result more interesting than the target harmonic. Those regions may produce additional local information, aiding the characterisation of conductivity variations.

The independent action of coil layers means, over the extended volume, this problem may benefit from a different inversion algorithm. One which does not implicitly assume a smooth current distribution would be advantageous. Practically though, careful consideration is necessary because inductance effects between individual coils will be magnified by larger discontinuities, likely detracting from the final field quality.

In conclusion, this arrangement of eighteen rectangular coils can produce a quasi-orthogonal set of at least twenty-nine unique magnetic profiles. It demonstrates that a significantly larger interrogation region can be used, though this detracts from the quality of the intended spherical harmonics. There are advantages to dynamically produced excitation patterns, which practically should be set by end-user requirements before calculations are performed.

5.2 Induction Systems - Scanning Applications & Three-Dimensional Current Sources

This section considers a transportable coil design for applications where the region of interest cannot be confined inside a system, limiting the overall control of the magnetic field. It also considers the possibility of utilising solid-geometry induction coils with a spatially varying current density. The intention was to establish whether these could benefit the accuracy of spherical harmonic field generation.

In large-scale problems like sub-surface geophysics, physical restrictions confine the induction system to a plane above the region of interest. Control over the functional form of the magnetic field is possible only close to the system, and profiles requiring variations along the z-axis become significantly harder to generate. These problems also require considerations of portability, with large-scale scanning preferring a minimal number of coils.

A simple topology is a planar design composed of an array of coils which may be laid flat and scanned across a survey region. Given the application of the superposition principle, careful experiment planning may allow fewer coils to gather data if intercoil induction effects are negligible, but this would increase the survey time. Functionally, generating a set of spherical harmonics requires a more nuanced design. Odd distributions, like a five-by-five grid, are well suited to odd harmonics, $n = 1, 3, 5 \ldots$, whereas even distributions

are ideal for the even harmonics, $n = 2, 4, 6 \ldots$. The number of components again links to the maximum amount of generatable variations and approximates the highest possible degree.

5.2.1 Array Composition

One way to address the odd/even problem is with a co-planar arrangement, inclusive of some small but negligible buffer between layers. Using a variable structure enables a system with both odd and even symmetries, which should permit a broader set of accurate spherical harmonic fields. A simple arrangement of this type would be a two-by-two grid on one plane with a single coil on a separate plane. Interestingly, this distribution has a relatively low condition number in Equation (3.23), even with many elements.

Another iterative process in terms of the Mean Field Fidelity enables a quick evaluation of where the quasi-optimal design may lie, with the first twenty-four harmonics, up to $n = 4$, $m = 4$, targeted in Figure 5.33.

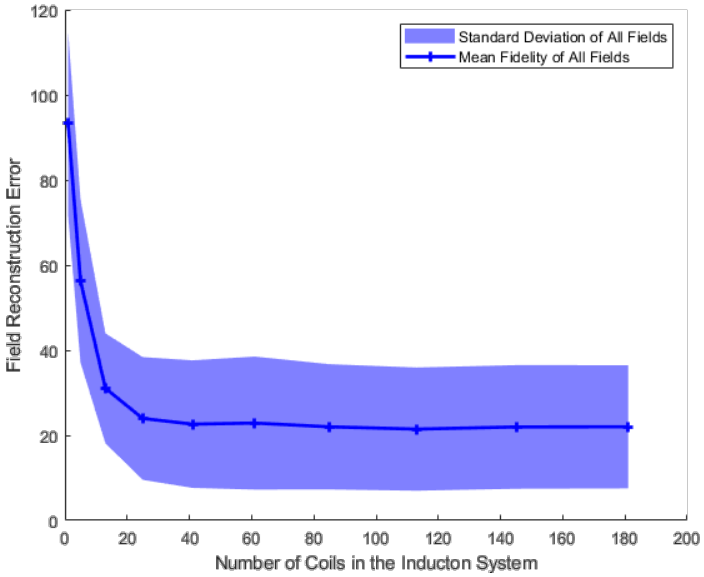


Figure 5.33: An iterative procedure exploring the ability for a planar square coil structure to reproduce a target set consisting of the first 24 spherical harmonic magnetic fields specified for Equation (3.23).

The Mean Field Fidelity plateaus rapidly, with little improvement apparent above the

use of forty-one coils - a coplanar structure with a lower layer of five-by-five and an upper four-by-four array of square coils. This, as well as the lack of variation in the standard deviation compared to the cylindrical setup, Figure 5.17, can be directly attributed to the form of the induction system well suited to the functional symmetries in the target set.

The negligible buffer between layers makes this design effectively planar, with all current density confined to an XY plane. This enables magnetic field variation along X or Y but introduces no additional Z control compared to a single current loop. Fields requiring Z-variations then become the dominant source of error, and control of the spatial field distribution is only possible up to depths approaching the size of the system, likely limiting investigations to very near-surface surveying in a geophysical setting.

From Figure 5.33, a quasi-optimal balance between maximising spherical harmonic field quality and minimising the number of coils is a thirteen-coil system. Two different models are then used to properly assess the potential of this induction system - Filament Current Sources and Solid Current Sources. These are intended to be roughly equivalent, with filament coils given a hundred turns and the solid coils designed to have a matching physical footprint. The driving potential for solid sources in Equation (3.29a) is scaled empirically, so the bulk current density resulting from Equation (3.28) is approximately equivalent.

To thoroughly compare and contrast these coil models, three magnetic profiles are presented for analysis - composed of the best and worst fields from Figure 5.39 and one additional result. The intention is to allow a fair discussion of the advantages, disadvantages and differences between the two coil models. Of interest is how the differences between the two models influence the final quality and composition of the targeted quasi-orthogonal magnetic field set, as well as the resulting magnetic field strength.

5.2.2 Three-Dimensional Current Sources

A system with three-dimensional current sources begins with a single component coil, Figure 5.34. As described in Section 3.5, a model uses a pair of coupled FEM problems to find the resulting magnetostatic field. Given a simple unit coil, the thirteen coil array

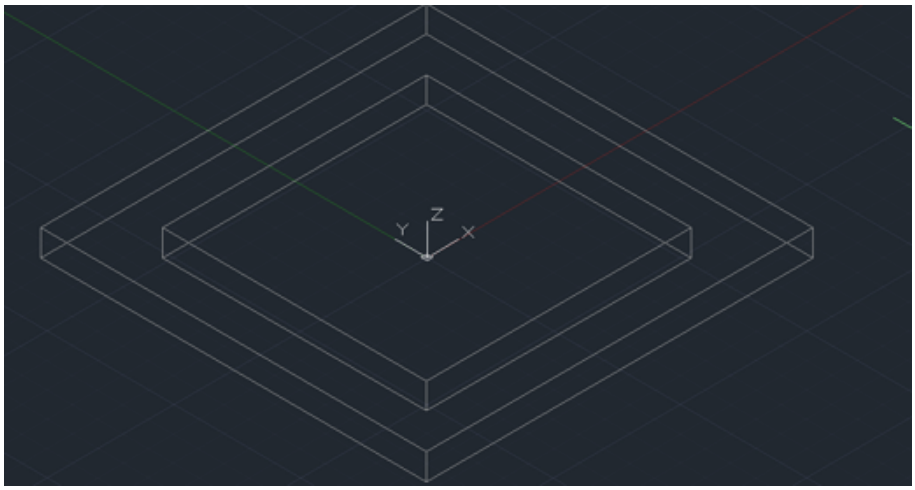


Figure 5.34: A three-dimensional current source used as the building block of the induction system. This is a hollow-cored square coil with outer side length of 2.5, inner side length of 1.5 and height (Z dimension) of 0.1 on normalised scales. This is then modelled as a solid block of copper through which a current is driven.

can then be constructed to match the source geometry in the filament coil setup. The resulting system is given in Figure 5.35.

To finalise an FEM model, a black-box simulation - i.e one contained in a volume large enough that the magnetic fields at its boundary can be assumed to be near-zero allowing the application of Dirichlet boundary conditions - is used with the associated geometry for the XZ plane given in Figure 5.36. A solution to Equation (3.23) results from mapping the induction system onto a two-dimensional setup, Figure 5.43, with the average dimensions of the three-dimensional array. The normalised current ratios are combined with the finite-element solution to Equation (3.26) on *cut coils* to specify the source current densities, leading to excitation patterns. The derivation of the driving potential comes from those ratios and Equation (3.30). Taking an XY view of a single coil, the value of d is then estimated empirically by assuming the average electron will take a roughly circular path around the coil as shown in Figure 5.37. This value is mesh-specific, defined for a

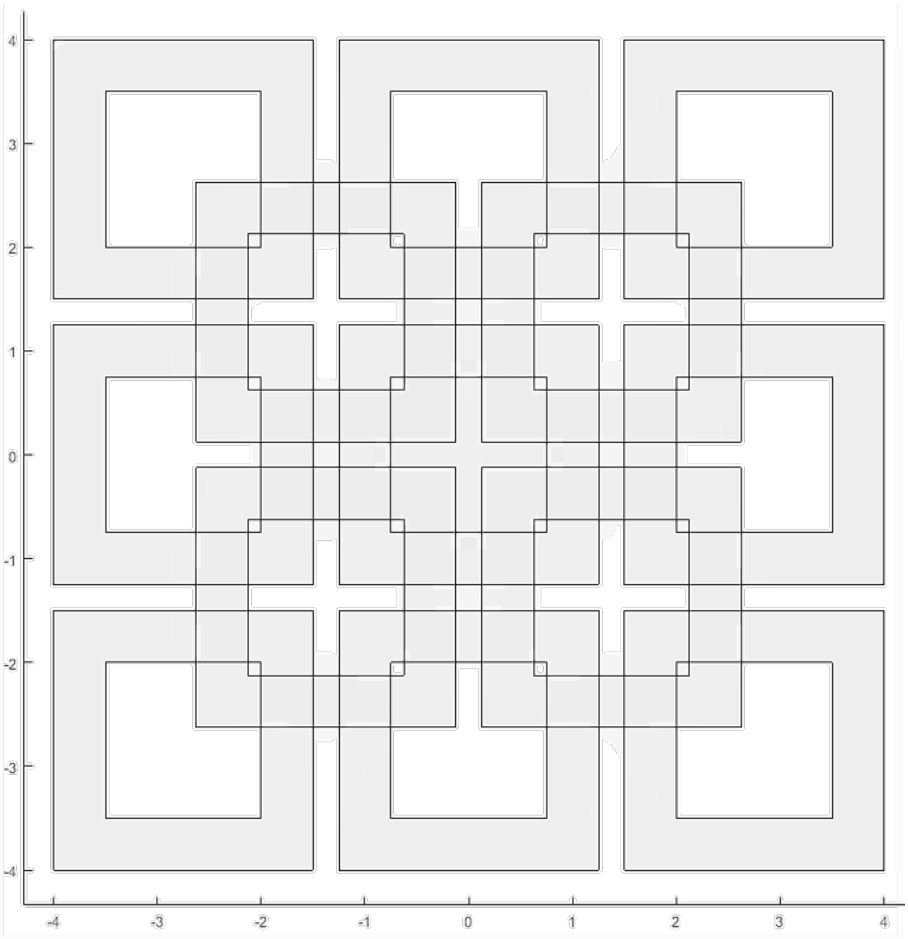


Figure 5.35: An XY view of the design constructed from thirteen solid unit coils. The inter-coil separation is at 10% of the outer side length. The two layers of coils are physically stacked, with the four coil layer placed directly on top of the nine coil layer with a total normalised height of 0.2.

bulk current density (cyan) equivalent to a reference filament coil with a hundred wire turns.

An exact excitation pattern, in terms of the magnitude of the electric field, for the optimised currents related to the linear $n = 1$, $m = 0$ magnetic field is given in Figure 5.38a. However, the relative contribution of each coil to the final magnetic field is ambiguous, and the direction of the current flow is unknown when plotting the magnitude. For ease of interpretation, it is easiest to use the average current density, weighted by the direction of current flow, Figure 5.38b.

Figure 5.36 contains fifteen conformal sub-domains meshed using the open-source GMSH

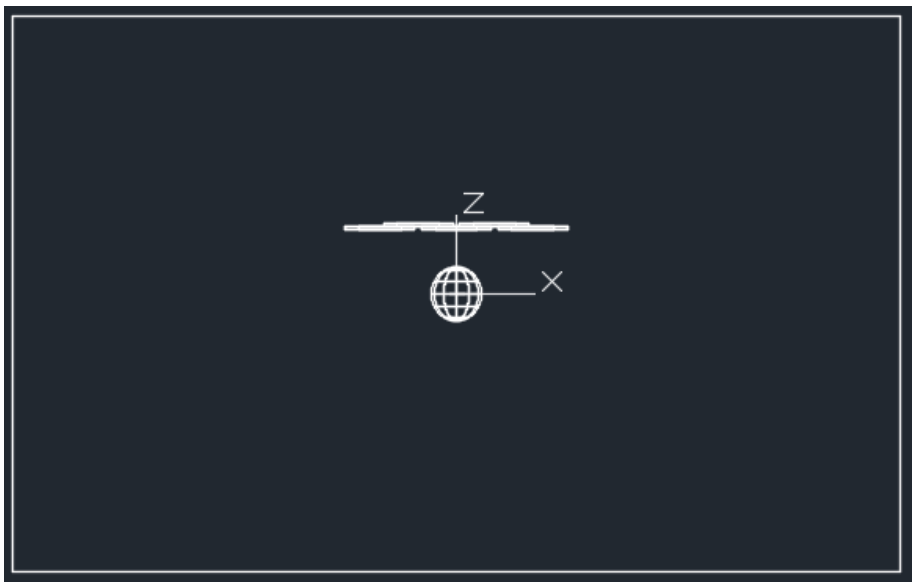


Figure 5.36: The complete geometry required to specify a 'black-box' magnetostatics simulation for the three-dimensional coil array. The spherical region specifies the field generation region for Equation (3.23), inside of which the analysis will assess the potential to produce a set of quasi-orthogonal magnetic fields.

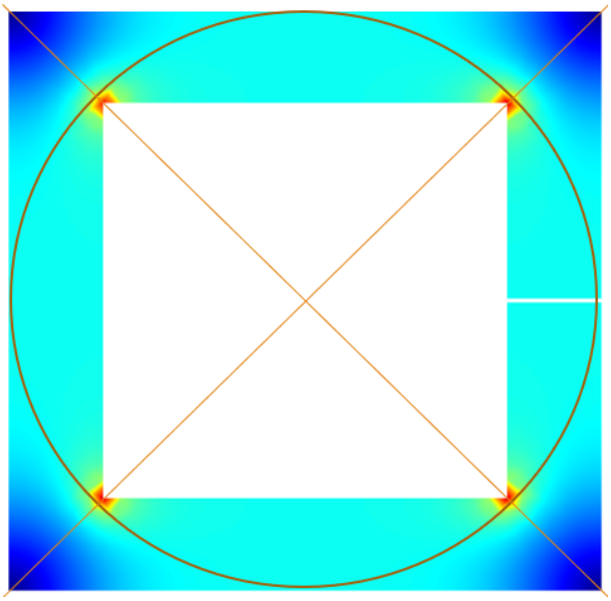
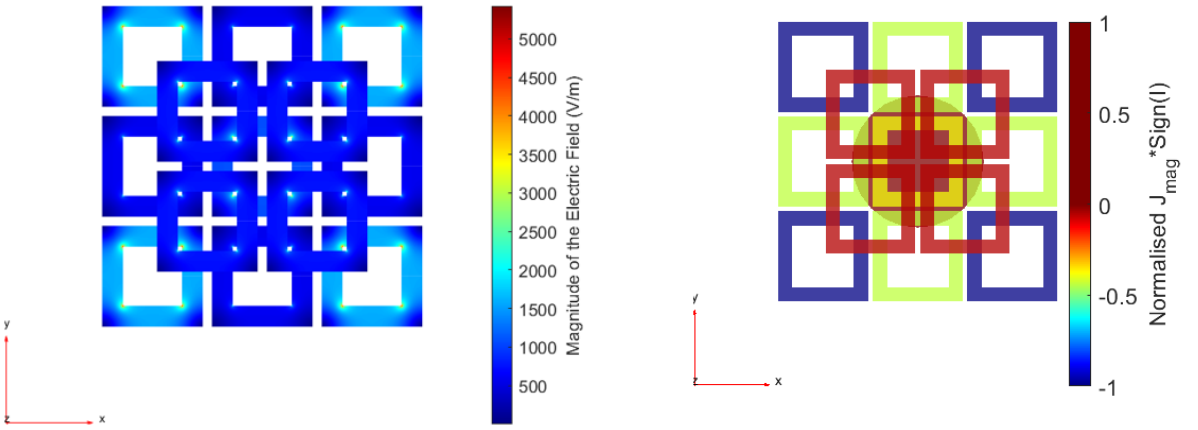


Figure 5.37: The path used to empirically estimate d in Equation (3.30) to ensure the generated magnetic field is approximately the same for solid and filament current sources, plotted over a *cut coil*. The colour map references current density, from minimal (dark blue) through average (cyan) to max (red).

software, which includes tools for fine-tuning and optimising the composition of the mesh. The spherical region allows the focusing of significant mesh refinements for computational ease.



(a) The Exact Excitation Pattern. (b) A Simplified Excitation Pattern, in A/m.

Figure 5.38: Two visualisations of the excitation pattern associated with the $n = 1$, $m = 0$ magnetic field for an array of solid copper coils.

To perform a consistent analysis the results at nodes within the prescribed spherical optimisation region are extracted and compared with the target magnetic fields. The distribution of fidelity values across the twenty-four target magnetic fields is in Figure 5.39 with a Mean Field Fidelity value of 33%

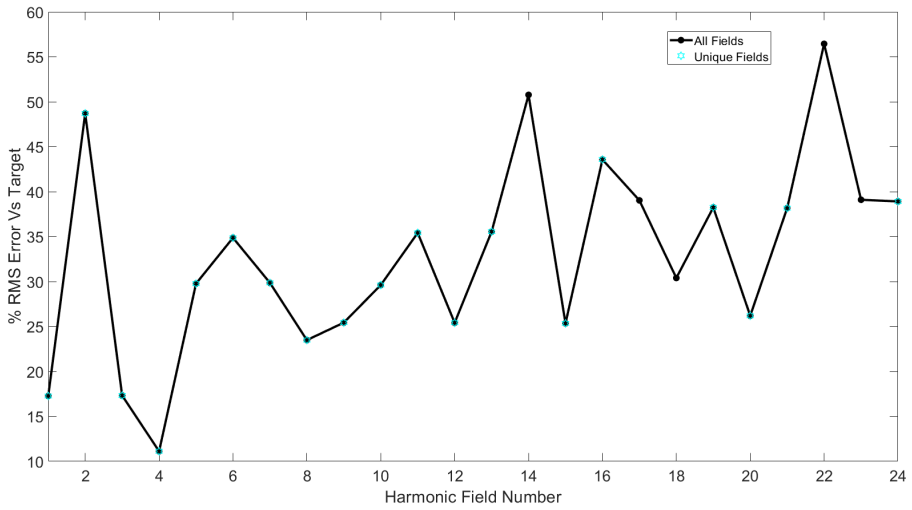


Figure 5.39: The distribution of RMS values for magnetic fields produced by a thirteen square coil induction system, with the unique profiles identified by cyan stars. The Mean Field Fidelity is 33% for all twenty-four profiles and is 30% for the unique data.

Recall the limited value in a lone RMS analysis of these fields. Using Equation (3.25) with a tolerance-based comparison leads to the conclusion that there are nineteen unique profiles from this thirteen-coil system. These are in highlighted Figure 5.39 with an ad-

justed Mean Field Fidelity of 30%. The non-unique fields can be grouped into four sets of fields with similar order, Table 5.1.

1:	$n = 2, m = 2,$	$n = 3, m = 2,$	$n = 4, m = 2.$
2:	$n = 3, m = -2,$	$n = 4, m = -2.$	
3:	$n = 3, m = -3,$	$n = 4, m = -3.$	
4:	$n = 3, m = 3,$	$n = 4, m = 3.$	

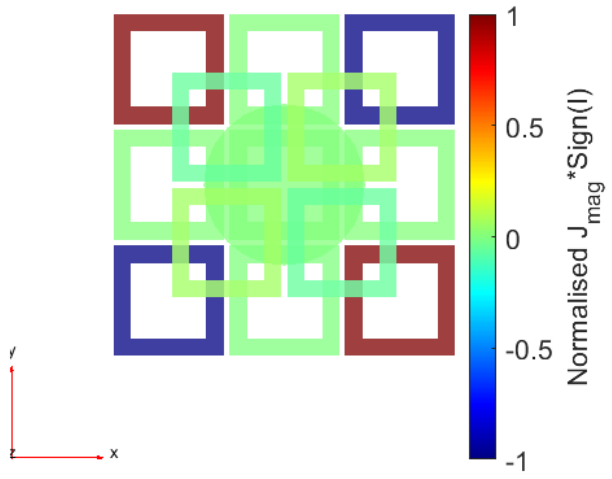
Table 5.1: Overlapping fields from the thirteen coil field generation problem.

The symmetries and the specific issues with generating fourth-degree SH field profiles match the theoretical expectations identified at the start of this section. For the three-dimensional array, the highest-quality field corresponds to the $n = 2$, $m = -2$ harmonic, Figure 5.40c, and the lowest-quality result is the $n = 1$, $m = 0$, Figure 5.41c. The other comparative field relates to the $n = 2, m = 1$ profile, Figure 5.42c.

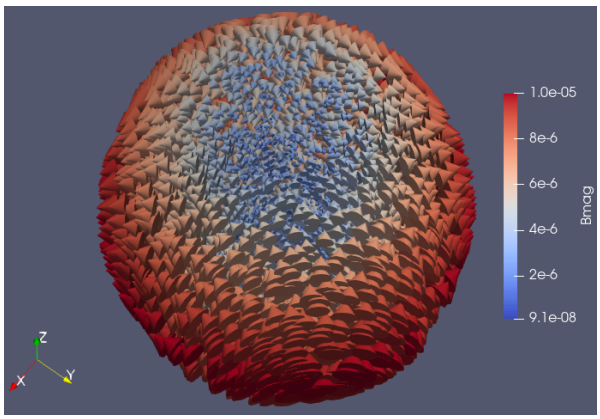
5.2.3 Filament Current Sources

The geometry of the comparative filament system, also used to find the excitation current distributions, is in Figure 5.43. Each is assigned one-hundred wire turns with currents in the normalised range of $[-1 \leq I \leq 1]A$. This analysis reuses the nodal coordinates from the spherical region associated with the three-dimensional array to allow a direct comparison. The distribution of fidelity values across the twenty-four target magnetic fields is in Figure 5.44, with a Mean Field Fidelity value of 29%. Applying the tolerance-based comparison to the results from Equation (3.25) identifies the same nineteen unique profiles, with a reduced Mean Field Fidelity of 26%.

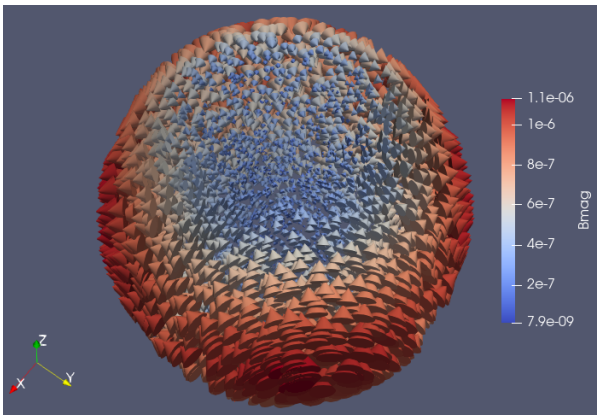
From the RMS values, the highest and lowest quality spherical harmonic fields differ using filament coils, with the highest quality reconstruction relating to $n = 1$, $m = 0$ at 6.1% RMS, Figure 5.47, and the lowest quality unique profile relating to $n = 4$, $m = -4$ at 53% RMS. Figure 5.46 gives the comparable result for the $n = 2$, $m = -2$ at 11% RMS, and Figure 5.47 gives the result for the $n = 2$, $m = 1$ at 33% RMS. For brevity, the excitation patterns are omitted from this filament coils section as the set of optimised *currents* are the same in both problems, only differentiated by the visualisation.



(a)



(b)

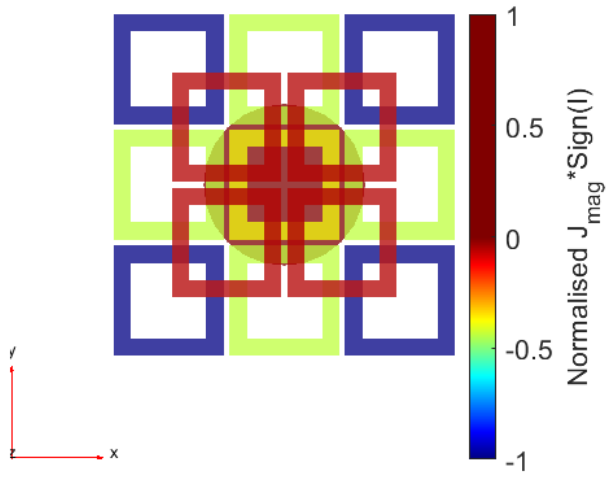


(c)

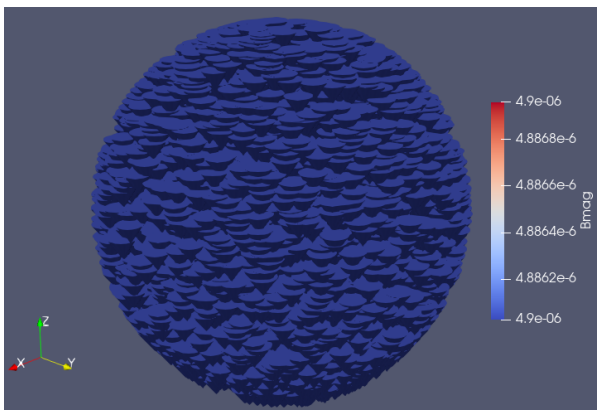
Figure 5.40: 5.40a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.40b) The target function related to the gradient of the real spherical harmonic n = 2 m = -2. 5.40c) The reconstruction from the excitation pattern with 11.1% RMS error, colour bar gives values in Tesla.

5.2.4 Comparison & Analysis

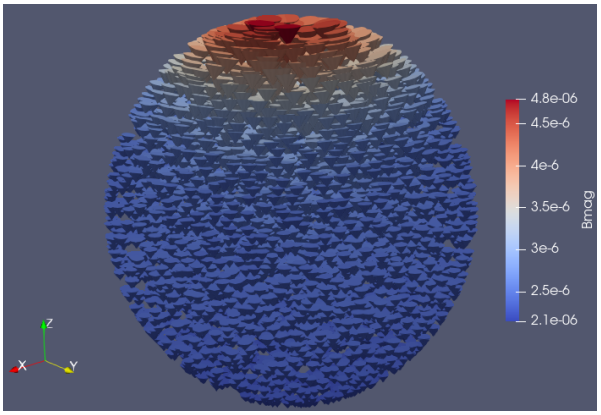
An initial comparison considers the RMS evaluation of field quality, Figures 5.39 and 5.44. The results are broadly consistent, with a slightly better Mean Field Fidelity from the filamentary system. The preference toward that system links well with the current density models. From Figure 5.37, each solid coil has some variance in the applied current density, which will impact the smoothness of the resulting magnetic field and introduce additional noise. Despite this, the RMS demonstrates marginal improvements for ten magnetic fields, eight of which are unique. These are usually directly-symmetric harmonics of odd order, with six paired results at $n = 1, m = \pm 1$, $n = 2, m = \pm 1$, $n = 3, m = \pm 3$.



(a)



(b)

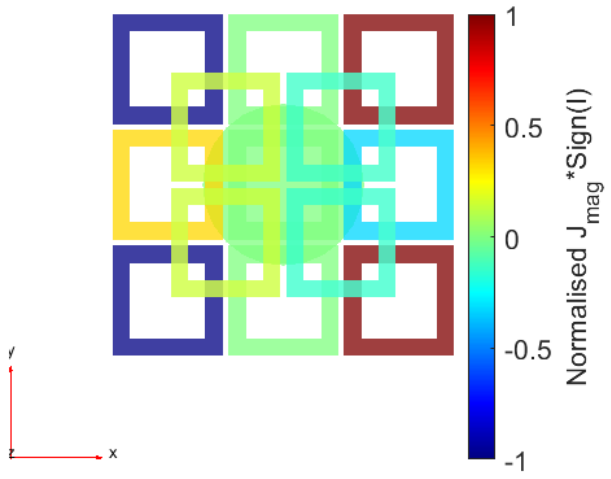


(c)

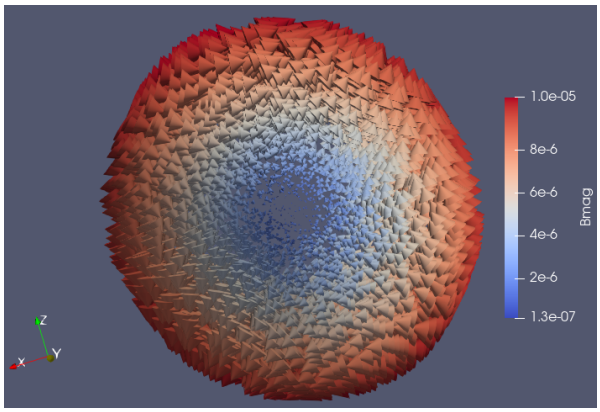
Figure 5.41: 5.41a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.41b) The target function related to the gradient of the real spherical harmonic n = 1 m = 0. 5.41c) The reconstruction from the excitation pattern with 48.7% RMS error, colour bar gives values in Tesla.

This potentially suggests limited improvements for certain functional symmetries. Numerical errors associated with the inexact FEM model may also impact these values when contrasted with the analytic filament solution, preventing a clear conclusion from being drawn.

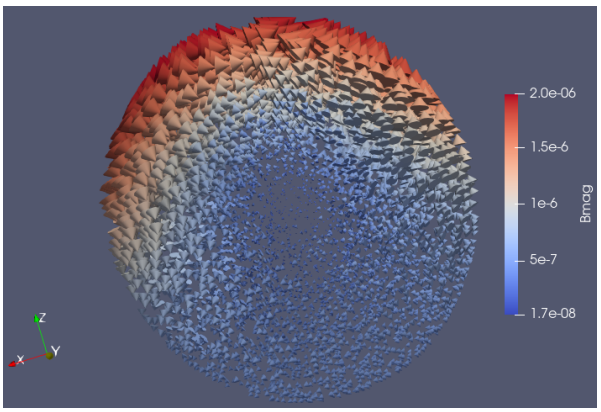
Contrasting the two versions of the $n = 2, m = -2$ field, Figures 5.40 and 5.46, the spatial distribution of the filamentary result is a good match, whereas, the solid system matches well in the outer-edges, but has a larger internal volume of weaker magnetic field. Both exhibit an excessive field gradient between inner and outer regions when compared



(a)



(b)



(c)

Figure 5.42: 5.42a) The excitation pattern in terms of normalised current distribution between the array of coils. 5.42b) The target function related to the gradient of the real spherical harmonic n = 4 m = -4. 5.42c) The reconstruction from the excitation pattern with 29.9% RMS error, colour bar gives values in Tesla.

to the target profile.

The distinguishing feature is the strength of the final magnetic field, which is approximately sixty times stronger throughout the volume of interest when produced by the solid inductors. In terms of practical implementation in MIT, this is advantageous because the associated eddy-currents produced in response to this particular magnetic field profile will be significantly stronger. Stronger eddy-currents have multiple potential implications including, but not limited to: allowing the detection of a greater range of conductivities with a specific magnetometer, the use of lower-frequency currents to minimise in-phase

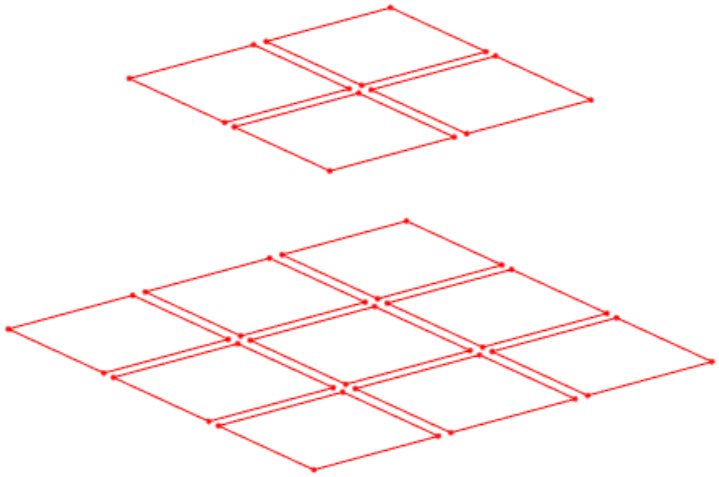


Figure 5.43: The geometry of the thirteen, co-planar square coil, filamentary, induction system. The aspect ratio is stretched along the Z-axis to ease interpretation, with the absolute separation between layers less than one-tenth the side length of an individual coil.

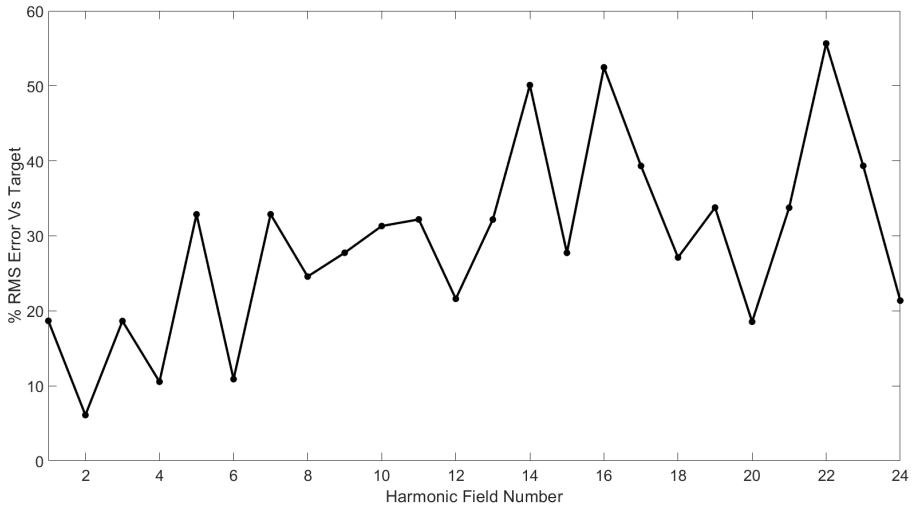
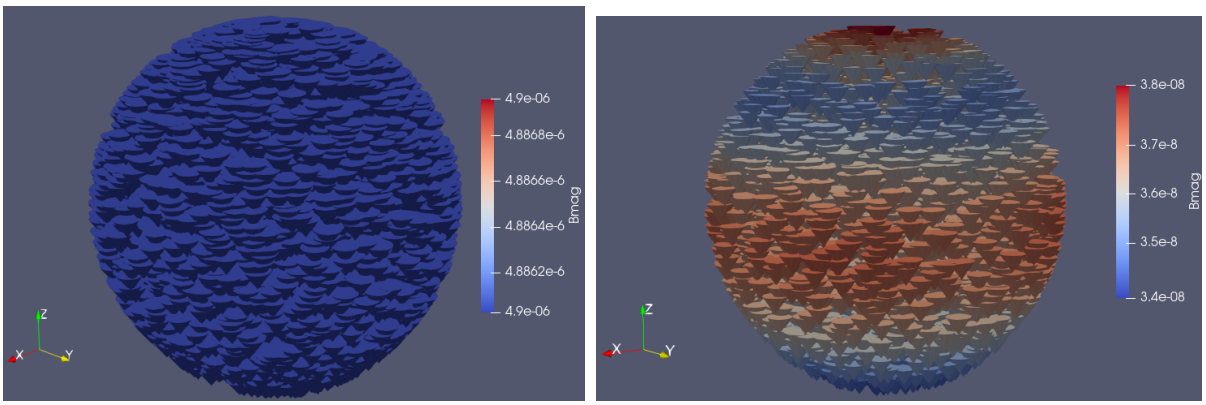


Figure 5.44: A plot of the various RMS values for each magnetic field produced by the filament version of this induction system. The Mean Field Fidelity for the complete data is 29%, which reduces to 26% for the nineteen unique field profiles, which match those in Figure 5.39.

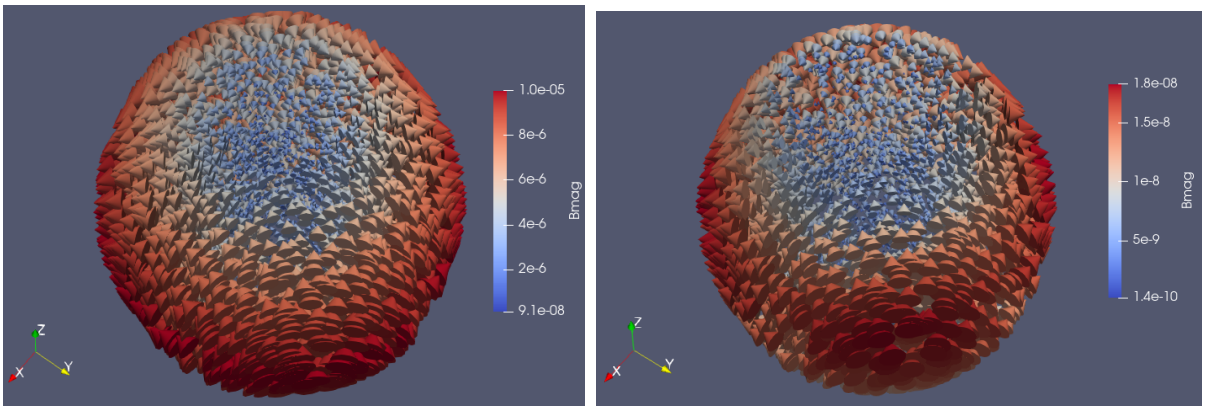
response contributions, and allowing the detection of conductive objects with a smaller volume.

The $n = 1$, $m = 0$ fields, Figures 5.41 and 5.45, are interesting because of the disparity



(a) (b)

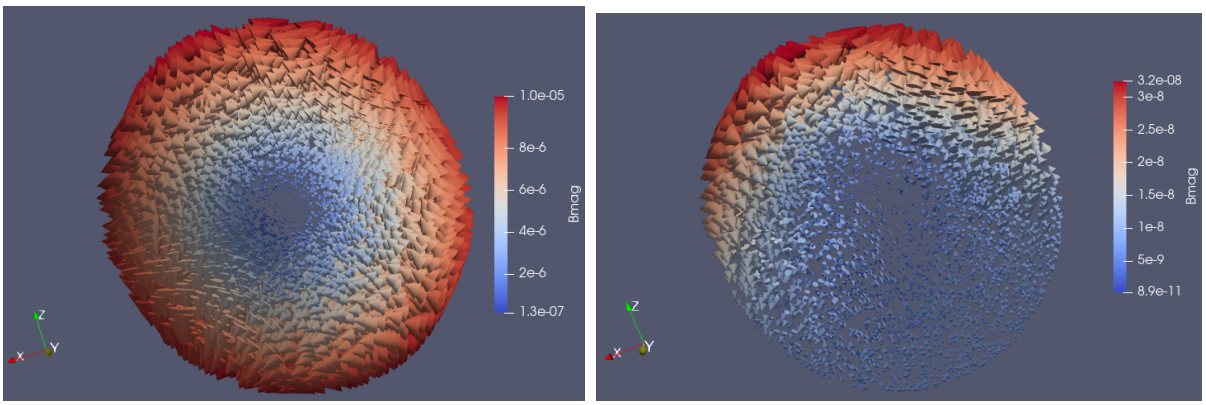
Figure 5.45: 5.45a) The target function related to the gradient of the real spherical harmonic n = 1 m = 0. 5.45b) The reconstruction from the excitation pattern with 6.1% RMS error, colour bar gives values in Tesla.



(a) (b)

Figure 5.46: 5.46a) The target function related to the gradient of the real spherical harmonic n = 2 m = -2. 5.46b) The reconstruction from the excitation pattern with 11 % RMS error, colour bar gives values in Tesla.

in RMS values. The difference arises from more significant variations in magnitude, with the solid coils displaying reduced control over the z-component at distances comparable to the size of a single inductor. There is similar behaviour for the $n = 2$, $m = 0$ result, suggesting the solid model has a noticeable unintended impact on the ability to produce magnetic fields with a dominant Z component. It is stronger than the filament result everywhere, ranging by a factor of between sixty and one-hundred-and-twenty times. The $n = 2$, $m = 1$ results, Figures 5.42 and 5.47, are also limited by control of the z-component and variations in magnitude. Interestingly the spatial distribution in the field is a good match throughout, but there is an obvious loss of strength with distance. The result is a



(a) (b)

Figure 5.47: 5.47a) The target function related to the gradient of the real spherical harmonic n = 2 m = 1. 5.47b) The reconstruction from the excitation pattern with 33 % RMS error, colour bar gives values in Tesla.

larger gradient than intended, which varies more significantly in the filament model. The solid coils still produce a universally stronger field, ranging by a factor of sixty to two hundred times.

In conclusion, differences are insignificant in the functional form of the fields produced by either current density model. Three-dimensional sources may have some minute functional advantages for specific targets, but the two-dimensional system makes higher-quality spherical harmonic approximations overall. The three-dimensional system does generate significantly stronger fields, ranging from one to two orders of magnitude. Nominally this allows for the detection of lower conductivity materials, the use of lower-frequency source currents, or the detection of smaller conductive objects.

5.3 Summary

Through this chapter an understanding of the proposed novel field generation philosophy is established, and the influence of various parameters investigated.

Section 5.1.1 provides some verification, highlighting that intuitive results can be reproduced. It further establishes that the RMS can give limited information on the field quality with respect to a target function, but not on whether a result is a unique member

of a (quasi-) orthogonal set. A combination of harmonic expansion into a set of basis functions and overlap integrals provides a robust framework for this assessment.

Section 5.1.2 establishes that increasing the number of coils substantially impacts the number and quality of generated fields. It identifies bounds on the precision of current control required from coil drivers in any practical design whilst demonstrating that λ is highly non-linear and incredibly influential on the orthogonality properties. There is also a trade-off between the fidelity (and subsequent orthogonality) of a generated profile and the field-per-unit current linking directly to choosing the current-norm as the stabilisation term in Equation 3.23.

Section 5.1.3 presents an idealistic design, with an average RMS across all 48 fields of just 2.2%. Overall this design successfully produces a larger set of higher-quality spherical harmonic magnetic fields compared to the equivalent system in Figure 5.6, and the improvement is generally quite significant. These provide a set of realistically generated high-accuracy fields, which are applied in the next chapter as the ideal excitations whilst evaluating the impact of orthogonality in the MIT inverse problem.

Overall, this chapter demonstrates that the algorithm in Equation 3.23 can successfully produce excitation patterns for various geometries. Many unique profiles are possible, even with a simple geometry, exceeding the number of coils. Increasing the number of coils is necessary to maximise orthogonality, and adaptive selection of the relative weighting of the current norm term (λ) is imperative and depends on all the inputs: target field, coil geometry, and excitation volume. The variability in any practical problem means λ should be selected quantitatively, with an L-curve-based method demonstrated to be versatile. A scannable system is also possible, but this sacrifices the control of the magnetic field along the Z-axis.

Chapter 6

Eddy-Currents & Inversion Problems In MIT

Through this chapter, we attempt to comprehensively verify the main features of the proposed approach to MIT. We demonstrate that including additional, spatially varying magnetic field profiles in an augmented approach to inversion can produce higher-quality images from inversion. Improving the orthogonality also has clear advantages to the image quality and minimises the computational effort (a proxy for inversion speed) required to produce them. This is true at high and low conductivities, even in less-than-ideal distributions, and at a high resolution where the problem becomes vastly underdetermined. Iterative algorithms are demonstrably superior to their traditional Tikhonov counterparts, which could not solve any of the problems considered. The specific choice of algorithm will always depend on the optimal balance between speed, inversion stability, and solution quality required for an application. We demonstrate the potential of numerous novel algorithms, including highlighting a disparity between the theoretical validity and practical usability of CGLS-type algorithms.

6.1 Application of Quasi-Orthogonal Magnetic Fields

This initial problem verifies the main hypothesised features of this novel approach to MIT. There are benefits to applying an increasing number of mutually unique magnetic field

landscapes, and improving the orthogonality of those landscapes maximises the available information leading to demonstrable improvements in the quality of image reconstructions. It shows iterative inversion algorithms are consistently superior to traditional Tikhonov counterparts and that improved orthogonality has further benefits in reducing the computational requirements needed to reach semi-convergence. Individual algorithms have specific advantages and combined ensure the improvements relate to the application of additional magnetic field profiles rather than any apriori assumptions.

A fixed forward problem is defined to compare results when the area is excited by twelve fields from the four coil induction system, Section 5.1.1, and an equivalently scaled, idealistic, sixty-six coil system, Section 5.1.3, targeting the same harmonics. Those have linear numbers $1-9, 17, 28, 30$, which the idealised system forms with less than 2% RMS.

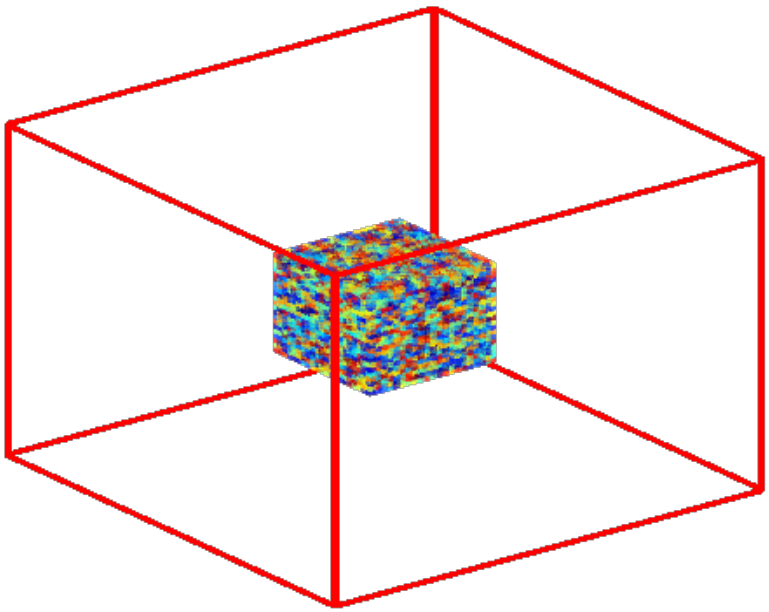


Figure 6.1: The discrete volume defined in the forward problem to generate synthetic signals.

This applies the approximations associated with Equation 3.65 and divides that volume into a $100 \times 100 \times 100$ grid of cubic voxels, recalling this covers a volume with dimensions one-third a coil side length, presented in Figure 6.1. The sensor choice is influential, so four arbitrary measurement positions are chosen relative and external to that imaging vol-

ume to formulate a contactless problem. These are presented in Figure 6.2, with all four

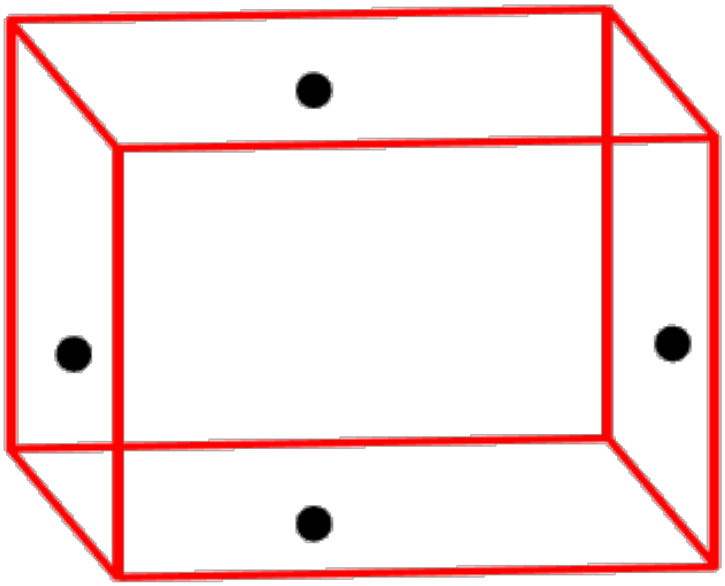
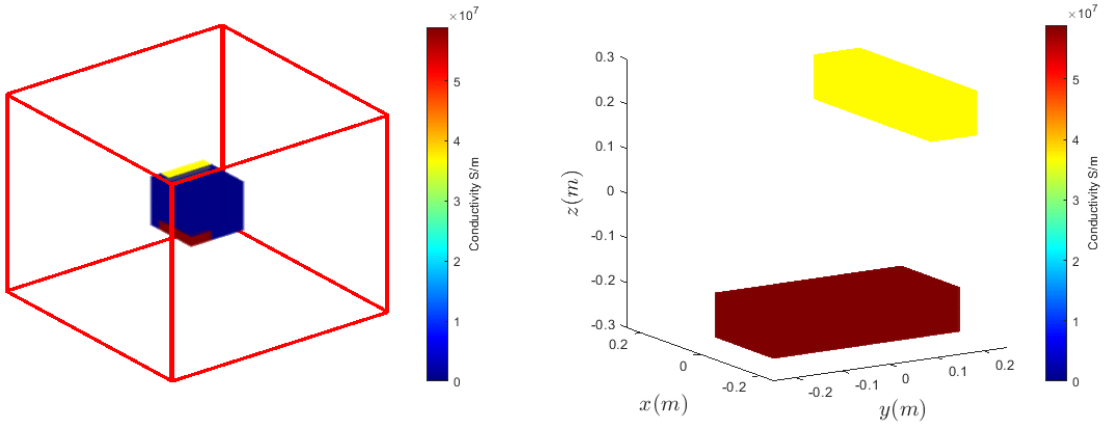


Figure 6.2: The arbitrary sensor distribution, as defined relative to the four coil system. Two sensors lie in-plane with the induction coils with length along the y-axis. The location of the other two sensors is above and below the induction system in-plane with the top/bottom of those coils.

skewed off-centre towards the viewpoint to minimise any bias relating to the commonly manipulated perfect axial symmetries. With sixty-six coils scaled to have an identical field generation volume, these sensors are contactless but located inside the cylindrical scanner bore. A specific conductivity profile representing a simple cluttered environment is specified over that volume, presented in Figure 6.3. These have normalised dimensions of; 0.2 in X, 0.4 in Y, 0.1 in Z for copper and; 0.4 in X, 0.1 in Y, 0.1 in Z for aluminium.

Simulation of the responses applies Equation 3.63 to this discrete problem for each coil system. With three-axial measurements, twelve fields, and four detectors, there are a maximum of one-hundred-and-forty-four responses. To avoid the inverse crime, these are combined with Gaussian white noise at one-percent of the signal magnitude, finalising a complete forward model.

Formulation of an inverse problem allows an evaluation of the influence varied magnetic field landscapes have on the ability to resolve the underlying conductivity profile. The



(a) The conductivity plotted relative to the induction system.

(b) An alternative view of the conductivity, setting empty voxels to be transparent.

Figure 6.3: The specified conductivity distribution consists of two separate objects made of copper (red, $\sigma = 5.89 \times 10^7$ S/m) and aluminium (yellow, $\sigma = 3.69 \times 10^7$ S/m), which have dimensions ranging from five to twenty times smaller than a single coil side length. In 6.3a, blue identifies the empty-space regions of the imaging volume as plotted relative to the four coil system.

limited data restricts the reasonably parameterised number of voxels so a sensitivity matrix calculation utilises Equation 3.65 for a dimensionally identical volume containing two-hundred-and-sixteen, or a $6 \times 6 \times 6$ grid of cubic voxels. Individually, these are twenty times smaller than a coil side length. Depending on the number of magnetic landscapes applied, this formulates an underdetermined inverse problem with two-hundred-and-sixteen parameters and between twelve and one-hundred and forty-four pieces of data.

Quantitative analysis can apply two tools from Chapter 4: the condition number and the information efficiency. The information efficiency uses a Tikhonov identity generalised inverse and a λ value selected via the L-curve, using the noisy signals from each forward model. The contrasting results for each system are given in Figures 6.4 and 6.5 respectively.

Inverse solutions then rely on selecting an appropriate algorithm, with the small size of these inverse problems allowing multiple to be applied rapidly. Tikhonov methods were unable to solve this problem for any number of magnetic fields, producing corrupt results typically spanning $\approx \pm 10^9$ S/m, even for $10^{-40} \leq \lambda \leq 10^{40}$ with all matrix combi-

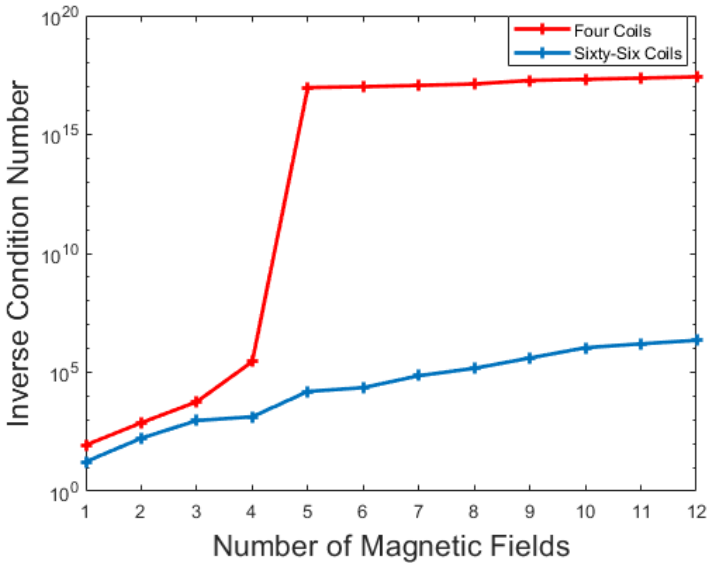


Figure 6.4: The variation in condition number as an increasing number of magnetic field landscapes are incorporated into an inversion problem attempting to resolve two-hundred-and-sixteen voxels. Each addition adds twelve measurements to the problem, defining an underdetermined problem with between twelve to one-hundred and forty-four pieces of data.

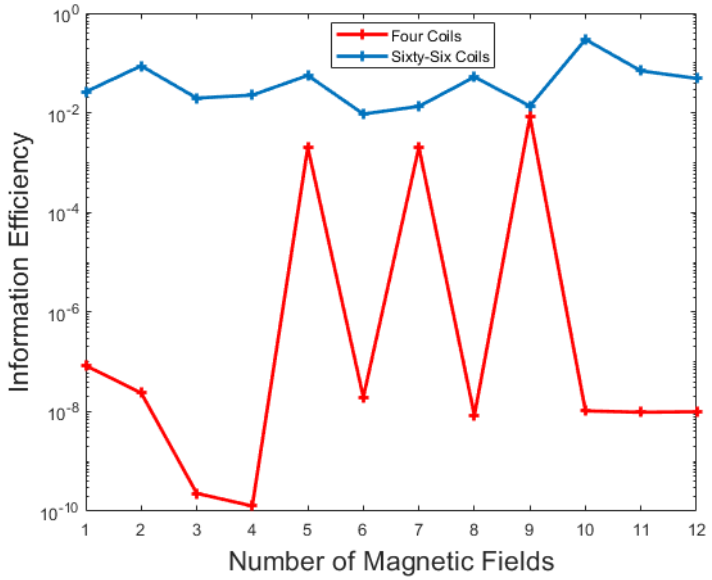


Figure 6.5: The variation in the information efficiency with an increasing number of magnetic field landscapes, as considered in an underdetermined inverse problem containing two-hundred-and-sixteen evenly distributed cubic voxels.

nations in Appendix E. This is one indication of the superiority of iterative algorithms, with basic choices producing an interpretable image with just twelve measurements.

Applying multiple algorithms serves several purposes, including demonstrations that:

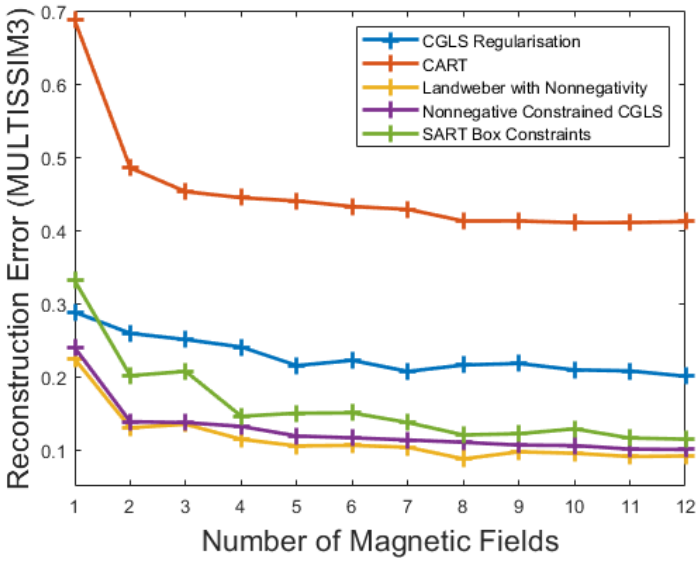
- The variation in image quality is robust, independent of the selected prior.
- There is a balance between speed, image quality, and solution stability.
- There are benefits to additional constraints.
- The algorithmic choice significantly impacts image quality.

Five algorithms are selected, the core details of which are in Chapter 4. The specific non-negative Landweber implementation uses a line-search to select an iteration-dependent relaxation parameter for convergence speed, whilst the SART algorithm uses physical box constraints of $0 \leq \sigma \leq 6.3 \times 10^7$ S/m. Non-negatively constrained CGLS refers to the NNFCGLS approach from [175]. From those algorithms, image quality results are in Figure 6.6, with sub-figures comparing the influence of increased field orthogonality.

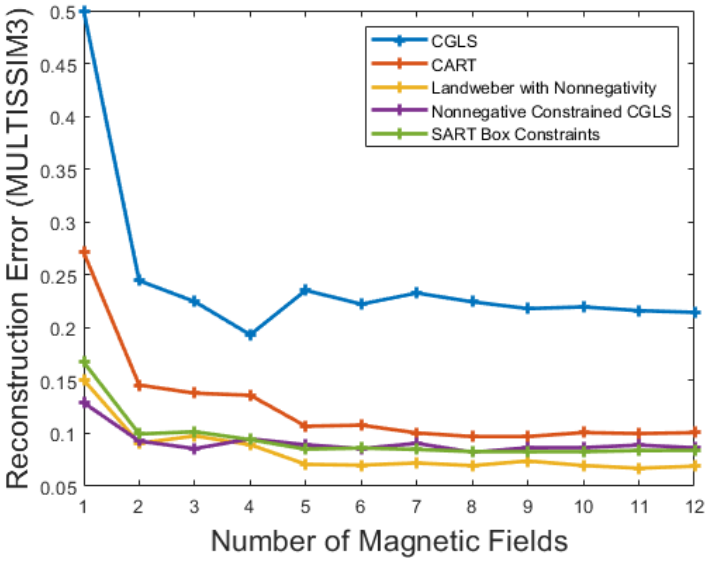
The number of included magnetic fields also influences the time taken to reach semi-convergence, which also varies between algorithms. The iteration results for four algorithms are in Figure 6.7, with the CART algorithm omitted because nearly all solutions utilise the same time-limited number of iterations.

For this and all subsequent inversion analyses, the MULTISSIM3 metric quantitatively assesses the images produced at all iterations, with the best image selected from this metric. This metric provides quantitative insights into trends and changes in these experiments but is not a substitute for a complete qualitative analysis of the resulting images. Results are, therefore, idealistic, bypassing the general requirement to apply a stopping rule. However, how necessary this is will be explicitly coupled with the algorithmic stability.

This problem is, contextually, information deficient, so the measurements are likely to be unable to completely capture the additional information provided by various field profiles, particularly as harmonic n increases and the number of spatial variations increases. The simulated eddy-currents from Equation (3.63) are easily-differentiable between ap-



(a) Four Coils

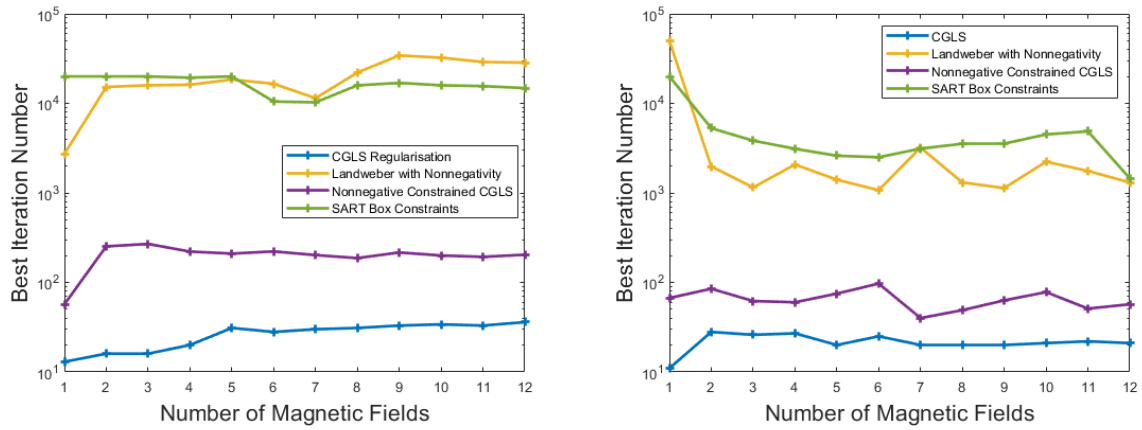


(b) Sixty-Six Coils

Figure 6.6: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, with increasing numbers of magnetic fields. The subplots compare identical models differentiated only by the quality of the generated spherical harmonic magnetic fields, demonstrating the influence of increased magnetic field orthogonality.

plied fields, indicating the capture of additional information.

Figure 6.4 demonstrates that the conditioning of the problem increases with the addition of magnetic fields, meeting expectations as the systems tends towards a well-determined problem. The variation is much smaller, and the system is significantly less sensitive to



(a) Four Coils (b) Sixty-Six Coils

Figure 6.7: The semi-convergent point, in terms of optimal iteration number for the specified problem from the MULTISSIM3 metric. Subplots compare magnetic fields with differing orthogonality properties, where the CART algorithm is omitted because only one of twenty-four results does not use the allowed 2500 iterations - Figure 6.7a, one magnetic field, CART uses fifty-seven.

noise when applying highly-orthogonal fields with increased measurement independence. The initial indication is, therefore, that orthogonal magnetic fields aid the ability to solve the inverse problem in MIT.

The condition number is not a physical parameter, so significant variation occurs when applying a vector-potential model such as Equation (3.58) instead of this flux model. From the first simulated field from four coils, there is a symmetric A_x, A_y distribution and a A_z component around 13 orders of magnitude smaller. The calculation of the singular values of the sensitivity matrix (applying Equation (4.3)) leads to a condition number of 4.3×10^{14}, whilst excluding the A_z contribution reduces this to 49. In reality, the A_z contribution approaches zero, compounding the issue by introducing singular values of zero, leading to an infinite condition number. A numerical quirk, exclusion of A_z has no influence on image quality even with 10% measurement noise, suggesting the condition number is a poor metric to solely quantify the physical invertibility of a particular system in this framework.

Figure 6.5 shows an increase in the average information content with the addition of mag-

netic fields, and that information content consistently increases in measurements from more orthogonal field profiles. Therefore, expanding the number of orthogonal magnetic fields improves the resolvability of the inverse problem. The consistently low values accentuate the understanding of the difficulties in producing a successful Tikhonov inversion. The oscillatory behaviour is unexpected with further testing demonstrating the preservation of shape, even when varying λ. This suggests some physical dependency which fluctuates with the number of fields, perhaps related to the limited number of measurements.

The quantitative image quality data in Figure 6.6 further demonstrates the value of applying spatially varying magnetic fields and increasing their orthogonality. Increasing the number of fields leads to a clear improvement in image quality for all five algorithms, and four of five algorithms demonstrate significant benefits to greater orthogonality. All five algorithms also use substantially fewer iterations when applying more orthogonal magnetic fields, implying inversions are produced faster with less computational effort, Figure 6.7. The improvement in speed is also typically more significant as the number of highly-orthogonal profiles increases. Individually each highlights how an appropriate choice of algorithm will further benefit reconstruction quality, though this is problem specific. Beyond image quality, the individual choice also has features of interest suited to different end-user requirements.

The error metric quantitatively illustrates trends and changes, but as previously suggested, when used exclusively is arbitrary. Therefore, investigating the ability to produce characterising, volumetric images of a conductivity distribution is incomplete without considering those images. For brevity, the inclusion of only a subset of the one-hundred and twenty image reconstructions composing Figure 6.6 is possible. Some aid the discussion of the strengths of individual algorithms, whilst the CART and Landweber algorithms with $1, 3, 6$ and 12 fields provide a visual illustration of the benefits of augmented use of quasi-orthogonal profiles. Each figure compares a perfect resolution-specific image (left)

to a reconstructed solution from inversion (right), using a matched colour map with a variable solution-dependent range. Each also suppresses noise below one percent of the peak (copper) conductivity and highlights object voxels through variable transparency to aid interpretation. The selected angle attempts to provide an idealised projection of a three-dimensional system however an interactive analysis (rotating, panning, varying transparency) was necessary for a detailed qualitative evaluation.

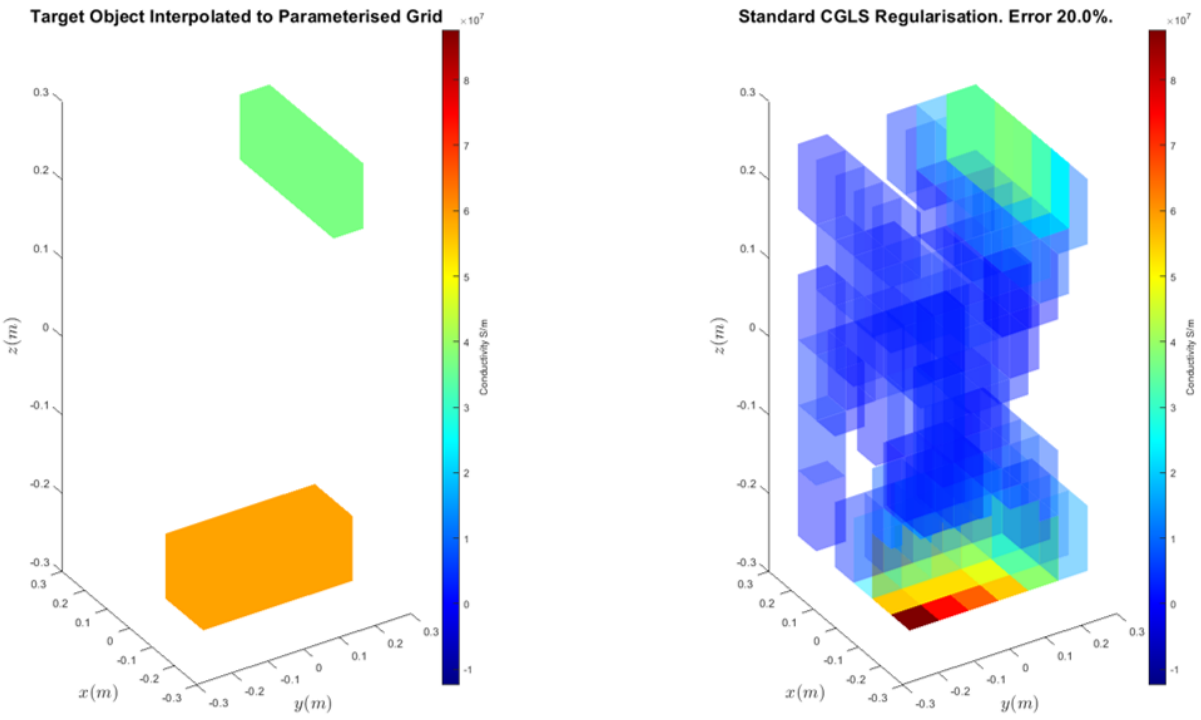


Figure 6.8: The best image reconstruction from the CGLS iterative reconstruction algorithm for a system which applies twelve harmonic-inspired magnetic fields generated by four coils.

CGLS is theoretically invalid due to the lack of a positive-definite sensitivity matrix. Empirically, this is a fast, easily implemented algorithm which produces substantially better images than comparable corrupt Tikhonov results. It does particularly well at characterising the object magnitudes, Figure 6.8, whilst using the fewest computations, Figure 6.7. This speed will be particularly advantageous for practical resolutions. However, images contain significant smoothing and sizeable infeasible negative conductivities around 20% of the expected peak magnitude, meaning poor reproduction of the majority of the volume.

The NNFCGLS algorithm, Figure 6.9, demonstrates the substantial benefits of includ-

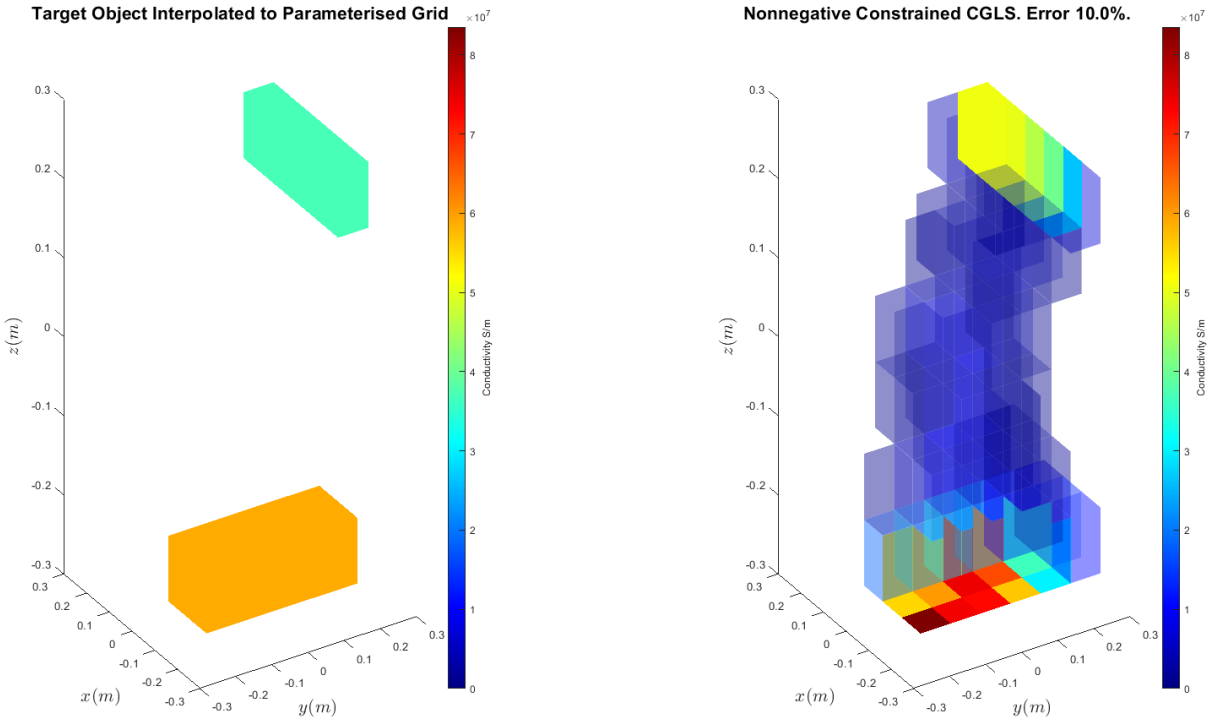


Figure 6.9: The best image reconstruction from the NNFCGLS iterative reconstruction algorithm for a system which applies twelve harmonic-inspired magnetic fields generated by four coils.

ing simple physical constraints balanced with increased computational requirements to produce the solution. For these images, the algorithm uses 204 iterations compared to 36 from the simple implementation, with a 10% improvement in image quality and elimination of unfeasible values. There is a reduction in smoothing effects, leading to better object localisation. Increased field orthogonality produced by the sixty-six coil system benefits both image quality, Figure 6.10, and computation speed, requiring just 67 iterations to converge. The IR tools implementation also contains a GCV stopping rule, preventing over-iteration and further minimising wasted computations.

The SART algorithm incorporates box-constraints preventing the over-estimations of conductivity prevalent in other solutions, Figure 6.11, enabling the best characterisation of the copper object for all solutions contained in Figure 6.6. Once again, increased orthogonality leads to a quicker solution, with semi-convergence falling from 14823 to 1455 iterations. This is true for all five algorithms, suggesting that as induced field inde-

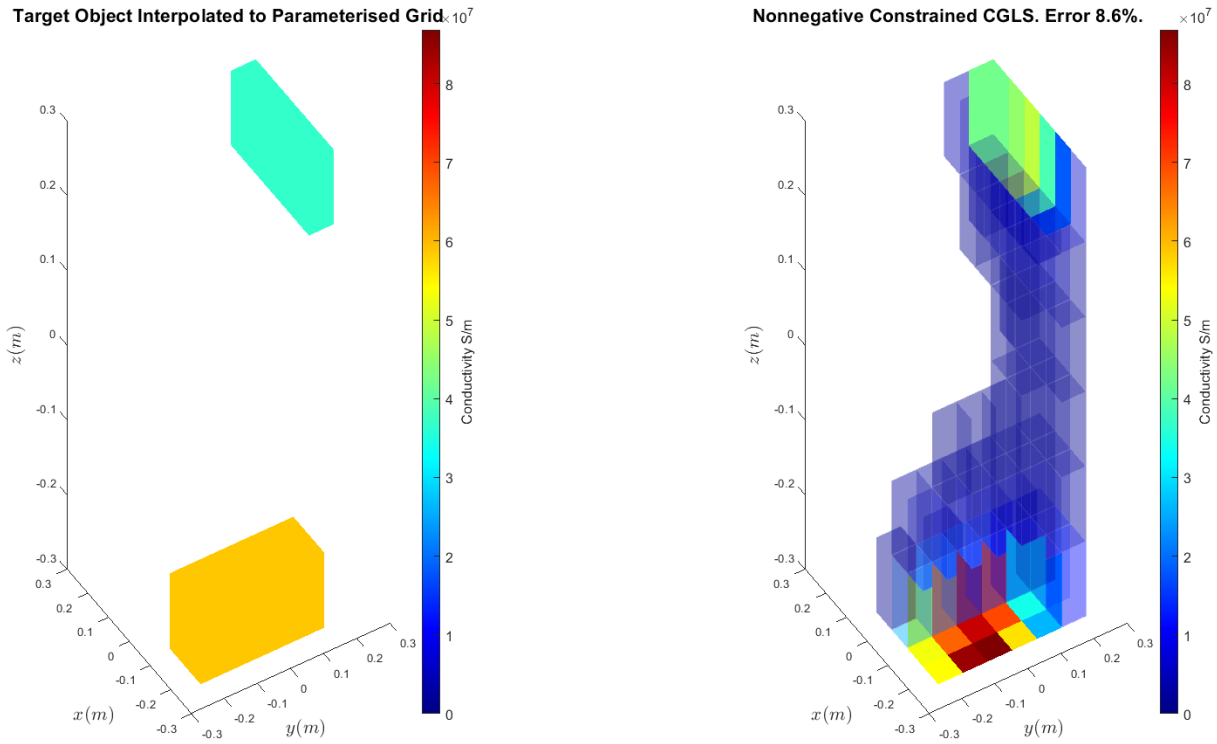


Figure 6.10: The best image reconstruction from the NNFCGLS iterative reconstruction algorithm for a system which applies twelve harmonic-inspired magnetic fields generated by sixty-six coils.

pendence increases, the solution space becomes more highly constrained, enabling faster convergence.

Theoretically, the CART algorithm is particularly well suited to this novel approach to MIT with orthogonal fields augmented into a single inverse problem owing to the hypothesised independence of measurements. The results for a single magnetic field are given in Figure 6.12 comparing one produced by four coils (top), sixty-six coils (bottom) followed by the related iterative behaviour, Figure 6.13. Contrasting the two solutions, Figure 6.12, a clear improvement in image quality is associated with the more orthogonal magnetic field, including enhanced object location, characterisation, and noise reduction. Neither object is localised well with the better image, Figure 6.12b, identifying a larger lower object and a smaller upper one - each with spatially varying conductivity.

The less orthogonal field is the only one of twenty-four CART solutions in Figure 6.6

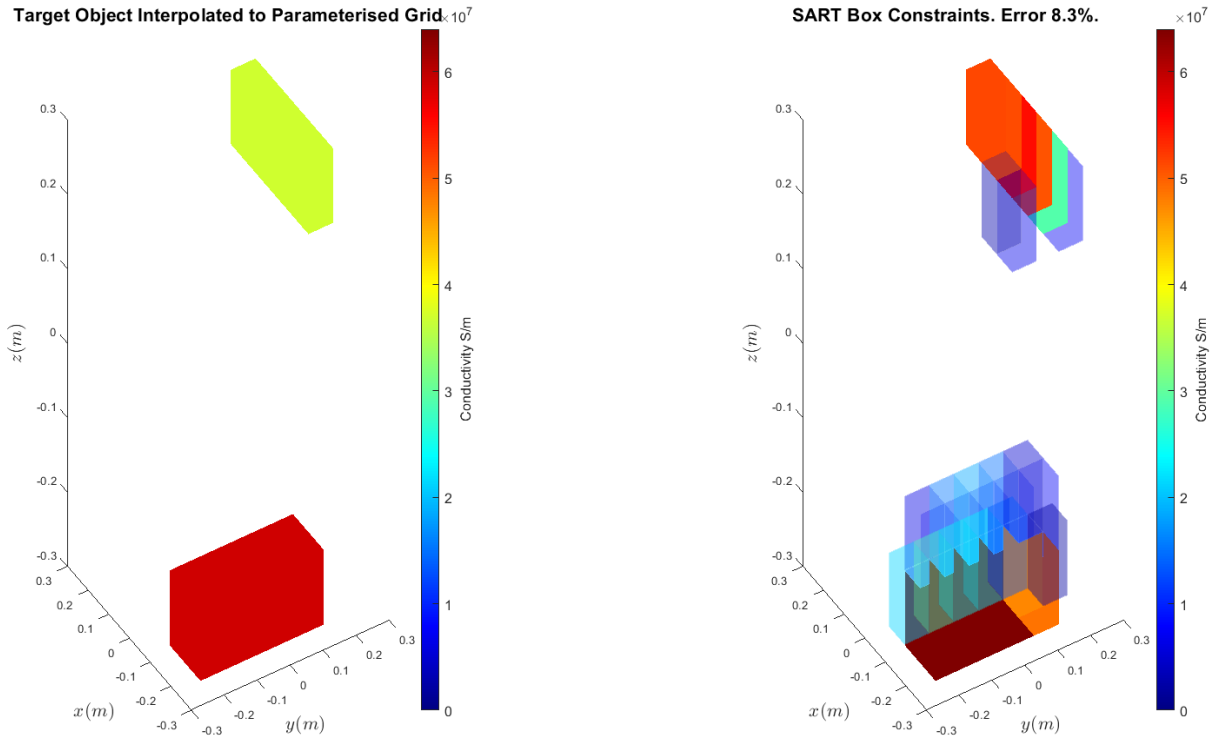
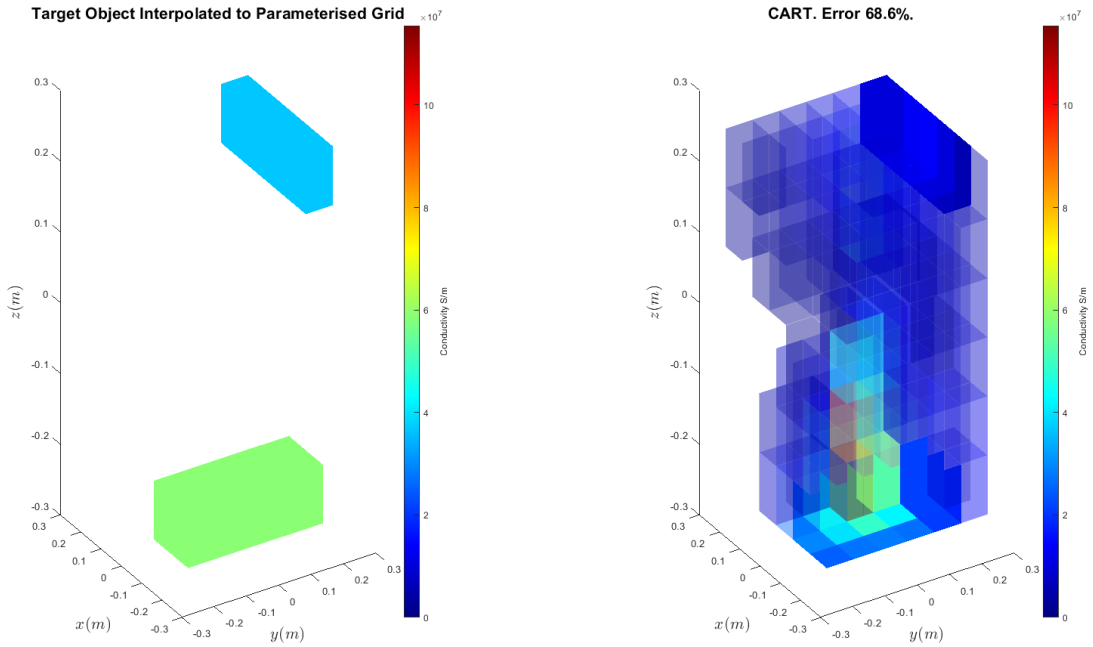


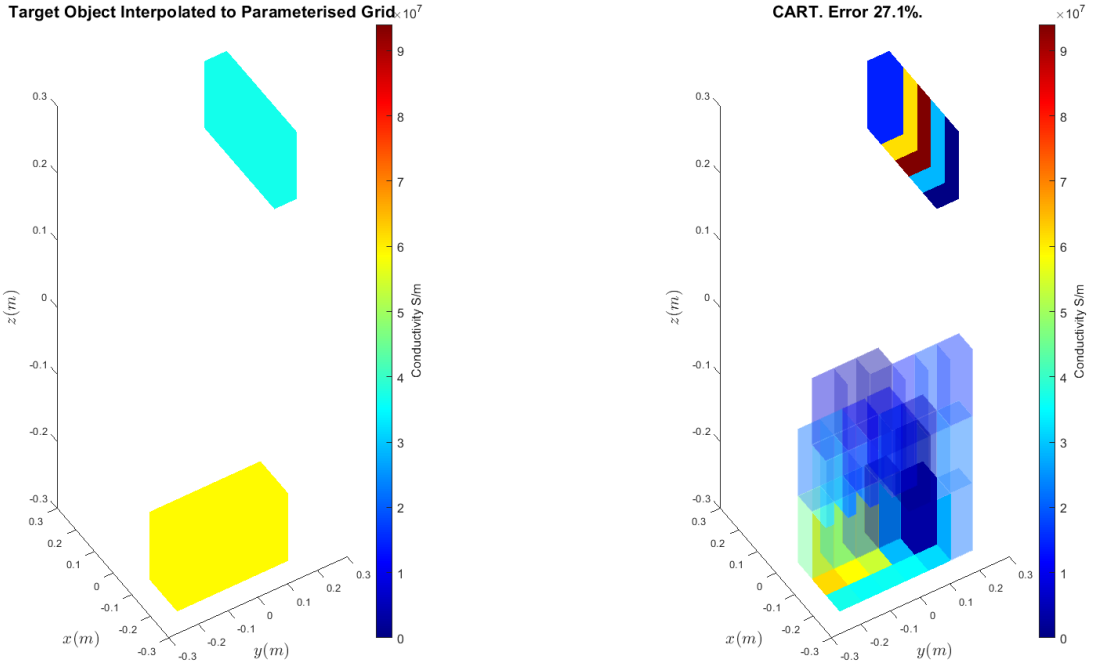
Figure 6.11: The best image reconstruction from the constrained SART iterative reconstruction algorithm for a system which applies twelve harmonic-inspired magnetic fields generated by sixty-six coils.

to pass semi-convergence in 2500 allowed iterations, Figure 6.13, but image quality falls little over the remaining iterations highlighting one key advantage of this algorithm. A comparable NNFCGLS approach with the same forward model decays rapidly away from the optimal result, Figure 6.14, making the practical selection of a near-optimal solution more difficult. Interestingly it appears the more orthogonal field from sixty-six coils improves the stability of the solving process. There is clearly a larger region of near-optimal solution, likely making it easier to select, and a slower decay to the point of semi-convergence when compared to the lower quality field from four coils.

The figure format is repeated sampling images from three, six, and twelve magnetic field profiles in Figures 6.19, 6.20, 6.21 respectively, with iterative behaviour provided concurrently in Figures 6.16, 6.17, 6.18. Individually each identifies consistently improved image quality when applying higher-quality spherical harmonic magnetic fields, true after any reasonable number of iterations, Figures 6.13, 6.16, 6.17, 6.18. The implication is



(a) Four Coils.



(b) Sixty-Six Coils.

Figure 6.12: The image reconstructions from the CART iterative reconstruction algorithm for a system which applies a single harmonic-inspired magnetic field, comparing four coil and sixty-six coil systems.

more orthogonal induced fields lead to increasingly independent eddy-responses, benefiting image quality as hypothesised. The shape of the blue curve in Figure 6.17 is also characteristic of this algorithm nearing semi-convergence, whilst all four orange results are yet to approach the first local minima. The implication is that further improvement

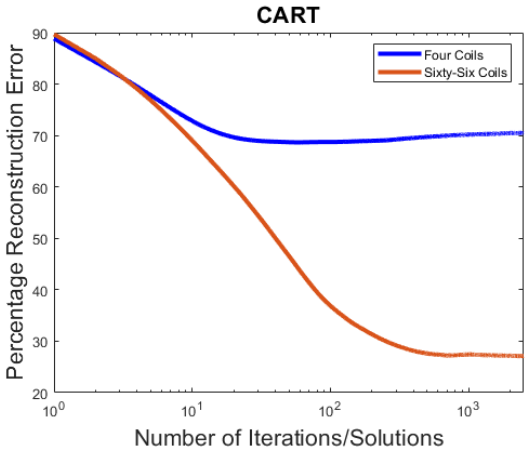


Figure 6.13: Iterative behaviour of the CART reconstruction algorithm for a system which applies a single harmonic-inspired magnetic field produced by two different induction systems.

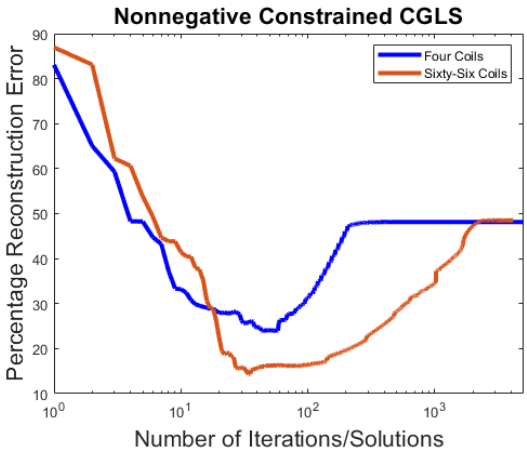


Figure 6.14: Iterative Behaviour for the NNFCGLS algorithm using a single-harmonic inspired magnetic field produced by two different induction systems.

in image quality is possible with additional computation, something verified in Figure 6.15, which runs the simulation from Figure 6.12b with six-hundred times more iterations (1.5×10^6) leading to a 12.3% improvement without reaching semi-convergence. The requirement is substantially more time, especially as CART has the slowest single iteration of the five selected algorithms. Practical use necessitates an early truncation, with 2500 iterations selected heuristically to roughly match the maximum allowed time for the other four algorithms.

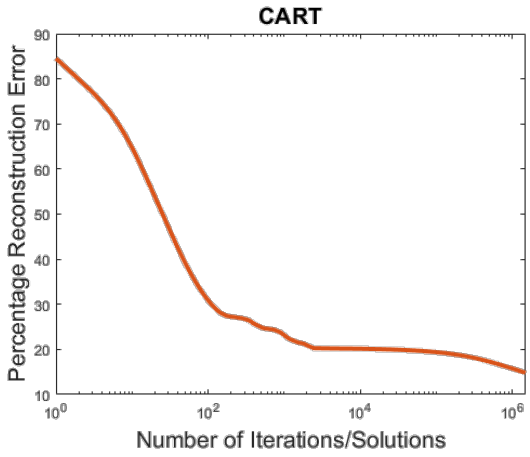


Figure 6.15: Iterative behaviour of the CART reconstruction algorithm for a system which applies a single harmonic-inspired magnetic field produced by sixty-six coils over 1.5×10^6 iterations.

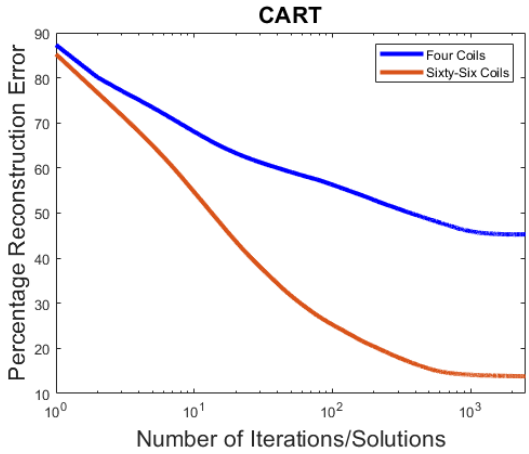


Figure 6.16: Iterative behaviour of the CART reconstruction algorithm for a system which applies three harmonic-inspired magnetic fields from two induction systems.

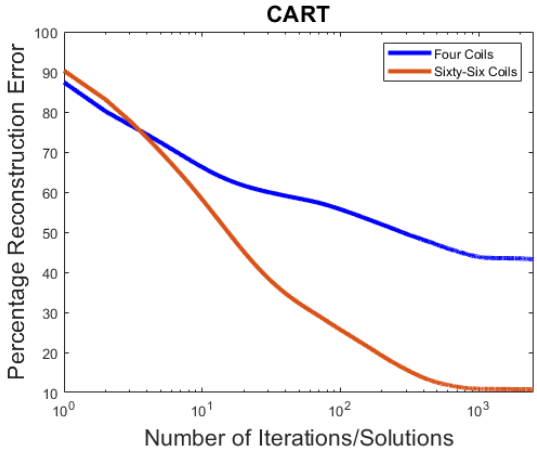


Figure 6.17: Iterative behaviour of the CART reconstruction algorithm for a system which applies six harmonic-inspired magnetic fields from two induction systems.

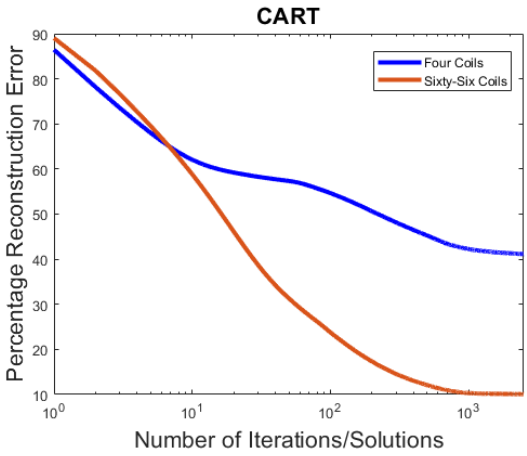
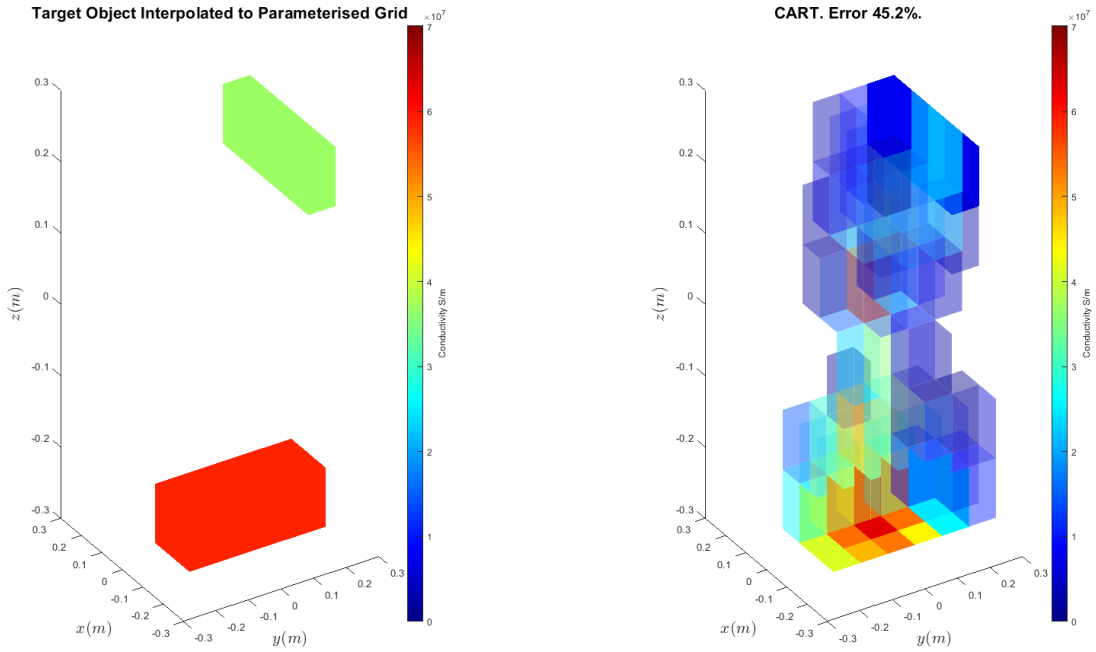
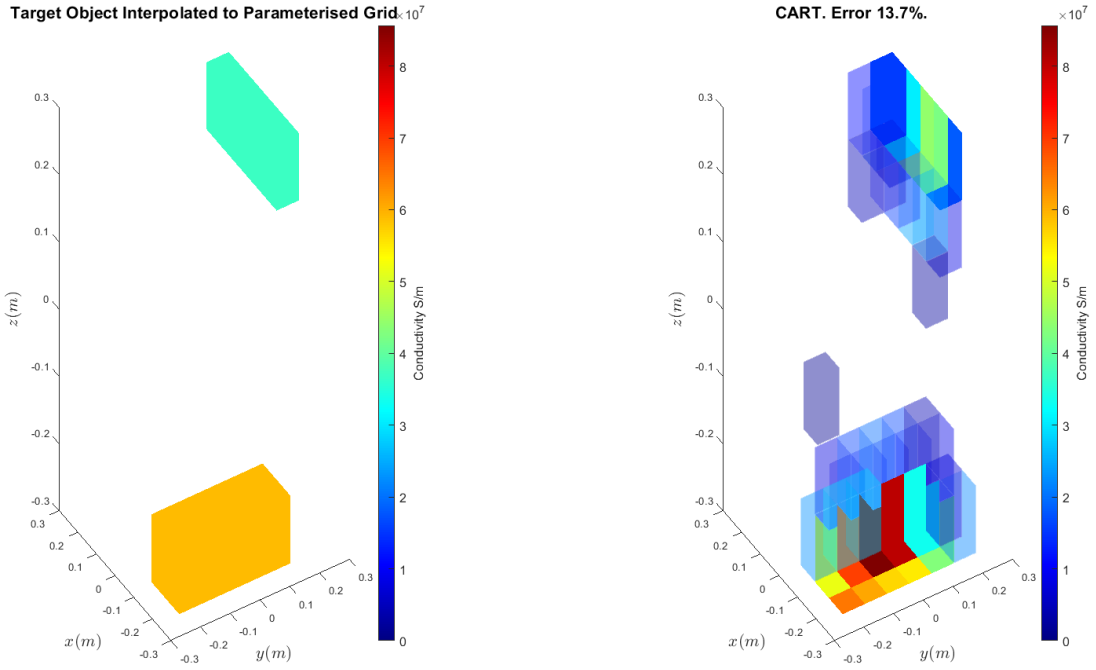


Figure 6.18: Iterative behaviour of the CART reconstruction algorithm for a system which applies twelve harmonic-inspired magnetic fields from two induction systems.

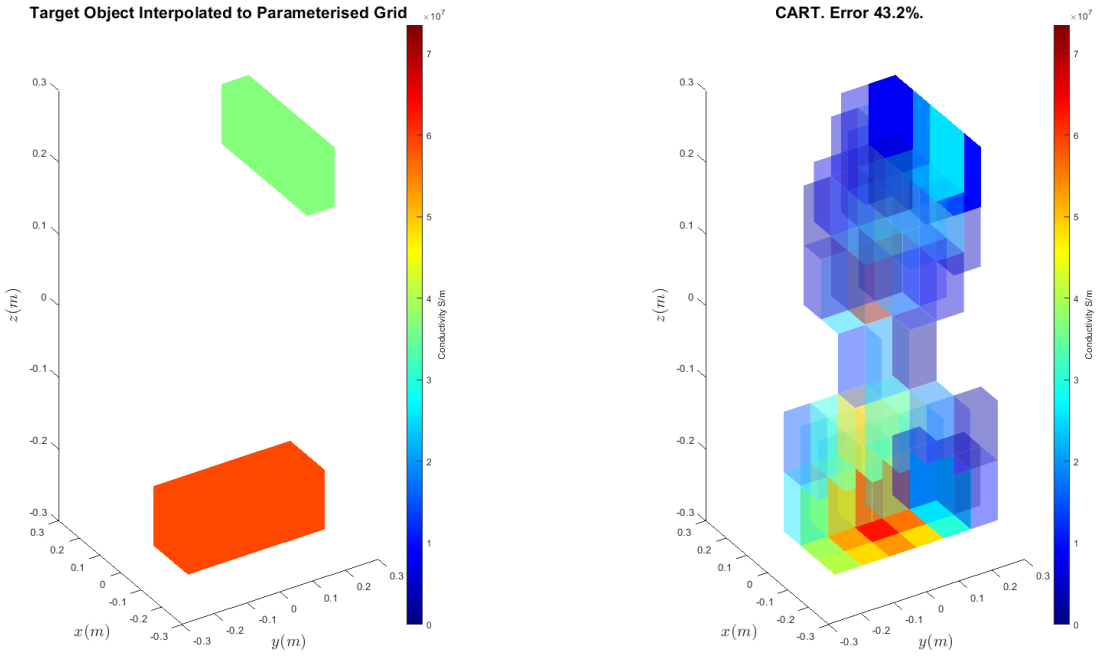


(a) Four Coils.

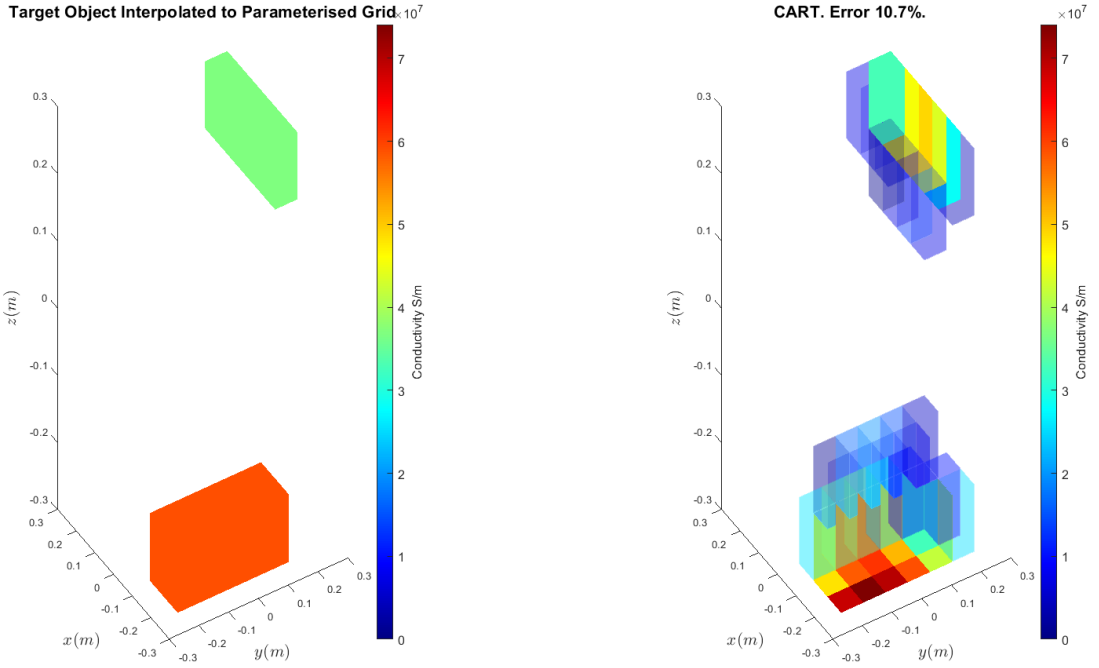


(b) Sixty-Six Coils.

Figure 6.19: The image reconstructions from the CART iterative reconstruction algorithm for a system which applies a three harmonic-inspired magnetic fields, comparing four coil and sixty-six coil systems.

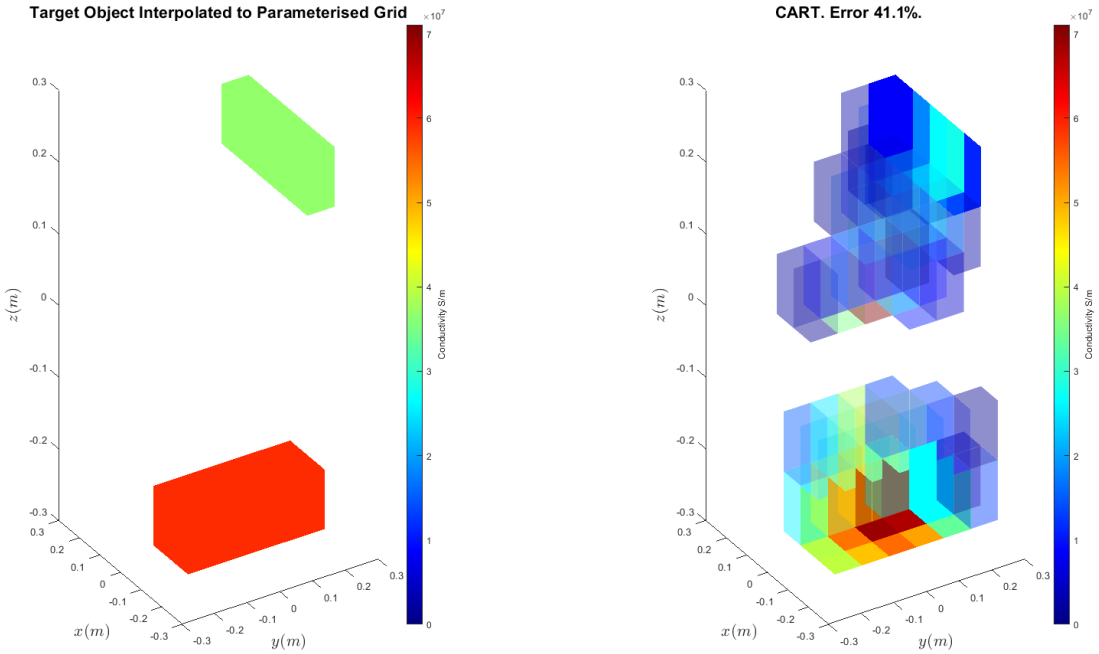


(a) Four Coils.

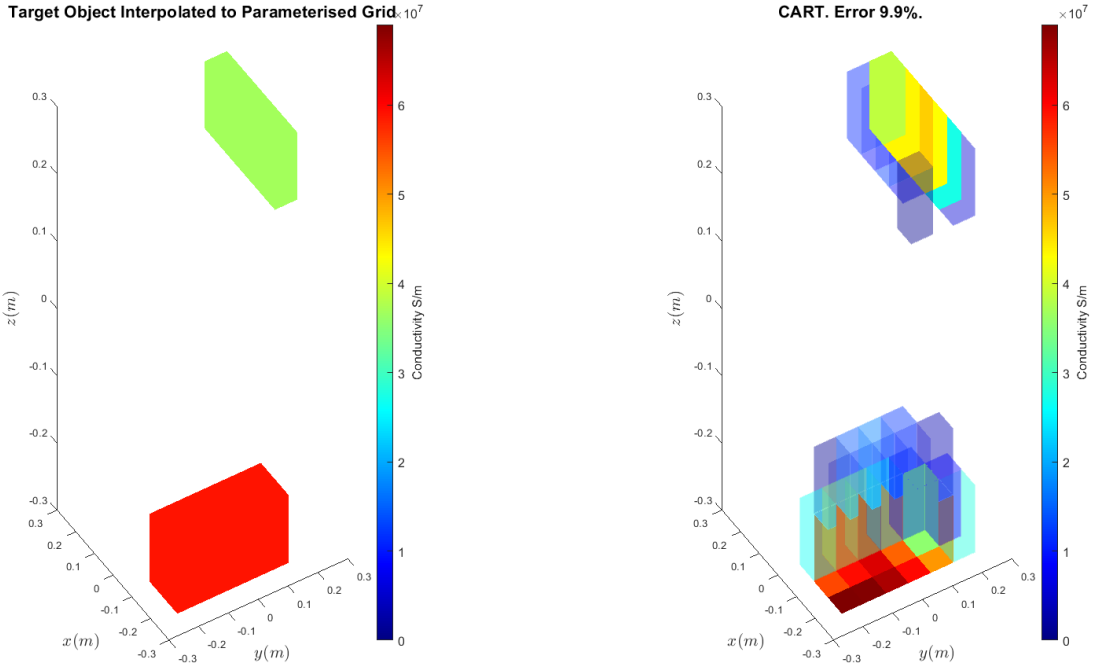


(b) Sixty-Six Coils.

Figure 6.20: The image reconstructions from the CART iterative reconstruction algorithm for a system which applies a six harmonic-inspired magnetic fields, comparing four coil and sixty-six coil systems.



(a) Four Coils.



(b) Sixty-Six Coils.

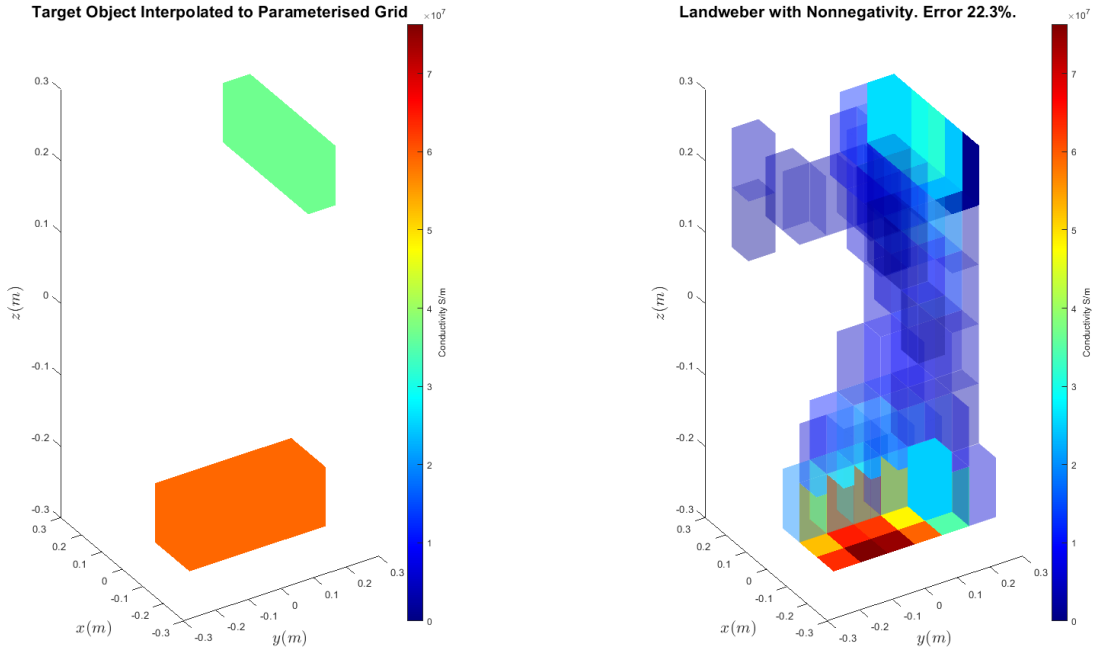
Figure 6.21: The image reconstruction and associated behaviour of the CART iterative reconstruction algorithm for a system which applies a twelve harmonic-inspired magnetic fields, comparing four coil and sixty-six coil systems.

All CART images using the four coil system contain some central phantom noise elements. Misinterpretation of these as an additional highly conductive object is not unlikely as their magnitude exceeds that of the real aluminium object. This phantom does not exist in any images associated with the sixty-six coil system. Contrasting the four Figures 6.12, 6.19, 6.20, 6.21, concurrently an increasing number of augmented fields produces demonstrable improvements in the accuracy of object location, characterisation of the target magnitude, and object uniformity, as well as a reduction in noise regardless of field quality. This is most significant between one and six fields, Figures 6.12 & 6.20, which reaffirms the approximately flat change in image quality seen in Figure 6.6 relating mostly to small reductions in noise and slight improvements in magnitude reproduction.

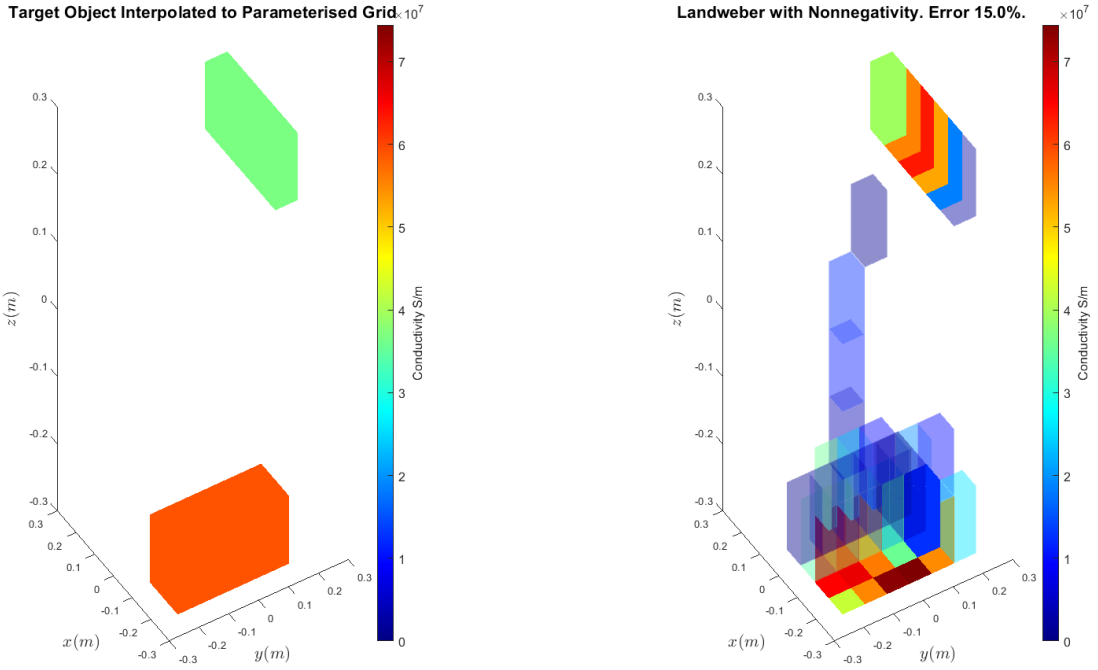
This suggests there is little benefit to the final six applied magnetic fields in either case. However, those profiles are of increasing spatial complexity. Viable mapping of those variations would require significantly more measurements and/or voxels, necessitating an inverse problem of higher complexity to characterise and assess their feasibility.

From Figure 6.6, the Landweber variant consistently produces the best image quality and is more advantageous as the orthogonality of the magnetic fields increases. For an end-user, this is contrasted in Figure 6.7 with a significant increase in computational effort compared to the NNFCGLS algorithm, suggesting it is ideal when computation time is less important.

For one magnetic field the images from the two induction systems are in Figure 6.22. Contrasted with Figure 6.12 the improvement is significant: there is a substantial reduction in noise, better object localisation, and more accurate magnitude characterisation. The four coils result has none of the high magnitude phantom voxels seen in the CART reconstructions and has a less significant overestimate of the peak conductivity. Comparing sub-figures 6.22a and 6.22b there is a clear improvement associated with the sixty-six coil system, comprising of less noise, and improved localisation of the two objects. The



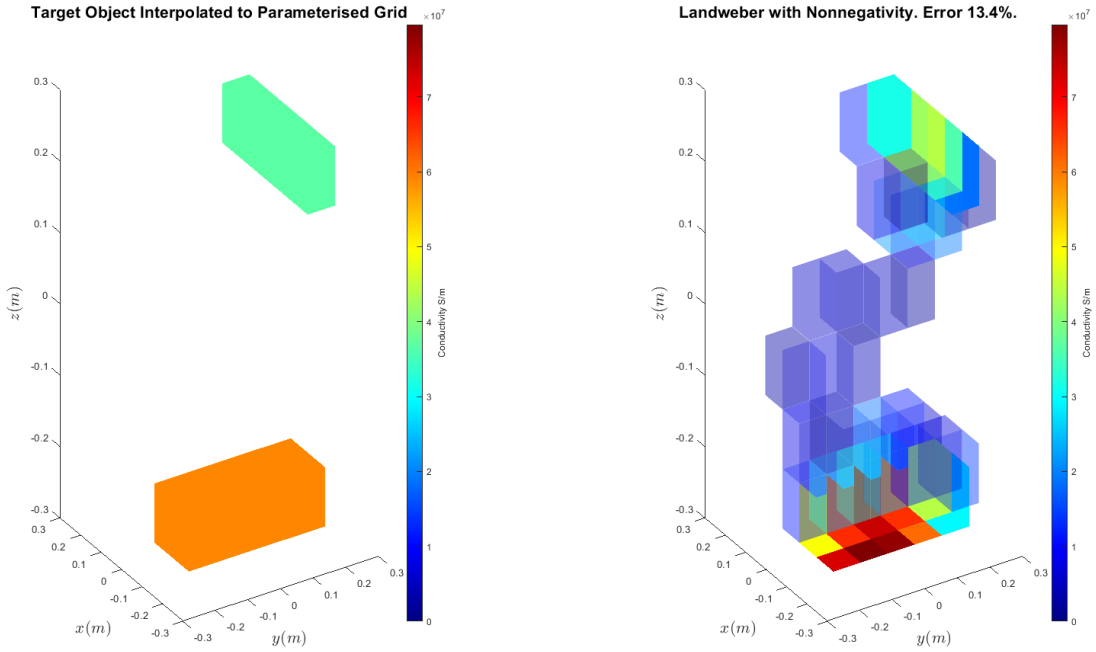
(a) Four Coils.



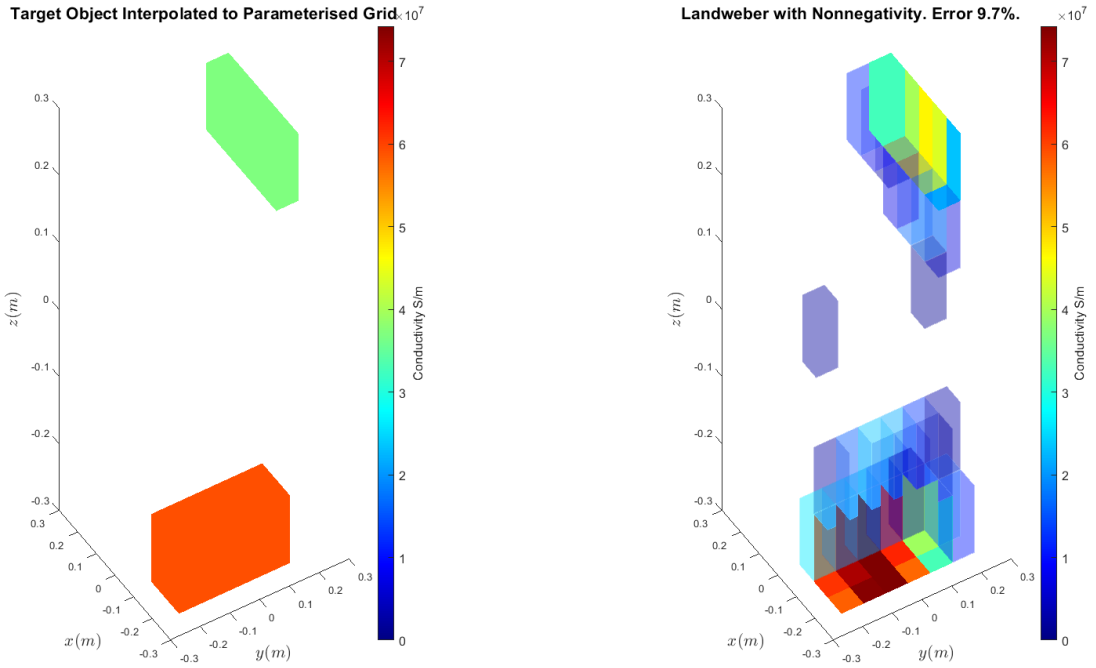
(b) Sixty-Six Coils.

Figure 6.22: The image reconstruction and associated behaviour of the constrained Landweber iterative reconstruction algorithm for a system which applies a single harmonic-inspired magnetic field, comparing four coil and sixty-six coil systems.

projection in Figure 6.15 obtains a similar error while requiring over thirty times more, individually slower, iterations than the 5×10^4 used to produce Figure 6.22b implying this Landweber variant is vastly more efficient.



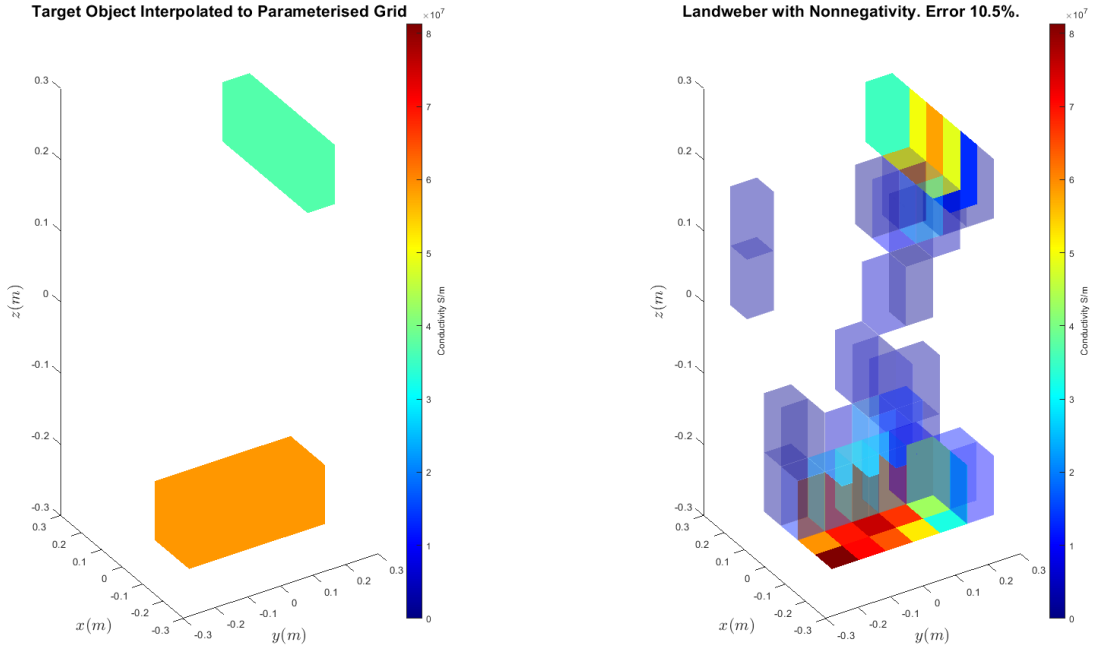
(a) Four Coils.



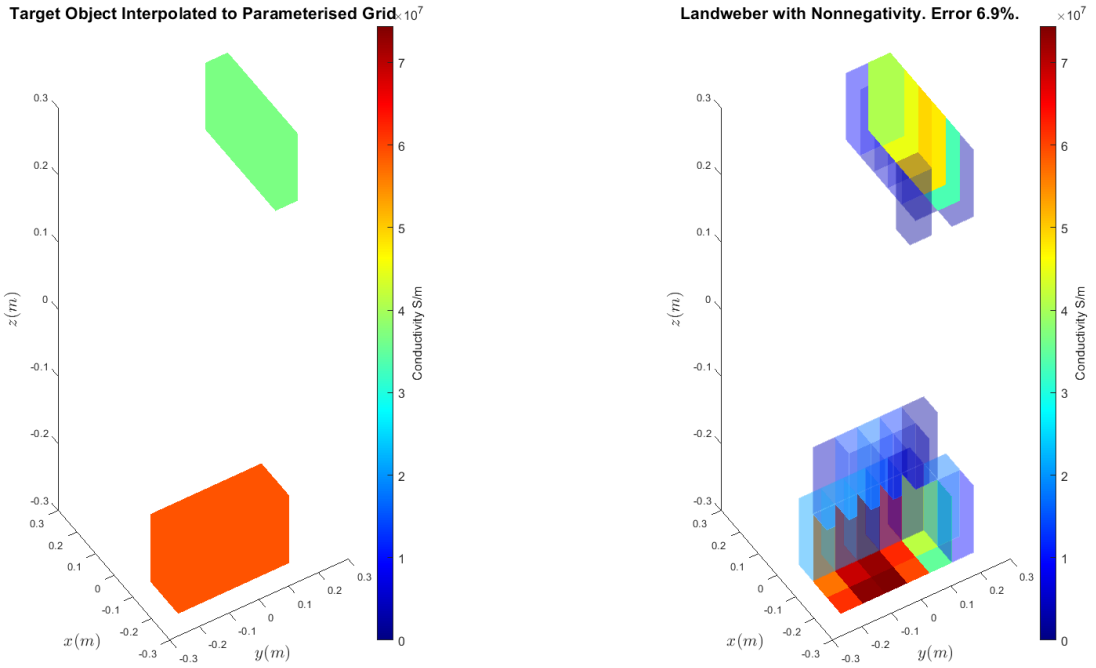
(b) Sixty-Six Coils.

Figure 6.23: The image reconstruction and associated behaviour of the constrained Landweber iterative reconstruction algorithm for a system which applies a three harmonic-inspired magnetic fields, comparing four coil and sixty-six coil systems.

As with the CART data, the resulting images for three, six, then twelve magnetic fields generated by both systems can be found in Figures 6.23, 6.24, and 6.25 respectively. Regardless of field quality, increasing the number of magnetic fields results in better-quality images with an obvious noise reduction. Both objects become increasingly uniform, suc-



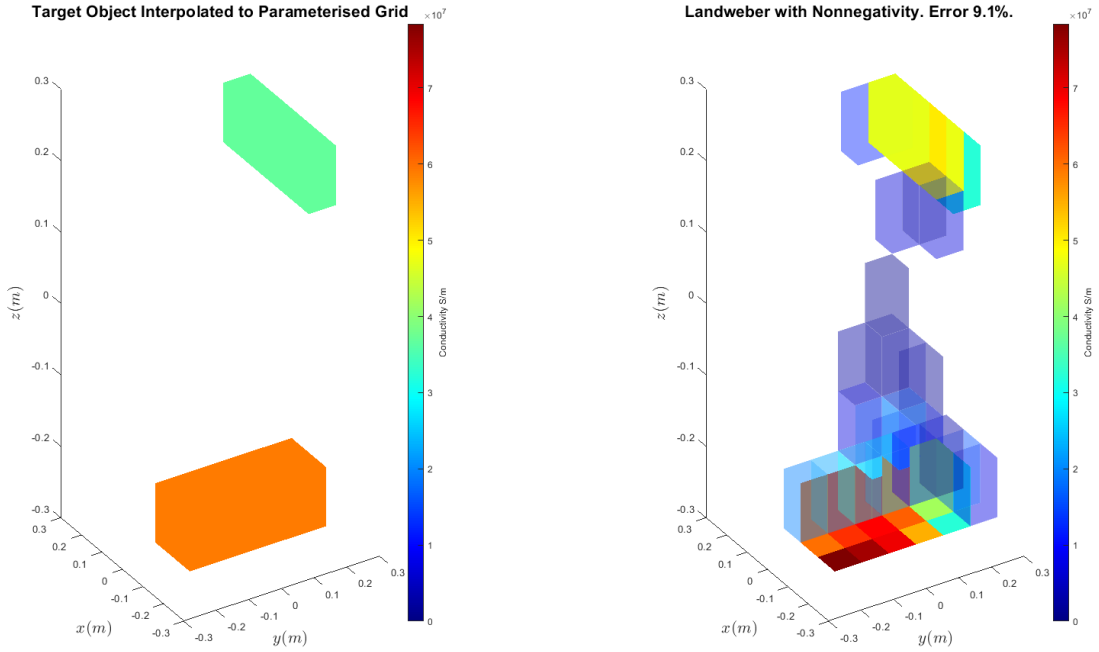
(a) Four Coils.



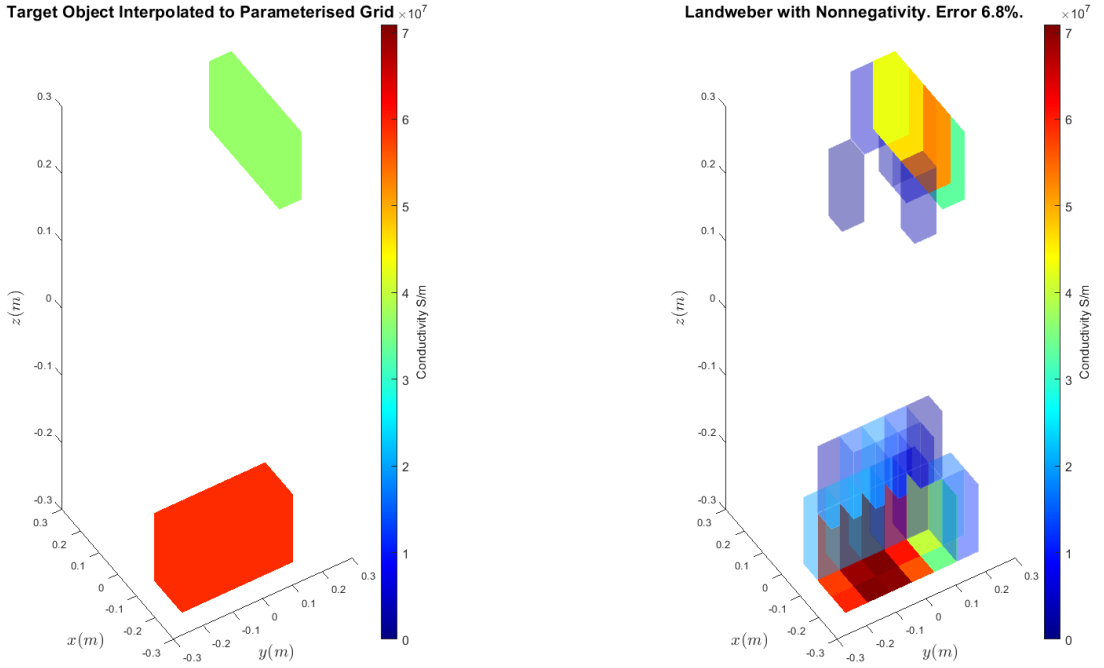
(b) Sixty-Six Coils.

Figure 6.24: The image reconstruction and associated behaviour of the constrained Landweber iterative reconstruction algorithm for a system which applies a six harmonic-inspired magnetic fields, comparing four coil and sixty-six coil systems.

cessfully approaching the target structure. The conductivity is consistently over-estimated with either system but is roughly proportionally correct. The main deviation from this is Figure 6.23, which produces the most accurate characterising image of the upper object, comparable to Figure 6.8 regardless of the fields' orthogonality. The more orthogonal



(a) Four Coils.



(b) Sixty-Six Coils.

Figure 6.25: The image reconstruction and associated behaviour of the constrained Landweber iterative reconstruction algorithm for a system which applies a twelve harmonic-inspired magnetic fields, comparing four coil and sixty-six coil systems.

fields persistently produce better images and usually have a less significant overestimation of the peak conductivity. The overestimate also reduces consistently as more profiles are added from the sixty-six coil system.

Sub-figures 6.24b and 6.25b are interesting because subjective interpretation implies twelve fields produce a slightly lower quality image. The twelve fields result contains fewer noise voxels and improved identification of the copper object but a lower quality characterisation of the upper aluminium object, highlighting a limitation associated with any attempts to quantitatively assess image quality. Other error metrics were considered but proved to be a worse match to the subjective assessment, particularly when evaluating iterative behaviour like in Figure 6.13.

Overall, this problem successfully verifies the main hypothesised features of our novel approach to MIT. Applying an increasing number of spatially unique magnetic fields as part of an augmented inversion problem leads to significant, demonstrable improvements in the quality of the final solutions. Increased field orthogonality further benefits the quality and minimises the computational requirements to reach semi-convergence, accelerating solution speed. Iterative inversion algorithms are superior to the traditional Tikhonov methods, which cannot solve any of the twenty-four sampled problems. Individual algorithms have specific advantages which may make them more suitable for various end-users. In combination, these verify that the variation associated with our MIT approach is related to the increasing augmentation of magnetic field profiles rather than the careful choice of apriori information. Individually, the results demonstrate an ability to image and characterise two dissimilar metallic objects in a cluttered environment. The ability to rapidly produce multiple, high-quality solutions also suggests there is potential to improve results further through post-processed combinations.

6.1.1 Varying The Conductivity Distribution

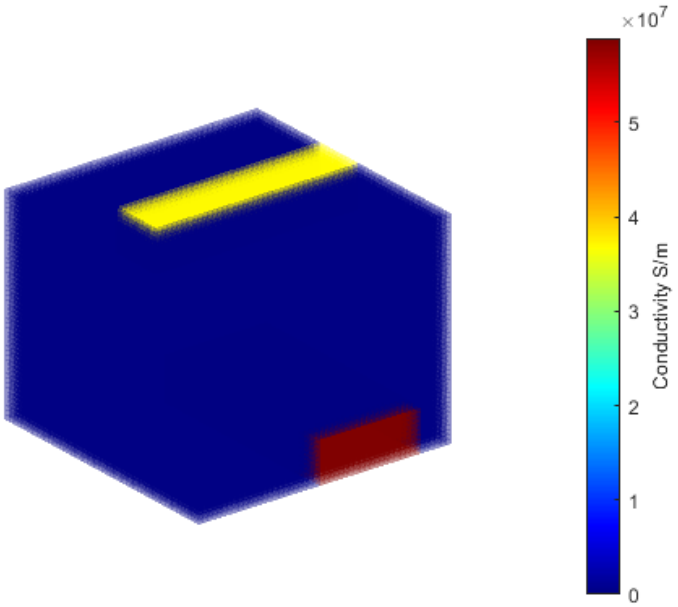
Given the success of the preceding analysis, it is worth considering the validity of this approach with different conductivity distributions, particularly one that has a poor functional overlap with the high magnitude components of the target spherical harmonic magnetic fields. The high-quality fields generated by sixty-six coils performed consistently better in the previous discussion, so evaluation uses those results exclusively. Through

the following two problems, we further demonstrate the benefits of applying orthogonal magnetic field landscapes, suggesting our proposal is robust even in less ideal conditions.

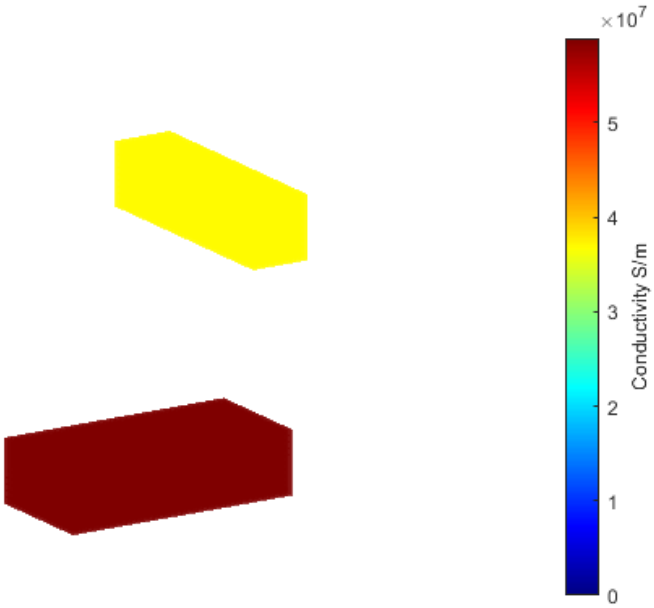
The existing sensitivity calculations from Equation 3.65 can be reused, with the location of the conductive blocks in the forward model adjusted to the positions shown in Figure 6.26. New signals are then produced with minimal computations using a simple matrix operation and combined with another one-percent white Gaussian noise profile. Maintaining a two-hundred-and-sixteen voxel resolution allows the reuse of the existing sensitivity matrix. This remains an intentionally information deficient problem, designed to highlight how quasi-orthogonal fields capture additional information through their augmented combination. The condition number is unchanged, Figure 6.4, and the information efficiency maintains a form which is visually indistinguishable from Figure 6.5.

The quantitative data concerning the variation in image quality with the number of applied magnetic fields for this conductivity profile is in Figure 6.27. The relative performance of individual algorithms mirrors that seen in Figure 6.6, with the constrained Landweber algorithm producing the best image quality. The CART algorithm again fails to converge in the allowed iterations, indicating that further improvements in image quality are probable, as demonstrated previously in Figure 6.15.

There are clear benefits to increasing the number of augmented fields regardless of the algorithm when varying from one to four. The quantitative variation then becomes more algorithm-dependent, but the best-performing algorithms continue to show improvements in image quality beyond this. The best qualitatively assessed image occurs using the Landweber algorithm with twelve magnetic fields, seen in Figure 6.28. The two objects can be clearly identified, and their dimensions are characterised with minimal background noise.



(a) The new conductivity distribution plotted from the angle in Figure 6.3a.



(b) Another view using variable transparency.

Figure 6.26: The second conductivity distribution considered, utilising the same copper (5.89×10^7 S/m) and aluminium (3.69×10^7 S/m) blocks at different locations within the imaging volume. Blue in 6.26a again identifies the empty-space regions of the imaging volume.

A further experiment considers another conductivity profile further from the ideal, Figure 6.29. The quantitative data concerning the variation in image quality with the number

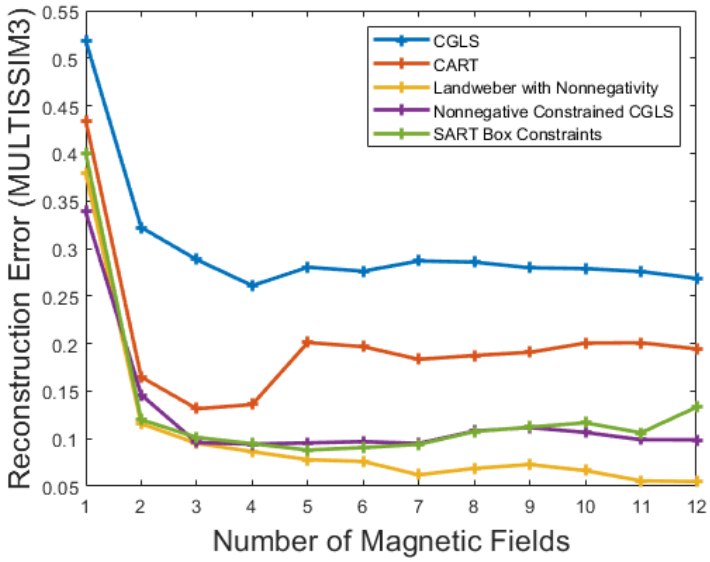


Figure 6.27: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, with increasing numbers of highly orthogonal magnetic fields produced by a sixty-six coil system.

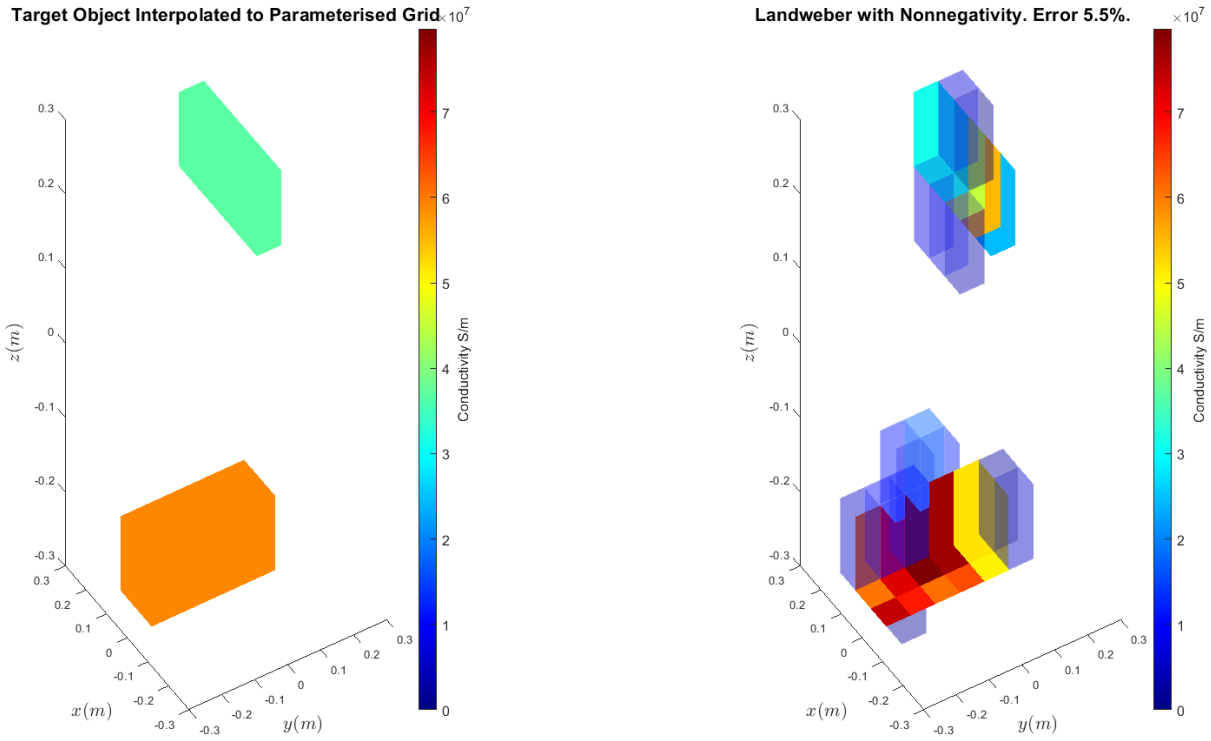
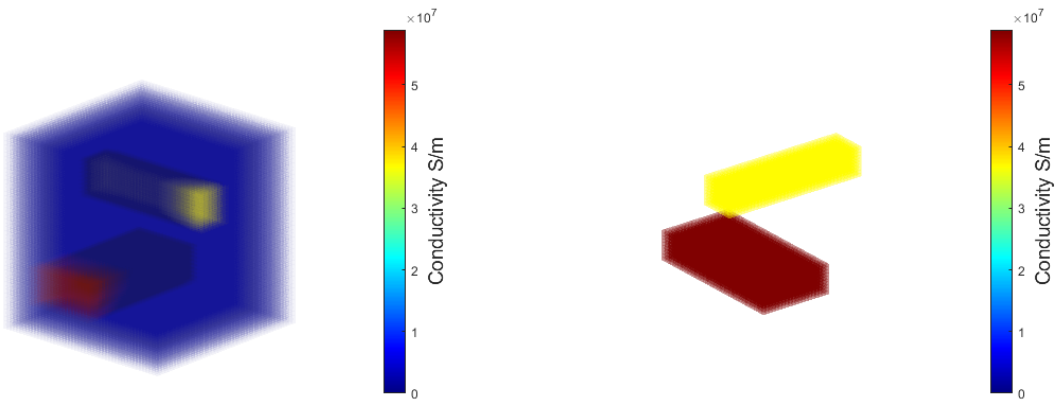


Figure 6.28: The best image reconstruction from the particular Landweber iterative reconstruction algorithm for a system which applies twelve harmonic-inspired magnetic fields generated by sixty-six coils to the second set of block positions.

of applied magnetic fields for this conductivity profile is in Figure 6.30.



(a) The new conductivity distribution. (b) Another view using variable transparency.

Figure 6.29: The second conductivity distribution considered, utilising the same copper (5.89×10^7 S/m) and aluminium (3.69×10^7 S/m) blocks at different locations within the imaging volume. Blue in 6.29a again identifies the empty-space regions of the imaging volume.

The relative performance of the algorithms is again similar, with an interchange between CART and CGLS similar to the four coils results for the first block position. The data

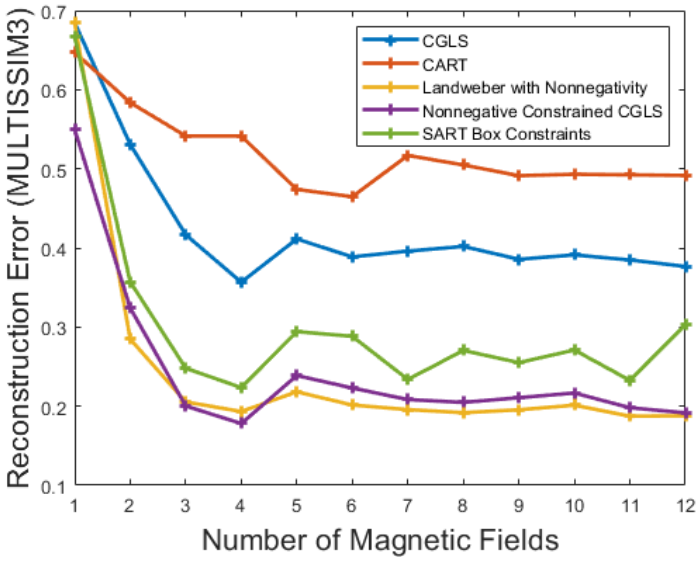


Figure 6.30: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, with increasing numbers of highly orthogonal magnetic fields produced by a sixty-six coil system for the least ideal set of block locations.

again affirms that there are quantifiable benefits to an increased number of orthogonal profiles, though this is less clear-cut than previous results. Partially this relates to the functional overlap between the profiles and the conductivity. The more complex fields

have a much weaker field in the conductive region, so the ensuing measurements are expected to capture less additional information. Quantitatively, the highest quality result is from the NNFCGLS algorithm with four fields with a MULTISSIM3 error of 17%, Figure 6.31. Subjectively, the twelve fields' result is of higher quality, having a lower overestimate of the copper conductivity, slightly fewer noise voxels, and suggesting more uniform conductivities in each block, Figure 6.32.

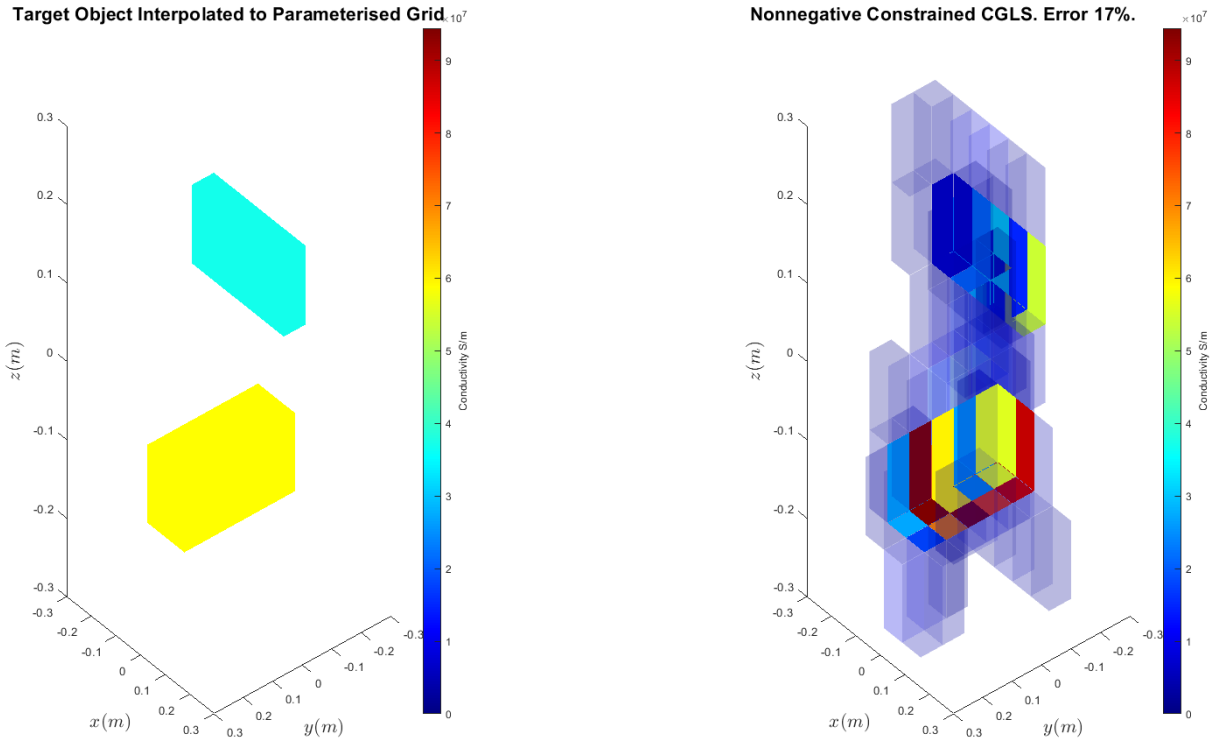


Figure 6.31: The best image reconstruction from the particular Landweber iterative reconstruction algorithm for a system which applies four harmonic-inspired magnetic fields generated by sixty-six coils to the third choice of block positions.

Overall one can conclude further success, with increasing field augmentation leading to demonstrable improvements in image quality. Given the object locations are less idealistic with respect to overlaps with the generated fields, this also suggests there is some robust benefit to the proposed orthogonal field landscapes even in a more general setting.

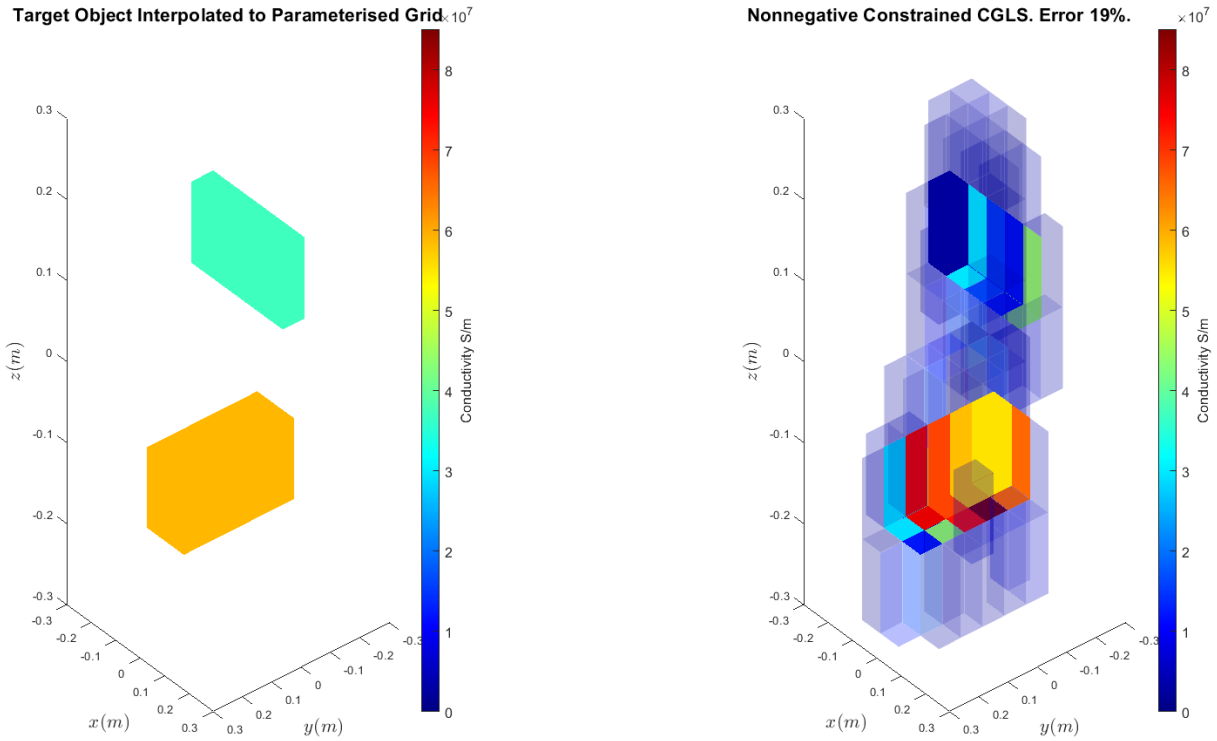


Figure 6.32: The best image reconstruction from the particular Landweber iterative reconstruction algorithm for a system which applies twelve harmonic-inspired magnetic fields generated by sixty-six coils to the third choice of block positions.

6.2 Optimising Sensor Positions

To this point, the choice of sensor positions has been entirely arbitrary. Practically, it is perhaps possible to identify an ideal set of sensor positions for the individual magnetic fields. However, in the augmented inversion problem, we consider measurements from numerous spatially unique profiles. Physically there will always be a trade-off between the ideal positions for each profile and those which are better for the system as a whole. Predicting the system-wide ideal locations is particularly challenging, so it is sensible then that some optimisation procedure could be particularly advantageous.

In optimisation, a local minima relates to a specific region of the search space, whilst a global minima is the minimum value for all possible parameters. Their equivalence is universally true for a quadratic polynomial, whereas a comparable quartic function has two minima for which the global minima is the lowest. MIT is likely to have numerous local minima, so algorithms from the MATLAB Global Optimisation toolbox are applied.

The process may minimise the condition number of a particular sensitivity matrix, the number of iterations to produce an inverted image of a specified quality or the overall error in image reconstruction. The latter is perhaps the most pertinent but is also extremely computationally expensive. Consider the process to produce one evaluation of image quality for a fixed system with precalculated excitation patterns: A forward model solves Equation 3.63 for all applied magnetic fields and sensor positions coupled to a particular conductivity profile, application of a noise profile occurs, then an inverse model solves Equation 3.65 at a target resolution, application of regularisation algorithms is necessary to produce solutions, then finally these are compared to some error metric.

For this reason, a proof of principle occurs using the four coil system with a target resolution of two-hundred-and-sixteen voxels, demonstrated in Figure 6.6a, varying the four sensor positions in an attempt to improve the best image. Figure 6.7a identifies the NNFCGLS algorithm to give an ideal balance between speed and image quality, with the best image produced with twelve magnetic fields. Non-linear constraints are imperative to prevent sensors from being placed inside the imaging region but within a practical range of the generating coils, excluding various common optimisation algorithms.

A comparison is made between two algorithms and their arbitrary counterpart. The Pattern Search algorithm is a good starting point because it provably converges to a solution with any constraint, even for non-smooth functions, but this is not guaranteed to be a global minimum. Non-linear constraints necessitate a particular variant, the Augmented Lagrangian Pattern Search algorithm, [207]. The Global Search algorithm aims to find a global minimum or identifies the best of many local minima depending on the problem, as described in [208].

The resulting optimised sensor positions from the Pattern Search are in Figure 6.33. The comparable Global Search positions are then in Figure 6.34. These deviate significantly

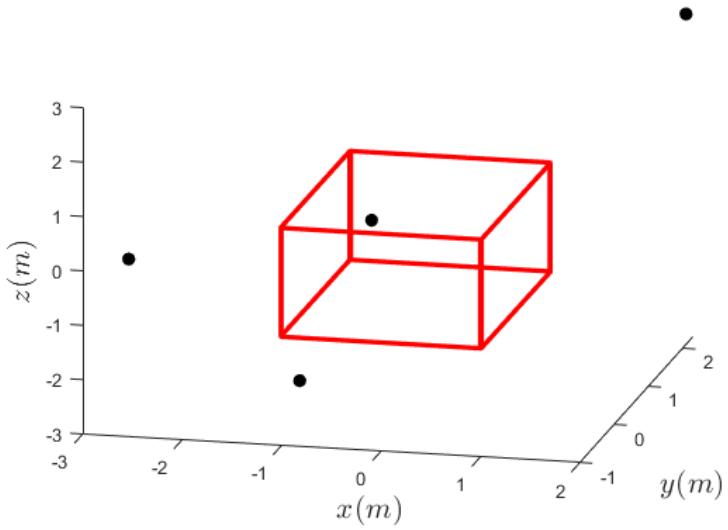


Figure 6.33: The optimised sensor distribution from the application of an Augmented Lagrangian Pattern Search algorithm, where non-linear constraints ensure sensors are located external to, but within a reasonable range of, the induction system.

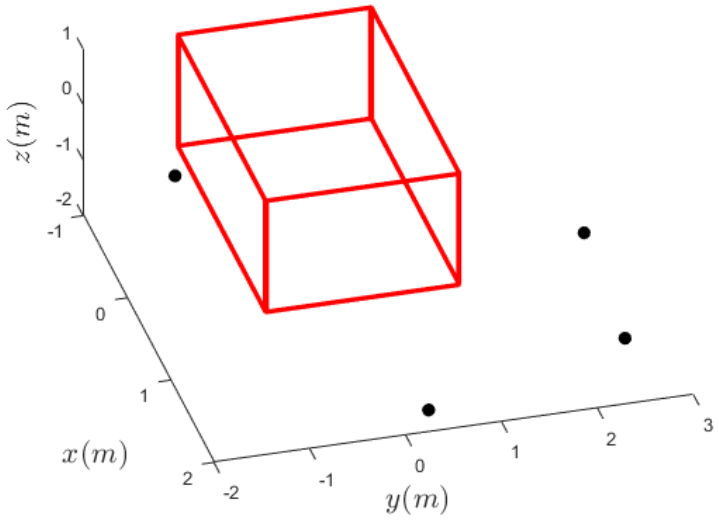


Figure 6.34: The optimised sensor distribution from the application of a Global Search algorithm, where non-linear constraints ensure sensors are located external to, but within a reasonable range of, the induction system.

from the original distribution, given in Figure 6.2. The likelihood of manually selecting these positions, especially given the relative dissimilarities, is low.

The related condition number data when varying between one and twelve fields is in Figure 6.35, with the optimised results demonstrating a systematically higher condition

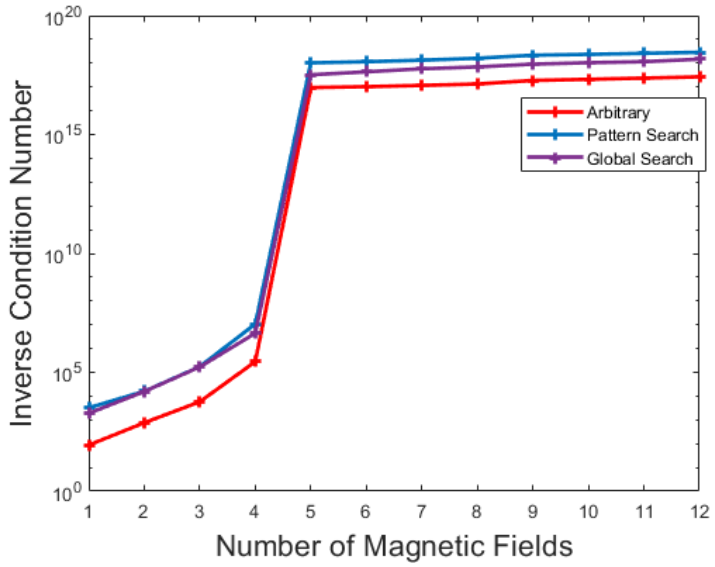


Figure 6.35: An overlay of the condition number values when varying the number of magnetic field landscapes comparing the arbitrarily chosen and Pattern Search optimised sensor positions.

number. This implies a greater sensitivity to noise, which appears to be an unintended detrimental consequence of this optimisation. The subsequent Information Efficiency val-

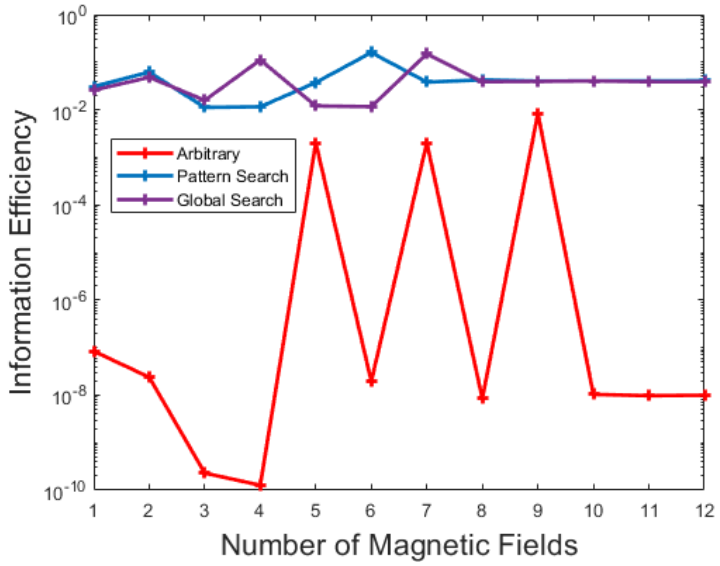


Figure 6.36: An overlay of the Information Efficiency values when varying the number of magnetic field landscapes comparing the arbitrarily chosen and Pattern Search optimised sensor positions.

ues are in Figure 6.36, with the optimised results indicating a substantial improvement in average model resolution regardless of the number of magnetic fields. The two optimised systems overlap significantly, with obvious differences only between four and seven fields.

These values indicate no clear preference for either of the optimised sensor distributions.

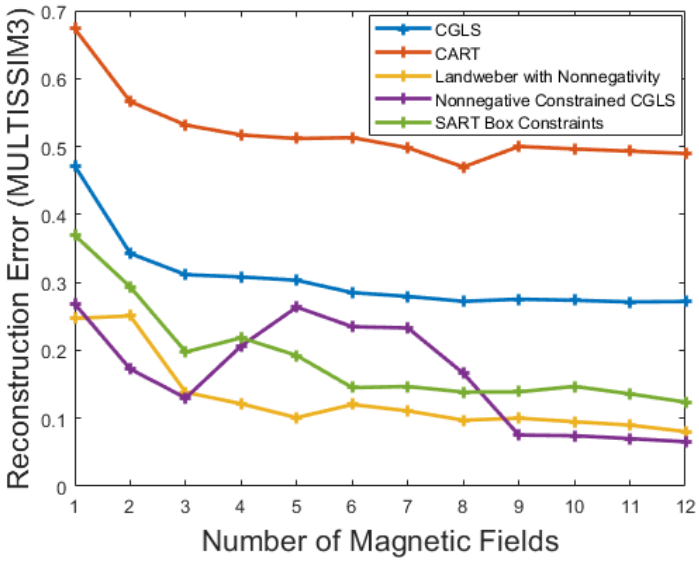


Figure 6.37: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, with increasing numbers of magnetic fields for sensor positions optimised by a Pattern Search.

The image quality data for the Pattern Search optimised sensors is in Figure 6.37, with the same experiment from the Global Search optimisation found in Figure 6.38. To highlight

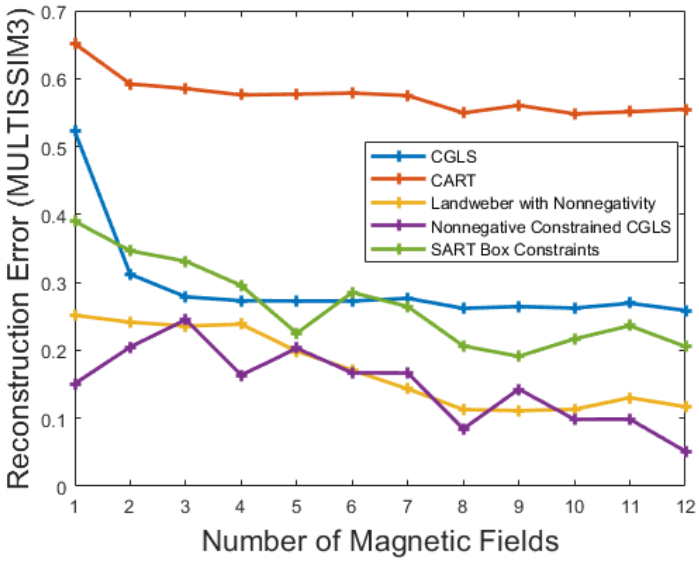


Figure 6.38: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, with increasing numbers of magnetic fields for sensor positions optimised by a Global Search.

the influence of these optimisations, an overlay compares the results for the NNFCGLS

algorithm in Figure 6.39. The optimised image reconstruction from the Pattern Search

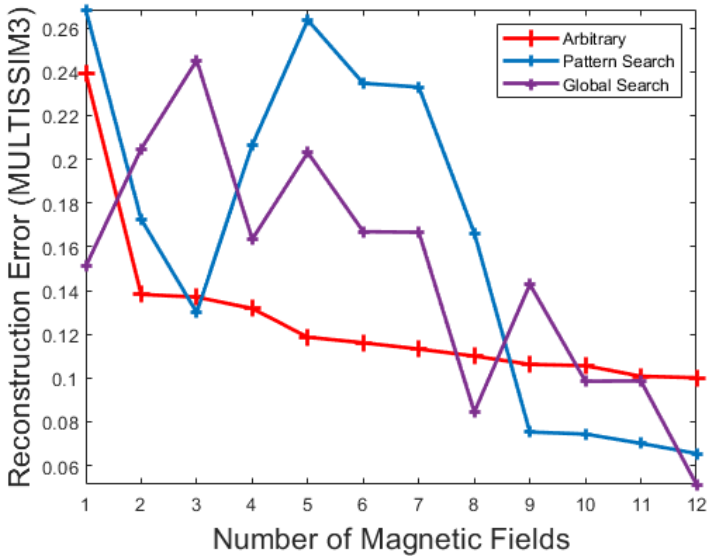


Figure 6.39: The variation in the MULTISSIM3 metric with the number of magnetic fields, for the three different sensor distributions when applying the NNFCGLS inversion algorithm utilised in the optimisation.

is in Figure 6.40. Compared to Figure 6.9, there is a substantial reduction in noise, with the blocks of conductivity characterised with greater uniformity. This image also demonstrates an improvement compared to the arbitrary sixty-six coils result, which was given in Figure 6.10, again concerning reduced noise and smoother object characterisation. The related optimised image from the Global Search sensor optimisation is in Figure 6.41. This demonstrates a further noise reduction compared to the Pattern optimised sensors, with additional improvements to object characterisation.

Reassuringly, regardless of algorithm or sensor positions, there is a clear trend that an increased number of augmented magnetic fields results in reduced image-reconstruction error, further validating the hypothesis. For this specific problem, both algorithms improve image quality, with the most significant change relating to the GlobalSearch algorithm. These match expectations given GlobalSearch intends to select the global minimum, and high-quality solutions already exist, leaving little space for significant improvements. The Pattern Search also demonstrates a slight enhancement to the Landweber solution with twelve magnetic fields, but all other algorithms perform worse than their arbitrary coun-

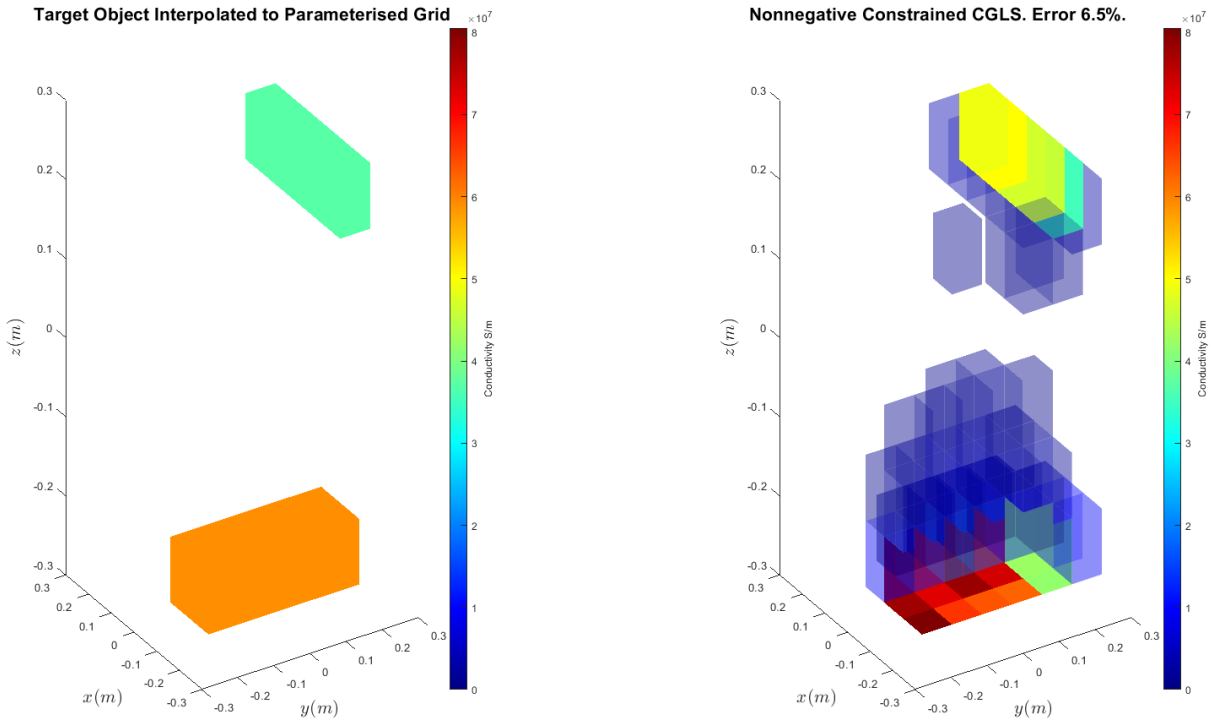


Figure 6.40: The pattern optimised image reconstruction for the NNFCGLS iterative reconstruction algorithm with the twelve fields generated by four coils.

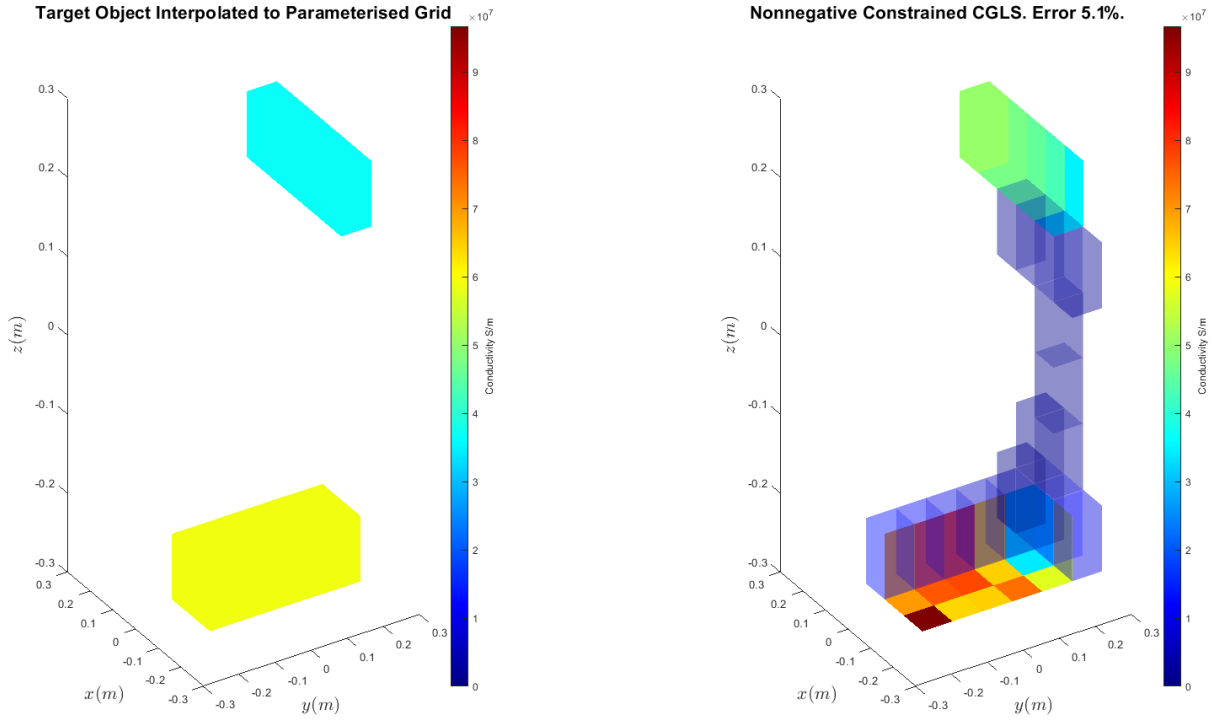


Figure 6.41: The globally optimised image reconstruction for the NNFCGLS iterative reconstruction algorithm with the twelve fields generated by four coils.

terpart. This is generally true for all five algorithms with fewer than twelve magnetic fields. The implication is that optimisation provides demonstrable improvements only on a specific fixed problem. In a limiting case such as fault analysis, this has inherent advantages. However, more frequently, this formulates sensor position optimisation as a dynamic, object-dependent problem requiring continual recalculation.

In future work, Multi-Objective optimisations may be beneficial, for example, minimisation of the reconstruction error and condition number, perhaps while maximising the strength of measured responses and the information efficiency. The Parento front should then highlight numerous solutions, with selection depending on an end-users specification. Other parameters, like resolution, the number of sensors, the size of the imaging volume, or the number/composition of field profiles, can also be optimised to improve inversion quality. Though, these are dependent on the specific goals and constraints of end-users.

To summarise, optimisations can successfully improve image quality in this generalised MIT problem, even for a well-solved problem. The results are computationally expensive and highly application specific. They are also unintuitive, and it would be unlikely they could be selected without the aid of an optimisation procedure. Given the implicit dependence on conductivity, these optimisations would almost certainly need to be dynamically calculated, which may limit their widescale applicability. Where the need for high-quality images exceeds the requirement for rapid image reconstruction, such as quality control, sensor position optimisation could be particularly effective.

6.3 Other Problems

6.3.1 Increasing The Number of Magnetic Fields

The sixty-six coil system utilised extensively up to this point can produce many more quasi-orthogonal fields than just twelve. To further validate the hypothesis, this section will examine whether applying additional magnetic profiles benefits image quality. The

initial problem, with the conductivity distribution in Figure 6.3, enables a direct comparison targeting a two-hundred-and-sixteen voxel image, but this time including between one and forty-eight magnetic fields. The new forward-modelled signals are again combined with one-percent white noise and solved with the same five algorithms.

The augmented application of magnetic fields is identical up to the ninth, then deviates from the previous model as they arise in the linear order of appearance associated with Figure 5.19. This new problem becomes well-determined with eighteen magnetic fields, tending to an over-determined problem with approximately 2.7 times more data than voxels with all forty-eight - a total of five-hundred and seventy-six unique measurements from four sensor positions.

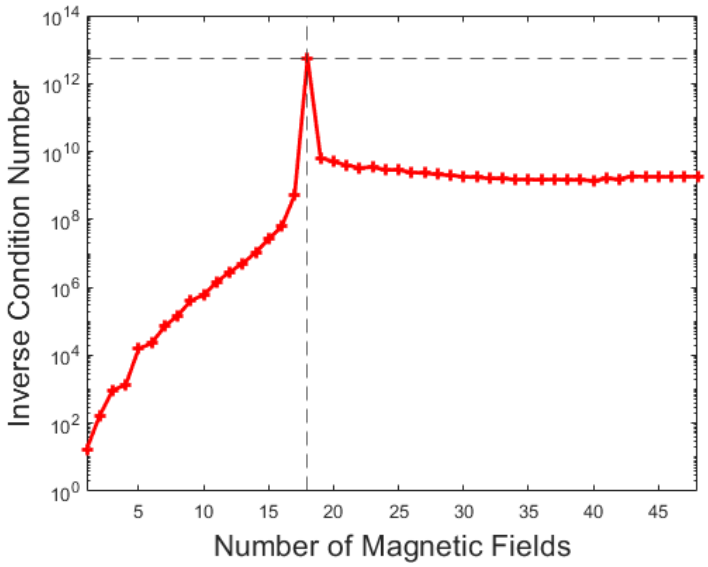


Figure 6.42: The variation in condition number as an increasing number of orthogonal magnetic field landscapes are augmented into an inversion problem which attempts to resolve two-hundred-and-sixteen voxels, with each step incorporating twelve additional measurements. The black dotted line marks where the problem becomes well-determined.

Qualitative assessment of this system is considered in terms of the condition number and the quality of image reconstructions. Figure 6.42 demonstrates that once the problem becomes well-determined, increasing the number of orthogonal magnetic field landscapes leads to a substantial fall in the condition number of the problem. The suggestion is

that the problem becomes less sensitive to noise when increasing the amount of data, potentially an additional advantage to increasing the number of orthogonal fields in an augmented inversion. Even at its maxima, the condition number is at least a factor of 10^4 lower than the four coils result in Figure 6.4.

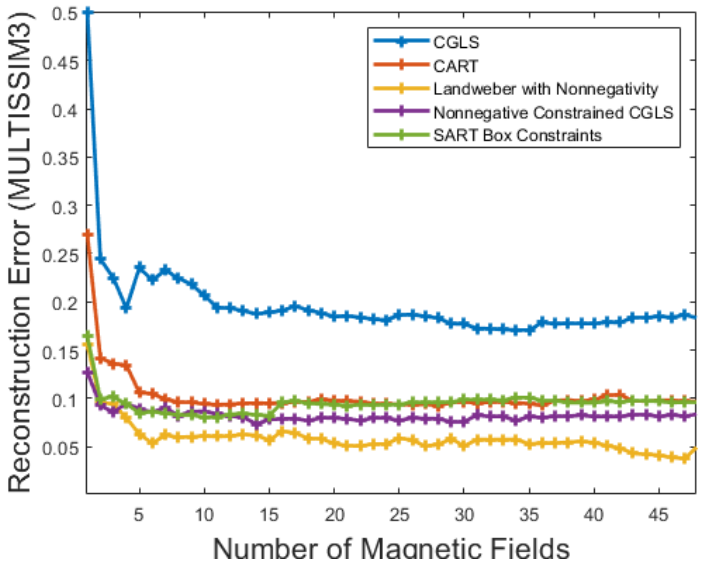


Figure 6.43: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, with increasing numbers of highly orthogonal magnetic fields up to the entire set produced by a sixty-six coil system.

Of paramount interest is the change in image quality, for which the results are in Figure 6.43. All five algorithms demonstrate some improvement beyond twelve magnetic field landscapes, the significance of which varies. The best performing Landweber algorithm exhibits continual improvement up to forty-seven applied fields, whilst the second best peaks with thirty-three.

The forty-seven fields image reconstruction is in Figure 6.44, which is contrasted with the comparable best image from the initial twelve fields in Figure 6.25b. This result is a smaller overestimate of the copper conductivity, with that object surrounded by fewer noise voxels. It also demonstrates a more accurate characterisation of the aluminium objects' conductivity, though it contains an additional noise voxel.

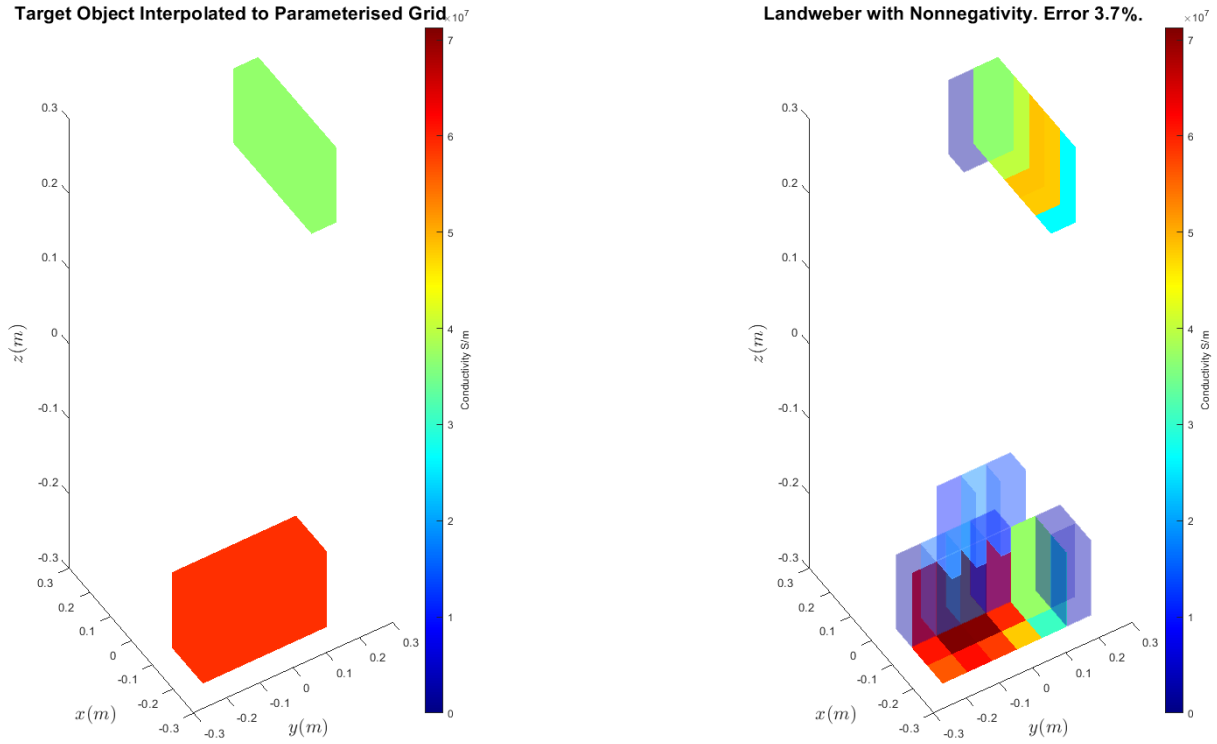


Figure 6.44: The best image reconstruction from the particular Landweber iterative reconstruction algorithm for a system which applies forty-seven harmonic-inspired magnetic fields generated by sixty-six coils.

This problem is solved particularly well with a limited number of fields, leaving negligible space for significant improvements with more data - especially with a 1% noise profile and large condition number. As previously discussed, it is unlikely that the resolution nor the specified number of measurements will be able to fully quantify the benefits of this approach to MIT. This becomes more pertinent for fields with higher n values, such as those in Figures 5.21 and 5.22, where a single voxel may incorporate, and average, many of the induced variations in the magnetic field.

Even for this very limited problem, the data support the hypothesis that increasing the number of quasi-orthogonal magnetic fields allows a quantifiable improvement in image quality. It also suggests the potential to reduce the sensitivity to noise in the inverse problem purely by varying the driving currents supplied to a modular system.

6.3.2 Increasing The Number of Measurements

Continuing with the sixty-six coil system, the other property we have yet to investigate is the significance of the number of measurements. Inversion theory suggests that additional data provide a closer approximation to a continuum and should reduce the sensitivity to noise of the inversion problem. This should, in turn, be beneficial to the final image quality. As previously demonstrated, the actual positions of individual sensors are incredibly influential on the result. Hence a future investigation may consider optimisation of a larger sensor array.

The obvious approach to considering more measurements is to simply take one at the

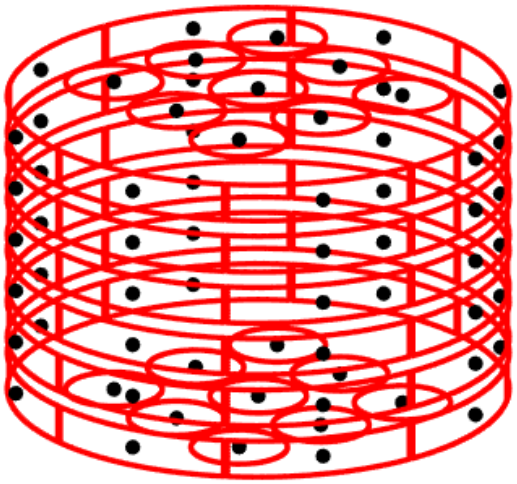


Figure 6.45: A system which uses sixty-six measurement positions, placed at the centre of each induction coil.

centre of each induction coil, shown in Figure 6.45. This has analogies with conventional MIT systems which may use off-coils as detectors or have detection coils placed close to excitation coils. The result is a problem with one-hundred and ninety-eight (sixty-six by three-axis) measurements per applied magnetic field. Maintaining the remaining parameters from the existing two-hundred-and-sixteen voxel inversion of the conductivity profile in Figure 6.3 this leads to an excess of data, with multiple fields reaching a $44:1$ ratio

when applying all forty-eight.

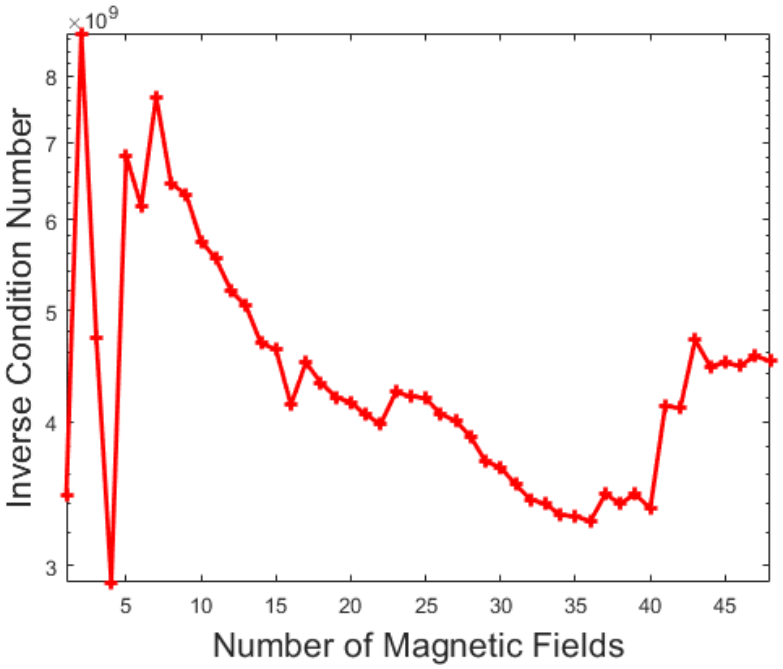


Figure 6.46: The variation in condition number when a significantly larger number of measurements are made of each orthogonal magnetic field.

For this over-determined problem, the condition number becomes significantly more stable but shows some improvement as the number of fields increases, Figure 6.46. It peaks roughly two orders of magnitude below the well-determined result in Figure 6.42, and is typically an order of magnitude better than the comparable over-determined results. The associated image quality data is in Figure 6.47. There is no significant change in the relative performance of the individual algorithms, though the NNFCGLS and SART algorithms now perform very similarly.

Separately, contrasting with Figure 6.43, all the algorithms tend to produce higher quality images, though the difference is marginally worse with CART. CGLS maintains the curve shape but performs substantially better (at least 5%) at all points. SART also performs consistently better, and the additional measurements remove the drop-off in image quality seen above fifteen augmented profiles. The best image occurs with thirteen profiles, but the variation in error is flat afterwards, peaking at just 0.3%.

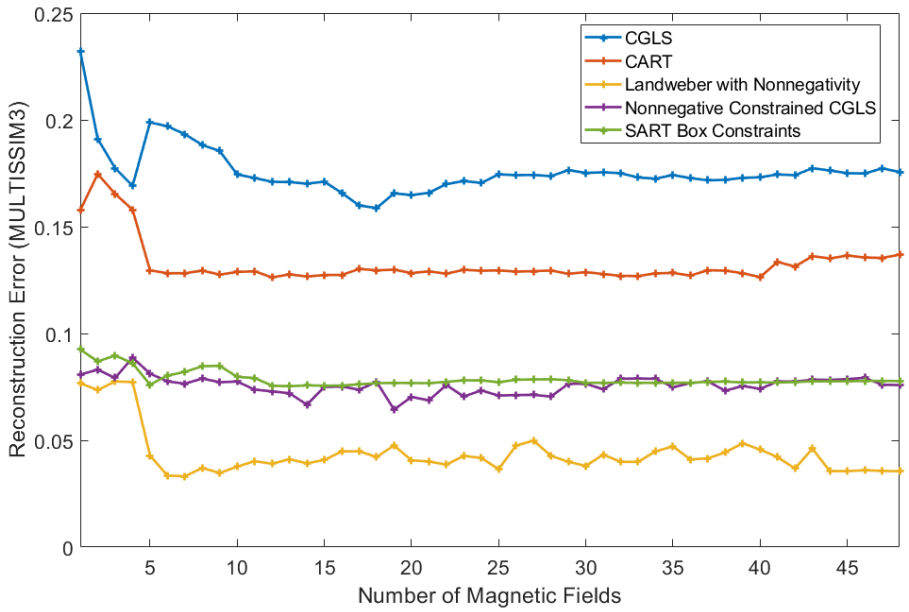


Figure 6.47: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, with increasing numbers of highly orthogonal magnetic fields from an idealised sixty-six coil system with the same number of measurements per profile.

Demonstrable improvements also transpire with the Landweber algorithm, where the best image has 3.3% error using seven profiles, with a similarly flat variation of 1.7% afterwards. One sample, using forty-three profiles, performs worse than its four measurement counterpart, but only by 0.3% on a solution which achieves significant accuracy. The highest-quality solution is in Figure 6.48.

Overall, this data demonstrates that increasing the number of measurements will improve the conditioning of the problem regardless of the number of field profiles. With more measurements, the conditioning also improves with an increasing amount of augmented magnetic fields. There is a consistent benefit to the final image quality, though the marginal improvements at points undoubtedly relates to the simplicity of the problem.

6.3.3 High Resolution/Low Conductivity with Novel Algorithms

The final set of problems demonstrates the breadth of the proposed approach to MIT. They consider significantly higher resolution inversions, targetting a characterisation at

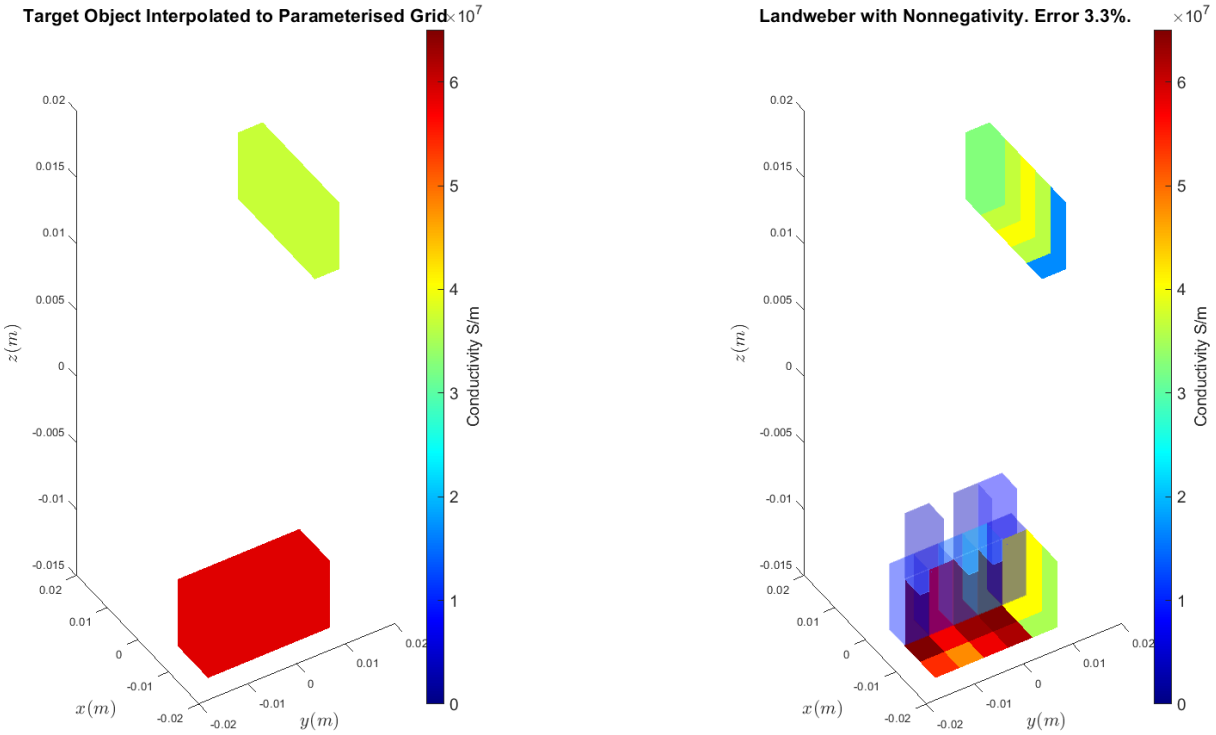
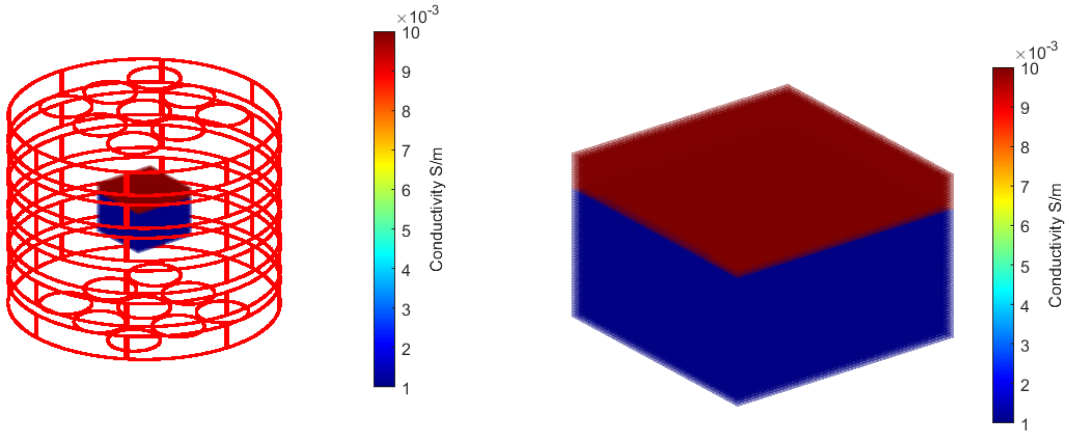


Figure 6.48: The best image reconstruction from the particular Landweber iterative reconstruction algorithm for a system which applies seven harmonic-inspired magnetic fields generated by sixty-six coils and measured at sixty-six points.

low conductivity, and several advanced, novel inversion algorithms. Consider the forty-eight profiles and four three-axis measurement positions studied in Section 6.3.1, which has a total of 576 measurements. A specific conductivity profile is in Figure 6.49, com-



(a) The conductivity plotted relative to the induction system.

(b) An alternative view of the conductivity.

Figure 6.49: The specified conductivity distribution consists of two separate layers (red, $\sigma = 1 \times 10^{-2}$ S/m and blue, $\sigma = 1 \times 10^{-3}$ S/m), both of very low conductivity similar to that in some geophysical problems.

pleting the forward model. It is a simple two-layer low conductivity structure where the entire imaging volume produces an eddy response.

The inverse model targets an image with 8000 voxels, or a $20 \times 20 \times 20$ grid, a significantly higher resolution. Individual voxels are of normalised side length 0.05, compared to $0.16\dot{7}$ in the previous problem. Against all forty-eight fields, this is approximately 14 voxels for each measurement. This problem also provides space to demonstrate the applicability of additional novel algorithms from Chapter 4 to MIT. As this problem is in a different conductivity range with an unrelated structure, the best-choice of algorithm may differ significantly.

A limited quantitative evaluation using fields $[1, 6, 12, 18, 24, 30, 36, 42, 48]$ is in Figure 6.50. The omission of a complete analysis with a variable number of magnetic fields builds on an assumption from the preceding data that using all forty-eight fields should produce something near the optimal solution. This links to significant increases in calculation time both in the sensitivity calculation and the inversion stage. There is a heavy weight towards the inversion step, with the TTLS algorithm requiring approximately ten minutes to evaluate and the CART algorithm requiring fifty minutes. For this problem, seventeen algorithms from the library could solve it with an error below 0.5%. These include: TTLS, NNFCGLS, Brakhage's ν-method, Non-negative Landweber, Enriched CGLS, Adaptively-constrained IRN and FISTA.

Two of the best results from the TTLS and first-derivative preconditioned CGLS algorithms are in Figure 6.51 and 6.52, with the latter giving the best available solution at 0.17% error. These highlight the issues associated with analytic error measurements for image reconstruction, even when using metrics designed for this class of problem. From the previous discussion, the error values quoted for both algorithms imply an incredibly high-quality result.

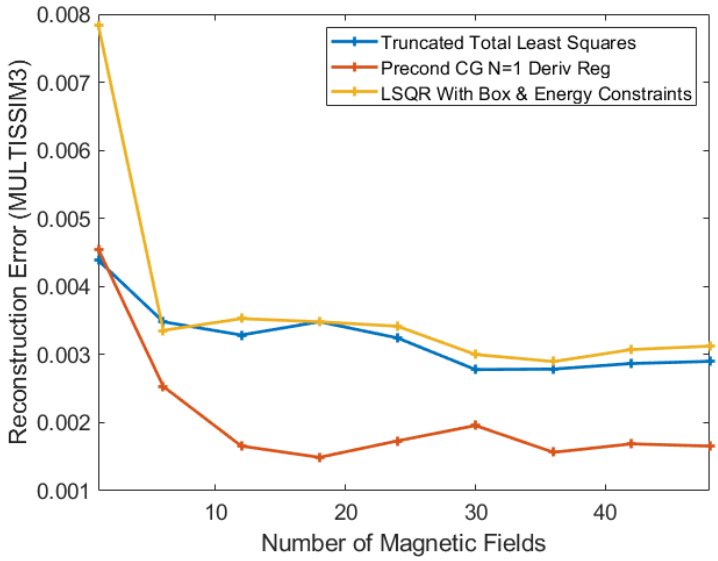
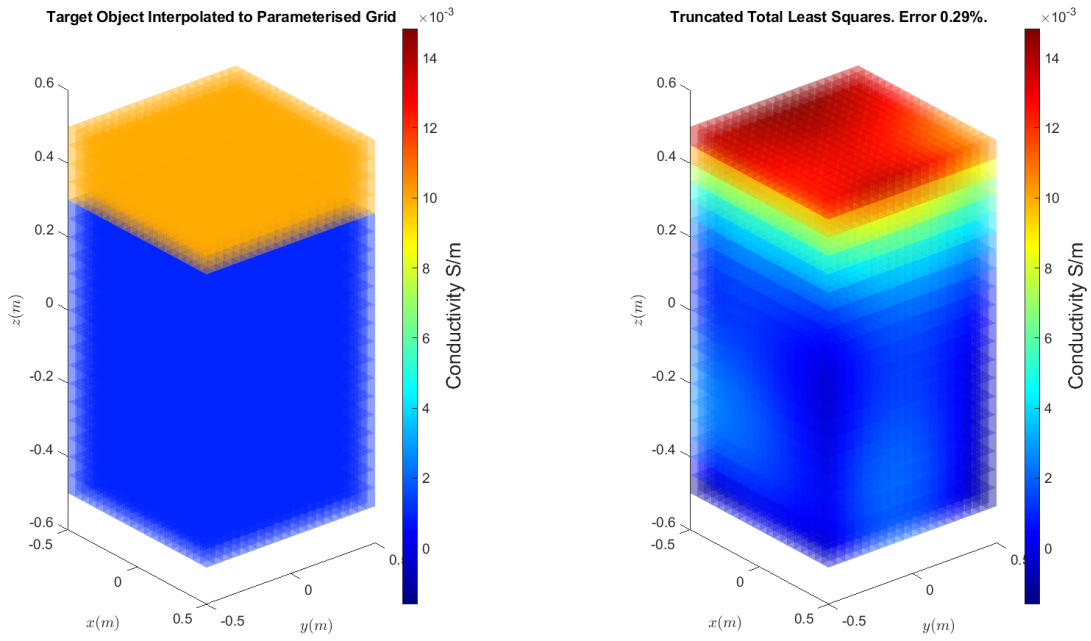


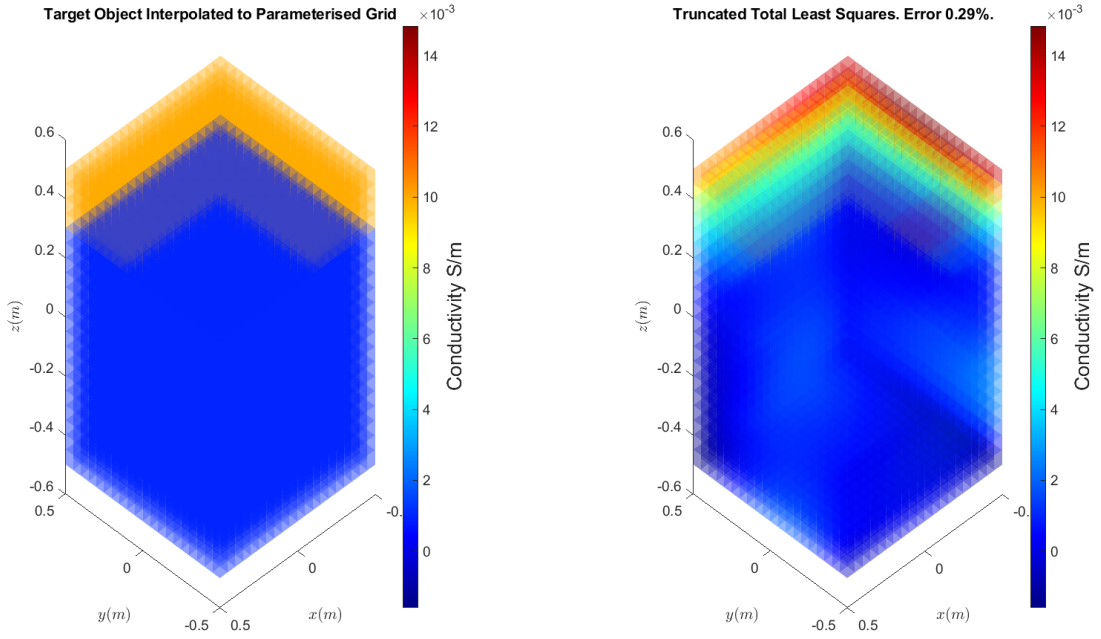
Figure 6.50: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, with increasing numbers of highly orthogonal magnetic fields from an idealised sixty-six coil system with four measurements per profile reconstructing a problem with eight-thousand parameters.

Yet human instinct when interpreting the visual results is that these images are much worse than the metric suggests. The RMS is similarly unhelpful, at 0.0016 and 0.0013 respectively. In part, these results will relate to the fact that there is an accurate characterisation of the vast majority of voxels. Much of the error is due to smoothing effects failing to reproduce the smaller, upper high conductivity layer. Characterisation of the dimensions of this layer is better with the CGLS variant, weighted against greater variations across the bulk of the lower layer. Both solutions also feature some small anomalously negative conductivities (darkest blue). These are in areas with no overlap, linking directly to numerical instabilities in the problem, so post-processing could filter these by taking a weighted combination of the two results.

Constrained Least Squares is a viable alternative algorithm producing a result with a slightly lower error metric than TTLS (0.31%) but no infeasible values. It uses physical box constraints and a non-negative energy constraint. The result is in Figure 6.53, with the second subfigure using variable transparency to highlight that this approach suggests some voxels of exactly zero conductivity. The variable transparency allows the visualisation of otherwise hidden voxels on the projected images. Careful selection of an



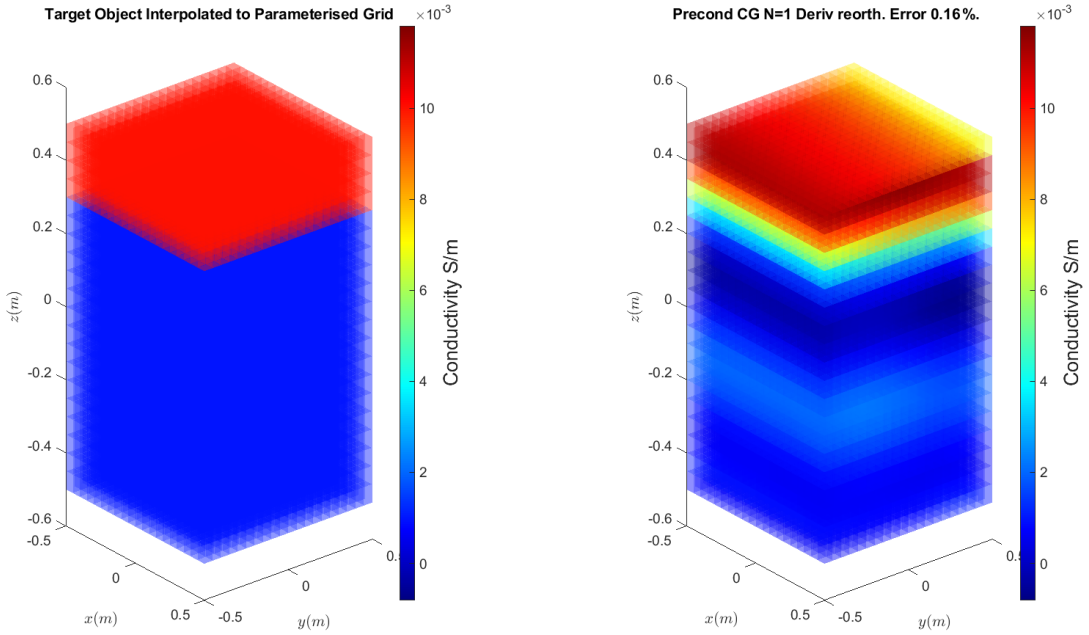
(a) A solution view.



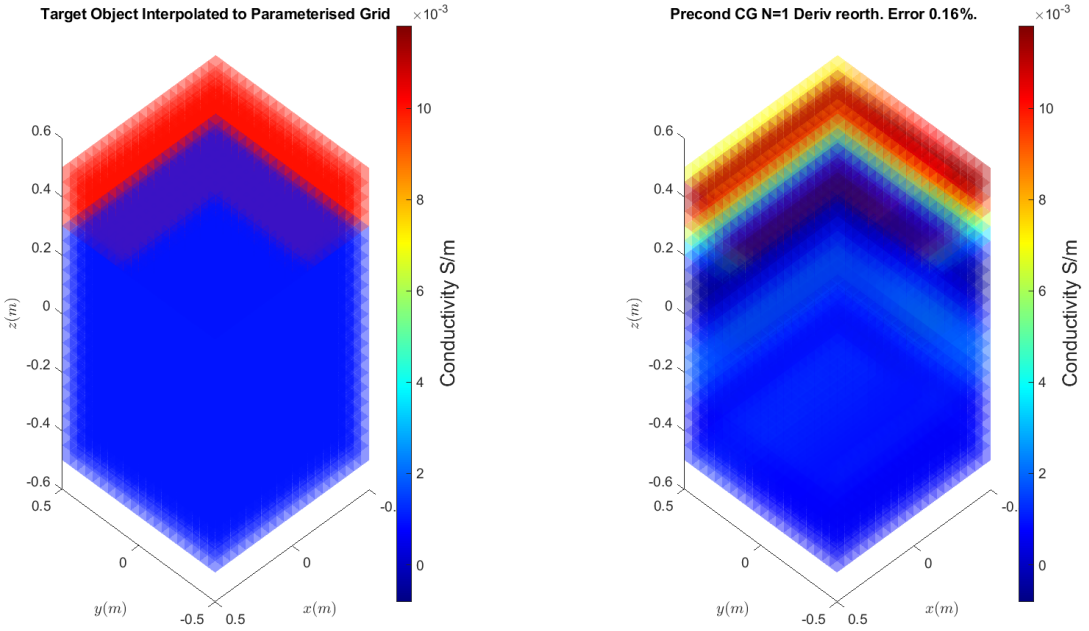
(b) An alternative view.

Figure 6.51: The best image reconstruction from the Truncated Total Least Squares (TTLS) algorithm for a system which applies forty-eight orthogonal magnetic fields measured at four locations. Subfigures compare diagonally opposing corners to aid comprehension.

appropriate positive energy constraint may further benefit solution quality but is difficult to estimate for an unknown problem. To the authors' knowledge, the application of TTLS and Constrained Least Squares to MIT is novel, and both have demonstrable value in this regime.



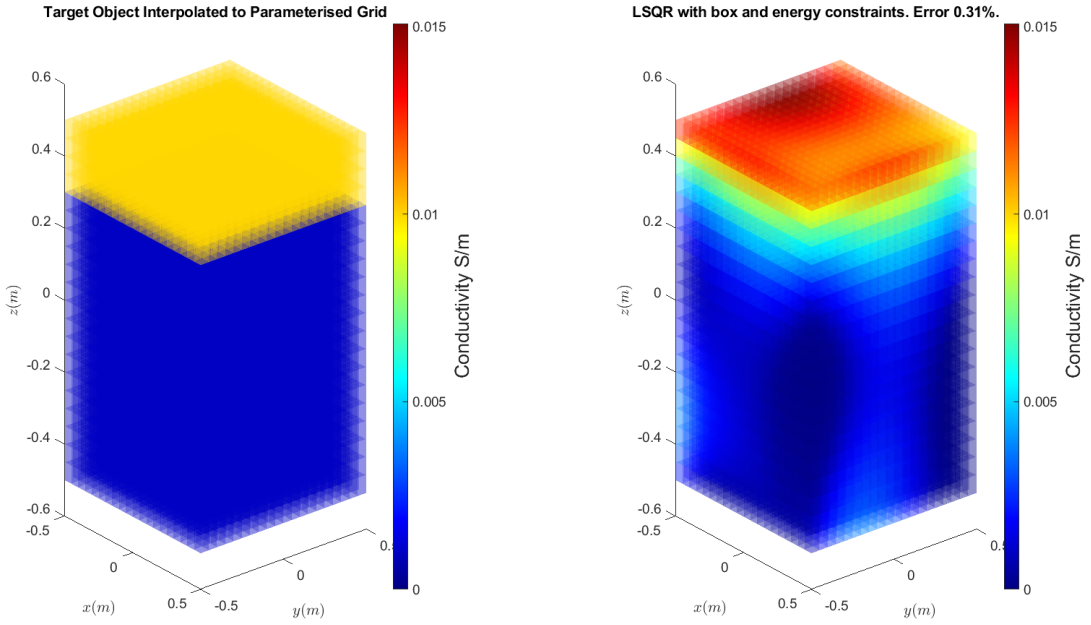
(a) A solution view.



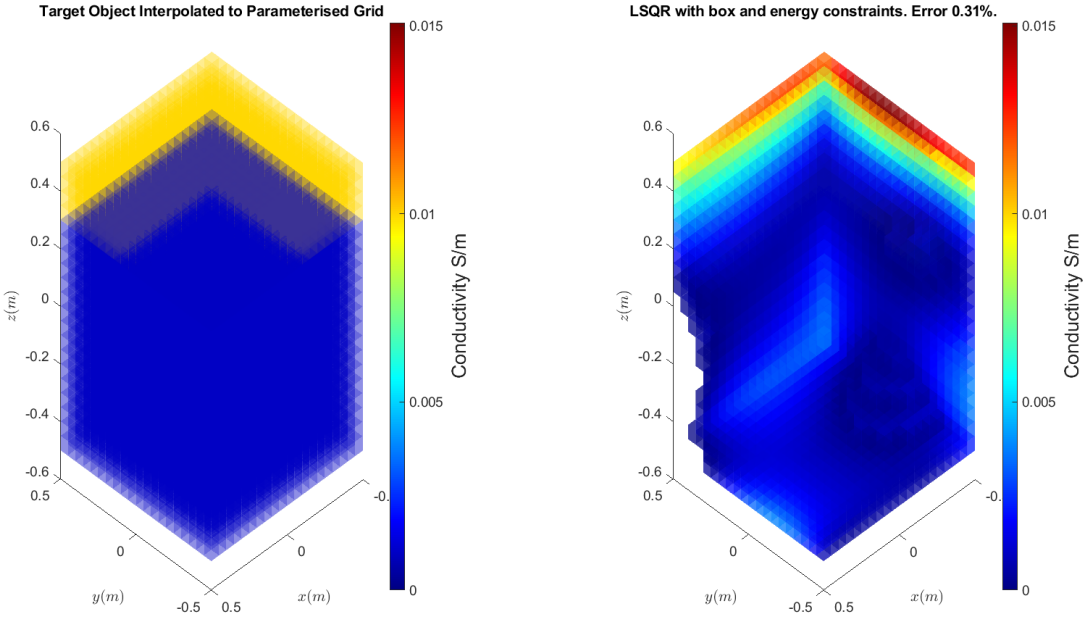
(b) An alternative view.

Figure 6.52: The best image reconstruction from a first-derivative preconditioned CGLS algorithm for a system which applies forty-eight orthogonal magnetic fields measured at four locations. Subfigures compare diagonally opposing corners to aid comprehension.

The conductivity profile is reasonably simplistic, so a second problem considers including an occlusive object. As seen in Figure 6.54, this is uniform along the y-axis and given the average conductivity of the two layers. The forward model again uses a consistent forty-eight fields from sixty-six coils with four measurement positions.



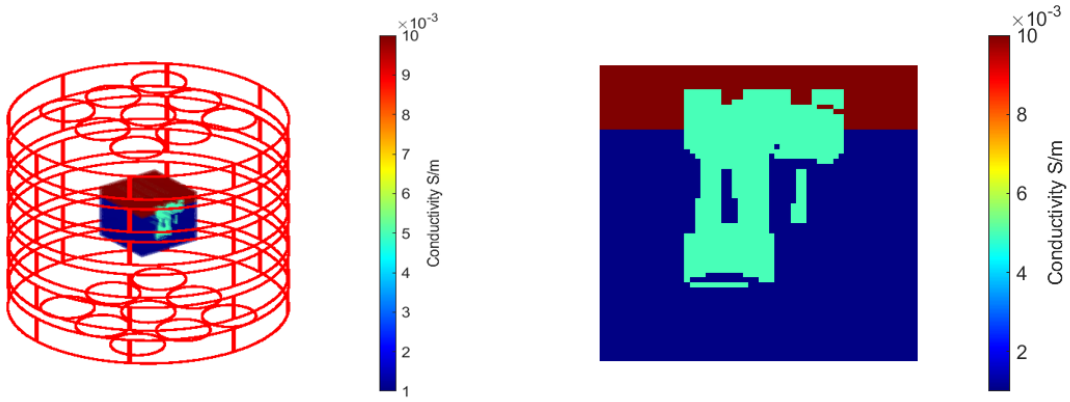
(a) A solution view.



(b) An alternative view.

Figure 6.53: The best image reconstruction from a Constrained Least Squares algorithm, using physical box limits (zero & silver) and a non-negative energy value, for a system which applies forty-eight orthogonal magnetic fields measured at four locations. Subfigures compare diagonally opposing corners to aid comprehension.

Sixteen algorithms meet the 0.5% threshold specified for an occlusion-containing distribution. Three of those are of particular interest, producing some of the best results whilst also being novel for MIT: FISTA, a variant of IRN (using adaptive box constraints



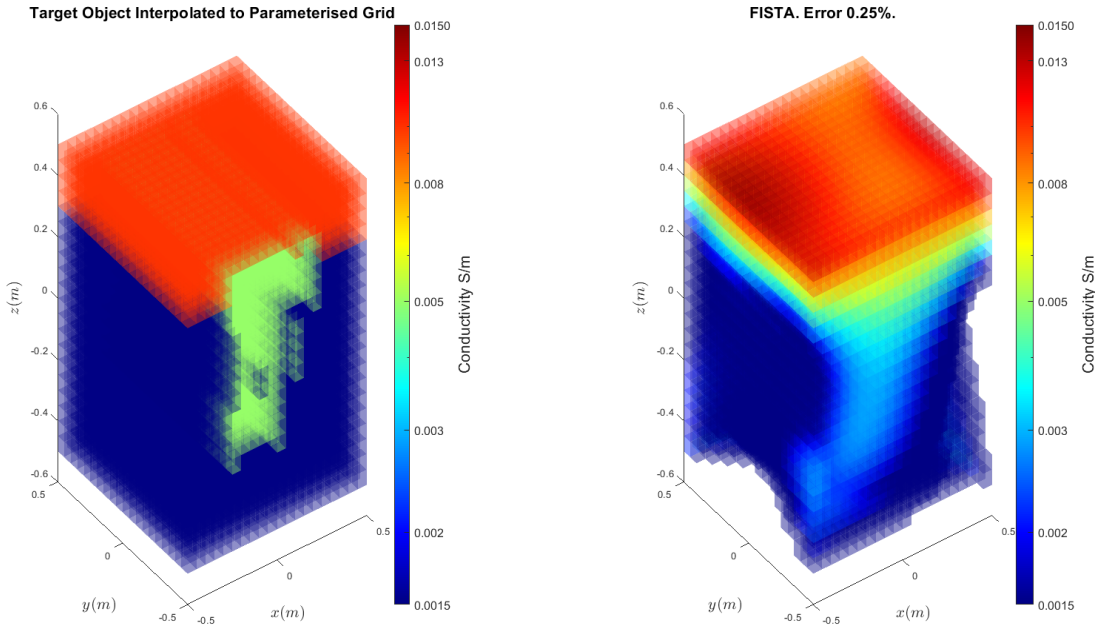
(a) A plot relative to the induction system. (b) An XZ view of the conductivity.

Figure 6.54: The specified conductivity distribution consists of two separate layers (red, $\sigma = 1 \times 10^{-}2$ S/m and blue, $\sigma = 1 \times 10^{-}3$ S/m) and an internal occlusion (green, $\sigma = 5 \times 10^{-}3$ S/m).

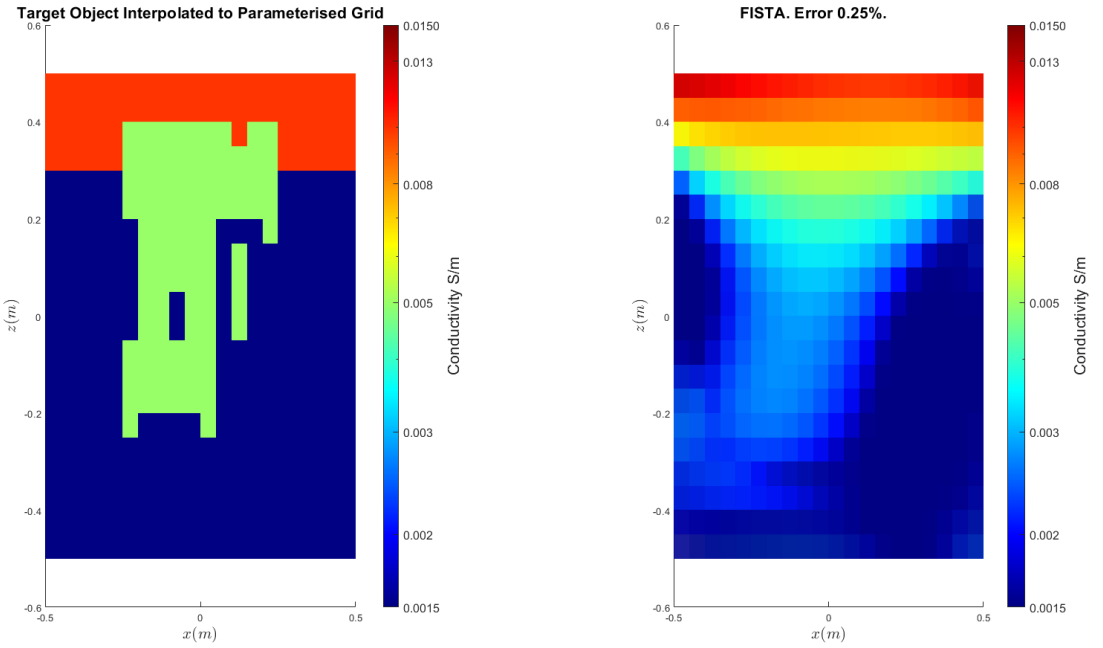
set at physical conductivity bounds, non-negativity and silver), and Brakhage's ν-method. The associated solutions are in Figures 6.55, 6.56, and 6.57. Qualitatively, the best result is from FISTA, followed by IRN and then the ν-method. All three results reproduce the two layers of conductivity reasonably well and identify an anomalous conductive occlusion embedded between them. That occlusion has roughly the expected symmetry along the Y-axis but underestimates the conductivity. In all cases, that object is an anomalous smudge rather than a precise characterisation of the individual features.

The FISTA and the IRN solutions feature some exact zeros, highlighted as transparent voxels. The implication is a poor characterisation of parts of the low conductivity layer, detracting from overall image quality. The zeros are similar between the two algorithms, which suggests that post-processed combinations of the two would provide little benefit. The ν solution contains only positive conductivity values with a minimal value of $1.3 \times 10^{-}3$, meaning it finds a more uniform lower layer without errant zero values but also over-estimates the conductivity of that area.

Other than minor variations in the image quality and particular noise voxels, the main reason for including these three similar results is differences in iterative behaviour, Figure 6.58. The ν algorithm continues to 5×10^5 iterations, but the error exceeds 50% after



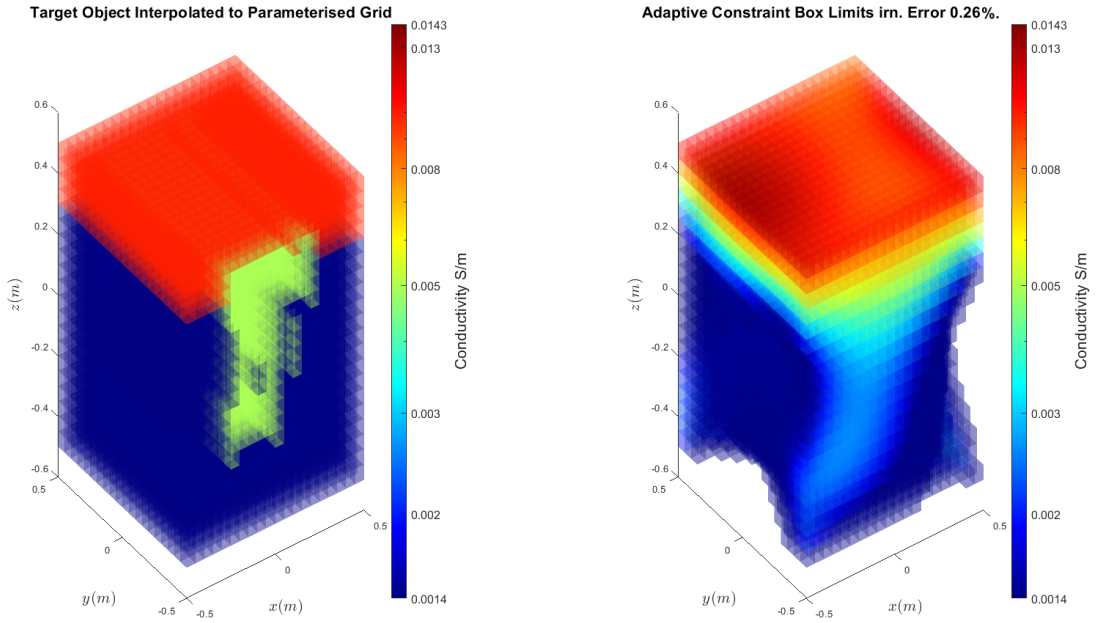
(a) A solution view.



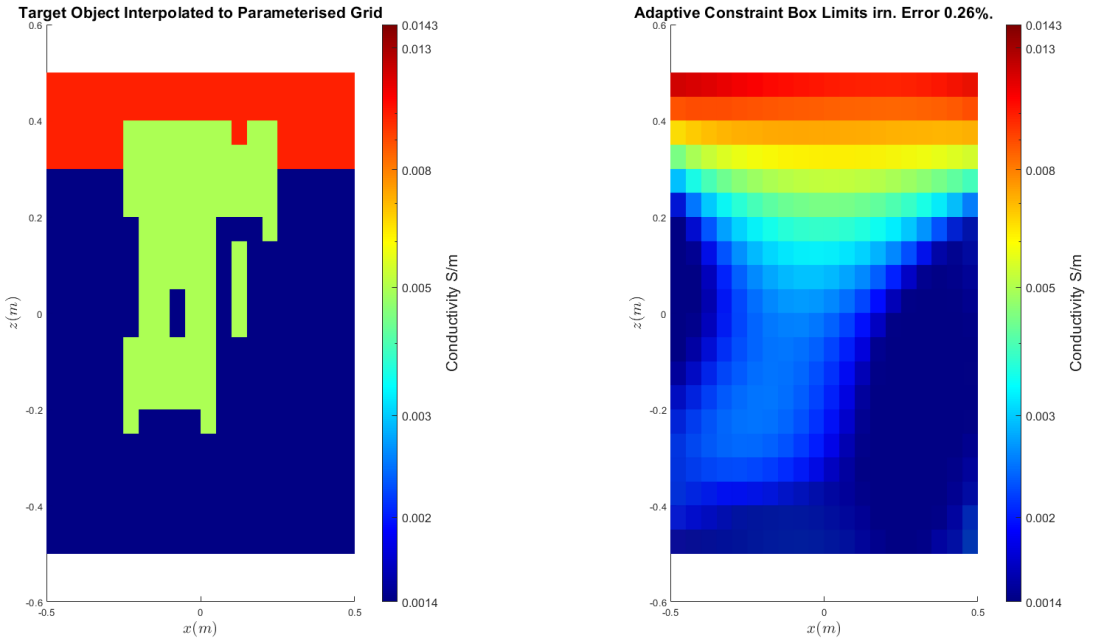
(b) An XZ view.

Figure 6.55: The best image reconstruction from a FISTA algorithm for a system which applies forty-eight orthogonal magnetic fields measured at four locations. Subfigures provide different viewpoints to aid interpretation.

1×10^4 iterations, and $\rightarrow$ 100% after 1×10^5 making it unstable. Total time was 846 seconds, with the optimal reached after 0.3s. The other two algorithms, FISTA and the



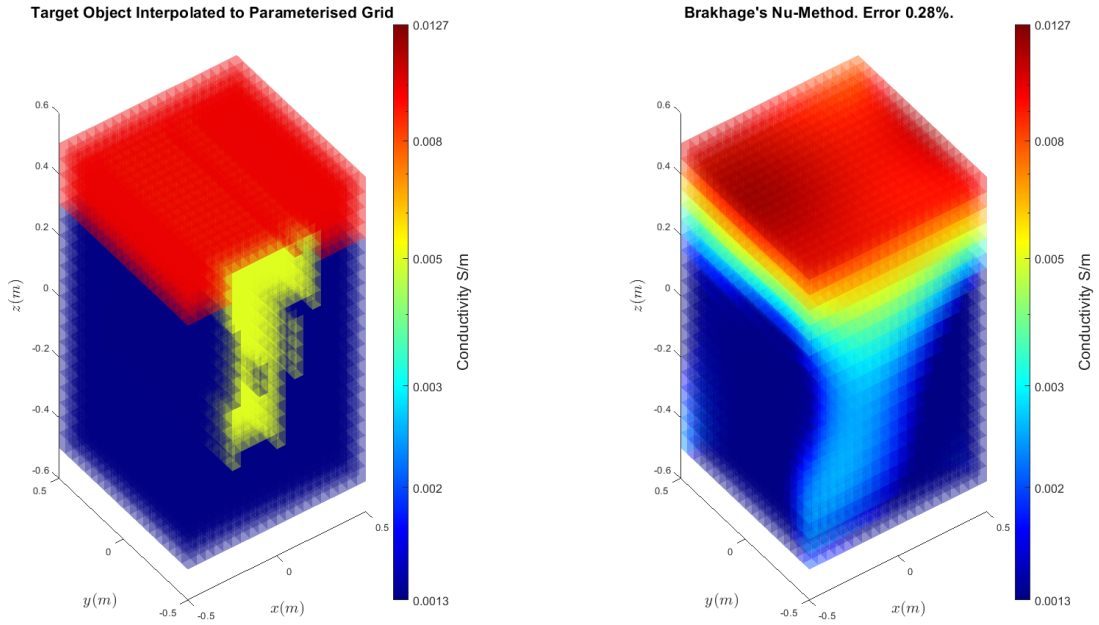
(a) A solution view.



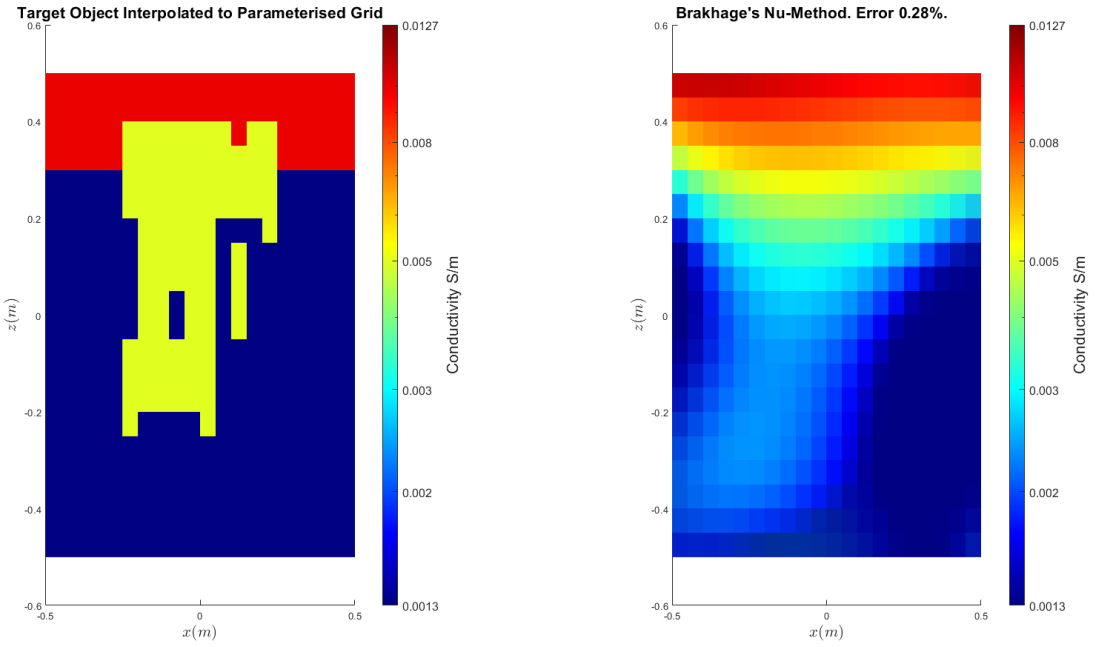
(b) An XZ view.

Figure 6.56: The best image reconstruction from a variant of the IRN algorithm for a system which applies forty-eight orthogonal magnetic fields measured at four locations. Subfigures provide different viewpoints to aid interpretation.

IRN variant, are automatically stopped by inbuilt stopping rules at 2.5×10^3 and 3.7×10^5 iterations with errors around 1.2%. FISTA took 9.3 seconds, with the optimal found after



(a) A solution view.



(b) An XZ view.

Figure 6.57: The best image reconstruction from Brakhage's ν-method for a system which applies forty-eight orthogonal magnetic fields measured at four locations. Subfigures provide different viewpoints to aid interpretation.

1.9 seconds, and the IRN variant ran for a total of 1.7×10^3 seconds, reaching an optimum after 222.5 seconds.

Like the CART algorithm, the IRN variant is much more stable after many iterations,

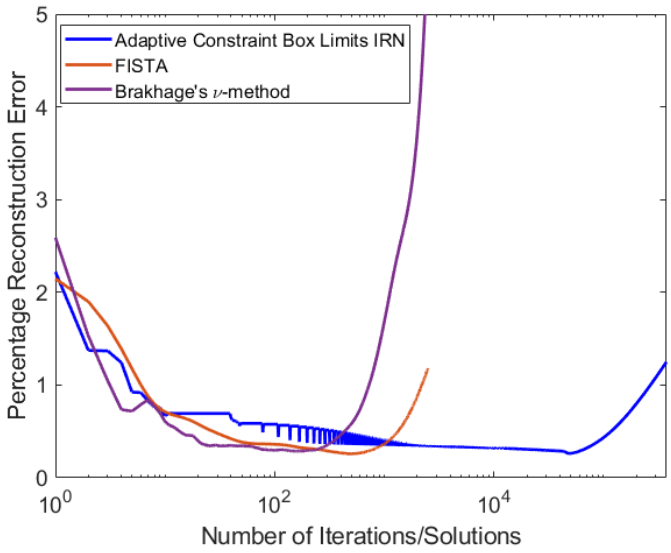


Figure 6.58: The contrasting iterative behaviour of the three sampled algorithms, as characterised by the MULTISSIM3 metric. The forward model uses the idealised sixty-six coil system, with four measurements per profile, to reproduce eight-thousand parameters for a two-layer conductivity distribution which also contains an occlusion.

making it more practical where inversion time is not a concern. The ν solution converges in the fewest iterations, which may be advantageous with increasing numbers of parameters.

Overall, these problems identify that forty-eight orthogonal fields can produce images of a filled conductive region at a significantly higher resolution, even with just four measurements per magnetic profile. These include reasonable characterisation of the conductivity values and the material structure. Identification and characterisation of an embedded conductive object is also possible, though precise feature imaging was not. Advanced inversion algorithms enable both inversions and appear to be novel in MIT. Future work could consider whether this is possible with a more reasonable number of measurements.

6.4 Future Work

There are a variety of investigations worth considering as research into orthogonal magnetic fields for MIT progresses. The conductivity distributions considered both here and throughout the literature are relatively simplistic. For security applications, such as baggage scanning, considerations of significantly more complex structures will be essential. This may include the shape of objects, the distribution and quantities of material throughout the scanner, the variety of possible properties incorporating inhomogeneity, and variability in the number of potentially overlapping items. Initial steps may consider a more cluttered environment, containing perhaps three or four different objects. Shielded objects, for example, uranium in a lead container would be of particular interest given conventional X-ray scanning is unable to image such objects.

Crucial work towards such goals will concern adapting the voxel meshing procedure. Conformal elements will become necessary, as will techniques such as adaptive mesh refinement to minimise the number of resolved parameters whilst maximising the quality in the inverse problem.

The vast variety of existing inversion algorithms is also worth further investigation. A notably interesting one is Enriched Least Squares. The augmented method of formulating a sensitivity matrix attempts to force spherical harmonic orthogonality into adjacent rows. Supplying a subspace with a related harmonic enrichment would sensibly benefit image quality and should positively impact solution speed.

A related formal study of algorithmic combination and post-processing techniques is also sensible. The main distinguishing feature between *good* solutions should relate to noise magnification, so combinations of algorithms are potentially a simple way to minimise error. These could be weighted by the model misfit - how well an individual solution solves the original system of equations.

Other work may consider optimisation of other free parameters like the resolution or dimensions of the image reconstruction volume and consider Multi-Objective functions as discussed in Section 6.2. One of specific interest concerns the composition of quasi-orthogonal magnetic field profiles in the augmented inverse problem. Figure 6.59 demon-

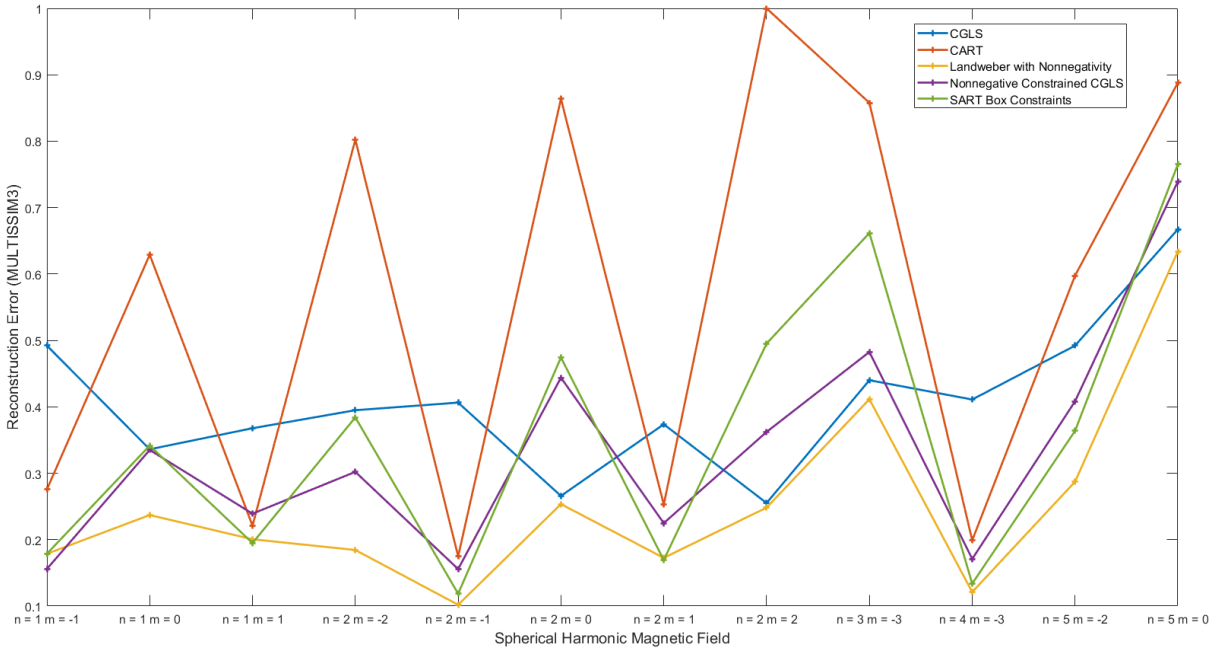


Figure 6.59: The variation in image quality, as characterised by the MULTISSIM3 metric for voxel-based imagery, when the twelve magnetic fields from sixty-six coils in the initial problem are applied separately. There is a clear correlation between the various algorithms which suggests functional interaction between fields and the conductive material allows for improved image quality.

strates that an augmentation approach to the test problem in Figure 6.6b provides higher quality solutions than the individual application of the component fields. But there is also variability in the quality of reconstructions produced by each. This likely relates to functional overlaps between the interrogation regions of specific profiles and conductive material, with a preference for areas of high magnitude.

An optimisation could identify the *best* profiles to apply, which has numerous potential benefits. Experimentally this could speed up the process, minimising data gathering time by reducing the number of time-dependent excitation patterns. There are also computational benefits, with fewer measurements translating to a smaller sensitivity matrix, decreasing the evaluation time of any inversion procedure. However, it may be that, as

with sensor positions, this would have to be a dynamic process for successful application to an unknown problem. Similar work may consider further improvements to the orthogonality of the generated profiles, the number of sensor positions, the strength and frequency of the applied fields, and the target resolution. Multi-objective optimisation could account for a number of these to produce an optimal system and perhaps consider reducing the overall sensitivity to noise of the final design in a single process.

6.5 Conclusion

In this chapter, we have verified the main hypothesised features of our proposed MIT methodology. Our augmented approach to inversion uses a system of spatially varying, relatively unique magnetic fields produced by the practical induction systems developed in the previous chapter.

We demonstrate that increasing the number of fields can improve the quality of image reconstructions significantly, reducing noise and allowing the accurate location and characterisation of objects. We illustrate this at high and low conductivities with an identical system, suggesting our approach is robust to variation.

Improving the orthogonality of the applied fields has consistent benefits for image quality, in some cases by up to 40% (see Figures 6.6 or 6.12 for example), alongside consistently reducing the computational effort (number of iterations) required to find those solutions. It may also improve the stability of the inversion algorithms (see Figure 6.14), making the practical selection of a near-optimal solution more probable in a realistic setting. Combined, this suggests multiplicative benefits to applying quasi-orthogonal magnetic field profiles.

Iterative algorithms are consistently superior to their Tikhonov counterparts, which cannot provide a solution. We consider the potential of various novel algorithms and highlight a disparity between the theory (which states CGLS-type algorithms are invalid)

and practical reality (where CGLS algorithms can produce high-quality results). Careful consideration of the balance between speed, stability, and image quality is imperative in algorithmic choice.

Future work should consider more complex conductivity environments and further optimisation, like identifying the best magnetic excitations to apply. A detailed analysis of state-of-the-art inversion algorithms may also be beneficial. With numerous solutions available, a comprehensive investigation into post-processing is also likely to improve image quality by eliminating some of the natural variation arising from an ill-posed problem.

Chapter 7

Summary

This body of work has explored a set of orthogonal magnetic fields which are theoretically ideal for magnetic-induced interrogation. These are linked to spherical harmonics, with a system-wide scalar potential defined for a particular harmonic, leading to an infinite set of mutually orthogonal inhomogeneous profiles. A technique proposes to produce multiple profiles from a modular induction system composed of arbitrary elements by independently varying the current driven to them, similar to common active magnetic-shielding approaches using matrix coils[209]. An estimate of the excitation patterns for arbitrarily distributed elements, and any physically feasible magnetic profile comes from a regularised inverse problem. There is also a framework to model three-dimensional sources with spatially varying current densities.

A discussion of the existing conventions associated with modelling the measured response, and calculating the sensitivity to variations in conductivity from a typical MIT system, reveals several limitations. These include an inability to account for inhomogeneous fields, difficulties associated with modelling magnetometer-based measurements, and a large square sensitivity matrix for each applied magnetic field. From a governing Helmholtz equation, a full-phasor derivation introduces a novel greens function approach to model the problem in a framework which supports the spherical harmonic magnetic profiles. Whilst computationally expensive, this produces a rectangular sensitivity matrix for a predetermined set of measurement positions, with augmentation enabling a single inver-

sion problem for many field profiles.

A further derivation provides simplified models for both the magnetic vector potential and magnetic flux. For an appropriate problem, this provides quick and easy calculation which is stable over a range of conductivities. The result is an ability to effortlessly estimate the eddy responses from, and sensitivity to, a variety of problems, with a rectangular sensitivity matrix suited to augmentation - perfect for rapid prototyping of new MIT systems.

Progression of this work then moves the considering inversion, specifically in terms of a linear, discrete, ill-posed problem, which, in principle, may apply to any of the preceding eddy-current models. The problem is typically ill-posed because there are infinite equally valid solutions to any specific set of measurements, and ill-conditioned, where extreme amplification of any tiny noise component occurs during inversion. Ill-conditioning leads naturally to the singular-value decomposition, which defines the condition number - a value which can quantify the noise sensitivity. The high sensitivity to noise also relates directly to a definition of the inverse crime in the context of MIT.

A solution to ill-posed/ill-conditioned problems is to apply regularisation. Functionally this works by introducing some additional information or assumptions about the properties of a result to select one which is both suitable and stable. An entire field of research is dedicated to these methods, though few have seen active application to MIT. The most common is Tikhonov regularisation which is simple and easy to implement, but the computation time scales dramatically with resolution. The SVD then identifies the function of regularisation is to introduce some filtering to the inverse problem, which leads instinctively to truncated inverse methods. A natural progression is then to iterative inversion algorithms where the regularisation parameter becomes the number of iterations. Different algorithms impart various assumptions about the sought solution, and these are typically semi-convergent.

Part of the difficulty of this type of inversion is that the regularisation parameter λ is problem specific and highly non-linear. Too little allows noise corruption to persist, but too much and the solution is a poor fit to the data. Realistically the target solution is unknown, and some automated selection criteria are required. Many possible methods exist, each with associated advantages and disadvantages, but a consistent feature of all existing approaches is that they select a sub-optimal solution.

Differentiation of individual algorithms beyond the regularisation scheme concerns the computational work involved in each iteration and their convergence properties. The convergence properties are particularly pertinent to selecting the optimal solution. In an ideal case of semi-convergence, the solution updates shrink, becoming smaller around this optimum. The result should then be relatively stable over many iterations making the accuracy of automated selection criteria less critical. The balance is, therefore, between the best possible image quality determined by the suitability of a regularisation scheme, desirable convergence qualities, and computational limitations. In reality, no single algorithm is ideal, motivating the availability of a library of relevant options.

The field design procedure can produce a set of many unique magnetic profiles, even from a simple modular system. With four square coils, twelve are possible, but these are poor representations of the orthonormal target set. Determining uniqueness is tricky, with many nuanced components involved in the final functional form of a magnetic field profile. RMS analysis and comparison of excitation patterns can provide some indication, but this is not rigorous. A robust framework seeks to evaluate this generally through a combination of the harmonic composition of individual fields, a tolerance approach to cross-comparison, and overlap calculations for additional verification.

Though it may be possible to produce quality spatial approximations of higher complexity harmonics, the sixth degree provides a good truncation point for this MIT investigation.

With thirty-eight coils targeting all harmonics up to and including that degree, the advantages of the excitation design procedure and the necessity of a robust framework to analyse uniqueness is much clearer. With this many elements, forty-six quasi-orthogonal fields of the spherical harmonic type are possible. These are typically a quality match to the spatial distribution but have limited success in producing intended magnitude variations.

To maintain the attained accuracy, high-precision current drivers with control to at least 0.1% of the peak amplitude are necessary. This is less critical for many low-degree harmonics, so where fewer fields will suffice, simpler experimental apparatus may be appropriate. The nonlinearity of the regularisation parameter significantly influences the quality of the generated harmonic approximations. Ignoring it is perilous, manual selection is slow, and analytic calculations are computationally expensive. Minute variations may cause a notable reduction in field quality and have a pronounced impact on the field produced per unit current. For MIT, this measure is imperative because it directly relates to how detectable a response field will be.

Further increases in the number of control elements lead to an idealistic setup which can produce forty-eight orthonormal profiles, targeting up to the sixth-degree harmonics. Most of these are quality representations of the target with below 2% RMS error. Simplified designs, such as eighteen rectangular elements, can also produce numerous fields. A significantly larger interrogation region is also feasible, but this detracts from the accuracy of the spherical harmonics. Dynamic calculations of the excitation patterns may compensate for this and are imperative when the interrogation region is varied.

Production of spherical harmonic-type fields is also possible with a co-planar system. The intention is to produce a scannable device, but there are limitations on harmonic quality due to an inability to generate significant variations along the z-axis. Improvements may be possible by varying the offset between layers or the side length of component coils to

provide greater control. Three-dimensional current sources have minimal influence on the functional form of the generated magnetic field but appear to provide a stronger magnetic field. However, the practicalities of driving current through a relevant device may negate any advantages.

The final chapter evaluates the influence of multiple mutually unique quasi-orthogonal magnetic fields when applied to MIT. There is an explicit dependence on an augmented inversion for all profiles and measurements enabled by the preceding derivations. Initial analysis establishes that mutually unique magnetic profiles of any form provide additional information, which improves the average resolution of each voxel and the quality of the resulting images.

Increasing the orthogonality of those profiles has a significant impact, drastically reducing the sensitivity to noise whilst substantially improving the voxel resolution and image quality. The computational effort required to produce the solution, regardless of the algorithm, is also reduced. These benefits to image quality also hold even when considering far less ideal conductivity distributions, implying the novel approach to MIT is robust. This problem also confirms the expected superiority of modern inversion algorithms over their traditional Tikhonov counterparts. Tikhonov methods cannot solve the forty-eight sampled problems (12 from four coils, 3×12 from sixty-six with various block locations).

The computational efficiency of the augmented sensitivity formulation also enables the optimisation of various parameters. Sensor positions are incredibly influential and can produce significant improvements in image quality. A Global-search algorithm performs better on the specified problem than a Pattern Search, but neither is particularly robust under perturbation. The specificity of the resulting optimisations means this process must be dynamic in practice, especially given their increased sensitivity to noise.

Increasing numbers of orthogonal magnetic fields lead to further quantitative benefits,

including consistent improvements in image quality. Once a problem becomes overdetermined, something vastly easier to achieve when it is possible to produce 48 unique measurements per detector, orthogonal excitation enables a reduction in noise sensitivity. Additional measurements are also intuitively beneficial, demonstrating that this excitation procedure is effective even with an information excess.

The final set of problems further demonstrates the versatility of this methodology. It functions well in a different, much lower conductivity regime and can produce higher-resolution images. Detection of an average conductivity occlusion is also possible, but with limited characterisation from four measurements. Six additional novel inversion algorithms perform best on these problems, each with distinguished convergence properties.

The presented research has achieved a number of the original aims. It is now possible to generate many quasi-orthogonal magnetic fields from any arbitrary distribution of modular current elements and consolidate them into a single memory-efficient inverse problem for MIT. The methodology appears robust and can produce high-quality three-dimensional characterising images with MIT through variable quasi-orthogonal spatial interrogation. It is possible to do this rapidly with augmented inversion, utilising various novel algorithms to maximise image quality and solution stability whilst minimising evaluation time. Future work should focus on experimental verification and then consider further optimisations, including minimising the number of applied fields and condition number, alongside maximising signal-to-noise, resolution, and image quality. Potential future applications are broad; enclosed excitation systems may produce safe, contactless medical images; interrogate suspect packages in an airport; or even detect the moisture content of food during manufacturing. Scannable designs could perform sub-surface geophysical image reconstruction or examine the contents of cargo containers.

Bibliography

[1] R. Streich, "Controlled-source electromagnetic approaches for hydrocarbon exploration and monitoring on land," *Survey Geophysics*, vol. 37, pp. 47–80, 9 2015.

[2] M. Rashed, "Rolling ball algorithm as a multitask filter for terrain conductivity measurements," *Journal of Applied Geophysics*, vol. 132, pp. 17 – 24, 2016. http://www.sciencedirect.com/science/article/pii/S0926985116301744.

[3] H. Huang, "Depth of investigation for small broadband electromagnetic sensors," *Geophysics*, vol. 70, pp. G135–G142, 10 2005.

[4] A. V. Christiansen, J. B. Pedersen, E. Auken, N. E. Søe, M. K. Holst, and S. M. Kristiansen, "Improved geoarchaeological mapping with electromagnetic induction instruments from dedicated processing and inversion," *Remote Sensing*, vol. 8, 12 2016.

[5] D. A. Robinson, "Advancing process-based watershed hydrological research using near-surface geophysics: a vision for, and review of, electrical and magnetic geophysical methods," *Hydrological Processes*, vol. 22, pp. 3604–3635, 3 2008.

[6] F. A. M. Santos, E. P. Almeida, R. Castro, R. Nolasco, and L. Mendes-Victor, "A hydrogeological investigation using em34 and sp surveys," *Earth Planets Space*, vol. 54, pp. 655–662, 2002.

[7] T. Dezert, Y. Fargierb, S. P. Lopesa, and P. Côte, "Geophysical and geotechnical methods for fluvial levee investigation: A review," *Engineering Geology*, vol. 260, 2019.

[8] T. R. Brosten, J. H. Bradford, J. P. McNamara, M. N. Gooseff, J. P. Zarnetske, W. B. Bowden, and M. E. Johnston, "Estimating 3d variation in active-layer thick-

ness beneath arctic streams using ground-penetrating radar," *Journal of Hydrology*, vol. 373, pp. 479–486, 5 2009.

[9] J. S. Whiteley, J. Chambers, S. Ihlemann, P. Wilkinson, and J. M. Kendall, "Recent developments in the direct-current geoelectrical imaging method," *Reviews of Geophysics*, vol. 57, pp. 106–145, 2 2019.

[10] M. S. Mansor, Z. Zakaria, I. Balkhis, R. Abdul Rahim, M. Fadzli Abdul Sahib, Y. Yunos, S. Sahlan, S. Bunyamin, K. h. Abas, M. Ishak, and K. A. Danapalasingam, "Magnetic induction tomography: A brief review," *Jurnal Teknologi*, vol. 73, pp. 91–95, 02 2015.

[11] A. Korzhenevskii and V. Cherepenin, "Magnetic induction tomography," *Journal of communication technology and electronics*, vol. 42, no. 4, pp. 469–474, 1997.

[12] H. Griffiths, "Magnetic induction tomography," *Measurement Science and Technology*, vol. 12, p. 1126, aug 2001.

[13] H. Giffiths, "Magnetic induction tomography," *Measurement Science and Technology*, vol. 12, pp. 1126–1131, 06 2001.

[14] L. Ma and M. Soleimani, "Magnetic induction tomography methods and applications a review," *Measurement Science and Technology*, vol. 28, p. 072001, 06 2017.

[15] K. Heil and U. Schmidhalter, "Theory and guidelines for the application of the geophysical sensor em38," *Sensors (Basel, Switzerland)*, vol. 19, no. 19, 2019.

[16] E. Auken, L. Pellerin, N. B. Christensen, and K. Sørensen, "A survey of current trends in near-surface electrical and electromagnetic methods," *Geophysics*, vol. 71, pp. 477–488, 09-10 2006.

[17] A. Revil, M. Karaoulis, T. Johnson, and A. Kemna, "Review: Some low-frequency electrical methods for subsurface characterization and monitoring in hydrogeology," *Hydogeology*, vol. 20, pp. 617–658, 2012.

[18] A. P. Aizebeokhai, "2d and 3d geoelectrical resistivity imaging: Theory and field design," *Scientific Reserch and Essays*, vol. 5, pp. 3592–3605, 12 2010.

[19] S. H. Ward and G. W. Hohmann, "Electromagnetic theory for geophysical applications," in *Electromagnetic Methods in Applied Geophysics: Voume 1, Theory*,

pp. 130–311, Society of Exploration Geophysicists, 1988.

[20] H. Maurer, D. E. Boerner, and A. Curtis, "Design strategies for electromagnetic geophysical surveys," *Inverse Problems*, vol. 16, pp. 1097–1117, 2000.

[21] M. Loke, J. Chambers, D. Rucker, O. Kuras, and P. Wilkinson, "Recent developments in the direct-current geoelectrical imaging method," *Journal of Applied Geophysics*, vol. 95, pp. 135–156, 2013.

[22] M. Gharibi, D. Laurence, and L. Bentley, "Resolution of 3-d electrical resistivity images from inversions of 2-d orthogonal lines," *Journal of Environmental and Engineering Geophysics - J ENVIRON ENG GEOPHYS*, vol. 10, 12 2005.

[23] R. M. D. Goncalves, *Hydrochemical water prediction (water quality) with transient electromagnetic soundings (TEM) =*. PhD thesis, Universidade de Lisboa (Portugal), Jan. 2012.

[24] W. Feng, "A new high-resolution electromagnetic method for subsurface imaging," 2016. http://hdl.handle.net/10150/612974.

[25] B. K. Sternberg, O. Krichenko, and S. L. Dvorak, "A new high-sensitivity subsurface electromagnetic sensing system: Part i—system design," *Journal of Environmental and Engineering Geophysics*, vol. 13, no. 3, pp. 247–261, 2008. httphttps://library.seg.org/doi/abs/10.2113/JEEG13.3.247.

[26] J. M. Legault, "Airborne electromagnetic systems – state of the art and future directions," *CSEG RECORDER*, vol. 40, 2015.

[27] F. A. M. Santos, E. P. Almeida, R. Castro, R. Nolasco, and L. Mendes-Victor, "A hydrogeological investigation using em34 and sp surveys," *Earth Planets Space*, vol. 54, pp. 655–662, 2002.

[28] L. Slater, "Near surface electrical characterization of hydraulic conductivity: From petrophysical properties to aquifer geometries—a review," *Survey Geophysics*, vol. 28, pp. 169–197, 2007.

[29] V. Rath, T. Radic, and Y. Krause, "Use of 3-d modeling in design of a new type of near-surface survey," *Three-Dimensional Electromagnetics (Geophysical Developments No 7)*, vol. 7, pp. 477–488, 01 1999.

[30] P. Tang, F. Chen, A. Jiang, W. Zhou, H. Wang, G. Leucci, L. de Giorgi, M. Sileo, R. Luo, R. Lasaponara, and N. Masini, "Multi-frequency electromagnetic induction survey for archaeological prospection: Approach and results in han hangu pass and xishan yang in china," *Surveys in Geophysics*, vol. 39, pp. 1285–1302, Nov 2018. httphttps://doi.org/10.1007/s10712-018-9471-5.

[31] P. Weidelt, "Fundamentals of geoelectrics," *Manual for exploring the underground of landfills and contaminated sites*, vol. 3, pp. 65–83, 1997.

[32] N. R. Edwards, K. L. Law, and J. M. DeLaurier, "On measuring the electrical conductivity of the oceanic crust by a modified magnetometric resistivity method," *Journal of Geophysical Research: Solid Earth*, vol. 86, no. B12, pp. 11609–11615, 1981.

[33] J. C. Goswami, M. Li, A. Abubakar, and T. M. Habashy, "Subsurface conductivity mapping using controlled source electromagnetic method," in *2017 IEEE International Symposium on Antennas and Propagation USNC/URSI National Radio Science Meeting*, pp. 2065–2066, July 2017.

[34] R. N. Edwards, "Marine controlled source electromagnetics: Principles, methodologies, future commercial applications," *Surveys in Geophysics*, vol. 26, pp. 675–700, 2005.

[35] S. Constable and L. J. Srnka, "An introduction to marine controlled-source electromagnetic methods for hydrocarbon exploration," *Geophysics*, vol. 72, no. 2, 2007.

[36] M. Wirianto, W. A. Mulder, and E. C. Slob, "A feasibility study of land csem reservoir monitoring in a complex 3-d model," *Geophysical Journal International*, vol. 181, no. 2, 2010.

[37] V. Puzyrev, E. Vilamajo, P. Queralt, J. Ledo, and A. Marcuello, "Three-dimensional modeling of the casing effect in onshore controlled-source electromagnetic surveys," *Surveys in Geophysics*, vol. 38, pp. 527–545, Mar 2017.

[38] K. Tietze, O. Ritter, and P. Veeken, "Controlled-source electromagnetic monitoring of reservoir oil saturation using a novel borehole-to-surface configuration," *Geophysical Prospecting*, vol. 63, no. 6, pp. 1468–1490, 2015.

[39] K. Tietze, O. Ritter, C. Patzer, P. Veeken, and M. Dillen, "Repeatability of land-based controlled-source electromagnetic measurements in industrialized areas and including vertical electric fields," *Geophysical Journal International*, vol. 218, pp. 1552–1571, 05 2019. https://doi.org/10.1093/gji/ggz225.

[40] S. Thomson, D. Fountain, and T. Watts, "Airborne geophysics – evolution and revolution," *Proceedings of Exploration 07: Fifth Decennial International Conference on Mineral Exploration*, pp. 19–37, 2007.

[41] P. Killeen, "Exploration trends and developments in 2013," *The Northern Miner*, 2014.

[42] D. Fountain, "Airborne electromagnetic systems - 50 years of development," *Exploration Geophysics*, vol. 29, no. 1-2, pp. 1–11, 1998. https://doi.org/10.1071/EG998001.

[43] A. Auken, T. Boesen, and A. V. Christiansen, "A review of airborne electromagnetic methods with focus on geotechnical and hydrological applications from 2007 to 2017," *Advances in Geophysics*, vol. 58, pp. 47–93, 2017.

[44] M. Darnet, C. Coppo, P. Reninger, P. Wawrzyniak, B. Bourgeois, and J. Girard, "Airborne and land-based controlled-source electromagnetic surveying in challenging electromagnetic environments," no. 1, pp. 1–5, 2017. https://www.earthdoc.org/content/papers/10.3997/2214-4609.201701357.

[45] M. V. Smirnova, M. Becken, C. Nittinger, P. Yogeshwar, W. Mörbe, R. Rochlitz, A. Steuer, S. Costabel, M. Y. Smirnov, and the DESMEX Working Group, "A novel semiairborne frequency-domain controlled-source electromagnetic system: Three-dimensional inversion of semiairborne data from the flight experiment over an ancient mining area near schleiz, germany," *GEOPHYSICS Technical Papers*, vol. 84, Sep 2019.

[46] J. F. Mason, "Crack-detectives foil aircraft failure: Detecting flaws in aircraft is based on nondestructive testing with new interpretive aids on the way," *IEEE Spectrum*, vol. 17, no. 2, pp. 47–53, 1980.

[47] D. J. Kyte, *Magnetic Induction Tomography and Techniques for Eddy-Current Imag-*

ing, Thesis for the Degree of Doctor of Philosophy, University of Surrey. PhD thesis, University of Surrey, January 1985. https://surrey.eprints-hosting.org/707/.

[48] B. J. Darrer, J. C. Watson, P. Bartlett, and F. Renzoni, "Toward an automated setup for magnetic induction tomography," *IEEE Transactions on Magnetics*, vol. 51, no. 1, pp. 1–4, 2015.

[49] H.-Y. Wei and M. Soleimani, "A magnetic induction tomography system for prospective industrial processing applications," *Chinese Journal of Chemical Engineering*, vol. 20, no. 2, pp. 406–410, 2012.

[50] H. Wei and M. Soleimani, "Theoretical and experimental evaluation of rotational magnetic induction tomography," *IEEE Transactions on Instrumentation and Measurement*, vol. 61, pp. 3324–3331, Dec 2012.

[51] W. T. Joines, Y. Zhang, C. Li, and R. L. Jirtle, "The measured electrical properties of normal and malignant human tissues from 50 to 900 mhz," *Medical physics*, vol. 21, no. 4, pp. 547–550, 1994.

[52] J. R. Feldkamp, "Single-coil magnetic induction tomographic three-dimensional imaging," *Journal of Medical Imaging*, vol. 2, no. 1, pp. 013502–013502, 2015.

[53] J. R. Feldkamp and S. Quirk, "Optically tracked, single-coil, scanning magnetic induction tomography," *Journal of Medical Imaging*, vol. 4, no. 2, pp. 023504–023504, 2017.

[54] S. Watson, H. C. Wee, H. Griffiths, and R. J. Williams, "A highly phase-stable differential detector amplifier for magnetic induction tomography," *Physiological Measurement*, vol. 32, pp. 917–926, jun 2011. https://doi.org/10.1088

[55] L. Marmugi and F. Renzoni, "Optical magnetic induction tomography of the heart," *Scientific Reports*, 2016. http://discovery.ucl.ac.uk/1477480/1/srep23962.pdf.

[56] L. M. Rushton, T. Pyragius, A. Meraki, L. Elson, and K. Jensen, "Unshielded portable optically pumped magnetometer for the remote detection of conductive objects using eddy current measurements," *Review of Scientific Instruments*, vol. 93, no. 12, p. 125103, 2022. https://doi.org/10.1063/5.0102402.

[57] P. D. D. Schwindt, S. Knappe, V. Shah, L. Hollberg, J. Kitching, L.-A. Liew, and J. Moreland, "Chip-scale atomic magnetometer," *Applied Physics Letters*, vol. 85, no. 26, pp. 6409–6411, 2004. https://doi.org/10.1063/1.1839274.

[58] H. Alkan and Çağlayan Balkaya, "Parameter estimation by differential search algorithm from horizontal loop electromagnetic (hlem) data," *Journal of Applied Geophysics*, vol. 149, pp. 77 – 94, 2018. http://www.sciencedirect.com/science/article/pii/S0926985117301052.

[59] "Standard guide for using the frequency domain electromagnetic method for subsurface investigations," 2008.

[60] J. Boaga, "The use of fdem in hydrogeophysics: A review," *Journal of Applied Geophysics*, vol. 139, pp. 36–46, 2 2017.

[61] J. McNeill, "Electromagnetic terrain conductivity measurement at low induction numbers," *Geonics. Ltd., Technical Note TN-6.*, 1980.

[62] A. MacDonald, J. Davies, and R. Peart, "Geophysical methods for locating groundwater in low permeability sedimentary rocks: examples from southeast nigeria," *Journal of African Earth Sciences*, vol. 32, no. 1, pp. 115 – 131, 2001. http://www.sciencedirect.com/science/article/pii/S0899536201900223.

[63] B. Spies and F. Frischknecht, "Electromagnetic sounding," *Electromagnetic Methods in Applied Geophysics, 2, part A, SEG*, pp. 285–425, 5 1991.

[64] D. C. Fraser, "The multicoil ii airborne electromagnetic system," *GEOPHYSICS*, vol. 44, no. 8, pp. 1367–1394, 1979.

[65] J. T. Smith, H. F. Morrison, L. R. Doolittle, and H.-W. Tseng, "Multi-transmitter multi-receiver null coupled systems for inductive detection and characterization of metallic objects," *Journal of Applied Geophysics*, vol. 61, 10 2005.

[66] A. Mester, E. Zimmermann, J. Kruk, H. Vereecken, and S. van Waasen, "Development and drift-analysis of a modular electromagnetic induction system for shallow ground conductivity measurements," *Measurement Science and Technology*, vol. 25, 04 2014.

[67] S. Delefortrie, D. Hanssens, T. Saey, E. V. D. Vijver, M. Smetryns, C. Bobe,

and P. D. Smedt, "Validating land-based fdem data and derived conductivity maps: Assessment of signal calibration, signal attenuation and the impact of heterogeneity," *Journal of Applied Geophysics*, vol. 164, pp. 179 – 190, 2019. http://www.sciencedirect.com/science/article/pii/S0926985117305748.

[68] J. Thiesson, P. Kessouri, C. Schamper, and A. Tabbagh, "Calibration of frequency-domain electromagnetic devices used in near-surface surveying," *Near Surface Geophysics*, vol. 12, no. 4, pp. 481–491, 2014. https://www.earthdoc.org/content/journals/10.3997/1873-0604.2014012.

[69] B. J. Minsley, B. D. Smith, R. Hammack, J. I. Sams, and G. Veloski, "Calibration and filtering strategies for frequency domain electromagnetic data," *Journal of Applied Geophysics*, vol. 80, pp. 56 – 66, 2012. http://www.sciencedirect.com/science/article/pii/S0926985112000286.

[70] P. Jerzy Wtorek and M. Anna Bujnowska, "Electromagnetic methods in geophysics," *Gdańsk University of Technology*, pp. 1–120, 2004.

[71] J. Guillemoteau and J. Tronicke, "Non-standard electromagnetic induction sensor configurations: Evaluating sensitivities and applicability," *Journal of Applied Geophysics*, vol. 118, pp. 15 – 23, 2015. http://www.sciencedirect.com/science/article/pii/S0926985115001354.

[72] M. Himi, J. Tapias, S. Benabdelouahab, A. Salhi, L. Rivero, M. Elgettafi, A. E. Mandour, J. Stitou, and A. Casas, "Geophysical characterization of saltwater intrusion in a coastal aquifer: The case of martil-alila plain (north morocco)," *Journal of African Earth Sciences*, vol. 126, pp. 136 – 147, 2017.

[73] D. Gomez-Ortiz, I. Blanco-Montenegro, J. Arnoso, T. Martin-Crespo, M. Solla, F. Montesinos, E. Vélez, and N. Sánchez, "Imaging thermal anomalies in hot dry rock geothermal systems from near-surface geophysical modelling," *Remote Sensing*, vol. 11, no. 6, 2019. http://search.proquest.com/docview/2304012765/.

[74] A. T. Selepeng, S. Sakanaka, and T. Nishitani, "3d numerical modelling of negative apparent conductivity anomalies in loop-loop electromagnetic measurements: a case study at a dacite intrusion in sugisawa, akita prefecture, japan," *Exploration*

Geophysics, vol. 48, no. 3, pp. 177–191, 2017. https://doi.org/10.1071/EG16027.

[75] C. von Hebel, S. Rudolph, A. Mester, J. A. Huisman, P. Kumbhar, H. Vereecken, and J. van der Kruk, "Three-dimensional imaging of subsurface structural patterns using quantitative large-scale multiconfiguration electromagnetic induction data," *Water Resources Research*, vol. 50, pp. 2732–2748, 3 2014.

[76] S. D. Carriere, K. Chalikakis, C. Ollivier, T. Heath, M. Mangin, J. Kempf, J. B. Randriatsitohaina, M. Gillon, T. Razakamanana, and J. Lapègue, "Sustainable groundwater resources exploration and management in a complex geological setting as part of a humanitarian project (mahafaly plateau, madagascar)," *Environmental Earth Sciences*, vol. 77, p. 734, Oct 2018. https://doi.org/10.1007/s12665-018-7909-1.

[77] R. Di Maio, S. Fais, P. Ligas, E. Piegari, R. Raga, and R. Cossu, "3d geophysical imaging for site-specific characterization plan of an old landfill," *Waste Management*, vol. 76, pp. 629–642, 2018.

[78] C. Brogi, J. Huisman, S. Pätzold, C. von Hebel, L. Weihermüller, M. Kaufmann, J. van der Kruk, and H. Vereecken, "Large-scale soil mapping using multi-configuration emi and supervised image classification," *Geoderma*, vol. 335, pp. 133 – 148, 2019. http://www.sciencedirect.com/science/article/pii/S0016706117315641.

[79] R. R. Frederiksen, A. V. Christiansen, S. Christensen, and K. R. Rasmussen, "A direct comparison of emi data and borehole data on a 1000ha data set," *Geoderma*, vol. 303, pp. 188 – 195, 2017. http://www.sciencedirect.com/science/article/pii/S0016706117302811.

[80] L. Pellerin and D. L. Alumbaugh, "Tools for electromagnetic investigation of the shallow subsurface," *The Leading Edge*, vol. 16, pp. 1631 – 1638, 1997.

[81] K. Sudduth, S. Drummond, and N. Kitchen, "Accuracy issues in electromagnetic induction sensing of soil electrical conductivity for precision agriculture," *Computers and Electronics in Agriculture*, vol. 31, pp. 239–264, 05 2001.

[82] D. Robinson, I. Lebron, S. Lesch, and P. Shouse, "Minimizing drift in electrical conductivity measurements in high temperature environments using the em-38,"

Soil Science Society of America, vol. 68, pp. 339–345, 2003a.

[83] S. J. H. Abdu, D.A. Robinson, "Comparing bulk soil electrical conductivity determination using the dualem-1s and em38-dd electromagnetic induction instruments," *Soil Science Society of America*, vol. 71, pp. 189–196, 2007.

[84] J. Lerssi, S. Niemi, and I. Suppala, "Gem-2 - new generation electromagnetic sensor for near surface mapping," vol. 2016, no. 1, 2016. https://www.earthdoc.org/content/papers/10.3997/2214-4609.201602051.

[85] K. Heil and U. Schmidhalter, "The application of em38: Determination of soil parameters, selection of soil sampling points and use in agriculture and archaeology," *Sensors*, vol. 17, 2017.

[86] J. Bonsall, R. Fry, C. Gaffney, I. Armit, A. Beck, and V. Gaffney, "Assessment of the cmd mini-explorer, a new low-frequency multi-coil electromagnetic device, for archaeological investigations.," *Archaeological Prospecting*, vol. 20, pp. 219–231, 2013.

[87] J. Huang, T. Koganti, F. M. Santos, and J. Triantafilis, "Mapping soil salinity and a fresh-water intrusion in three-dimensions using a quasi-3d joint-inversion of dualem-421s and em34 data," *Science of The Total Environment*, vol. 577, pp. 395 – 404, 2017.

[88] M. Dabas, A. Anest, J. Thiesson, and A. Tabbagh, "Slingram EMI Devices for Characterizing Resistive Features Using Apparent Conductivity Measurements: check of the DualEM-421S Instrument and Field Tests," *Archaeological Prospection*, vol. 23, no. 3, pp. 165–180, 2016. https://hal.sorbonne-universite.fr/hal-01376280.

[89] U. Stockmann, J. Huang, B. Minasny, and J. Triantafilis, "Utilizing a dualem-421 and inversion modelling to map baseline soil salinity along toposequences in the hunter valley wine district," *Soil Use and Management*, vol. 33, no. 3, pp. 413–424, 2017. https://onlinelibrary.wiley.com/doi/abs/10.1111/sum.12352.

[90] M. A. Vallée, R. S. Smith, and P. Keating, "Metalliferous mining geophysics — state of the art after a decade in the new millennium," *GEOPHYSICS*, vol. 76, no. 4, pp. W31–W50, 2011.

[91] F.-X. Simon, A. Sarris, J. Thiesson, and A. Tabbagh, "Mapping of quadrature magnetic susceptibility/magnetic viscosity of soils by using multi-frequency emi," *Journal of Applied Geophysics*, vol. 120, pp. 36 – 47, 2015. http://www.sciencedirect.com/science/article/pii/S0926985115001883.

[92] J. Thiesson, A. Tabbagh, F.-X. Simon, and M. Dabas, "3d linear inversion of magnetic susceptibility data acquired by frequency domain emi," *Journal of Applied Geophysics*, vol. 136, pp. 165 – 177, 2017. http://www.sciencedirect.com/science/article/pii/S0926985116304694.

[93] F.-X. Simon, M. Pareilh-Peyrou, S. Buvat, A. Mayoral, P. Labazuy, K. Kelfoun, and A. Tabbagh, "Quantifying multiple electromagnetic properties in emi surveys: A case study of hydromorphic soils in a volcanic context – the lac du puy (france)," *Geoderma*, vol. 361, p. 114084, 2020. http://www.sciencedirect.com/science/article/pii/S0016706119304483.

[94] T. R. Andersen, S. E. Poulsen, P. Thomsen, and K. Havas, "Geological characterization in urban areas based on geophysical mapping: A case study from horsens, denmark," *Journal of Applied Geophysics*, vol. 150, pp. 338–349, 2018.

[95] Q. Xiaodong, W. Chen, F. Guangyou, and Y. Hejun, "Detecting anomaly targets using handheld frequency domain electromagnetic system," *Sensors and Actuators A: Physical*, vol. 268, pp. 155 – 163, 2017.

[96] M. Zhang, C. G. Farquharson, and C. Liu, "Response characteristics of gradient data from the frequency-domain controlled-source electromagnetic method," *Journal of Applied Geophysics*, vol. 172, p. 103873, 2020.

[97] S. A. Bakr and D. Pardo, "A multi-domain decomposition-based fourier finite element method for the simulation of 3d marine csem measurements," *Computational Geosciences*, vol. 21, pp. 345–357, Jun 2017.

[98] W. Zhu, Y. Liu, and X. Zhu, "Design of frequency-domain controlled source for electromagnetic prospecting based on multi-frequency resonance," in *2017 IEEE 2nd Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*, pp. 1819–1822, Dec 2017.

[99] J. Macnae, "Geophysical prospecting with electric fields from an inductive source," *University of Toronto*, 11 1981.

[100] J. Macnae and R. J. Irvine, "Inductive source resistivity: A tool for outlining silicification in gold exploration," *Exploration Geophysics*, vol. 19, no. 4, pp. 471–480, 1988. https://doi.org/10.1071/EG988471.

[101] J. MacNae, P. McGowan, and Y. Lamontagne, "Resistivity mapping using inductive sources," *Exploration Geophysics*, vol. 20, no. 1-2, pp. 47–50, 1989. https://www.tandfonline.com/doi/abs/10.1071/EG989047.

[102] D. Moghadas, F. André, J. H. Bradford, J. van der Kruk, H. Vereecken, and S. Lambot, "Electromagnetic induction antenna modelling using a linear system of complex antenna transfer functions," *Near Surface Geophysics*, vol. 10, no. 3, pp. 237–247, 2012. https://www.earthdoc.org/content/journals/10.3997/1873-0604.2012002.

[103] F. Lavoué, J. van der Krak, J. Rings, F. André, D. Moghadas, J. Huisman, S. Lambot, L. Weiherrnüller, J. Vanderborght, and H. Vereecken, "Electromagnetic induction calibration using apparent electrical conductivity modelling based on electrical resistivity tomography," *Near Surface Geophysics*, vol. 8, no. 6, pp. 553–561, 2010. https://www.earthdoc.org/content/journals/10.3997/1873-0604.2010037.

[104] T. B. Candansayar ME, "Two-dimensional joint inversion of radio magnetotelluric and direct current resistivity data," *Geophys Prosp*, vol. 56, pp. 737–749, 2008.

[105] E.-K. H. Monteiro Santos FA, "Comparative study of local versus global methods for 1d joint inversion of direct current resistivity and time domain electromagnetic data," *Near Surface Geophysics*, vol. 8, pp. 135–143, 2010.

[106] P. A. Bala J., "Validation of joint inversion of direct current and electromagnetic measurements," *Acta Geophys*, vol. 58, pp. 135–143, 2010.

[107] P. Sentenac, V. Benes, V. Budinsky, H. Keenan, and R. Baron, "Post flooding damage assessment of earth dams and historical reservoirs using non-invasive geophysical techniques," *Journal of Applied Geophysics*, vol. 146, pp. 138–148, 9 2017.

[108] N. Goldshleger, O. Shamir, U. Basson, and E. Zaady, "Frequency domain electromagnetic method (fdem) as a tool to study contamination at the sub-soil layer,"

Geosciences, vol. 9, no. 9, 2019. https://www.mdpi.com/2076-3263/9/9/382.

[109] Y. Xue, Z. Tang, Q. Hu, and J. Drahota, "Using the electromagnetic induction survey method to examine the depth to clay soil layer (bt horizon) in playa wetlands," *Journal of Soils and Sediments*, vol. 20, pp. 556–570, Jan 2020. https://doi.org/10.1007/s11368-019-02390-8.

[110] F. Cassel and S. Sharma, "Delocalization of salt solution in a semiarid farmland topsoil," *Water, Air, & Soil Pollution*, vol. 229, p. 254, Jul 2018. https://doi.org/10.1007/s11270-018-3905-z.

[111] P. McLachlan, J. Chambers, S. Uhlemann, and A. Binley, "Geophysical characterisation of the groundwater–surface water interface," *Advances in Water Resources*, vol. 109, pp. 302 – 319, 2017. http://www.sciencedirect.com/science/article/pii/S0309170817304463.

[112] B. Narayanasamy and F. Luo, "A survey of active emi filters for conducted emi noise reduction in power electronic converters," *IEEE Transactions on Electromagnetic Compatibility*, vol. 61, pp. 2040–2049, Dec 2019.

[113] W. Indrasari, W. Srigutomo, M. Djamal, and R. N. S, "Development of an ac power source for csem method using full-bridge switching configuration," *AIP Conference Proceedings*, vol. 1656, no. 1, p. 040008, 2015. https://aip.scitation.org/doi/abs/10.1063/1.4917115.

[114] "20 shielding tips and tricks," *Appliance Design*, vol. 66, pp. 14–17, 01 2018. http://ezproxy.nottingham.ac.uk/login?note=https://search.proquest.com/docview/2103547502?accountid=8018.

[115] F. Luo, D. Dong, D. Boroyevich, P. Mattavelli, and S. Wang, "Improving high-frequency performance of an input common mode emi filter using an impedance-mismatching filter," *IEEE Transactions on Power Electronics*, vol. 29, pp. 5111–5115, Oct 2014.

[116] L. Fang, D. Boroyevich, P. Mattevelli, and N. Gazel, "A comprehensive design for high power density common mode emi inductor," in *2011 IEEE Energy Conversion Congress and Exposition*, pp. 1861–1867, Sep. 2011.

[117] F. Luo, S. Wang, F. Wang, D. Boroyevich, N. Gazel, Y. Kang, and A. C. Baisden, "Analysis of cm volt-second influence on cm inductor saturation and design for input emi filters in three-phase dc-fed motor drive systems," *IEEE Transactions on Power Electronics*, vol. 25, pp. 1905–1914, July 2010.

[118] Y. Maillet, R. Lai, S. Wang, R. Burgos, and D. Boroyevich, "High-density emi filter design for dc-fed motor drives," *Power Electronics, IEEE Transactions on*, vol. 25, pp. 1163 – 1172, 06 2010.

[119] S. Wang, F. C. Lee, W. G. Odendaal, and J. D. van Wyk, "Improvement of emi filter performance with parasitic coupling cancellation," in *2005 IEEE 36th Power Electronics Specialists Conference*, pp. 1780–1786, June 2005.

[120] Shuo Wang, F. C. Lee, J. D. van Wyk, and J. D. van Wyk, "A study of integration of parasitic cancellation techniques for emi filter design with discrete components," *IEEE Transactions on Power Electronics*, vol. 23, pp. 3094–3102, Nov 2008.

[121] P. P. Regtien, "3 - uncertainty aspects," in *Sensors for Mechatronics* (P. P. Regtien, ed.), pp. 35 – 56, Oxford: Elsevier, 2012. http://www.sciencedirect.com/science/article/pii/B9780123914972000030.

[122] S. Wang, Y. Y. Maillet, F. Wang, D. Boroyevich, and R. Burgos, "Investigation of hybrid emi filters for common-mode emi suppression in a motor drive system," *IEEE Transactions on Power Electronics*, vol. 25, pp. 1034–1045, April 2010.

[123] X. Wang, Z. Fu, Y. Wang, W. Wang, W. Liu, and J. Zhao, "A wide-frequency constant-amplitude transmitting circuit for frequency domain electromagnetic detection transmitter," *Electronics*, vol. 8, no. 6, 2019. httphttps://www.mdpi.com/2079-9292/8/6/640.

[124] E. Penovi, S. Maestri, R. García Retegui, N. Wassinger, and M. Benedetti, "Event-based control system suitable for high-precision pulsed current source applications with improved switching behavior," *IEEE Transactions on Industrial Informatics*, vol. 11, pp. 987–996, Aug 2015.

[125] D. Zhou, S. Yang, and Y. Tang, "Model-predictive current control of modular multilevel converters with phase-shifted pulse width modulation," *IEEE Transactions*

on Industrial Electronics, vol. 66, pp. 4368–4378, June 2019.

[126] H. Xiao, Y. Ma, Y. Lv, T. Ding, S. Zhang, F. Hu, L. Li, and Y. Pan, "Development of a high-stability flat-top pulsed magnetic field facility," *IEEE Transactions on Power Electronics*, vol. 29, pp. 4532–4537, Sep. 2014.

[127] P. E. Wannamaker, "Anisotropy versus heterogeneity in continental solid earth electromagnetic studies: Fundamental response characteristics and implications for physicochemical state," *Surveys in Geophysics*, vol. 26, pp. 733–765, Nov 2005.

[128] C. J. Kohnke, "Electromagnetic effects of steel-cased wells," *Colorado School of Mines*, pp. 1–74, 2016.

[129] F. Roméo and D. I. Hoult, "Magnet field profiling: Analysis and correcting coil design," *Magnetic Resonance in Medicine*, vol. 1, no. 1, pp. 44–65, 1984. https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910010107.

[130] P. J. Hobson, *Magnetic Field Design for Quantum Technologies*. PhD thesis, University of Nottingham, 07 2022.

[131] P. Silvester, "High-order polynomial triangular finite element for potential problems," *International Journal of Engineering Science - INT J ENG SCI*, vol. 7, pp. 849–861, 08 1969.

[132] N. Demerdash, T. NEHL, and F. FOUAD, "Finite element formulation and analysis of three dimensional magnetic field problems," *Magnetics, IEEE Transactions on*, vol. 16, pp. 1092 – 1094, 10 1980.

[133] N. Demerdash, F. Fouad, T. Nehl, and O. Mohammed, "Three dimensional finite element vector potential formulation of magnetic fields in electrical apparatus," *Power Apparatus and Systems, IEEE Transactions on*, vol. PER-1, pp. 4104 – 4111, 09 1981.

[134] M. Sadiku, "A simple introduction to finite element analysis of electromagnetic problems," *IEEE Transactions on Education*, vol. 32, no. 2, pp. 85–93, 1989.

[135] M. Sadiku, A. Shadare, and S. Musa, "Computational electromagnetics," *International Journal of Engineering Research*, vol. 6, pp. 414–416, 08 2017.

[136] I. The MathWorks, *Partial Differential Equation Toolbox*. Natick, Massachusetts,

United States, 2021. https://www.mathworks.com/help/pde/.

[137] D. J. Griffiths, *Introduction to electrodynamics; 3rd ed.* Upper Saddle River, NJ: Prentice-Hall, 1999. https://cds.cern.ch/record/611579.

[138] M. Soleimani, "Sensitivity maps in three-dimensional magnetic induction tomography," *Insight - Non-Destructive Testing and Condition Monitoring*, vol. 48, no. 1, pp. 39–44, 2006. https://www.ingentaconnect.com/content/bindt/insight/2006/00000048/00000001/art00007.

[139] M. Soleimani, W. Lionheart, A. Peyton, X. Ma, and S. Higson, "A three-dimensional inverse finite-element method applied to experimental eddy-current imaging data," *IEEE Transactions on Magnetics*, vol. 42, no. 5, pp. 1560–1567, 2006.

[140] D. Dyck, D. Lowther, and E. Freeman, "A method of computing the sensitivity of electromagnetic quantities to changes in materials and sources," *IEEE Transactions on Magnetics*, vol. 30, no. 5, pp. 3415–3418, 1994.

[141] F. Ihlenburg and I. Babuška, "Finite element solution of the helmholtz equation with high wave number part i: The h-version of the fem," *Computers & Mathematics with Applications*, vol. 30, no. 9, pp. 9–37, 1995. https://www.sciencedirect.com/science/article/pii/089812219500144N.

[142] S. Macintyre, "A portable low noise low frequency three-axis search coil magnetometer," *IEEE Transactions on Magnetics*, vol. 16, no. 5, pp. 761–763, 1980.

[143] S. Tumanski, "Induction coil sensors—a review," *Measurement Science and Technology*, vol. 18, pp. R31–R46, jan 2007.

[144] Q. Zhao, G. Chen, J. Hao, K. Xu, and W. Yin, "Numerical approach for the sensitivity of a high-frequency magnetic induction tomography system based on boundary elements and perturbation method," *Measurement Science and Technology*, vol. 24, no. 7, p. 074004, 2013.

[145] G. Sinha and S. Prabhu, "Analytical model for estimation of eddy current and power loss in conducting plate and its application," *Reviews of Modern Physics - REV MOD PHYS*, vol. 14, 06 2011.

[146] J. Nagel, "Induced eddy currents in simple conductive geometries due to a time-

varying magnetic field," *IEEE Antennas and Propagation Magazine*, vol. PP, pp. 1–1, 12 2017.

[147] J. D. Jackson, *Classical electrodynamics*. American Association of Physics Teachers, 1999.

[148] R. G. Brown, *Electrodynamics II*. Duke University Lecture Notes, 2007. https://www.lulu.com/content/1144184?page=1&pageSize=4.

[149] M. Wang, *Industrial tomography: systems and applications*. Elsevier, 2015.

[150] H. Dierckx, K. Blauwe, M. Meirvenne, and H. Verschelde, "Inversion of electromagnetic induction data using a 2d geophysical response function," 05 2018.

[151] J. Hadamard, "Lectures on cauchy's problem in linear partial differential equations, yale univ," *Press. New Haven*, 1923.

[152] P. C. Hansen, *Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion*. SIAM, 1998.

[153] V. Klema and A. Laub, "The singular value decomposition: Its computation and some applications," *IEEE Transactions on Automatic Control*, vol. 25, no. 2, pp. 164–176, 1980.

[154] J. Todd, "The condition of a certain matrix," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 46, no. 1, p. 116–118, 1950.

[155] P. C. Hansen, "Regularization tools version 4.0 for matlab 7.3," *Numerical algorithms*, vol. 46, no. 2, pp. 189–194, 2007.

[156] D. L. Colton, R. Kress, and R. Kress, *Inverse acoustic and electromagnetic scattering theory*, vol. 93. Springer, 1998.

[157] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, vol. 2, pp. 1398–1402, Ieee, 2003.

[158] A. Tikhonov, V. Arsenin, and F. John, *Solutions of Ill-posed Problems*. Halsted Press book, Winston, 1977. https://books.google.co.uk/books?id=ECrvAAAAMAAJ.

[159] A. Y. Ng, "Feature selection, l 1 vs. l 2 regularization, and rotational invariance," in *Proceedings of the twenty-first international conference on Machine learning*, p. 78,

2004.

[160] A. Stothers, A. Davie, and I. Gyongy, *On the Complexity of Matrix Multiplication*. University of Edinburgh, 2010.

[161] P. C. Hansen, "Truncated singular value decomposition solutions to discrete ill-posed problems with ill-determined numerical rank," *SIAM Journal on Scientific and Statistical Computing*, vol. 11, no. 3, pp. 503–518, 1990.

[162] S. R. Arridge, "Optimisation and inverse problems in imaging," *Department of Computer Science, University*, 2014.

[163] C. F. Van Loan, "Generalizing the singular value decomposition," *SIAM Journal on numerical Analysis*, vol. 13, no. 1, pp. 76–83, 1976.

[164] C. C. Paige and M. A. Saunders, "Towards a generalized singular value decomposition," *SIAM Journal on Numerical Analysis*, vol. 18, no. 3, pp. 398–405, 1981.

[165] P. C. Hansen, "Regularization, gsvd and truncatedgsvd," *BIT numerical mathematics*, vol. 29, no. 3, pp. 491–504, 1989.

[166] M. Donatelli and L. Reichel, "Square smoothing regularization matrices with accurate boundary conditions," *Journal of Computational and Applied Mathematics*, vol. 272, pp. 334–349, 2014.

[167] M. Ekstrom and R. Rhoads, "On the application of eigenvector expansions to numerical deconvolution," *Journal of Computational Physics*, vol. 14, no. 4, pp. 319–340, 1974.

[168] R. D. Fierro, G. H. Golub, P. C. Hansen, and D. P. O'Leary, "Regularization by truncated total least squares," *SIAM Journal on Scientific Computing*, vol. 18, no. 4, pp. 1223–1241, 1997.

[169] K. Miller, "Least squares methods for ill-posed problems with a prescribed bound," *SIAM Journal on Mathematical Analysis*, vol. 1, no. 1, pp. 52–74, 1970.

[170] P. C. Hansen and D. P. O'Leary, "The use of the l-curve in the regularization of discrete ill-posed problems," *SIAM journal on scientific computing*, vol. 14, no. 6, pp. 1487–1503, 1993.

[171] H. Scharfetter, K. Hollaus, J. Rosell-Ferrer, and R. Merwa, "Single-step 3-d im-

age reconstruction in magnetic induction tomography: Theoretical limits of spatial resolution and contrast to noise ratio," *Annals of biomedical engineering*, vol. 34, no. 11, pp. 1786–1798, 2006.

[172] M. Schmidt, "Graphical model structure learning with l1-regularization," *University of British Columbia*, 2010.

[173] P. C. Hansen and M. Saxild-Hansen, "Air tools—a matlab package of algebraic iterative reconstruction methods," *Journal of Computational and Applied Mathematics*, vol. 236, no. 8, pp. 2167–2178, 2012.

[174] P. C. Hansen and J. S. Jørgensen, "Air tools ii: algebraic iterative reconstruction methods, improved implementation," *Numerical Algorithms*, vol. 79, no. 1, pp. 107–137, 2018.

[175] S. Gazzola, P. C. Hansen, and J. G. Nagy, "Ir tools: a matlab package of iterative regularization methods and large-scale test problems," *Numerical Algorithms*, vol. 81, no. 3, pp. 773–811, 2019.

[176] C. Lawson and R. Hanson, "Chapter 23 p.121," *Solving Least-Squares Problems. Upper Saddle River, New Jersey: Prentice-Hall*, 1974.

[177] M. R. Hestenes and E. Stiefel, "Methods of conjugate gradients for solving," *Journal of research of the National Bureau of Standards*, vol. 49, no. 6, p. 409, 1952.

[178] S. Fan, "An introduction to krylov subspace methods," *arXiv preprint arXiv:1811.09025*, 2018.

[179] D. S. Watkins, *The matrix eigenvalue problem: GR and Krylov subspace methods*. SIAM, 2007.

[180] H. Brakhage, "On ill-posed problems and the method of conjugate gradients," in *Inverse and ill-posed Problems*, pp. 165–175, Elsevier, 1987.

[181] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst, *Templates for the solution of linear systems: building blocks for iterative methods*. SIAM, 1994.

[182] H. A. Van der Vorst, "Bi-cgstab: A fast and smoothly converging variant of bi-cg for the solution of nonsymmetric linear systems," *SIAM Journal on scientific and*

Statistical Computing, vol. 13, no. 2, pp. 631–644, 1992.

[183] S. Kaczmarz, "Approximate solution of systems of linear equations: Bulletin international de l'acad 'e mie polonaise des sciences et des lettres," 1937.

[184] R. Gordon, R. Bender, and G. T. Herman, "Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray photography," *Journal of Theoretical Biology*, vol. 29, no. 3, pp. 471–481, 1970. https://www.sciencedirect.com/science/article/pii/0022519370901098.

[185] T. Strohmer and R. Vershynin, "A randomized kaczmarz algorithm with exponential convergence," *Journal of Fourier Analysis and Applications*, vol. 15, no. 2, pp. 262–278, 2009.

[186] Å. Björck and T. Elfving, "Accelerated projection methods for computing pseudoinverse solutions of systems of linear equations," *BIT Numerical Mathematics*, vol. 19, no. 2, pp. 145–163, 1979.

[187] P. Eggermont, G. Herman, and A. Lent, "Iterative algorithms for large partitioned linear systems, with applications to image reconstruction," *Linear Algebra and its Applications*, vol. 40, pp. 37–67, 1981. https://www.sciencedirect.com/science/article/pii/0024379581901397.

[188] T. Elfving, P. C. Hansen, and T. Nikazad, "Convergence analysis for column-action methods in image reconstruction," *Numerical Algorithms*, vol. 74, 03 2017.

[189] S. J. Wright, "Coordinate descent algorithms," *Math. Program.*, vol. 151, p. 3–34, jun 2015.

[190] T. Hastie, "Fast regularization paths via coordinate descent," in *The 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Denver*, vol. 2009, 2008.

[191] G. Cimmino, "Cacolo approssimato per le soluzioni dei systemi di equazioni lineari," *La Ricerca Scientifica (Roma)*, vol. 1, pp. 326–333, 1938.

[192] C. D. Meyer, *Matrix analysis and applied linear algebra*, vol. 71. Siam, 2000.

[193] R. Hartshorne, *Algebraic Geometry*. Graduate Texts in Mathematics, Springer, 1977.

[194] A. Beutelspacher, B. Albrecht, and U. Rosenbaum, *Projective geometry: from foundations to applications*. Cambridge University Press, 1998.

[195] R. P. Stanley *et al.*, "An introduction to hyperplane arrangements," *Geometric combinatorics*, vol. 13, no. 389-496, p. 24, 2004.

[196] Y. Censor, D. Gordon, and R. Gordon, "Component averaging: An efficient iterative parallel algorithm for large and sparse unstructured problems," *Parallel computing*, vol. 27, no. 6, pp. 777–808, 2001.

[197] Y. Censor, T. Elfving, G. T. Herman, and T. Nikazad, "On diagonally relaxed orthogonal projection methods," *SIAM Journal on Scientific Computing*, vol. 30, no. 1, pp. 473–504, 2008.

[198] C. R. Vogel, *Computational methods for inverse problems*. SIAM, 2002.

[199] S. Gazzola, M. E. Kilmer, J. G. Nagy, O. Semerci, and E. L. Miller, "An inner–outer iterative method for edge preservation in image restoration and reconstruction," *Inverse Problems*, vol. 36, p. 124004, Dec 2020. http://dx.doi.org/10.1088/1361-6420/abb299.

[200] S. Gazzola, C. Meng, and J. Nagy, "Krylov methods for low-rank regularization," 2019.

[201] M. Schmidt, "Least squares optimisation with l1-norm regularization, class cs542b: Non-linear optimization by m. friedlander," 2005. https://www.cs.ubc.ca/~schmidtm/Documents/2005_Notes_Lasso.pdf/.

[202] T. Günther, "Inversion methods and resolution analysis for the 2d/3d reconstruction of resistivity structures from dc measurements," 2004.

[203] P. B. Wilkinson, M. H. Loke, P. I. Meldrum, J. E. Chambers, O. Kuras, D. A. Gunn, and R. D. Ogilvy, "Practical aspects of applied optimized survey design for electrical resistivity tomography," *Geophysical Journal International*, vol. 189, no. 1, pp. 428–440, 2012.

[204] P. B. Wilkinson, S. Uhlemann, J. E. Chambers, P. I. Meldrum, and M. H. Loke, "Development and testing of displacement inversion to track electrode movements on 3-d electrical resistivity tomography monitoring grids," *Geophysical Journal In-*

ternational, vol. 200, no. 3, pp. 1566–1581, 2015.

[205] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[206] A. Herrera-May, L. Aguilera-Cortés, P. Garcia, and E. Manjarrez, "Resonant magnetic field sensors based on mems technology," *Sensors (Basel, Switzerland)*, vol. 9, pp. 7785–813, 10 2009.

[207] A. R. Conn, N. I. Gould, and P. Toint, "A globally convergent augmented lagrangian algorithm for optimization with general constraints and simple bounds," *SIAM Journal on Numerical Analysis*, vol. 28, no. 2, pp. 545–572, 1991.

[208] Z. Ugray, L. Lasdon, J. Plummer, F. Glover, J. Kelly, and R. Martí, "Scatter search and local nlp solvers: A multistart framework for global optimization," *INFORMS Journal on computing*, vol. 19, no. 3, pp. 328–340, 2007.

[209] N. Holmes, M. Rea, R. M. Hill, J. Leggett, L. J. Edwards, P. J. Hobson, E. Boto, T. M. Tierney, L. Rier, G. R. Rivero, V. Shah, J. Osborne, T. M. Fromhold, P. Glover, M. J. Brookes, and R. Bowtell, "Enabling ambulatory movement in wearable magnetoencephalography with matrix coil active magnetic shielding," *NeuroImage*, vol. 274, p. 120157, 2023.

[210] G. Giovannetti, "Comparison between circular and square loops for low-frequency magnetic resonance applications: theoretical performance estimation," *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, vol. 46B, no. 3, pp. 146–155, 2016. https://onlinelibrary.wiley.com/doi/abs/10.1002/cmr.b.21343.

[211] M. E. Rudd and J. Craig, "Optimum spacing of square and circular coil pairs," *Review of Scientific Instruments*, vol. 39, no. 9, pp. 1372–1374, 1968.

[212] A. Gil, J. Segura, and N. M. Temme, "Numerical methods for special functions," 2007.

[213] S. Chandrasekhar, "Radiative transfer," pp. 61 and 64,65, 1960.

[214] F. B. Hildebrand, "Introduction to numerical analysis," pp. 327–330, 1987.

[215] R. Merwa, K. Hollaus, P. Brunner, and H. Scharfetter, "Solution of the inverse problem of magnetic induction tomography (MIT)," *Physiological Measurement*, vol. 26, pp. S241–S250, mar 2005. https://doi.org/10.1088/0967-3334/26/2/023.

[216] M. Zolgharni, P. Ledger, D. Armitage, D. Holder, and H. Griffiths, "Imaging cerebral haemorrhage with magnetic induction tomography: numerical modelling," *Physiological measurement*, vol. 30, no. 6, p. S187, 2009.

[217] S. Watson, R. Williams, W. Gough, and H. Griffiths, "A magnetic induction tomography system for samples with conductivities below 10 s m- 1," *Measurement Science and Technology*, vol. 19, no. 4, p. 045501, 2008.

[218] B. Dekdouk, *Image reconstruction of low conductivity material distribution using magnetic induction tomography*. PhD thesis, The University of Manchester (United Kingdom), 2011.

[219] L. Reichel and Q. Ye, "Simple square smoothing regularization operators," *Electronic Transactions on Numerical Analysis*, vol. 33, p. 63, 2009.

[220] M. Han, X. Cheng, and Y. Xue, "Comparison with reconstruction algorithms in magnetic induction tomography," *Physiological measurement*, vol. 37, no. 5, p. 683, 2016.

[221] C. Cohen-Bacrie, Y. Goussard, and R. Guardo, "Regularized reconstruction in electrical impedance tomography using a variance uniformization constraint," *IEEE transactions on medical imaging*, vol. 16, no. 5, pp. 562–571, 1997.

[222] J. W. Woods, "Chapter 7 - image enhancement and analysis," in *Multidimensional Signal, Image, and Video Processing and Coding (Second Edition)* (J. W. Woods, ed.), pp. 223–256, Boston: Academic Press, second edition ed., 2012.

[223] M. E. Hochstenbach, S. Noschese, and L. Reichel, "Fractional regularization matrices for linear discrete ill-posed problems," *Journal of Engineering Mathematics*, vol. 93, no. 1, pp. 113–129, 2015.

Appendix A

Spherical Harmonic Look-Up Tables

The real, or tesseral, spherical harmonic functions utilised to generate magnetic fields in this thesis are given by Figure A.1-A.4, followed by a tabulated list of the first 24 magnetic fields.

$$
\begin{aligned}
Z_{00} &= \frac{1}{\sqrt{4\pi}} \\
\hline
Z_{11}^{s} &= \sqrt{\frac{3}{4\pi}}[y/r] \\
Z_{10} &= \sqrt{\frac{3}{4\pi}}[z/r] \\
Z_{11}^{c} &= \sqrt{\frac{3}{4\pi}}[x/r] \\
\hline
Z_{22}^{s} &= \frac{1}{4}\sqrt{\frac{15}{\pi}}[2xy/r^2] \\
Z_{21}^{s} &= \frac{1}{2}\sqrt{\frac{15}{\pi}}[yz/r^2] \\
Z_{20} &= \frac{1}{4}\sqrt{\frac{5}{\pi}}[(3z^2 - r^2)/r^2] \\
Z_{21}^{c} &= \frac{1}{2}\sqrt{\frac{15}{\pi}}[xz/r^2] \\
Z_{22}^{c} &= \frac{1}{4}\sqrt{\frac{15}{\pi}}[(x^2 - y^2)/r^2] \\
\hline
Z_{33}^{s} &= \sqrt{\frac{35}{32\pi}}[(3x^2y - y^3)/r^3] \\
Z_{32}^{s} &= \sqrt{\frac{105}{16\pi}}[2xyz/r^3] \\
Z_{31}^{s} &= \sqrt{\frac{21}{32\pi}}[y(5z^2 - r^2)/r^3] \\
Z_{30} &= \sqrt{\frac{7}{16\pi}}[z(5z^2 - 3r^2)/r^3] \\
Z_{31}^{c} &= \sqrt{\frac{21}{32\pi}}[x(5z^2 - r^2)/r^3] \\
Z_{32}^{c} &= \sqrt{\frac{105}{16\pi}}[(x^2 - y^2)z/r^3] \\
Z_{33}^{c} &= \sqrt{\frac{35}{32\pi}}[(x^3 - 3xy^2)/r^3]
\end{aligned}
$$

Figure A.1: The real spherical harmonics up to n = 3, m = 3 from [Source].

$$
\begin{aligned}
Z^s_{44} &= \frac{3}{16}\sqrt{\frac{35}{\pi}}[4(x^3y - xy^3)/r^4] \\
Z^s_{43} &= \frac{3}{8}\sqrt{\frac{70}{\pi}}[(3x^2y - y^3)z/r^4] \\
Z^s_{42} &= \frac{3}{8}\sqrt{\frac{5}{\pi}}[2xy(7z^2 - r^2)/r^4] \\
Z^s_{41} &= \frac{3}{4}\sqrt{\frac{5}{2\pi}}[yz(7z^2 - 3r^2)/r^4] \\
Z_{40} &= \frac{3}{16}\frac{1}{\sqrt{\pi}}[35z^4 - 30z^2r^2 + 3r^4)/r^4] \\
Z^c_{41} &= \frac{3}{4}\sqrt{\frac{5}{2\pi}}[xz(7z^2 - 3r^2)/r^4] \\
Z^c_{42} &= \frac{3}{8}\sqrt{\frac{5}{\pi}}[(x^2 - y^2)(7z^2 - r^2)/r^4] \\
Z^c_{43} &= \frac{3}{8}\sqrt{\frac{70}{\pi}}[(x^3 - 3xy^2)z/r^4] \\
Z^c_{44} &= \frac{3}{16}\sqrt{\frac{35}{\pi}}[(x^4 - 6x^2y^2 + y^4)/r^4]
\end{aligned}
$$

Figure A.2: The real spherical harmonics for n = 4 from [Source].

$$
\begin{aligned}
Z_{55}^{s} &= \sqrt{\frac{693}{512\pi}}[(5x^4y - 10x^2y^3 + y^5)/r^5] \\
Z_{54}^{s} &= \sqrt{\frac{3465}{256\pi}}[4(x^3y - xy^3)z/r^5] \\
Z_{53}^{s} &= \sqrt{\frac{385}{512\pi}}[(3x^2y - y^3)(9z^2 - r^2)/r^5] \\
Z_{52}^{s} &= \sqrt{\frac{1155}{64\pi}}[2xy(3z^3 - zr^2)/r^5] \\
Z_{51}^{s} &= \sqrt{\frac{165}{256\pi}}[y(21z^4 - 14z^2r^2 + r^4)/r^5] \\
Z_{50} &= \sqrt{\frac{11}{256\pi}}[(63z^5 - 70z^3r^2 + 15zr^4)/r^5] \\
Z_{51}^{c} &= \sqrt{\frac{165}{256\pi}}[x(21z^4 - 14z^2r^2 + r^4)/r^5] \\
Z_{52}^{c} &= \sqrt{\frac{1155}{64\pi}}[(x^2 - y^2)(3z^3 - zr^2)/r^5] \\
Z_{53}^{c} &= \sqrt{\frac{385}{512\pi}}[(x^3 - 3xy^2)(9z^2 - r^2)/r^5] \\
Z_{54}^{c} &= \sqrt{\frac{3465}{256\pi}}[(x^4 - 6x^2y^2 + y^4)z/r^5] \\
Z_{55}^{c} &= \sqrt{\frac{693}{512\pi}}[(x^5 - 10x^3y^2 + 5xy^4)/r^5]
\end{aligned}
$$

Figure A.3: The real spherical harmonics for n = 5 from [Source].

$$
\begin{aligned}
Z_{66}^{s} &= \frac{231}{64}\sqrt{\frac{26}{231\pi}}[(6x^5y - 20x^3y^3 + 6xy^5)/r^6] \\
Z_{65}^{s} &= \sqrt{\frac{9009}{512\pi}}[(5x^4y - 10x^2y^3 + y^5)z/r^6] \\
Z_{64}^{s} &= \frac{21}{32}\sqrt{\frac{13}{7\pi}}[4(x^3y - xy^3)(11z^2 - r^2)/r^6] \\
Z_{63}^{s} &= \frac{1}{32}\sqrt{\frac{2730}{\pi}}[(3x^2y - y^3)(11z^3 - 3zr^2)/r^6] \\
Z_{62}^{s} &= \frac{1}{64}\sqrt{\frac{2730}{\pi}}[2xy(33z^4 - 18z^2r^2 + r^4)/r^6] \\
Z_{61}^{s} &= \frac{1}{8}\sqrt{\frac{273}{4\pi}}[yz(33z^4 - 30z^2r^2 + 5r^4)/r^6] \\
Z_{60} &= \frac{1}{32}\sqrt{\frac{13}{\pi}}[(231z^6 - 315z^4r^2 + 105z^2r^4 - 5r^6)/r^6] \\
Z_{61}^{c} &= \frac{1}{8}\sqrt{\frac{273}{4\pi}}[xz(33z^4 - 30z^2r^2 + 5r^4)/r^6] \\
Z_{62}^{c} &= \frac{1}{64}\sqrt{\frac{2730}{\pi}}[(x^2 - y^2)(33z^4 - 18z^2r^2 + r^4)/r^6] \\
Z_{63}^{c} &= \frac{1}{32}\sqrt{\frac{2730}{\pi}}[(x^3 - 3xy^2)(11z^3 - 3zr^2)/r^6] \\
Z_{64}^{c} &= \frac{21}{32}\sqrt{\frac{13}{7\pi}}[(x^4 - 6x^2y^2 + y^4)(11z^2 - r^2)/r^6] \\
Z_{65}^{c} &= \sqrt{\frac{9009}{512\pi}}[(x^5 - 10x^3y^2 + 5xy^4)z/r^6] \\
Z_{66}^{c} &= \frac{231}{64}\sqrt{\frac{26}{231\pi}}[(x^6 - 15x^4y^2 + 15x^2y^4 - y^6)/r^6]
\end{aligned}
$$

Figure A.4: The real spherical harmonics for n = 6 from [Source].

n, m	Harmonic Number	Field Component Bx	Field Component By	Field Component Bz
1,-1	1	0	$\sqrt{\frac{3}{4\pi}}$	0
1,0	2	0	0	$\sqrt{\frac{3}{4\pi}}$
1,1	3	$\sqrt{\frac{3}{4\pi}}$	0	0
2,-2	4	$\sqrt{\frac{15}{5\pi}}$ y	$\sqrt{\frac{15}{5\pi}}$ x	0
2,-1	5	0	$\sqrt{\frac{15}{4\pi}}$ z	$\sqrt{\frac{15}{4\pi}}$ y
2,0	6	$\frac{1}{4}\sqrt{\frac{5}{\pi}}$ -2x	$\frac{1}{4}\sqrt{\frac{5}{\pi}}$ -2y	$\frac{1}{4}\sqrt{\frac{5}{\pi}}$ 4z
2,1	7	$\sqrt{\frac{15}{4\pi}}$ z	0	$\sqrt{\frac{15}{4\pi}}$ x
2,2	8	$\sqrt{\frac{15}{5\pi}}$ x	- $\sqrt{\frac{15}{4\pi}}$ y	0
3,-3	9	$\frac{1}{4}\sqrt{\frac{35}{2\pi}}$ 6xy	$\frac{1}{4}\sqrt{\frac{35}{2\pi}}$ $(3x^2 - 3y^2)$	0

Figure A.5: The first nine spherical harmonic magnetic fields.

3,-2	10	$\frac{1}{2}\sqrt{\frac{105}{\pi}}$ yz	$\frac{1}{2}\sqrt{\frac{105}{\pi}}$ xz	$\frac{1}{2}\sqrt{\frac{105}{\pi}}$ xy
3,-1	11	$\frac{1}{4}\sqrt{\frac{21}{2\pi}}$ -2xy	$\frac{1}{4}\sqrt{\frac{21}{2\pi}}(4z^2 - x^2 - 3y^2)$	$\frac{1}{4}\sqrt{\frac{21}{2\pi}}$ 8zy
3,0	12	$\frac{1}{4}\sqrt{\frac{7}{\pi}}$ -6xz	$\frac{1}{4}\sqrt{\frac{7}{\pi}}$ -6yz	$\frac{1}{4}\sqrt{\frac{7}{\pi}}(6z^2 - 3x^2 - 3y^2)$
3,1	13	$\frac{1}{4}\sqrt{\frac{21}{2\pi}}(4z^2 - 3x^2 - y^2)$	$\frac{1}{4}\sqrt{\frac{21}{2\pi}}$ -2xy	$\frac{1}{4}\sqrt{\frac{21}{2\pi}}$ 8zx
3,2	14	$\frac{1}{4}\sqrt{\frac{105}{\pi}}$ 2xz	$\frac{1}{4}\sqrt{\frac{105}{\pi}}$ -2yz	$\frac{1}{4}\sqrt{\frac{105}{\pi}}(x^2 - y^2)$
3,3	15	$\frac{1}{4}\sqrt{\frac{35}{2\pi}}(3x^2 - 3y^2)$	$\frac{1}{4}\sqrt{\frac{35}{2\pi}}$ -6yx	0
4,-4	16	$\frac{3}{4}\sqrt{\frac{35}{\pi}}(3x^2y - y^3)$	$\frac{3}{4}\sqrt{\frac{35}{\pi}}(x^3 - 3xy^2)$	0
4,-3	17	$\frac{3}{4}\sqrt{\frac{35}{2\pi}}$ 6xyz	$\frac{3}{4}\sqrt{\frac{35}{2\pi}}(3x^2z - 3y^2z)$	$\frac{3}{4}\sqrt{\frac{35}{2\pi}}(3x^2y - y^3)$
4,-2	18	$\frac{3}{8}\sqrt{\frac{5}{\pi}}$ -2y(3x² + y² - 6z²)	$\frac{3}{8}\sqrt{\frac{5}{\pi}}$ -2x(x² + 3y² - 6z²)	$\frac{3}{8}\sqrt{\frac{5}{\pi}}$ 24xyz
4,-1	19	$\frac{3}{4}\sqrt{\frac{5}{2\pi}}$ -6xyz	$\frac{3}{4}\sqrt{\frac{5}{2\pi}}(4z^3 - 3x^2z - 9y^2z)$	$\frac{3}{4}\sqrt{\frac{5}{2\pi}}(12z^2y - 3x^2y - 3y^3)$

Figure A.6: Spherical harmonic magnetic fields 10-19.

4,0	20	$\frac{3}{16}\sqrt{\frac{1}{\pi}}(12x^3 + 12xy^2 - 48xz^2)$	$\frac{3}{16}\sqrt{\frac{1}{\pi}}(12y^3 + 12yx^2 - 48yz^2)$	$\frac{3}{16}\sqrt{\frac{1}{\pi}}(12z^3 + 12zx^2 - 48zy^2)$	
4,1	21	$\frac{3}{4}\sqrt{\frac{5}{2\pi}}(4z^3 - 9x^2z - 3y^2z)$	$\frac{3}{4}\sqrt{\frac{5}{2\pi}}$ -6xyz	$\frac{3}{4}\sqrt{\frac{5}{2\pi}}(12xz^2 - 3x^3 - 3xy^2)$	
4,2	22	$\frac{3}{8}\sqrt{\frac{5}{\pi}}(12xz^2 - 4x^3)$	$\frac{3}{8}\sqrt{\frac{5}{\pi}}(-12yz^2 + 4y^3)$	$\frac{3}{8}\sqrt{\frac{5}{\pi}}(12zx^2 - 12zy^2)$	
4,3	23	$\frac{3}{4}\sqrt{\frac{35}{2\pi}}(3x^2z - 3y^2z)$	$\frac{3}{4}\sqrt{\frac{35}{2\pi}}$ -6xyz	$\frac{3}{4}\sqrt{\frac{35}{2\pi}}(x^3 - 3xy^2)$	
4,4	24	$\frac{3}{16}\sqrt{\frac{35}{\pi}}(4x^3 - 12xy^2)$	$\frac{3}{16}\sqrt{\frac{35}{\pi}}(4y^3 - 12yx^2)$	0	

Figure A.7: Spherical harmonic magnetic fields 20-24.

Appendix B

Biot-Savart Law & Numerical Validation Problems

This section begins by finding an analytic solution to Equation (3.8) at some arbitrary point in space, $\vec{r}$, due to a finite-length wire segment with unspecified orientation carrying some current I. For numerical ease this is completed in the Cartesian coordinate system, though this does lead to a relatively cumbersome expression.

For a segment extending from $[x'_1, y'_1, z'_1]$ to $[x'_2, y'_2, z'_2]$ comprehension of the resulting expression is aided by defining some additional terms of the form, $\tilde{x_1} = x - x'_1$ and $\tilde{x_2} = x - x'_2$. The term $\tilde{x_1}$ is then a geometrical relationship linking a measurement

coordinate x to a related point on the current carrying segment, x_1'. This leads to,

$$B_x(\vec{r}) = \frac{\mu_0 I}{4\pi}\Big(- \big(- \tilde{y_1}\big(\frac{\tilde{z_1}}{(\tilde{x_1}^2 + \tilde{y_1}^2)(\tilde{x_1}^2 + \tilde{y_1}^2 + \tilde{z_1})^{\frac{1}{2}}} - \frac{\tilde{z_2}}{(\tilde{x_1}^2 + \tilde{y_1}^2)(\tilde{x_1}^2 + \tilde{y_1}^2 + \tilde{z_2})^{\frac{1}{2}}}\big)\big) \tag{B.1}$$

$$+\big(- \tilde{z_1}\big(\frac{\tilde{y_1}}{(\tilde{x_1}^2 + \tilde{z_1}^2)(\tilde{x_1}^2 + \tilde{y_1}^2 + \tilde{z_1})^{\frac{1}{2}}} - \frac{\tilde{y_2}}{(\tilde{x_1}^2 + \tilde{z_1}^2)(\tilde{x_1}^2 + \tilde{y_2}^2 + \tilde{z_1})^{\frac{1}{2}}}\big)\big)\Big),$$

$$B_y(\vec{r}) = \frac{\mu_0 I}{4\pi}\Big(- \big(- \tilde{x_1}\big(\frac{\tilde{z_1}}{(\tilde{x_1}^2 + \tilde{y_1}^2)(\tilde{x_1}^2 + \tilde{y_1}^2 + \tilde{z_1})^{\frac{1}{2}}} - \frac{\tilde{z_2}}{(\tilde{x_1}^2 + \tilde{y_1}^2)(\tilde{x_1}^2 + \tilde{y_1}^2 + \tilde{z_2})^{\frac{1}{2}}}\big)\big) \tag{B.2}$$

$$+\big(- \tilde{z_1}\big(\frac{\tilde{x_1}}{(\tilde{y_1}^2 + \tilde{z_1}^2)(\tilde{x_1}^2 + \tilde{y_1}^2 + \tilde{z_1})^{\frac{1}{2}}} - \frac{\tilde{x_2}}{(\tilde{y_1}^2 + \tilde{z_1}^2)(\tilde{x_2}^2 + \tilde{y_1}^2 + \tilde{z_1})^{\frac{1}{2}}}\big)\big)\Big),$$

$$B_z(\vec{r}) = \frac{\mu_0 I}{4\pi}\Big(- \big(- \tilde{x_1}\big(\frac{\tilde{y_1}}{(\tilde{x_1}^2 + \tilde{z_1}^2)(\tilde{x_1}^2 + \tilde{y_1}^2 + \tilde{z_1})^{\frac{1}{2}}} - \frac{\tilde{y_2}}{(\tilde{x_1}^2 + \tilde{z_1}^2)(\tilde{x_1}^2 + \tilde{y_2}^2 + \tilde{z_2})^{\frac{1}{2}}}\big)\big) \tag{B.3}$$

$$+\big(- \tilde{y_1}\big(\frac{\tilde{x_1}}{(\tilde{y_1}^2 + \tilde{z_1}^2)(\tilde{x_1}^2 + \tilde{y_1}^2 + \tilde{z_1})^{\frac{1}{2}}} - \frac{\tilde{x_2}}{(\tilde{y_1}^2 + \tilde{z_1}^2)(\tilde{x_2}^2 + \tilde{y_1}^2 + \tilde{z_1})^{\frac{1}{2}}}\big)\big)\Big).$$

The analytic solution to Equation (3.17) is found by simple differentiation of this expression with respect to I. A related expression can also be found for the Magnetic Vector Potential by solving Equation (3.55) in the same manner. Combination of this expression with a superposition derived summation over multiple wire-segments allows the calculation of both the magnetic flux from, and the flux-sensitivity with respect to current due to, an arbitrarily shaped (closed-loop) current source at any point in space. This concept can be further extended to find the field produced by multiple such sources acting as a single induction system.

B.0.1 Validation 1: Square Loop

A simple numerical verification problem concerns a simple square current loop of side length S, for which the on axis flux solution for a coil in the XY plane is given by[210],

$$B_z = \frac{\mu_0 I}{2\pi}\frac{S^2}{[z^2 + \frac{S}{2}^2]\sqrt{z^2 + 2(\frac{S}{2})^2}}. \tag{B.4}$$

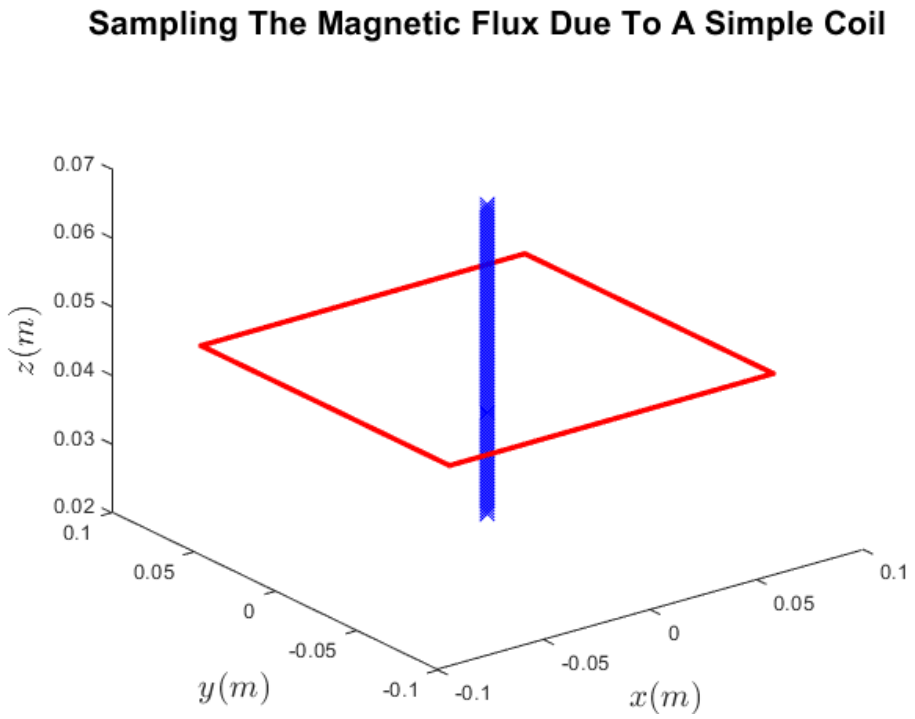


Figure B.1: The square current source modelled, and the related sampling points for comparison to an analytic result in the first validation problem. The scales are normalised for convenience.

The geometry, consisting of four wire segments, is generated and Equation (B.1) is combined with Equation (3.18) to find the associated magnetostatic field. The results for this implementation are given in Figure B.2 for completeness and confirms this analytic technique holds exactly for the test problem and has been correctly implemented for a simple geometry.

This is programmed in MATLAB which is well suited to matrix calculations, but poorly suited to looped computations. For this reason, noting the arbitrary case of individual inductors (i), measurement/simulation positions (j) and wire elements in an individual inductor (k), it is convenient to repeat the various vectors containing the related data into three dimensional arrays of size $[i, j, k]$. For example, the number of measurement positions is a vector of size $[1xj]$ which is then repeated i times in dimension 1 and k times in dimension 3. This reduces the arbitrary problem, with three dimensional simula-

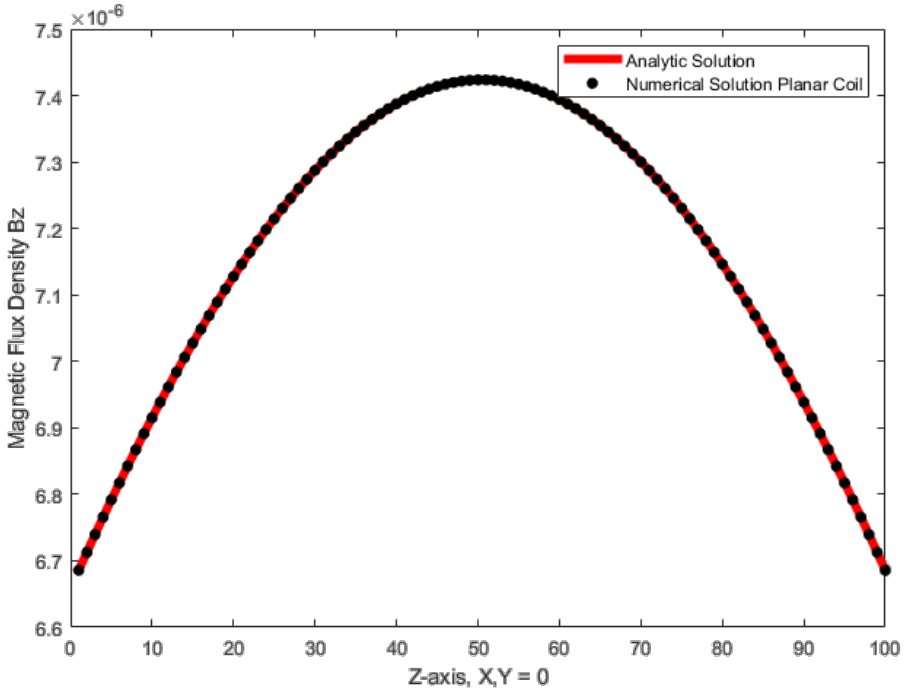


Figure B.2: A comparison between the numerical results from a square current source constructed of four individual wire segments and the analytic result for the same square current loop.

tions of $[B_x, B_y, B_z]$, to a series of matrix-vector operations which can be solved effectively instantaneously provided the set of matrices stay within the available memory.

B.0.2 Validation 2: Curved Geometries

A problem of more interest is the extension of the application of this equation to curved geometries, for example circular coils which are convenient because they technically have an infinite number of sides. For a circular coil of radius R orientated along the XY axis, the on axis magnetic flux due to the current source is given by,

$$B_z = \frac{\mu_0}{4\pi} \frac{2\pi R^2 I}{(z + R^2)^{3/2}}. \tag{B.5}$$

The visual form of a circle as approximated by varying numbers of straight-segments can be seen in Figure B.3. Using the same numerical approach of Equation (B.1) combined with Equation (3.18), the results for those geometries are given in Figure B.4.

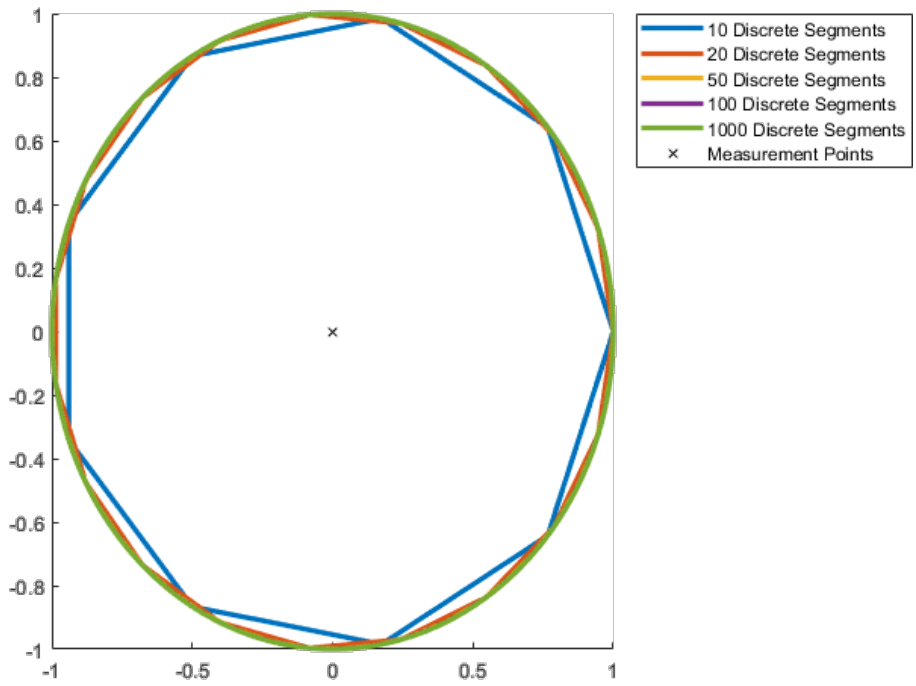


Figure B.3: The various approximations to a circular coil as constructed with a varying number of individual line-segments. Scales are normalised.

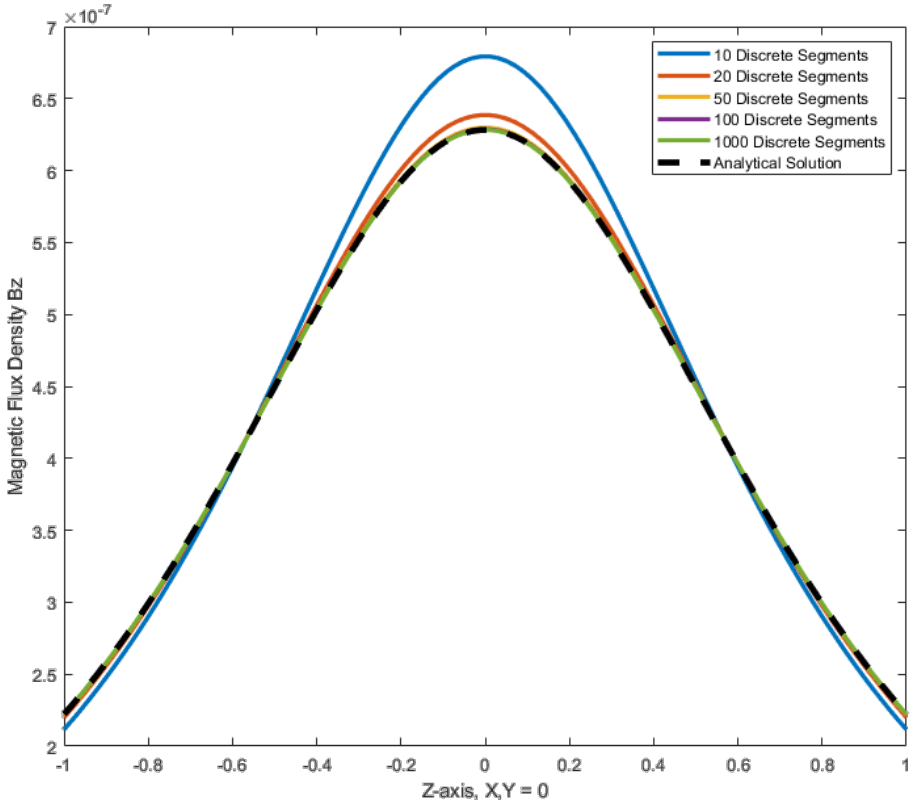


Figure B.4: A comparison between the numerical results for various approximately circular current sources constructed of a varying number of segments and the analytic result for a perfect circular current loop.

The behaviour matches the intuitive conclusion from the visualisation of approximations to a circular coil. With a small number of elements there is some deviation between the numerical and analytic solution, tending to an exact answer with around 10,000 segments. This leads to an additional caveat in the simulation of curved geometries requiring some minimal amount of discretisation.

B.0.3 Validation 3: Least-Squares Optimised Currents

Knowing from Appendices B.0.1-B.0.2 that the basic physics is simulated correctly, Equation (3.23) should be able to reproduce known behaviour given multiple current sources. The obvious candidate is the well studied Helmholtz pair, which has optimal separations of half the diameter and 0.5445 times the side length for circular and square coil pairs respectively[211].

A rendering of two square inductors of this type is provided in Figure B.5. This pair

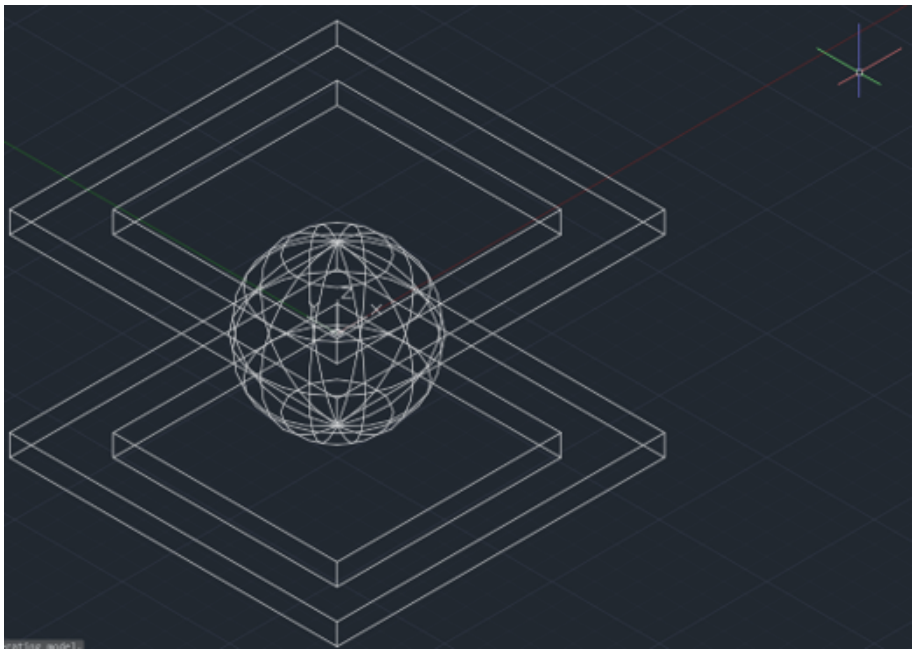


Figure B.5: A three dimensional rendering of a Helmholtz pair. Here the spherical region represents the volume of interest in the optimisation problem - the region over which the coils should have good control of the form of the magnetic field.

is the perfect test candidate because equal currents produce a uniform on-axis magnetic - corresponding directly with the $n = 1$ targets, which differ only by rotations onto each

Cartesian axis. Opposite currents then produce an anti-helmholtz field which relates to the $n = 2$, $m = 0$ target, leading to six test-cases demonstrating the implementation of Equation (3.23).

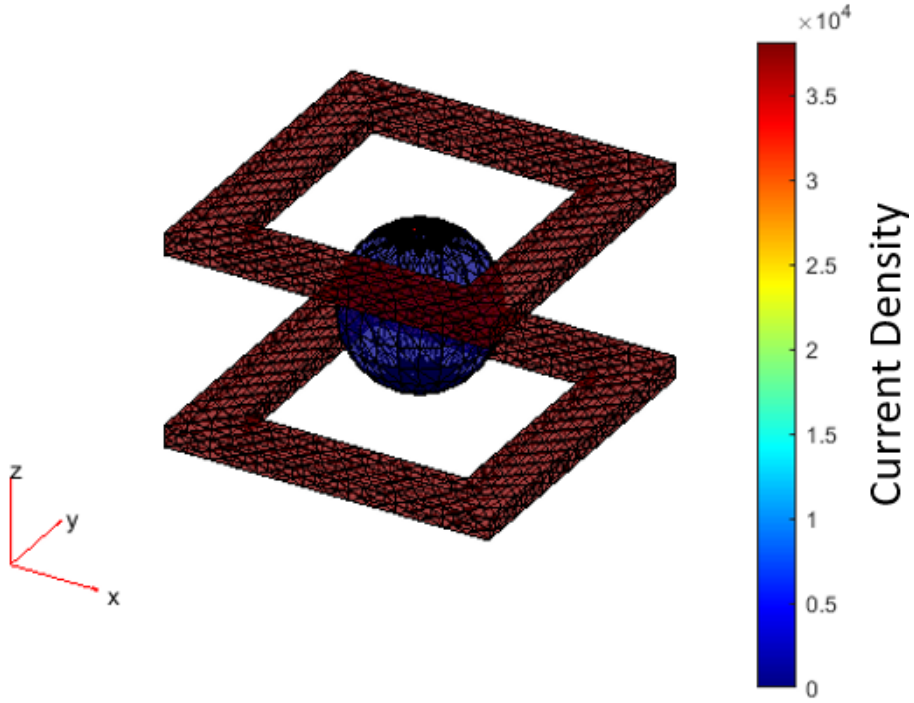


Figure B.6: A depiction of the currents derived from the optimisation procedure for the Helmholtz test case with the induction system orientated in the XY plane. The intention is to produce a uniform B_z field, and so this result matches intuition with identical currents through each coil.

The test results for the $n = 1$, $m = 0$ case are depicted in Figure B.6, matching expectations up to thirteen decimal places of precision. Repeating this process along the three possible orientations then confirms the algorithm can correctly target $\vec{B}_x$, $\vec{B}_y$ and $\vec{B}_z$.

Appendix C

Numerical Solutions to Integral Expressions

Two distinct mathematical approaches to modelling MIT have been considered. The first formulates the problem as a partial differential equation and is well suited to the finite-element approach. The second uses some integral expression, for which a different numerical approach is often necessary where analytic solutions do not exist.

A simple, easy to comprehend, approach to numerical integration is the midpoint rule. This can be used widely, especially in problems such as Equation (3.23) where the numerical accuracy of the final solution is not essential. For an arbitrary one-dimensional function, $f(x)$, divided into n_{mp} segments over the interval [a,b],

$$G(x) = \int_a^b f(x) = \sum_{\kappa=1}^{n_{mp}-1} (x_{\kappa+1} - x_\kappa) f\big(0.5(x_{\kappa+1} + x_\kappa)\big), \tag{C.1}$$

with each segment then of length $b - a/n_{mp}$. In essence, this represents an approximation of the curve, $f(x)$ in term of piece-wise constant functions. The main advantage of this is it is a quick, easily implemented method with accuracy limited by the number of discrete points used. However for high accuracy in three-dimensional problems this often requires a significant number of points.

More accurate discrete approximations of an integral expression are useful to reduce the number of computations required to solve a model. The composite Simpsons rule is another well-known example that, in the analogy, considers approximating the function in terms of piece-wise quadratic functions. Requiring an even number of intervals, n_{simp}, this becomes,

$$G(x) = \int_a^b f(x) = \frac{h}{3} \sum_{\kappa=1}^{n_{simp}/2} \Big[f\big(x_{2\kappa-2}\big) + 4f\big(x_{2\kappa-1}\big) + f\big(x_{2\kappa}\big) \Big], \tag{C.2}$$

with $h = \frac{(b-a)}{n_{simp}}$, $x_\kappa = a + \kappa h$ and $\kappa = 0, 1, ..., n$. Both of these numerical integration approaches can be described as quadrature rules, or weighted sums of function values at some points within the integration domain.

This brings us to an important class of quadrature rule - the Gaussian quadrature rule. An n point Gaussian quadrature rule is a numerical approach used to find an exact result when integrating polynomials of degree $2n - 1$. This is achieved by specific choices of weights and nodes, such,

$$G(x) = \int_{-1}^{1} f(x)dx = \sum_{\kappa=1}^{n} w_\kappa f(x_\kappa). \tag{C.3}$$

The most common is known as Gauss-Legendre quadrature, in which case the weights become,

$$w_\kappa = \frac{2}{(1 - x_\kappa^2)[P_n'(x_\kappa)]}, \tag{C.4}$$

and the nodes are then given by the κ^{th} roots of the associated Legendre polynomial such,

$$P_n(x_\kappa) = 0. \tag{C.5}$$

More often than not, an integration problem of interest involves some arbitrary bound $[a, b]$ rather than the standard interval $[-1, 1]$, requiring a change of interval,

$$\int_a^b f(x)dx = \int_{-1}^{1} f(\frac{b-a}{2}\Omega + \frac{b+a}{2})\frac{\partial x}{\partial \Omega}d\Omega, \tag{C.6}$$

with $\frac{\partial x}{\partial \Omega} = \frac{b-a}{2}$. The n-point Gaussian quadrature rule then becomes,

$$\int_a^b f(x)dx \approx \frac{b-a}{2} \sum_{\kappa=1}^{n} w_\kappa f(\frac{b-a}{2}\Omega_\kappa + \frac{b+a}{2}), \tag{C.7}$$

leaving the need to be able to calculate appropriate nodes, x_κ, and weights, w_κ, for a general problem. The Golub-Welsch algorithm is applied, which involves expressing a three-term recurrence relation resulting from the orthogonal set of (associated Legendre) polynomials in matrix form, and then finding the eigenvalues of the resulting tridiagonal matrix[212]. The root finding problem is thereby reduced to a simple eigenvalue problem which can be solved relatively efficiently in a matrix-based programming language.

Beyond the obvious convenience of highly accurate solutions to integration problems, Gaussian quadrature rules are also useful due to their flexibility. For example, a geophysical application of Equation (3.46) may consider a semi-infinite bound, leading to a problem of the form,

$$\int_0^\infty e^{-x} f(x)dx, \tag{C.8}$$

which is well suited to the application of the generalised Laguerre polynomials, also known as Gauss-Laguerre quadrature[213]. Analogously, problems over an infinite bound maybe expressed in the form,

$$\int_{-\infty}^\infty e^{-x^2} f(x)dx, \tag{C.9}$$

for which Hermite polynomials, and the associated Gauss-Hermite quadrature becomes the ideal numerical method[214]. Application of a form of Gauss quadrature to functions with endpoint singularities may find Jacobi polynomials, and the associated Gauss-Jacobi quadrature useful, whilst Chebyshev polynomials and Chebyshev–Gauss quadrature are useful for functions of the form $\sqrt{1-x^2}$.

Appendix D

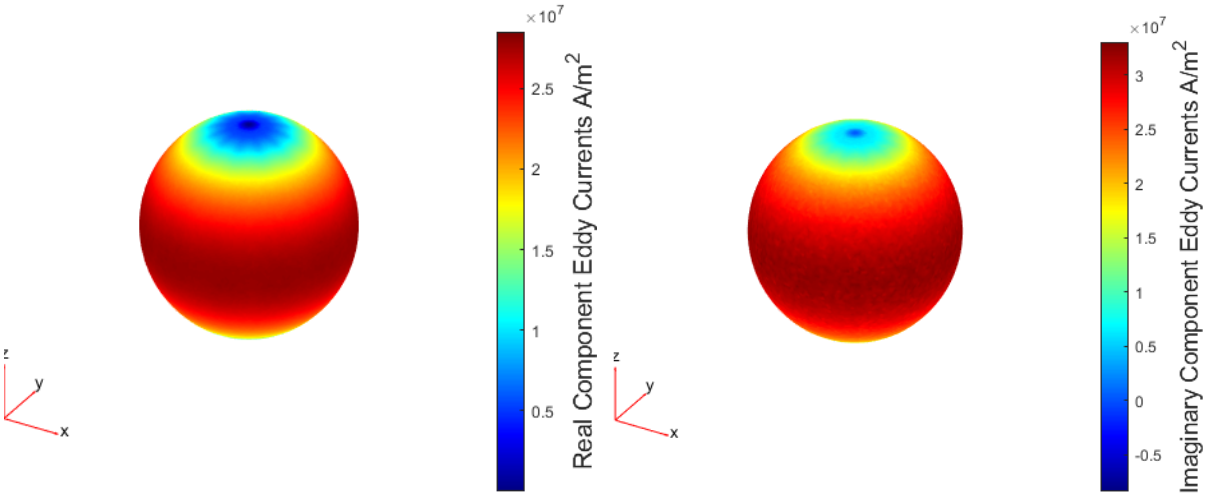
Frequency Domains and the Skin Depth

This appendix helps to identify appropriate uses of the approximate model developed in Equation 3.59 by exploring the impact of frequency on the skin effect.

D.1 Helmholtz Equation

An initial discussion begins with verification that the vector Helmholtz expression, Equation (3.41), is correctly implemented and can be extract directly from the reference paper [146] for the eddy-currents induced in a conducting sphere. Specifically a homogeneous sinusoidally oscillating magnetic flux of 0.1T is applied along the Z axis at 10kHz across a highly conductive sphere at 5MS/m.

Much of the literature suggests solving this problem requires edge-FEM, or some similar complex approach which is difficult and time-consuming to initiate. This is verified by the fact that a linear tetrahedral mesh generated using Delaunay triangulation frequently fails to produce the expected results. Applying NETGEN optimisation from the open-source GMSH software can overcome this issue, allowing fairly high-quality solutions to be found with linear nodal elements, Figures D.1a and D.1b. The reference problem samples the eddy-current distribution explicitly for a line in the Y-axis, with an appropri-



(a) Real (b) Imaginary

Figure D.1: The eddy-current distribution in a highly conductive sphere due to a homogeneous sinusoidally oscillating $10kHz$ magnetic field at $0.1T$ induced along the Z axis.

ate analytic solution which is provided in Figure D.2. The inaccuracies near the edges of the sphere could be rectified in a number of ways: by applying an edge-based FEM, using higher order nodal elements, or reducing the operating frequency therefore minimising the influence of the skin effect.

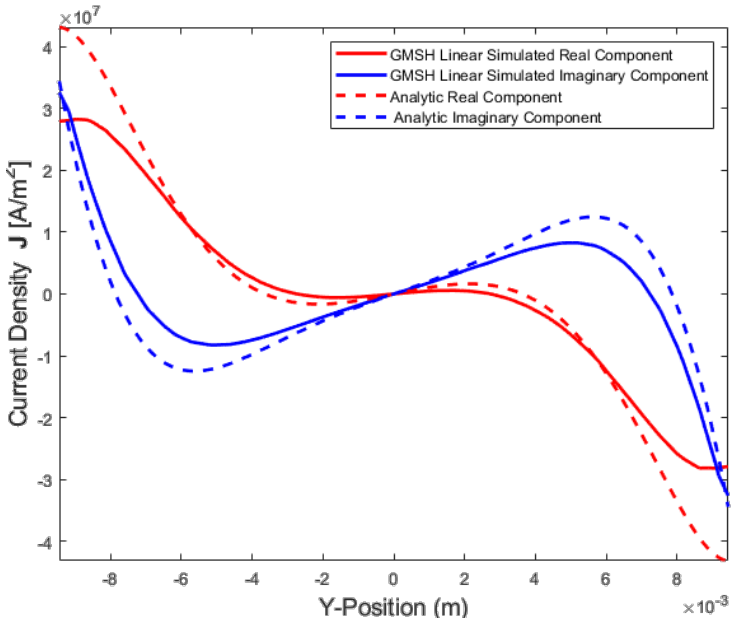


Figure D.2: A comparison between the analytic and nodal finite-element solution to the Helmholtz equation for a highly conducting sphere experiencing a homogeneous sinusoidally oscillating 0.1T induced field at 10KHz.

D.2 A Conductive Sphere

An discussion of the influence of the frequency of the impressed magnetic field, and the skin effect on the ability to image a region of material with passive electric properties can begin by varying the validation system in Figure D.1. For a highly conductive sphere in a uniform magnetic field at some quasi-static frequency, recall the skin depth can be estimated as,

$$\delta = \sqrt{\frac{2}{\omega\mu\sigma}}.$$

Initially, this is approximately 2.3mm - one fifth of the conductive sphere. The skin effect is then clearly reflected in the response profile of the sphere, Figure D.2 with a multi-phased response concentrated in the upper layer of the sphere. The response strength decays approaching the centre of the object, meaning central regions are effectively not interrogated by the applied magnetic field. For a three-dimensional imaging problem, this is clearly a problem because there is little information about the central regions in any measured response.

The skin-depth can be used to estimate a frequency which should allow thorough three-dimensional interrogation. From a rough estimate this corresponds to 540Hz to ensure radial penetration of the field, which has a notable effect on the distribution of eddy-currents - Figure D.3. The majority of the response is out-of-phase, with a simple linear profile containing most of the conductivity information. This may simplify experimentation because this component can be measured in the presence of a significantly stronger impressed magnetic field. Perfect characterisation must account for both components, with purely out-of-phase characterisation likely to underestimate the conductivity due to the 'missing' contributions from the in-phase response.

Doubling the skin depth to the diameter of the sphere leads to an approximate frequency of 135Hz and the eddy-current distribution given in Figure D.4. Physically this implies an impressed field which completely penetrates the sphere without dissipation.

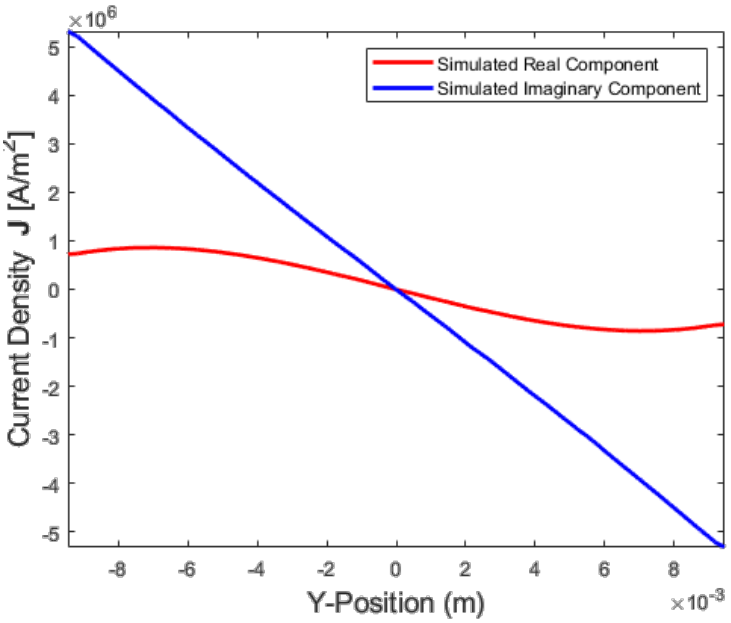


Figure D.3: The eddy-current distribution in a highly conductive sphere due to a homogeneous sinusoidally oscillating 540Hz magnetic field at 0.1T impressed along the Z axis.

The in-phase component is now effectively negligible, hence all in-phase measurements

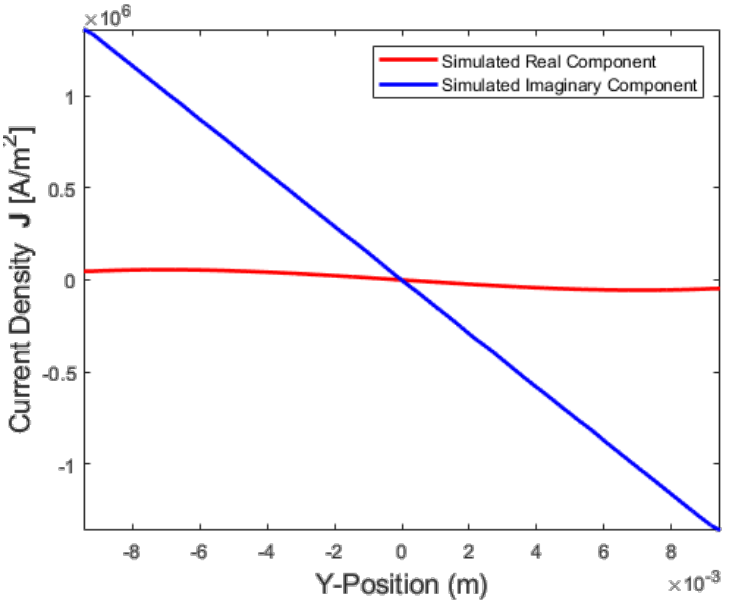


Figure D.4: The eddy-current distribution in a highly conductive sphere due to a homogeneous sinusoidally oscillating 135Hz magnetic field at 0.1T impressed along the Z axis.

can be directly attributed to the impressed field provided by the induction system. The out-of-phase component has then seen a linear reduction proportional to the frequency of the impressed field, and contains all information about the conductivity of the interrogated region. Experimentally both frequencies have a linear out-of-phase response, with

a balance sought between producing a detectable response and maintaining a minimal in-phase contribution - something most easily achieved by increasing the strength of the impressed field.

D.3 A Two Object System

The skin-effect is a ubiquitous problem but unfortunately becomes significantly more complex in systems with multiple objects, even with a simple impressed field. Take a slightly more complicated system with a same 5MS/m spherical object suspended at the centre of a box with a side-length equivalent to 32 radii - depicted in Figure D.5. Using

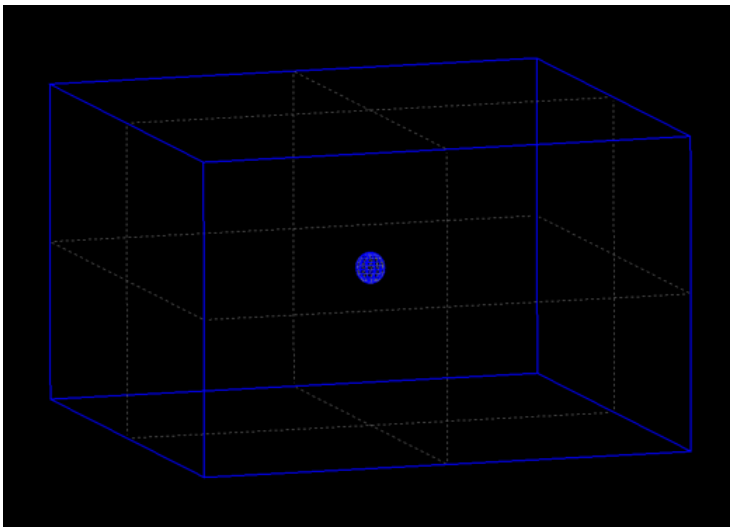


Figure D.5: A rendering of the geometry of the sphere suspended in a box test system.

the same homogeneous z-field at 0.1T and 10KHz, the box is assumed to be filled with seawater with a conductivity of 4.8S/m. A naive interpretation would assume then that the six orders of magnitude conductivity contrast would make the spherical object easily identifiable yet in reality this is far from true as seen in Figure D.6.

There is effectively zero eddy-current penetration inside the spherical object, with just a small peak in the upper surface layer. A similar effect is seen in the conductive box, with minimal response from roughly fifty-percent of the volume and, where the vast-majority of the systems' response is concentrated in the upper layer. The phase contributions also change directions between the skin-effect seen in each object, and the roughly equal but

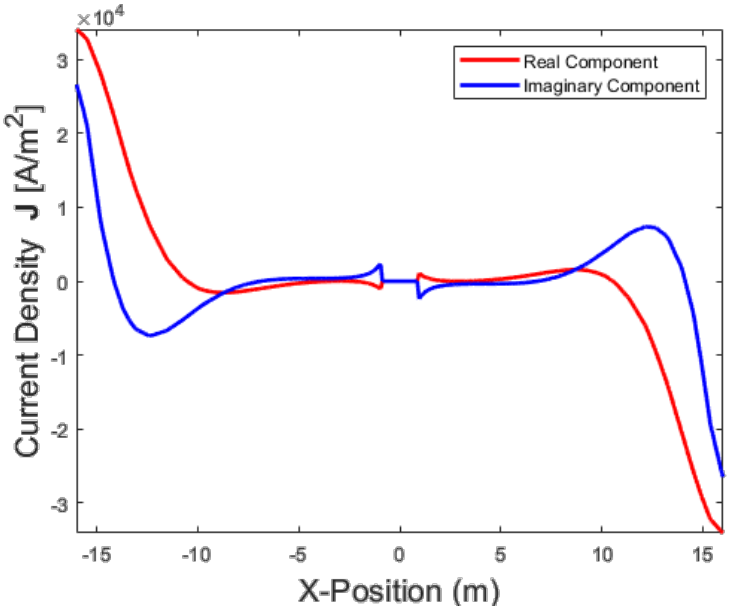


Figure D.6: The eddy-current distribution in a system consisting of a highly conductive sphere suspended at the centre of a seawater box due to a homogeneous sinusoidally oscillating 10KHz magnetic field at 0.1T impressed along the Z axis and sampled on a line at Y,Z = 0.

opposite composition of the high-conductivity response means any measurements of an amalgamated total response would be likely to render that object undetectable. The relative magnitudes are more easily understood by accounting for the volumetric nature of eddy-currents, but the complex phase relationship is not.

This can also help to identify the limitations of the simple estimation of the skin-depth in Equation (2.2). The highest conductivity is the sphere, and an MIT system wishes to ensure complete penetration of the whole domain. Taking σ to be that of the sphere and δ to be the size of the domain leads to an impressed field frequency of just 8.28×10^{-5}Hz. This is effectively a DC magnetic field which is apparent in the subsequent distribution of eddy-currents given in Figure D.7. Relating this to the flux theoretically detected, this would require a measurement of order 1×10^{-7}T in the presence of a 0.1T impressed field. Taking the limit where this incredibly weak response can be detected accurately (and measurements can be made inside the metallic sphere), there is still an additional issue for MIT. The eddy response is completely localised to the highly conductive object, with essentially no interrogation of the much larger box in which it is embedded, leading to an

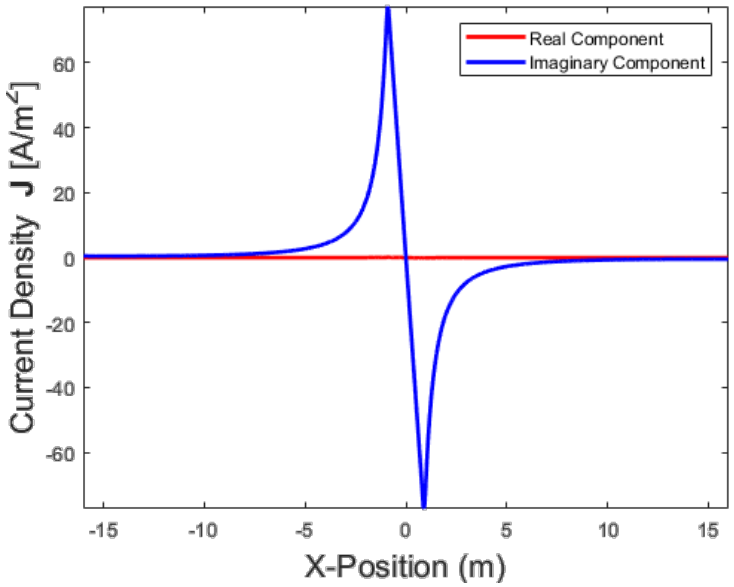


Figure D.7: The eddy-current distribution in a system consisting of a highly conductive sphere suspended at the centre of a seawater box due to a homogeneous 8.28×10^{-5}Hz magnetic field at 0.1T impressed along the Z axis and sampled on a line at Y,Z = 0.

MIT system which can only detect a small fraction of the volume it intends to image - far from ideal.

The balance to be sought then is the offset between fully inducing eddy-currents in the region of interrogation, and generating a responses which are significant enough to be detected. Most frequencies require considerations of the full-phase relationship for accurate characterisation. At 0.06Hz - $\delta \approx$ spheres' radius - response is approximately equally distributed between the in and out-of-phase contributions. The resulting response field is however much stronger making it significantly easier to detect.

Detection of both regions of conductivity can be achieved with careful choice of the induction frequency, however characterisation beyond the geometry of the objects is a difficult task. To do so, the frequency of the impressed field must be high enough that a response is generated from the bulk volume, but not too high that the skin effect dominates as in Figure D.6. In this system, a frequency roughly in the range 250Hz to 2KHz should be sufficient to detect both regions with responses in a manner equivalent to the naive physical expectation. Though both regions are detected through this frequency range,

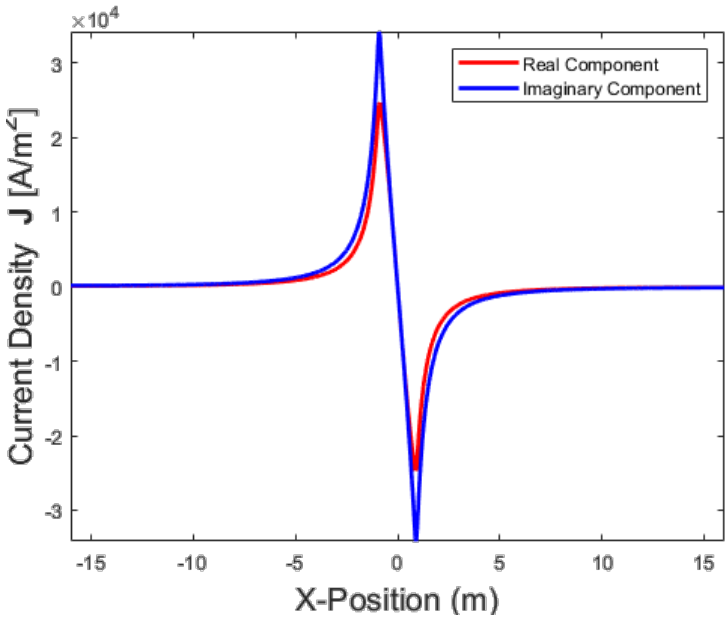
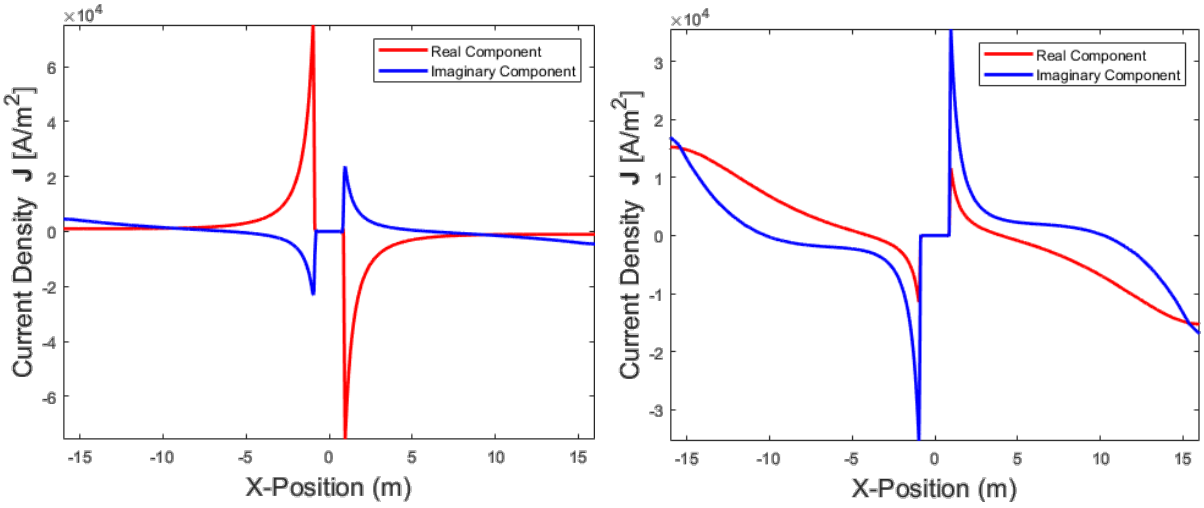


Figure D.8: The eddy-current distribution in a system consisting of a highly conductive sphere suspended at the centre of a seawater box due to a homogeneous 0.06Hz magnetic field at 0.1T impressed along the Z axis and sampled on a line at Y,Z = 0.



(a) 250Hz (b) 2KHz

Figure D.9: The eddy-current distribution in a system consisting of a highly conductive sphere suspended at the centre of a seawater box due to a homogeneous magnetic field at 0.1T with two frequencies impressed along the Z axis and sampled on a line at Y,Z = 0.

there is no flux penetration inside the sphere in either case. Thorough characterisation would likely require the application of multiple frequencies, for example 2KHz could provide information about the seawater box then 0.06Hz could ensure penetration of the conductive sphere.

There are also resonances which may be useful to provide further information with respect to what materials are present, which is identified in Figure D.9 by the direction-change in the magnitude between the two frequencies. Similar behaviour occurs in the out-of-phase component which can be seen by comparing Figure D.9 and D.8. Searching for resonances using frequency-sweeps is likely to find practical application in establishing an appropriate frequency for a particular MIT interrogation problem.

Appendix E

Regularisation, Preconditioning, and Prior-conditioning Matrices

Whilst the improved MIT formulations discussed in this thesis can vary significantly from a traditional implementation, the physics involved and subsequent inverse problem are similar enough that the same types of regularisation matrices have a good chance of being successful. As a result, the first set of regularisation matrices to be discussed are taken from existing literature in MIT.

The most common is the identity matrix[215] [216] [217] [218],

$$\mathbf{L} = \mathbf{I} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}, \tag{E.1}$$

of dimensions $(\kappa \times \zeta)$ for our existing set of equations. The influence of this regularisation approach has already been discussed in the previous section.

Beyond the identity matrix, the most common choice of $\mathbf{L}$ is plausibly a banded ap-

proximation to the Y^{th} order derivative operator derived from a finite-difference stencil. The first order matrix in banded form, $\mathbf{L_1}$, is given by,

$$\mathbf{L_1} = \begin{bmatrix} -1 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & -1 & 1 & \cdots & 0 \\ \vdots & \vdots & & \ddots & \ddots & \\ 0 & 0 & 0 & 0 & -1 & 1 \end{bmatrix}, \tag{E.2}$$

with dimensions of $(\kappa - 1 \times \kappa)$. This matrix has a uniform null-space, and so a constant component of the solution to the inverse problem this is applied to will not be suppressed by subsequent regularisation.

Often, particularly in the iterative schemes discussed later, a square regularisation matrix is required. To avoid inhibiting the smoothing properties of the matrix, the best approach is to add some small perturbation ($\tilde{\epsilon}$) to form a square matrix for example,

$$\tilde{\mathbf{L_1}} = \begin{bmatrix} -1 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & -1 & 1 & \cdots & 0 \\ \vdots & \vdots & & \ddots & \ddots & \\ 0 & 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & \tilde{\epsilon} \end{bmatrix}, \tag{E.3}$$

where the perturbation is appended as an additional row, with a prepended/inserted row also possible. This forces the related element of the solution to have some magnitude related to the size of the perturbation - i.e a small perturbation leads to a small conductivity in the respective element in an MIT problem.

Beyond the first-derivative operator, the second order derivative operator, or Laplace

operator, in banded form is,

$$\mathbf{L_2} = \begin{bmatrix} -1 & 2 & -1 & 0 & \cdots & 0 \\ 0 & -1 & 2 & -1 & & \vdots \\ \vdots & & \ddots & \ddots & \ddots & 0 \\ 0 & 0 & 0 & -1 & 2 & -1 \end{bmatrix}, \tag{E.4}$$

with dimensions of $(\kappa - 2 \times \kappa)$. In this case, the null-space means this regularisation approach does not affect constant or linear vectors in the solution space[166].

One approach to produce a square regularisation matrix for the $\mathbf{L_2}$ matrix is by the addition of boundary conditions. For Dirichlet conditions this becomes,

$$\tilde{\mathbf{L_2}} = \begin{bmatrix} 2 & -1 & 0 & 0 & \cdots & 0 \\ -1 & 2 & -1 & 0 & \cdots & 0 \\ 0 & -1 & 2 & -1 & & \vdots \\ \vdots & & \ddots & \ddots & \ddots & 0 \\ 0 & 0 & 0 & -1 & 2 & -1 \\ 0 & 0 & 0 & 0 & -1 & 2 \end{bmatrix}, \tag{E.5}$$

whilst Neumann conditions lead to,

$$\tilde{\mathbf{L_2}} = \begin{bmatrix} 1 & -1 & 0 & 0 & \cdots & 0 \\ -1 & 2 & -1 & 0 & \cdots & 0 \\ 0 & -1 & 2 & -1 & & \vdots \\ \vdots & & \ddots & \ddots & \ddots & 0 \\ 0 & 0 & 0 & -1 & 2 & -1 \\ 0 & 0 & 0 & 0 & -1 & 1 \end{bmatrix}, \tag{E.6}$$

with combinations of the two also possible depending on the specific problem[166]. In MIT Dirichlet conditions are often appropriate in medical or defence applications, where the measured response can be confined to a small known region of interest. Neumann condi-

tions may be more appropriate in a geophysics application, where the conductivity profile is typically continuous. Both banded derivatives, Equation (E.2) and Equation (E.4), can be found directly applied to MIT in [218] with analysis of their effectiveness on a system which produces a 'simple' magnetic field.

Anti-reflective boundary conditions can be applied through the use of zero padding at each end of the matrix in Equation (E.4). This then preserves the continuity of both the solution and its first derivative at the boundary and should be a more effective smooth matrix than the alternatives when the desired conductivity profile has a significant linear component[166].

Another approach to modifying Equation (E.4) to generate a square regularisation matrix is through the use of a circulant extension - well suited to problems with periodic solutions[219]. For problems with desired solutions containing both periodicity and linear growth, taking a weighted average of the anti-reflective boundary conditions and circulant extensions - the 'weighted average of circulant and zero-padded regularisation operators[219]' - can be an ideal modification. Here the null space then contains discretizations of both linear and slowly oscillating functions.

An alternative way to approximate the Laplace regularisation operator in Equation (E.4) for MIT based regularisation is through the construction of a nearest-neighbours matrix $\mathbf{N}$[215] [220] [171] with $\mathbf{L}^\top\mathbf{L} = \mathbf{N}$ and,

$$N_{ij} = \begin{cases} n_{ele}, & \text{for } i = j, \\ -1, & \text{for i,j neighbours ,} \\ 0, & \text{otherwise ,} \end{cases} \tag{E.7}$$

where n_{ele} is the number of elements which share a common face. The exact form of this matrix is dependent upon the form of discretisation used and is noticeably simpler for structured grids.

For both $\mathbf{L_1}$, $\mathbf{L_2}$, and all their variants, the null-space constrains the problem towards some increasingly smooth result, plus whatever other apriori information is included in the null-space of $\mathbf{L}$. These matrices are mostly sparse and this can be utilised to minimise the storage space when κ gets large, assuming the inverse procedure is capable of handling sparse matrices. The package developed by Hanson in [155] allows higher order derivatives to be experimented with as well which can be useful for some highly non-linear problems.

Another explicit regularisation approach well suited to MIT is one which includes a variance uniformization constraint. In this approach the regularisation matrix is chosen to account for the strong attenuation of electromagnetic phenomena within media as the distance from the source increases. Recalling the SVD of the sensitivity matrix, $M_{\zeta\kappa} = \mathbf{U}\eta\mathbf{V}^\top$, the regularisation is then chosen such, $\mathbf{L}^\top\mathbf{L} = \mathbf{V}\mathbf{D}\mathbf{V}^\top$, with the $\mathbf{D}$ matrix defined by,

$$\mathbf{D} = \frac{\eta}{\sqrt{c}} - \eta^2 \tag{E.8}$$

where c is some additional scalar regularisation parameter[215] [221].

There are a broad number of additional approaches which fall into this category of Regularisation Matrices, and including them all is an impossible task. For this reason, the reader is referred to some relevant papers for others which were implemented, tested, and found to be applicable to at least one MIT inversion problem over the course of this PhD - noting that the sheer volume of algorithms implemented means it is effectively impossible to include data for all of them.

Two algorithms which focus on edge-detection, or the reconstruction of sharp boundaries in high-contrast MIT such as metals in air are the Total Variation approach [139] and Sobel regularisation [222]. A more state-of-the-art approach considers the application of Fractional Regularisation Matrices[223], this was found to be particularly relevant in full-phased inverse problems, where the sensitivity matrix is some combination of real

and imaginary components such as those derived from Equation (3.41).

The majority of the methods discussed so far can be implemented relatively conveniently by making use of the package already widely cited throughout this thesis in [155], with some simple modifications required to properly implement the extensions and boundary conditions on various regularisation matrices.