3D QUANTITATIVE INTERPRETATION OF ARCHAEO-MAGNETIC SURVEYS: APPLICATION OF MATHEMATICAL MODELLING TO DETERMINE DEPTHS AND PHYSICAL CHARACTERISTICS OF BURIED MATERIALS

Thesis submitted for the degree of

Doctor of Philosophy

at the University of Leicester

by

Samuel Cheyney MGeol (Leicester)

Department of Geology

University of Leicester

2012

3D quantitative interpretation of archaeomagnetic surveys: application of mathematical modelling to determine depths and physical characteristics of buried materials

Samuel Cheyney

Abstract

High-resolution total-field magnetic data can be collected rapidly and relatively cheaply over large archaeological sites due to recent advances in data collection. However, interpretation of these datasets still generally comprises a sequence of data correction and filtering operations prior to a 2D visual interpretation based on pattern recognition. In contrast, current developments in aero-magnetic interpretation have led to several tools for identifying location, shape and depth information of anomalous sources.

Information regarding the horizontal location of sources can be obtained from derivative-based methods such as the horizontal gradient magnitude, tilt-angle and theta-map, however these methods typically fail for archaeological purposes due to the high noise content of these datasets. Here it is shown that using pseudogravity data with these techniques overcomes the problem of noise amplification that has previously hampered archaeological implementation.

Ideally, what it desired from the survey is information regarding the physical properties of causative features, rather than their magnetic responses. To achieve this, 3D inversion of the dataset can be used to produce a magnetic susceptibility model of the subsurface. It is demonstrated here that significant changes to the routine procedure employed during inversion of aero-magnetic data are required to avoid large discrepancies between the true and modelled depths of archaeological features. Here, a workflow is proposed that demonstrates how to process and model magnetic data for archaeological interpretation. A number of different stages are identified, dependant on the level of interpretation required, from enhancing the data image to full 3D quantitative results.

Greater interpretation of the magnetic data will aid the planning of higher resolution surveys such as GPR or resistivity, or to position targeted excavations. In some situations the detailed 3D geophysical model may supply sufficient archaeological information to avoid excavation, thus reducing cost and risk to the archaeological material.

Acknowledgements

There are many people who have contributed to allow this thesis to be completed. Firstly, I would like to thank the supervisors, Ian Hill, Stewart Fishwick, Neil Linford and Chris Leech, who allowed me the freedom to take the project in the direction I wished, yet were always available to give advice and guidance. Also I would like to acknowledge the Natural Environment Research Council (NERC) and Geomatrix for providing the funding without which none of this would have been possible.

A lot of this thesis is due to data from English Heritage, who allowed me to assist on site with their fieldwork and slowly acquire their knowledge in the evenings. Therefore, as well as Neil, this thesis owes a lot to the help and support of Paul Linford, Louise Martin and Andy Payne.

At Leicester, there are many more people to thank. Peter Maguire and Richard England are accountable along with Ian for enthusing me with the prospect of being a geophysicist in the first place, and Stephen Grebby and Sam Matthews deserve credit for attempting to be pseudo-geophysicists whenever I've had trouble throughout this project. Mohib Mohibullah Mohibullah Kakar Khan has endured wind, rain and snow with me, all within our little office, and has never once stopped smiling which has been an inspiration!

I would like to thank Geraint Jones, Skander Hili and David Andréis for patiently teaching me inverse methods, Linux and Matlab at OHM, all of which contributed largely to this project, and especially to Tim Tylor-Jones and Nick McArdle who convinced me I could undertake it during long walks along Aberdeen beach.

A big thank you must go to Matt Garthwaite, Ian Jarrett, Laura Balderstone and Alice Barkell for encouragement and distraction in just the right measure, with the added note to Matt for his part in the skills exchange between Latex and Matlab.

Becky Williams deserves a special mention because she's a very special person, and it would be impossible for me to express my thanks and love in just a few lines.

Finally, the biggest thank you of all must go to mum, dad and Hannah who persuaded me not to quit after a month of being an undergraduate, and I'm sure must be continually amazed that I am still here 11 years later. I owe everything to your unquestionable support, and for listening to me, even when I have nothing to say.

P.s. Ian, I hope this thesis wasn't the cause of your early retirement. Enjoy the sailing ;-)

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Magnetic surveying is a commonly used tool for first-pass non-invasive archaeological surveying, and is often used to target areas for more detailed geophysical investigation, or excavation. Quick and routine processing of magnetic datasets mean survey results are typically viewed as 2D greyscale maps and the shapes of anomalies are interpreted in terms of likely archaeological structures. This technique is simple, but ignores some of the information content of the data.

The data collected using dense spatial sampling with modern precise instrumentation are capable of yielding numerical estimates of the depths to buried structures, and their physical properties. The magnetic field measured at the surface is a superposition of the responses to all anomalous magnetic susceptibilities and remanence in the subsurface, and is therefore capable of revealing a 3-dimensional model of the causative magnetic properties. The application of mathematical modelling techniques to very near surface surveys such as for archaeology is quite rare, however similar methods are used in regional scale mineral exploration surveys.

The aim of this project is to achieve 3D quantitative interpretation of archaeo-magnetic total field datasets. This is likely to include a workflow for the rapid generation of information regarding the spatial position of anomalous bodies, which can be used as a constraint on further characterisation of the subsurface magnetic properties. To achieve this the following objectives must be achieved:

1. An evaluation of methods developed for use with large-scale aero-magnetic datasets, for transforming the total field dataset into an easily interpretable form. This should investigate any problems associated with the low signal-to-noise ratio, and high sampling density typically encountered with archaeo-magnetic datasets.

2. Produce a 3D model of the subsurface magnetic properties, which is consistent with the characteristics of typical archaeological structures and comparable to models of other geophysical properties and excavation information.

3. Investigate, where possible, the ability to achieve the above goals with minimal user input, to avoid biasing the final result based on individual operator preferences.

Part of the project has included a placement with English Heritage to observe the data acquisition techniques, and initial data processing typical of an archaeo-magnetic survey. This has included fieldwork resulting in several of the datasets used as case studies here.

This thesis is presented as a series of chapters examining the suitability of existing and novel approaches to deriving quantitative information from datasets, and any adaptations that are required to work effectively with the low signal-to-noise ratio typical of archaeo-magnetic datasets.

1.1 Chapter summary

Chapter 2. Magnetic methods are amongst the oldest geophysical techniques, with the first recordings of the Earth's field made in the 16th Century (Telford *et al.*, 1990). Since then, magnetic techniques have developed to be used today for piecing together the evidence for plate tectonics, searching for sedimentary basins and mapping igneous intrusions, minerals deposits, civil engineering targets and archaeology. Chapter 2 introduces the theory behind magnetic methods, introduces key terminology, causes of anomalous magnetisation as well as discussing

the equipment, methods of data collection and processing that is required to produce archaeo-magnetic datasets.

Chapter 3. The primary information typically required from magnetic datasets is the lateral position of any anomalous material. This information is generally derived manually from visual analysis of the data, and is often the full extent of the interpretation. The nature of magnetic signals mean that the peak anomaly will rarely be positioned directly above the source, due to the inclination of the inducing field. A number of methods are increasingly adopted in aero-magnetic surveys, which can be used to highlight source body parameters in potential field data. In Chapter 3, methods involving the derivatives of the total magnetic field are used in association with advanced image processing techniques to outline the edges of anomalous bodies more accurately.

Chapter 4. In a similar way to how the derivative-based methods are used to highlight the edge positions of anomalies, they can be used to provide estimates for the depth to sources. This is typically based on the attenuation rate with distance of the magnetic field from structures with simple geometries. In Chapter 4, these methods are evaluated, in particular the applicability of simple source geometries to typical archaeological structures.

Chapter 5. To produce a full 3D model of the causative apparent susceptibilities that result in the total field anomaly measured at the surface inverse modelling is required. Inverse modelling techniques have inherent ambiguity due to the nature of the mathematical 'inverse problem'. Often, although a good fit to the recorded values can be obtained, the final model will be non-unique and may be heavily biased by the starting model provided. Chapter 5 investigates the methodology required to invert synthetic archaeo-magnetic datasets with MAG3D, a code developed at the University of British Columbia for the modelling of aero-magnetic data. The different assumptions that are required for archaeological purposes are tested.

Chapter 6. Ideally, a subsurface model can be created that accounts for all the observations measured at the surface, however a proportion of the the magnetic field measured at any point

will be caused by magnetic bodies far from the survey area. In Chapter 6, the processing required to create a dataset suitable for modelling with a finite-sized mesh are investigated. This include methods for separating the features of interest from other trends, and assessing the presence of remanent magnetisation that would invalidate key assumptions used in the inverse modelling.

Chapter 7. The observations made in the previous chapters are used in Chapter 7 to produce a methodology for 3D quantitative interpretation of archaeo-magnetic datasets that is applied to three case studies. The first case study presented is from the NSGG Test Site 1, Leicester, a geophysical test site, where the depth and position of the buried features is known. The second dataset is from Silchester Roman Town, Hampshire, which has been an area of archaeological interest for many years, and therefore a wealth of geophysical and excavation information is available for comparison. The third case study is from Brading, Isle of Wight, which is representative of a site where currently the magnetic data is the only information available regarding the subsurface properties. This is typical of a situation where more thorough analysis of the magnetic data can aid interpretation and be used to target subsequent investigation. Where appropriate the susceptibility model is compared to existing information and data relating to the buried features.

Chapter 8. The conclusions for the project are presented in Chapter 8, where the aims and objectives of the project are revisited, along with a proposed work-flow for processing and interpreting total field archaeo-magnetic datasets. Further developments of this work are also proposed.

Appendices Some chapters have contributed heavily to articles published in peer-review journals. A large proportion of Chapter 3 was published in Archaeological Prospection in May 2011 (Cheyney *et al.*, 2011), and a copy can be found in Appendix A. Sections of Chapter 6 have contributed to a paper currently in press with Near Surface Geophysics and can be found in Appendix B.

Throughout the thesis, where data processes have been carried out, they have been performed by scripts written in MATLAB, interpreted from the original paper or equation, and presented on CD in Appendix C. The exceptions to this are Euler Deconvolution, which was performed in Geosoft Oasis Montaj, and magnetic inversion, which was conducted using the UBC-GIF's MAG3D.

Chapter 2

Magnetic surveying in archaeological prospection: scientific principles and literature review

2.1 Introduction

Magnetic surveying has become the most widely used method of geophysical prospecting over archaeological sites (see Gaffney & Gater (2003); Linford (2006); Aspinall *et al.* (2008)). It has been applied to a wide range of targets, in many differing geological settings, and has produced results where very small deviations in the Earth's magnetic field can be used to detect and provide evidence for previous land uses that have long since disappeared from the surface.

In this chapter the theory behind magnetic methods, the key terminology, and the causes of anomalous magnetisation will be discussed, as well as a review of the equipment used for archaeological prospection and the methods of data collection and processing that are typically applied to archaeo-magnetic datasets.

2.2 Fundamental terms

2.2.1 Magnetising force

Magnetic materials can be considered to consist of magnetic poles, which may either be north-seeking or south-seeking, as they tend to move towards the Earth's north or south poles respectively. For two poles with pole strengths p_1 and p_2, the force between the poles is given by:

$$\overrightarrow{F} = -\frac{\mu_0 p_1 p_2}{4\pi\mu_R r^2}\hat{r} \tag{2.1}$$

where μ_0 is the magnetic permeability of free space ($4\pi \times 10^7\ WbA^{-1}m^{-1}$), μ_R is the relative magnetic permeability of the medium separating the poles, r is the distance separating the poles, and $\hat{r}$ is the unit radial vector. If the pole types are alike the force is repulsive and negative, and if they are opposites the force will be attractive, and positive.

2.2.2 Magnetic field

A magnet can be thought of as consisting of two poles of opposite polarity. Surrounding a magnet is a magnetic field, which is the space though which the influence of the magnet exists. Within this, free magnetic dipoles tend to align their axes parallel to the field. The magnetic field strength can be quantified by the magnetic flux density which is calculated by:

$$\overrightarrow{B} = \frac{\mu_0 p}{4\pi\mu_R r^2}\hat{r} \tag{2.2}$$

In the SI system of units, B is measured in teslas (T), however this unit is very large relative to measurements in archaeo-magnetics, which are often presented in nanoteslas (nT). The magnetic field is related to the magnetising force, typically referred to as H, by the permeability of free space μ_0

$$\overrightarrow{B} = \mu_0 \overrightarrow{H} \tag{2.3}$$

where H has the units Am^{-1}. The magnetic flux density can be modified in the presence of a medium and is often presented as:

$$\overrightarrow{B} = \mu_R \mu_0 \overrightarrow{H} \tag{2.4}$$

μ_R may be a material constant, or a function based on the magnetising field strength.

2.2.3 Magnetic potential

The magnetising force is the negative gradient of the magnetic potential (V), which is the scalar value representing the work done against the magnetic field to bring a unit magnetic pole to that point from infinity. The magnetic potential is defined as:

$$V = \frac{\mu_0 p}{4\pi \mu_r r} \tag{2.5}$$

It is possible to calculate the component of the magnetic field in any direction from the magnetic potential, by the partial derivative in that particular direction.

2.2.4 Magnetic moment

The magnetic moment quantifies the strength and orientation of the dipolar magnetic field generated by a pair of magnetic poles. Where the poles are of strength p, separated by $\overrightarrow{l}$, it is defined as:

$$\overrightarrow{m} = p\,\overrightarrow{l} \tag{2.6}$$

where the magnetic moment, $\overrightarrow{m}$, is a vector directed from the negative pole to the positive pole, with the units Am^2.

2.2.5 Intensity of induced magnetisation

Magnetisation can be induced into a material when it is placed within a magnetic field. This occurs as magnetic poles are created on the faces perpendicular to the direction of the magnetising field, and hence tends to the direction of the magnetising field. The intensity of induced magnetisation is expressed as:

$$\overrightarrow{M} = \frac{\overrightarrow{m}A}{l} \tag{2.7}$$

where m is the induced magnetic moment, A is the area of the pole face, and l is the distance between the two poles. The units for the magnetisation are Am^{-1}, which are the same as that for the magnetising force H. The magnetisation and magnetising force are therefore related by a dimensionless quantity:

$$\overrightarrow{M} = k\overrightarrow{H} \tag{2.8}$$

where the dimensionless quantity k is the magnetic susceptibility. Although unit-less, magnetic susceptibility is often displayed with the postfix SI to show the system of units used. It should be noted that there is a difference between the susceptibility in the SI (International Standard) system of units, and susceptibility in the emu (Electromagnetic Units) of $1\ emu = 4\pi\ SI$.

2.2.6 Magnetic susceptibility and permeability

The magnetic susceptibility is a constant value and does not depend on the quantity of material analysed. It is related to the relative magnetic permeability μ_R by

$$\mu_R = 1 + k \tag{2.9}$$

The relative permeability is also dimensionless, and in a vacuum or free air is equal to 1. Therefore from Equation 2.9 it can be seen that the value of susceptibility in a vacuum or free air is zero.

2.3 The Earth's magnetic field

The major component of the magnetic field measured close to the Earth's surface is that due to the field generated within the Earth itself. Convection currents within the liquid outer core move around the inner core. Both the inner and outer core consist of iron and nickel, and therefore both are good electrical conductors. This motion around the inner core generates a field that is often approximated by a dipole situated towards the centre of the Earth, yet offset from the spin axis of the Earth by around 11°. The field can be measured at the surface ranging from 25000 nT at the geomagnetic equator to 65000 nT at the geomagnetic poles (Hansen *et al.*, 2005). In the United Kingdom, the Earth's field is around 48500 nT.

In reality, the field is much more complicated than a simple dipole model, with many distortions caused by the complex movement of material. The Earth's magnetic field is also known to vary with time due to changes in the convection currents in the core (Telford *et al.*, 1990). The field is modelled empirically by data from satellite and geomagnetic observatories to produce the International Geomagnetic Reference Field (IGRF). The IGRF can be used to provide regional corrections for some large scale surveys, however it has a minimum nominal wavelength of around 3000 km. A base station or further processing steps are therefore required to account for smaller scale deviations. For archaeological purposes, the IGRF can be used to identify the inclination, declination and magnitude of the inducing field for use when reducing a total field dataset to an anomaly dataset or carrying out processing steps such as reduction to the pole and the pseudogravity transformation. The secular changes to the IGRF are modelled by extrapolating observations, therefore the IGRF is revised every five years (Finlay *et al.*, 2010).

2.3.1 External magnetic fields

The geomagnetic field also includes a component that originates from sources external to the Earth. Diurnal variations occur due to solar radiation ionising the upper layers of the atmosphere, and can typically cause anomalies of the order of a few tens of nT throughout the day. Around the equator the diurnal variations can be larger, generating anomalies greater than 100 nT around a 1000 km wide zone called the equatorial electrojet. In polar regions similar variations are observed around the auroral electrojet. Magnetic storms caused by sunspot and solar flare activities can also generate anomalies of hundreds of nT, with much shorter periods.

2.4 Types of magnetic domain orientations

Materials can have electric charges, due to the particular configuration of electrons and protons in their atomic structure. Electrons exist in sub-shells around the central nuclei. A circulating

current can be generated by electrons in the sub-shells which in turn create a magnetic dipole and resultant magnetic field. This field however, is small in comparison to the field generated by the spinning nature of the electron sub-shells. The electrons in the sub-shells create magnetic moments that are aligned in two anti-parallel directions. No two electrons can have the same quantum state, therefore electrons pair up, with one of each pair spinning in each of the two directions. Each sub-shell can only accept a certain number of electrons, and when full, there is an even number of electrons, each paired up. The alignment of the magnetic moments is therefore even in both directions, and there is no resultant magnetic field. Where an odd number of electrons exist in a sub-shell a resultant magnetic moment is produced, however this is usually in a random direction which cancels throughout the material to produce a net zero magnetisation.

In the presence of an external field, the spin axis of some of the electrons align parallel to the external field. The motion of the electron produces a small current that creates a magnetic moment in the opposite direction to the external field. This effect is known as *diamagnetism*, and appears as a negative magnetisation or susceptibility. The effect of diamagnetism generally contributes little to the magnetisation of a material as the effect is small in comparison to other types of magnetisation. Water is an example of a diamagnetic material, where the other types of magnetisation are not large enough to overshadow the diamagnetism, and often exhibits a small negative susceptibility. The electron spin also has an effect on the magnetic field. For the paired electrons, this in itself has no effect on the magnetic field, yet for the unpaired electrons, the spin moments partially align with the external field and enhance it, resulting in a magnetisation in the direction of the applied field and a positive susceptibility. This is called *paramagnetism*. When taken out of the external field, thermal effects randomly orient these alignments, and the paramagnetism often disappears.

All minerals are diamagnetic, and many are paramagnetic, however the strength of these two fields is very small, and insignificant in relation to the strength of fields generated when surrounding atoms interact with each other. Where there are several unpaired electrons in overlap-

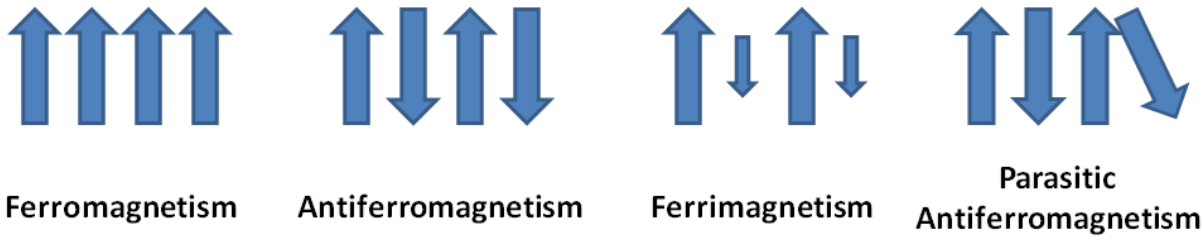

Figure 2.1: Types of magnetic domain orientations. Ferromagnetic and ferrimagnetic materials can have strong magnetic fields in the direction of the inducing field. While antiferromagnetic materials have no magnetic field, small inperfections often cause parasitic antiferromagnetic fields.

ping shells, the dipoles created can be coupled with exchange energy to the surrounding atoms, and produce magnetic domains. These domains can either be parallel or anti-parallel depending on the composition of the mineral (Figure 2.1). *Ferromagnetic* minerals, are where the alignment of the magnetic domains are parallel to each other, and in the same direction, causing a high, positive magnetisation and susceptibility. While several typical elements have ferromagnetic properties (iron, cobalt, nickel), these are often only when they are found in their pure form. The most common types of ferromagnetic materials are those containing ferric (Fe^{3+}) or ferrous (Fe^{2+}) iron.

In other materials, the magnetic domains align themselves parallel and anti-parallel, so the effect is self-cancelling and the magnetisation and susceptibility are close to zero. These are called *antiferromagnetic* materials. Pure haematite (Fe_2O_3) is an example of antiferromagnetism. Antiferromagnetic materials are rare, as often impurities are present, and the two opposite domains do not exactly cancel. This is called *parasitic antiferromagnetism*. Other materials can include both ferromagnetic and antiferromagnetic alignments. These are called *ferrimagnetic* materials. The alignment of the magnetic domains is still parallel and anti-parallel, however they are no longer perfectly balanced, which results in a net magnetisation. The magnetic fields generated from ferrimagnetic materials can be very high, with magnetite (Fe_3O_4) being a common example.

2.5 Remanent magnetisation

Temperature often affects the magnetisation of a material. For diamagnetic and paramagnetic materials, rising temperature results in lowering the susceptibility of the material as the thermal agitation causes a randomisation of the magnetic moments. Ferromagnetic and ferrimagnetic materials behave differently. The bonding of the magnetic domains is strong enough in these materials to withstand the effect of the rise in temperature, until a critical point, called the Curie temperature is reached. At the Curie temperature, the thermal energy is great enough to overcome the binding of the magnetic domain structure, and thus the material behaves as a paramagnetic material. This temperature will vary depending on the type and magnetic grain size of the mineral.

There are several types of remanent magnetisation. Thermoremanent magnetisation (TRM) can be found when ferromagnetic or ferrimagnetic materials are heated to beyond their Curie temperatures. Above the Curie temperature the magnetic domains will align with the ambient field and upon cooling back to below the Curie temperature the material may retain some properties of the ambient magnetic field that it was situated within. TRM is typical in igneous rocks, where they take on a TRM in the direction of the Earth's magnetic field in which they are formed, and thus allowed the identification of spreading margins and the discovery of plate tectonics.

Heating may also result in a change in the chemical composition in a material below the Curie temperature. This may also result in a remanent magnetisation becoming present, and is known a thermochemical remanent magnetisation (TCRM).

2.5.1 Hysteresis loop

When placed within an external magnetic field, the magnetisation of diamagnetic and paramagnetic materials is linear to the external field, and therefore displays a constant susceptibility.

When the external field is removed, the magnetisation of the material returns to zero. This is not the case with ferromagnetic materials, which do not always behave linearly.

When a ferromagnetic material is placed within a strengthening external magnetic field, the magnetisation rises, until the saturation magnetisation is reached. At this point, despite the external field continuing to get stronger, the magnetisation of the material remains the same, and the differential susceptibility becomes equal to zero. At this point, all of the magnetic dipoles are aligned with the external field, and thus, the magnetisation cannot be strengthened. If the external field is then removed, the magnetisation will not return to zero, as the material will have taken on a isothermal remanent magnetisation (IRM). IRM may also be caused by lightening strikes, which results in a very short, yet intense exposure to a strong magnetic force.

An example hysteresis curve is shown in Figure 2.2. A material with no remanent magnetisation is placed within the strong magnetic field. This is shown by the black line, and causes the magnetisation of the material to reach the saturation magnetisation. Upon removal of this field (the blue line until H=0) the material keeps a magnetisation despite no longer being placed within a field (marked as IRM). The strength of the magnetic field in the opposite direction required to remove this magnetisation is termed the coercive force. The saturation magnetisation can also be achieved in the opposite direction, and when the field is removed, an IRM exists in the opposite direction (red line).

2.6 Archaeological magnetisation

Sub-surface archaeological features can be detected using magnetometry, as they have either remanent magnetisation, induced magnetisation, or both, that allows a contrast in magnetisation to exist between the feature of interest and the surrounding environment. There are many reasons why anthropogenic land use could result in an anomalous magnetisation. Remanent magnetisation can be locked into archaeological materials such as kilns, ovens, floors that are

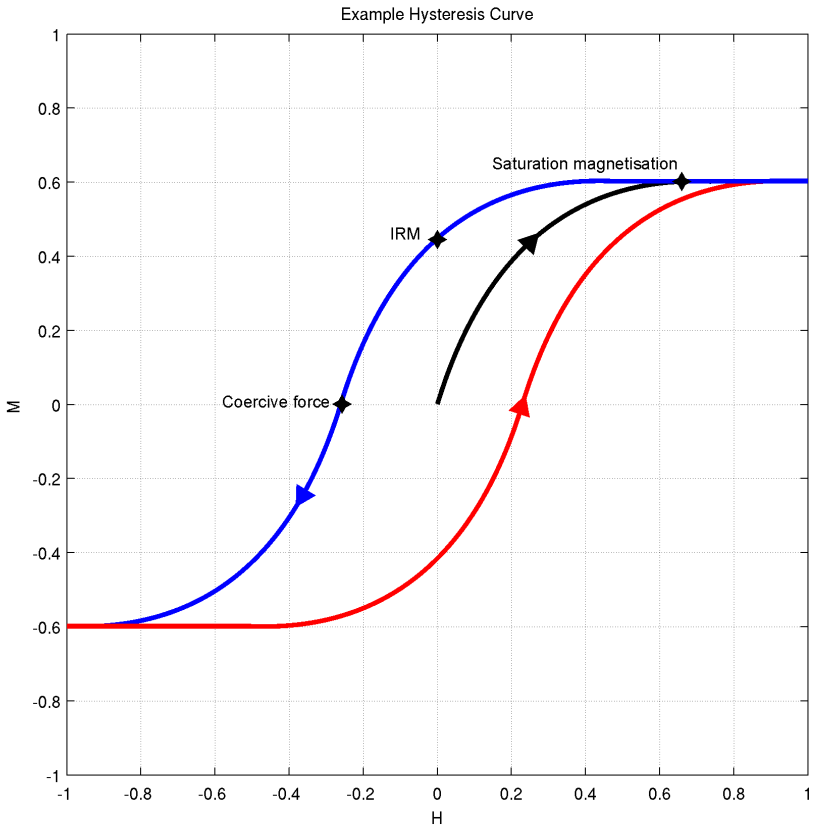

Figure 2.2: Example hysteresis curve for a material that is placed into a magnetic field until the saturation magnetisation is reached (black line), then the field is removed and reversed (blue line), and finally, placed again into a positive magnetic field.

baked, and fired bricks. Heating these materials results in a high thermal remanent magnetisation (TRM). Where the material is high in iron oxides, the heating process causes paramagnetic behaviour above the Curie temperature, and when subsequently cooled the magnetic moments are aligned in the direction of the Earth's magnetic field at the time. In the cases of kilns and floors, this will have remained and can be used to date the time of heating (Clark *et al.*, 1988). As clay bricks are usually moved and arranged into structures after being fired, the alignment of adjacent bricks is often different, and the combined magnetisation of a building is not as strong as an anomaly from a kiln or baked floor (Bevan, 1994; Hesse *et al.*, 1997).

Induced magnetisation also plays a large part in the ability to detect archaeological features. Anthropogenic uses of the land has been shown to often lead to an enhanced magnetic susceptibility in the upper soil level. This is due to the presence of high concentrations of iron oxides, typically haematite (α-Fe_2O_3), magnetite (Fe_3O_4) and maghemite (γ-Fe_2O_3). As haematite is anti-ferromagnetic, and therefore has a very low magnetic susceptibility, enhancement of the subsurface involves converting this haematite into the more magnetic forms of iron oxide, magnetite and maghemite. Haematite and maghemite exist in a fully oxidised state (Fe^{3+}) and have the same chemical makeup, yet different arrangements of ions in the crystal lattice. Magnetite, however, is partially oxidised, where some ions are Fe^{3+}, and others are Fe^{2+}.

Anthropogenic enhancement of the subsurface comes from reduction of haematite that results in the more magnetic form magnetite, and thus raising the magnetic susceptibility of the soil. Reduction can be caused by a number of reasons, that are well laid out in Aspinall *et al.* (2008). They are:

- Burning vegetation or other surface fires results in a lack of oxygen in the subsoil and hence a reducing effect on haematite resulting in an increase of magnetite (Le Borgne, 1955).
- Microbially mediated by the decay of organic waste (Linford, 2004). In this scenario bacteria create the reducing conditions.

- Similarly, Fassbinder *et al.* (1990) discovered a special type of magnetotactic bacteria which are able to create small magnetite crystals from the iron oxides in the soil. This is an effect that is typically seen in decaying wood, resulting in enhanced soil around post holes and other foundation features.

- The addition of magnetic material, such as broken pottery, brick fragments, metallic debris from machinery etc, to the soil (Weston, 2002).

- Natural enhancement of the soil's magnetic susceptibility during pedogenesis. This is a process that occurs in-situ without the presence of micro-organisms (Maher & Taylor, 1988).

Once the soil has experienced reducing conditions and haematite is converted to magnetite, it is often likely that it will undergo re-oxidisation. Re-oxidisation of magnetite tends not to result in a return to haematite, but instead a conversion to maghemite, which itself is also ferromagnetic and therefore has a high susceptibility leading the material to retain it's higher magnetisation.

The majority of the processes listed above require the presence of organic matter in order to aid the conversion to a higher susceptible form. Organic matter in the sub-surface is often found in greater abundances in the upper soil level, hence the top soil generally becomes much more susceptible then deeper soils, which can have iron oxides washed out by persistent water-logging of the soil. This results in a contrast in susceptibility that may become detectable should the magnetic stratigraphy become disturbed by anthropogenic processes.

2.7 Types of anomaly

Magnetic anomalies measured at the surface are indicative of a contrast between the feature and the susceptibility of the surrounding soil. The magnitude of these anomalies will therefore depend on both the susceptibility of both the original soil as well as the feature of interest. In

an area where the soil is low in iron oxides, anomalies may be much smaller in amplitude than areas where the soil contains a high concentration of iron oxides. For the same type of feature, the anomaly and subsequent susceptibility contrast will not necessarily reveal information about the type of material that has generated it.

Data from archaeomagnetic surveys typically reveals both positive and negative anomalies, as both positive and negative contrasts to the surrounding soil susceptibility can be present. Positive contrasts are typical in areas where ditches or other intrusions have been cut into top soil that has undergone enhancement of susceptibility. This is demonstrated in Figure 2.3, reproduced from Linford (2006). Here the ditch cuts through the highly susceptible top soil into a lower magnetised soil. Both types of soil become mixed, or the excavated soil becomes enhanced though further anthropogenic activity at the site, and once abandoned it is this soil that in-fills the ditch creating a positive contrast against the original soil. Subsequent sedimentation of soil on the site is likely to be as a uniform covering layer, resulting in a buried positive magnetic susceptibility source.

Apparent negative readings in magnetic anomaly datasets, can be formed in a similar way. Here the infill material is of low magnetic susceptibility compared to the surrounding soil, thus producing an apparent low reading. This is commonly found with foundations and walls made on a low susceptibility rock such as limestone, or where the ditches have been cut into high susceptibility soils and infilled with low-susceptibility alluvium (Aspinall *et al.*, 2008).

2.8 Magnetometers

Magnetometry is the most commonly used geophysical technique applied to archaeological site investigation, due to its passive nature and ability to cover large areas relatively quickly. Several different types of magnetometers have been developed, each having their advantages and disadvantages. Some magnetometers such as the proton precession and alkali vapour, output a scalar

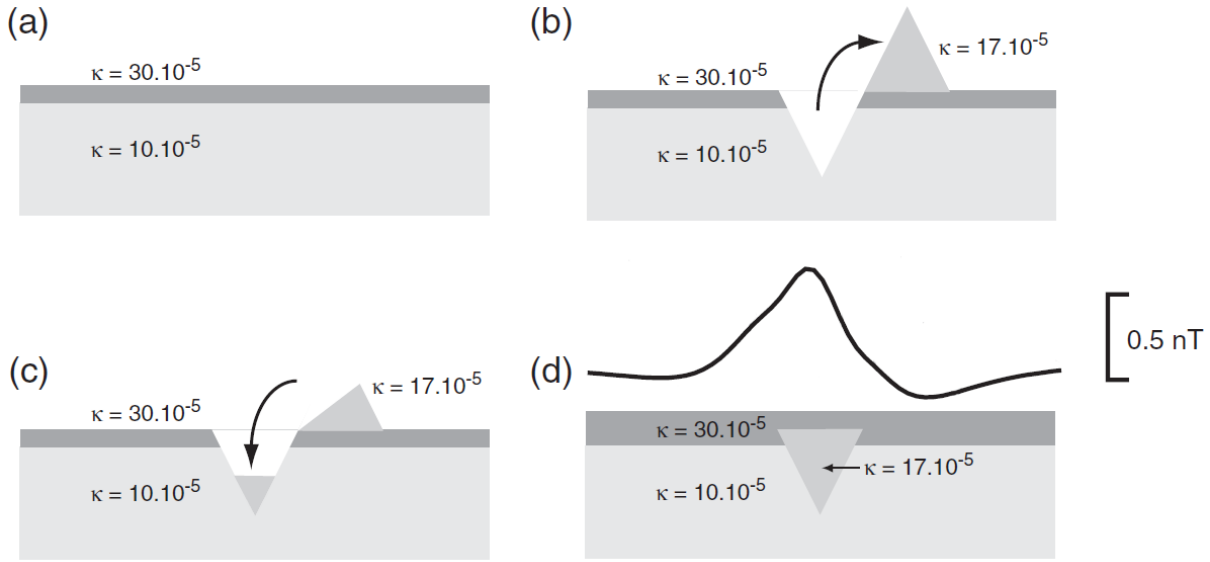

Figure 2.3: Example of the formation of a positive archaeological anomaly. (a) A ditch is cut through a magnetically enhanced top soil, into a lower susceptible ground. (b) During the excavation the enhanced top soil gets mixed with the lower susceptible soil. (c) The ditch becomes filled using the excavated mixture with now has a higher susceptibility to the lower soil. (d) Subsequent sedimentation is even across the site hence having no effect on the amplitude of the anomaly. The anomaly is generated from the contrast in susceptibility between the infill and surrounding area. From Linford (2006).

measurement of the total field, while others, such as the fluxgate and more recent development, the SQUID magnetometer, output the vector of the field parallel with the instruments alignment.

2.8.1 Fluxgate gradient magnetometer

Fluxgate magnetometers were first developed during the second world war as an instrument for detecting submarines. They typically consist of two identical cores of ferromagnetic material which have a very high permeability at low magnetic fields. The cores are separated at a distance usually around 0.5 m, and wound with primary and secondary coils. The wiring is arranged so that the cores are in opposition when they are mounted parallel. The primary coils are then connected in series and energised by low frequency alternating currents, which are strong enough to ensure that both coils are sufficiently magnetised to reach saturation. Due to the arrangement of the wiring the saturation is in opposite polarities. The secondary coils are

connected to a differential amplifier and the output consists of the difference between the two signals in the secondary coils.

In the absence of an external magnetic field the voltages in the two secondary coils cancel each other out, and the response is zero. However, in the presence of an external field, saturation occurs earlier in one of the coils than the other, and the induced voltages are out of phase. This causes the output of the two secondary coils to consist of voltage pulses proportional to the amplitude of the external field in the direction of the axis of the cores.

As fluxgate magnetometers do not measure the absolute magnetic field, calibration of each instrument is required. Other problems such as thermal drift can occur as the magnetic properties of the cores alter as temperature varies. Fluxgate magnetometers continue to be abundant in use by archaeologists, and typically record the accuracy of the vertical gradient component of the field to around 0.1-1 nT.

2.8.2 Proton precession magnetometer

Proton precession magnetometers consist of a bottle containing a fluid that is rich in hydrogen nuclei such as a high-purity decane, a primary coil and a secondary coil. The two coils may be combined into a single coil. The magnetometer works on the basis that the nuclei have a net magnetic moment, which coupled with their spin results in a precession around the axis of a magnetic field.

A measurement is taken by creating a polarising magnetic field using the primary coil, which is a solenoid around the bottle of protons. The field is much larger than the Earth's magnetic field (around 5-10 mT) and directed approximately at 90° to it. This results in the polarisation of the protons roughly normal to the Earth's field. When the current is switched off, the protons realign and precess about the Earth's field at an angular velocity that is proportional to the ambient

magnetic field. This induces an alternating voltage in the secondary coil with a frequency proportional to the field strength, which can be used to calculate the total magnetic field.

This process can take longer than a second to produce a reading, therefore proton precession magnetometers cannot be used for fast, continuous readings. One advantage to proton precession magnetometers is that they are not so susceptible to drift compared to a fluxgate magnetometer, and can produce readings of the total field accurate to within 0.1-1 nT.

2.8.3 Alkali vapour magnetometer

The alkali vapour magnetometer works on the principal of optical pumping. The alkali vapour is typically caesium, which is preferred due to the single electron in the outer shell of the atoms. Optical pumping works by illuminating the caesium-vapour cell by a light whose filament also consists of caesium. The emitted light is at the correct wavelength for incident photons to be absorbed by the vapour, and the outer electrons with an initial energy level called *A1*, become excited to higher energy levels, *B*. Electrons with the energy level *B* are unstable and thus decay to a lower energy level, which can be either the same as before, *A1* or a different, higher level called *A2*. This process is repeated and eventually the number of electrons with energy *A2* becomes overpopulated at the expense of electrons with energy *A1*. This causes the cell to stop absorbing light and it turns from opaque to transparent.

The difference between the energy levels *A1* and *A2* is proportional to the strength of the ambient magnetic field. The total magnetic field strength is obtained through a process of depolarisation, where the frequency of high frequency radio waves that trigger the transition of the elections back to *A1* is determined.

Alkali vapour magnetometers cannot read continuously, yet they are capable to repeating this procedure rapidly, achieving sample rates of up to 20 readings per second and sensitivity of 0.01 to 0.05 nT typical with modern equipment.

2.8.4 SQUID magnetometer

A recent development in magnetometers has seen the use of superconducting quantum interference device (SQUID) magnetometers for use in the field. Previously this technology has been used to make high precision measurements within laboratory conditions. They work on the principle that electrical conductivity of some metals increases as temperature falls, and at very low temperatures the metal becomes superconducting, with no resistance to electric current. When a magnetic flux is passed through the plane of the superconducting loop, a current is generated that opposes the applied flux. By measuring this current, the vector component of the Earth's field aligned with the sensor's axis can be measured.

The downside to using SQUIDs in the field is the need to keep the metal in a superconducting state. Two types of magnetometer have been developed. Low-temperature SQUIDs are typically made with niobium and the SQUID is required to be kept at liquid helium temperatures of 4°K, (-269°C). More recently high-temperature SQUIDs have been developed that contain complex compounds required to be kept at liquid nitrogen temperatures (77°K, -196°C) in the field. The advantage of SQUID magnetometers however, is the extremely high sample rate that can be obtained, capable of several thousand measurements per second, and the accuracy of around 0.00001 nT that can be achieved.

2.9 Data collection

The survey design of archaeo-magnetic surveys can differ from other ground based magnetometry surveys, due to particular characteristics of the target bodies. As a guide for aero-magnetic surveys, Reid (1980) states that the sample spacing should be less than the depth of the target. As archaeological targets can be buried in the very near subsurface this leads to a requirement for very small sample intervals. For detailed surveys an in-line sample interval no greater than

0.25 m is recommended by English Heritage (2008). Rapid data collection with small sampling intervals is becoming commonplace through the use of continually recording systems. Fluxgate systems manufactured by Geoscan, Bartington and Foerster can all be used to record repetitively while the operator walks along the survey line. The data is recorded automatically in a data logger, so the operator is not required to stop at each recording station. Recently towed systems have been developed that consist of multiple magnetometers mounted on a platform and towed at distance behind a vehicle. Such systems have been developed using caesium vapour magnetometers (Hill *et al.*, 2004) and fluxgate magnetometers (Dabas, 2010). The data shown in this thesis have been collected using English Heritage's magnetometer cart (Linford *et al.*, 2007), where four caesium vapour magnetometers are placed 0.5 m apart, normal to the survey lines, and recording is conducted at 10 Hz while the cart is pulled along the survey line. Each pull of the cart results in four lines of data being collected with an approximate spatial resolution of 0.5×0.125 m. The cart can be seen in Figure 2.4.

Sources of noise in archaeo-magnetic surveys will be discussed in Section 2.10, however a typical method applied to ground based, and aero-magnetic surveys in order to improve the signal to noise ratio of the data is to place the sensor at a height above the surface, in order to attenuate the signal of very near surface features (Hansen *et al.*, 2005). This is not an appropriate methodology with archaeological surveys, as the targets themselves are often of very small amplitude, and shallowly buried, therefore recording as close to the ground surface as possible is required to get the strongest response.

Data have historically been positioned by placing a survey grid relative to permanent features in the area, using lines, tapes or spray paint to indicate the positioning. Increasingly, Differential Global Positioning System (GPS) is used for locating the edges of survey grids, and more recently portable real-time systems are being used for the positioning of individual survey points. These systems do not require the operator to follow a fixed grid system, however this is generally recommended, due to any directional sensitivity that may exist with the sensors. The direction of data collection is often determined by any physical barriers present at the site,

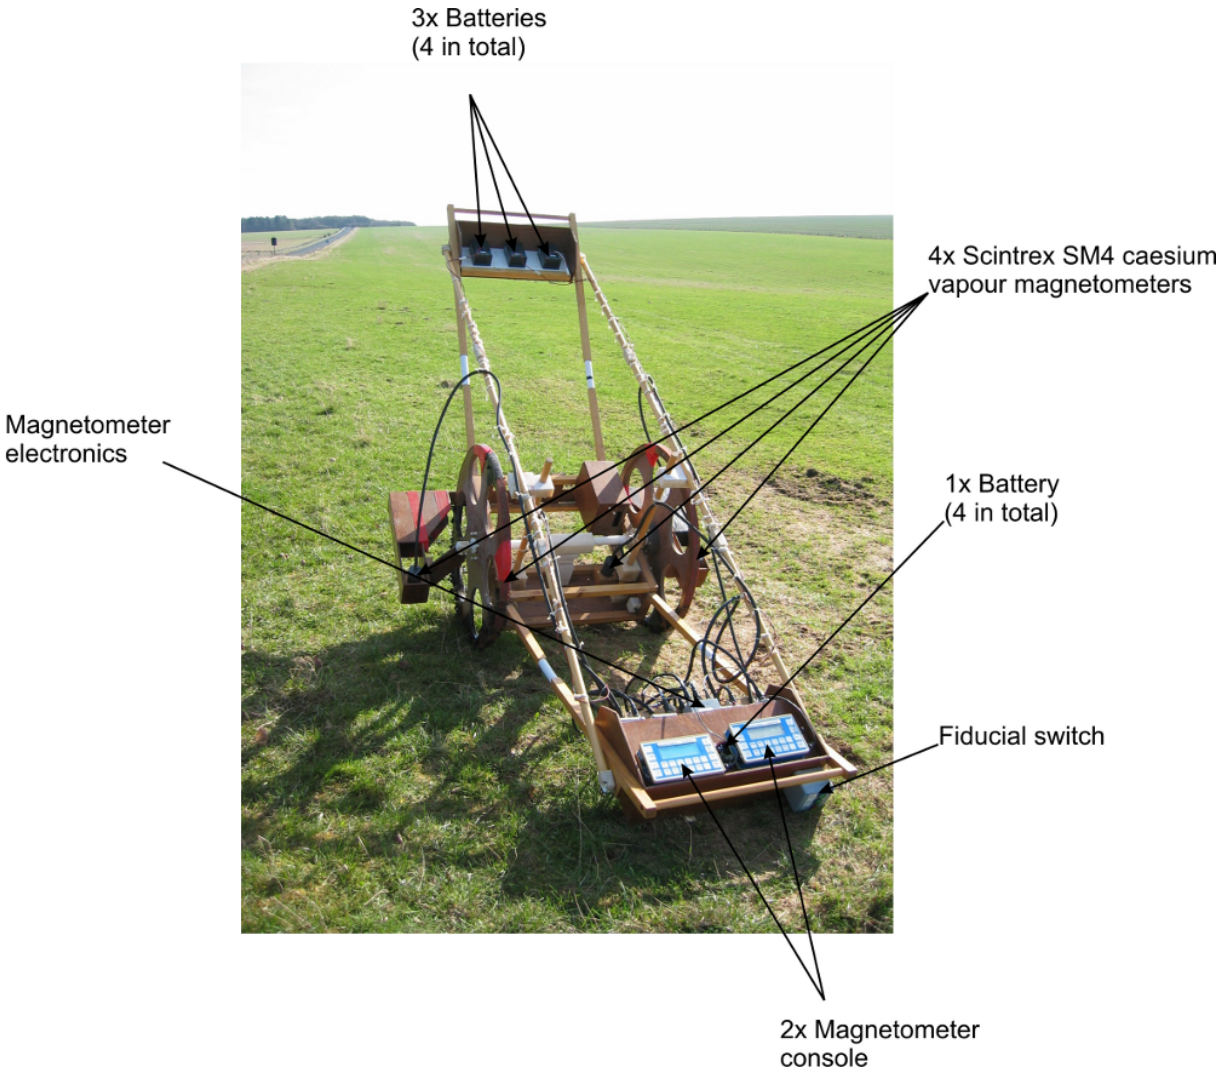

Figure 2.4: Magnetometer cart developed by English Heritage. The cart consists of four caesium vapour magnetometers mounted 0.5 m apart. Recording is conducted at 10 Hz as the survey is progressing resulting in a spatial resolution of $\sim 0.5 \times 0.125$ m.

however generally a north-south direction is preferred in order to collect the highest resolution across the asymmetric shape of anomalies in an inclined magnetic field.

Gradiometry surveys are commonplace in archaeological prospection. The often stated advantages with gradiometry surveys, over surveys that measure the total field is that due to the rapid falloff rates of magnetic anomalies ($1/r^3$ for a dipole source), magnetic gradiometers set up to measure the vertical gradient of the field register only anomalies in the near surface. An added advantage to this is that there is not a requirement for a base station magnetometer, as any diurnal variations will affect each sensor equally at the same time. Fluxgate magnetometers are always gradiometers. Proton precession and alkali vapour magnetometers, on the other hand, have the advantage of recording the intensity of the total field, but not its direction. Where two or more sensors are available it is possible to connect the sensors into a gradiometry configuration, however it has been demonstrated by Tabbagh (2003) that equivalent datasets can be generated by upward continuing the total field data and calculating the difference between the original and upward-continued data. Therefore Tabbagh (2003) suggests, where several sensors are available it is often preferable to use them to cover the survey area quicker by using them in a horizontal configuration, and to calculate the vertical gradient numerically. A similar observation is made by Linford *et al.* (2007) where a base station magnetometer is sacrificed in order to utilise the sensor to increase the coverage of the survey area. Corrections for the diurnal variations are subsequently carried out using median low-pass filters both parallel and perpendicular to the traverse direction.

2.10 Sources of noise

Noise in archaeo-magnetic datasets can be categorised into two distinct groups. Firstly, there is instrumental noise, and secondly there are unwanted, yet measured, magnetic variations.

Instrumental noise can be produced by the electronic circuitry of the instrument, and drift of the instrument due to its sensitivity to temperature. A study in Linford *et al.* (2007) shows fluxgate instruments produce higher instrument noise than caesium vapour magnetometers, however this type of noise is generally a lot smaller than the other factors that cause undesired magnetic readings that may mask some of the archaeological responses. Further instrument noise, such as a zero-offset error which is likely to be a constant value, and calibration errors can be reduced by ensuring the instrument is correctly set up.

Sometimes cultural signals from the surrounding area contribute to give an accurate reading of the magnetic field at the observation point, however are often unwanted in the interpretation of the dataset, and may also be considered noise. These measurements may be random or coherent. Random noise will be caused by a moving object in, or nearby the survey area which cause time-dependant variations in the signal. These are typically caused by passing traffic, people or electric fences (Aspinall *et al.*, 2008). Power lines and cathodically protected pipelines will also cause time-dependant variations in the reading that cannot be repeated, and therefore make interpretation of anomalies of interest difficult (Gharibi & Pedersen, 2000).

Other cultural noise is repeatable, even if undesired. These generally relate to fixed objects such as buildings, stationary vehicles, fences, underground pipes and small metallic debris that may be scattered across the site such as coins, nuts, bolts and other pieces of machinery.

Scollar *et al.* (1990) provide a good discussion about how the geology, soil conditions and atmospheric effects can also be classed as noise in an archaeo-magnetic survey. Subsurface rocks may produce anomalies in the magnetic data, particularly those formed from igneous materials such as dykes or sills, or the infilling of palaeochannels. Soils deriving from igneous rocks are likely to contain rapid variations in susceptibility dependant upon the minerals present, which will give the whole area a background noise level. Variations in micro-topography and infill of undulations caused by land usage for farming are also likely to produce anomalies that will be superimposed onto the archaeological response. Diurnal variations as discussed above

are required to be removed from the data either through use of a gradiometer, base station or advanced processing techniques.

The operator of the magnetic survey is also liable to introduce unwanted signals into the dataset should care not be taken while conducting the survey. It is important that the operator refrains from carrying or wearing ferrous materials, and that care is taken to ensure instrument alignment, positioning and sensor height are kept consistent for all observations. Varying sensor height with walking stride is particularly a problem with continuously recording instruments, as is the proximity to the towing vehicle in towed surveys.

2.11 Processing steps

Typical processing techniques applied to archaeo-magnetic data are discussed in detail by Aspinall *et al.* (2008) and English Heritage (2008), and a summary is presented here. Typically the processing steps can be divided into two separate types. The first steps relate to correcting the data for artefacts that become introduced due to survey procedures, and the second stage is used to enhance the corrected signal for interpretation purposes.

2.11.1 Correcting data

Zero-mean gridding

Instrument drift is particularly a problem associated with surveys conducted with a fluxgate magnetometer, and it is often advised that the operator returns to a certain point to repeat the reading frequently. This allows a linear interpolation of the difference to be applied to all the data collected between returns to the zero point. Instrument drift in unprocessed datasets

is often illustrated by a gradual variation as the survey grid progresses, often resulting in an artificial slope across the area.

Where the survey has been conducted by surveying several, separate smaller grids of data, which are then combined into the final dataset, offsets are often are located at the edges of the smaller grids which cause a 'patched' look in the final dataset. This can be caused by using different zero points, or by surveying an area with differently calibrated instruments. Typically a shift will be applied to each grid of data by ensuring that the resultant mean of the grid is equal to zero. An alternative to this is to analyse the edges, where data on adjacent grids are compared and the required correction made (Haigh, 1992).

De-striping

Mis-alignment can also exist between individual traverses. This is common should the survey be conducted in a zig-zag orientation where directional sensitivity affects the sensor, or if multiple sensors with different baselines have been used. In a similar way to the matching of grids, each traverse is shifted so that the mean of each survey line is equal to zero. It may be necessary to remove outliers while calculating the mean to avoid causing offsets for lines that have particularly high or low amplitude anomalies along them. A danger of this processing step is that it can potentially remove real features that exist in the direction of the traverse, such as pipelines or ditches. Should these features not be desired in the dataset, Gaffney & Gater (2003) suggest purposely aligning the survey grid along these features so that this process can be used to remove them from the final dataset.

De-staggering

The final process of correcting the data is often a de-staggering technique, which correct for displacement errors along each survey line. Often when a survey is conducted in a continuously

recording data collection mode, and a zig-zag survey line pattern is performed, the topography will cause the operator to walk at a different speed along adjacent survey lines. This is highlighted in the unprocessed data where linear anomalies crossing the traverse appear to have a staggered appearance. When this is the case, cross correlation can be used to align adjacent traces (Ciminale & Loddo, 2001).

2.11.2 Enhancing data

De-spiking

After the corrections have been applied, it is common for the data to be enhanced to aid interpretation. Noise associated with small ferrous material in the immediate subsurface is occasionally the target for the survey, for example, over battlefield sites. However, often these anomalies just mask the signal of interest. De-spiking algorithms are presented by Ciminale & Loddo (2001) and Scollar *et al.* (1990), and generally involve identifying values which differ from the mean or median value by more than a specified amount, being removed and replaced with a value closer to the average.

Bandpass filtering

The Fourier transformation can be used to convert magnetic field data, from a function in the space domain, into the wavenumber domain (often referred to as the Fourier domain), or vice versa. Many transformations of magnetic data are typically conducted in the Fourier domain of the signal. The Fourier transform is defined as (Blakely, 1995):

$$F(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-i(k_x + k_y)} dx dy \tag{2.10}$$

and the inverse Fourier transform is:

$$f(x,y) = \frac{1}{4\pi^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F(k_x,k_y)e^{i(k_x+k_y)}dk_x dk_y \tag{2.11}$$

where *f(x,y)* are the data in the space domain, *F(k_x, k_y)* is the function in the Fourier domain, and k_x and k_y are inversely related to wavelength, by $k_x = \frac{2\pi}{\lambda_x}$ and $k_y = \frac{2\pi}{\lambda_y}$. The Fourier transform is a complex function and can be separated into two parts. The amplitude spectrum is defined as:

$$|F(k)| = \sqrt{ReF(k)^2 + ImF(k)^2} \tag{2.12}$$

and the phase spectrum:

$$\Theta(k) = atan\left(\frac{ImF(k)}{ReF(k)}\right) \tag{2.13}$$

The amplitude function can be squared to give the energy-density spectrum $|F(k)|^2$, which from here on in is used to represent the Fourier domain in this thesis. The Fourier representation of the signal is often used to apply bandpass filters to the dataset. High-pass filters can be used to remove long wavelength regional trends that occur from high susceptibility geological sources, fences and other large items, from the data. This allows short wavelength anomalies likely to be produced from the shallow subsurface to be better visualised. Alternatively, low-pass filters can be used to smooth the data and remove the often high amplitude, short-wavelengths that may be masking weaker features that extend across the area. This approach is often employed to remove near-surface, cultural noise from the signal.

2.12 Displaying data

Interpretation of archaeo-magnetic surveys is generally performed by visual analysis of a graphical display of the data. Numerous methods for presenting the data have been devised over time, and the most common ones are presented in Figure 2.5, for a synthetic anomaly.

The most obvious method for visualising the data is to plot the recorded values against distance for each profile collected. When a grid of data is collected, trace plots can be created by offsetting the profile vertically for each progressive survey line. An example can be seen in Figure 2.5a. This technique derives from the late 1960s when the need to easily visualise data in the field was becoming a reality through the use of portable X-Y plotters (Scollar *et al.*, 1986). Trace plots have several advantages as they show the actual profiles of the collected data, rather than a plan view. Unfortunately interpretation becomes difficult where many survey lines are presented together, and small deviations can be difficult to identify in the presence of spikes or large amplitude anomalies.

More recently, trace plots developed into three-dimensional plotting techniques, which display the data as a wireframe mesh, or surface over a topography related to the magnitude of the value observed in the field. An example of a 3D surface plot is shown in Figure 2.5b.

Contour plots are created by defining certain threshold levels, and then linking data of equal value together across the survey area. An example is presented in Figure 2.5c. Where contour plots are coloured, everything between each threshold level is given the same colour. This inevitably leads to some level of smoothing of the data dependant upon the contour frequency chosen by the user. Contour plots can be ideal for delineating large amplitude anomalies, however they can lead to smaller amplitude anomalies being missed during the interpretation.

Dot-density plots became the standard way for plotting grids of magnetic data in the 1960's (see Scollar *et al.* (1986) and references within). The method involves dividing the survey area

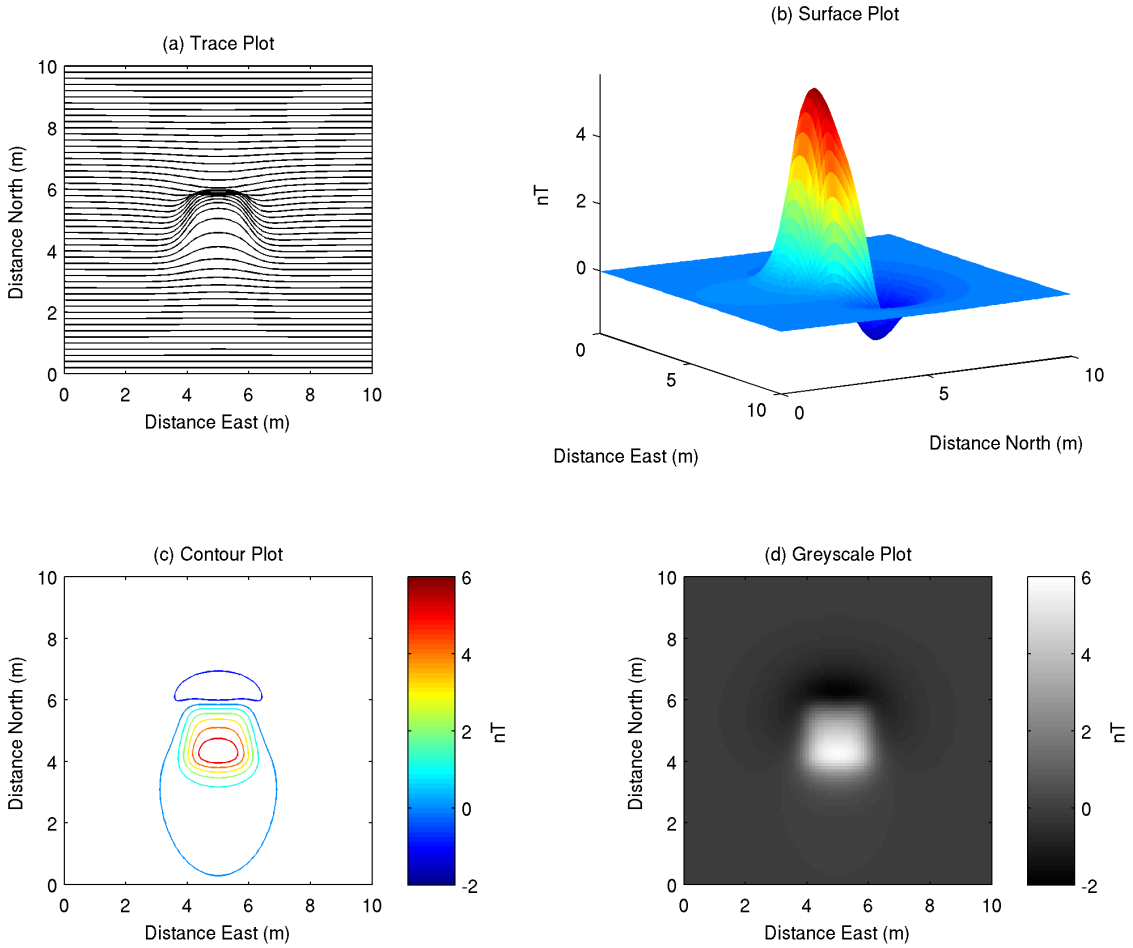

Figure 2.5: Different methods typically applied for the visualisation of archaeo-magnetic datasets. (a) Trace plot, showing alternate traces. (b) 3-D surface plot. (c) Contour plot. (d) Greyscale plot.

into squares with each square representing a data point. The square is subsequently filled with a number of black dots proportional to the amplitude of the magnetic field at that point. As printer techniques advanced, this type of plot has developed into the greyscale plot (Figure 2.5d), where instead of filling each square with a number of dots, a shade of grey is chosen to represent the magnitude of the magnetic field at that point. This technique has the advantage over contour plots in that the value of each data point can be identified.

2.13 Advanced processing/interpretation

Magnetic surveys are becoming increasingly common on archaeological sites due to the amount of data that can be collected rapidly in a non-invasive manner. Due to their relatively low cost compared to excavation they commonly provide the only dataset that covers an entire archaeological site, which can then be used to target other surveys and excavations to areas of interest. Typically interpretation is done visually using optimised images of the raw data, which when dealing with large datasets can be time-consuming and subjective. There are, however, some more advanced processing and interpretation steps that can be implemented to improve the clarity of the dataset, or produce a transformation that will aid interpretation.

2.13.1 Reduction to the pole

Reduction to the pole (RTP) is a commonly implemented transformation, used to simplify magnetic data collected at different latitudes. A magnetic anomaly will not appear directly above the causative body unless the magnetisation and ambient field are both directed vertically. In a non-vertical field, the magnetic response will be skewed, rather than the symmetrical response that would be obtained if data were collected in a vertical field. An example of this is shown in Figure 2.6a. Reduction to the pole removes the asymmetry caused by inclination by performing a phase shift to make the data appear as if they were collected at the north magnetic pole (Fig-

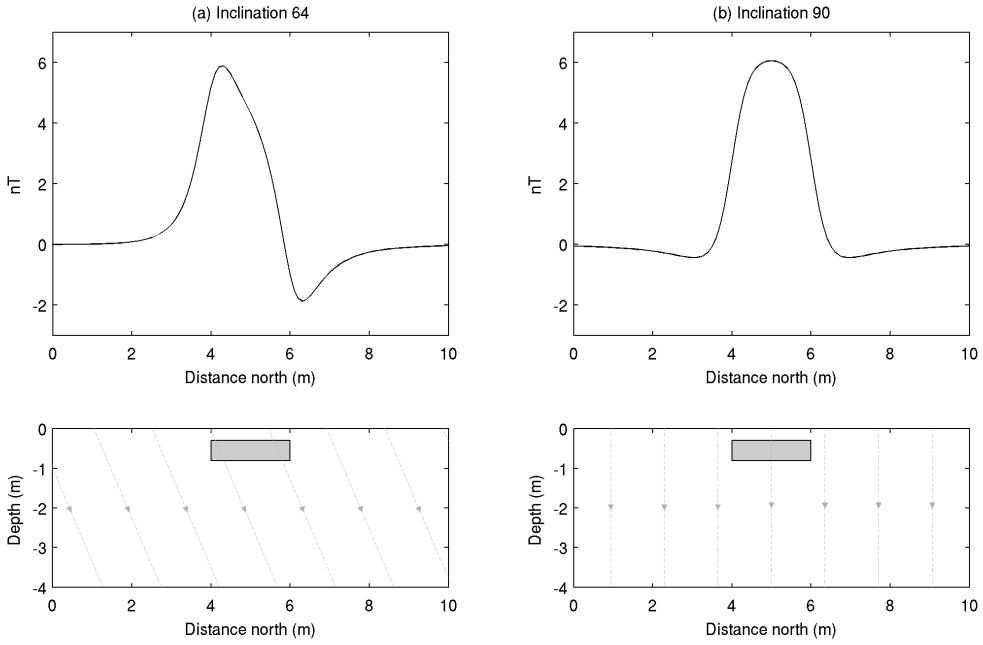

Figure 2.6: Example of reduction to the pole. (a) The response to a prism recorded in an inclined field of 64°. (b) The same anomaly reduced to the pole. Note the symmetrical response, which eases interpretation.

ure 2.6b). This has the effect of shifting anomalies laterally to be centred over their sources, and simplifies interpretation.

2.13.2 Derivative-based methods

Various derivative based methods have been developed recently to aid the interpretation of magnetic data. They work on the fundamental basis that the magnetic field is attenuated rapidly as the distance between the source and observation is increased. By analysing the gradients of the magnetic field, inferences can be made regarding the location of edges, and depth to subsurface magnetic bodies. Their use is gaining popularity in aero-magnetic regional and mineral exploration surveys. The horizontal gradient magnitude (Cordell & Grauch, 1985), analytic signal (Nabighian, 1972), theta map (Wijns *et al.*, 2005) and tilt angle (Miller & Singh, 1994) are all shown to derive information regarding the horizontal location of anomalies, and Source Parameter Imaging (Thurston & Smith, 1997) and Euler deconvolution (Reid *et al.*, 1990) have

been demonstrated to produce depth estimates to certain types of subsurface structures. Despite this they are rarely used in archaeological survey interpretation. This is likely to be due to the particular challenges posed due to the low amplitude and high-wavenumber content of archaeomagnetic surveys, where features of interest are often only slightly above the noise level. A thorough analysis of these methods is made in Chapters 3 and 4.

2.13.3 Inversion of magnetic data

The methods discussed above can be used to obtain certain parameters relating to the magnetic source bodies buried in the subsurface. This may be in the form of a lateral position, depth estimate, or an indication of source type. However, these methods do not produce a 3-dimensional model of the subsurface, which is the ideal result for an archaeologist attempting to understand the history of the site, and a geophysicist looking to compare the physical properties to other data-types.

Calculating the magnetic field that would be observed by an anomalous susceptibility distribution in the subsurface is relatively easily accomplished based on the theory of potential fields. The inverse process of this, calculating the subsurface distribution of susceptibility from the total field data observed in the field is a much more complicated procedure, and has led to formation of several different methods.

Limitations in the computer power required for inversion, and the problems associated with identifying the 'most-likely' model from the many solutions that will produce a good fit to the observations led to some novel approaches. The simplest approaches to this problem involved obtaining a small number of properties to a largely pre-defined problem. Powell (1967) produced a model of a dyke, and a step, or fault model. The optimisation problem was then solved to adjust the model parameters to alter the pre-defined model so that the response produced an acceptable fit to the observations. A similar approach can be seen with archaeological datasets for pit-houses (Herwanger *et al.*, 2000). Here, the susceptibility contrast between the

infill and the surrounding material was measured in a laboratory to produce a constraint on the final model. The pit houses were represented by a series of vertical prisms representing the measured susceptibility. The inversion was then used to adjust the length of these prisms in order to produce a model response that was a good fit to the observations. These techniques use the inversion to adjust a pre-defined model structure, and therefore the input structure and susceptibility contrasts have a large effect on the final model.

A neural network approach can be seen applied to archaeological data by Bescoby *et al.* (2006) for investigating the depths to limestone walls, and for a range of ditch-like scenarios by Sheen (1998). Neural networks work by using training datasets to teach the algorithm the response to synthetic, or known scenarios, and then applying this to field data. Similarities between the real data and the training datasets are used to make inferences to the source body parameters. These approaches have been seen to be quite successful, however the quality of the output will be determined by the training data that is used, and its similarity to the types of structures present in the real data.

Where the target is unknown, the subsurface is required to be represented by a less prejudiced structure. Eder-Hinterleitner *et al.* (1995) did this by forming a subsurface consisting of a 3-dimensional grid of dipoles positioned 0.5 m apart. The inversion was used to allocate each dipole one of the four chosen possible dipole moments, in order to produce a model of ring ditches in Lower Austria. The results produced very encouraging results, however the final model will still be heavily influenced by the values allowed for the dipole moments.

Dittmer & Szymanski (1995) used a slightly different approach to Eder-Hinterleitner *et al.* (1995), and instead of representing the subsurface with a grid of dipoles, used a mesh of cells. Each cell was allowed to account for one value of susceptibility, without a finite number of individual values for susceptibility being defined. The technique was limited to 2-dimensional models, due to the computational limitations. Despite this, Dittmer & Szymanski (1995) demon-

strated how this approach could be used to successfully model highly magnetised synthetic and real-case dyke features in the UK.

Generalised inversion schemes involving the division of the subsurface into a 3-dimensional mesh were developed initially for mineral exploration purposes (Li & Oldenburg, 1996; Pilkington, 1997). These approaches have the advantage that they can still be implemented in situations where no *a priori* information is available, and use constraints such as a requirement for a smoothly varying susceptibility distribution, and non-negative susceptibilities to identify a final model from the inversion. The principles in the Li & Oldenburg (1996) paper, were later adapted by the University of British Columbia Geophysical Inversion Facility (UBC-GIF) to produce a commercially available program called MAG3D (UBC-GIF, 2005). Li & Oldenburg (1996) employ this method to model copper-gold porphyry deposits in central British Columbia, and Williams (2008) shows how the addition of geological constraints can be used to refine the resultant model from the Agnew-Wiluna greenstone belt and Perseverance nickel deposit in Australia. MAG3D has also been implemented in archaeological situations by Piro *et al.* (2007). Here, chamber tombs at the Sabine Necropolis of Colle del Forno in Rome, Italy, were modelled using the vertical gradient of the magnetic field. The success of the method using the vertical gradient was noted with caution by the authors, and the method should not strictly be applicable without adapting the mathematics behind the forward modelling routine of the inversion, to account for the vertical component rather than the total field. Piro *et al.* (2007) also noted that modelling both positive and negative susceptibility contrasts would be a problem, as the code placed a constraint on the recovered susceptibilities to only allow modelling of positive susceptibility contrasts with the surrounding material. A development to the code, to relax this constraint has recently been added (UBC-GIF, 2005). Another example of an archaeological implementation can be seen in Sheriff *et al.* (2010), for investigation of late 19th century, early 20th century building foundations in Yellowstone, USA. The nature of these anomalies result in recordings of hundreds of nT, and related large susceptibility contrasts to the surrounding area. The suitability of inverting archaeo-magnetic datasets using MAG3D will

be investigated in Chapter 5, following analysis of more standard processing and interpretation techniques.

Chapter 3

Identifying lateral extents of archaeo-magnetic targets

3.1 Introduction

Several techniques are becoming increasingly popular for quickly deriving the surface locations of the edges of subsurface bodies. These can help to rapidly identify the locations and shapes of anomaly sources of interest from large spatial surveys. While implementation is becoming increasingly common in mineral exploration surveys, they remain rarely applied to archaeo-magnetic surveys. Recent examples of use of the horizontal gradient magnitude can be seen by Sheriff *et al.* (2010), locating foundations over a historic town in the USA, whilst theta map and tilt angle methods have been applied by Bueyueksarac *et al.* (2008) to synthetic and real data from a Bronze Age burial ground in Turkey. Although the data are collected and processed in fundamentally the same way in mineral and archaeological surveys, the spatial size, contrast in magnetisation, and signal-to-noise ratio of archaeological surveys can provide problems when implementing these techniques. The use of the analytic signal, involved in calculating the theta map, has become more typical with Tabbagh *et al.* (1997) and Godio & Piro (2005) providing good examples, however the problem caused by amplifying noise in derivative results, particularly second-order derivatives is demonstrated by Milea *et al.* (2010).

Here, these derivative-based methods will be examined for their suitability to be applied to archaeo-magnetic datasets, and an investigation into the alterations required to the methodologies for implementation with archaeo-magnetic datasets will be made.

3.2 Synthetic data

For this study the synthetic response to a compound anomalous body was generated. The anomaly shown in Figure 3.1a consists of two connected parts, in order to test the sensitivity to features of differing depths and magnetic response. The top of the shallower part of the body is positioned 0.3 m below ground level, while the deeper part is 0.8 m below ground level. The base of both parts is situated at 1.2 m below ground level. The anomaly has a susceptibility of 0.001 SI, and is situated within a magnetic field of 37.5 Am^{-1}, causing a magnetisation intensity of 0.0375 Am^{-1}. The total field response has been calculated at a height of 0.2 m above ground level, at a sampling density of 0.1×0.1 m using the method described by Rao & Babu (1993). To replicate a dataset reduced to the pole, the response has been generated with a magnetisation inclination of 90°. This represents a well-sampled, noise free, dataset which is typical of sampling resolutions used in archaeo-magnetic survey, and can be seen in Figure 3.1b.

To test the effects of noise on this study, a zero-mean symmetric Gaussian distribution was generated with a standard deviations of 0.1, 0.2 and 0.3 nT. A histogram of the distribution of noise for the 0.1 nT scenario is presented in Figure 3.1c, alongside the synthetic data with the noise added in Figure 3.1d. For the 0.2 nT and 0.3 nT scenarios, the noise and synthetic datasets are shown in Figures 3.1e and 3.1g, and Figures 3.1f and 3.1h respectively.

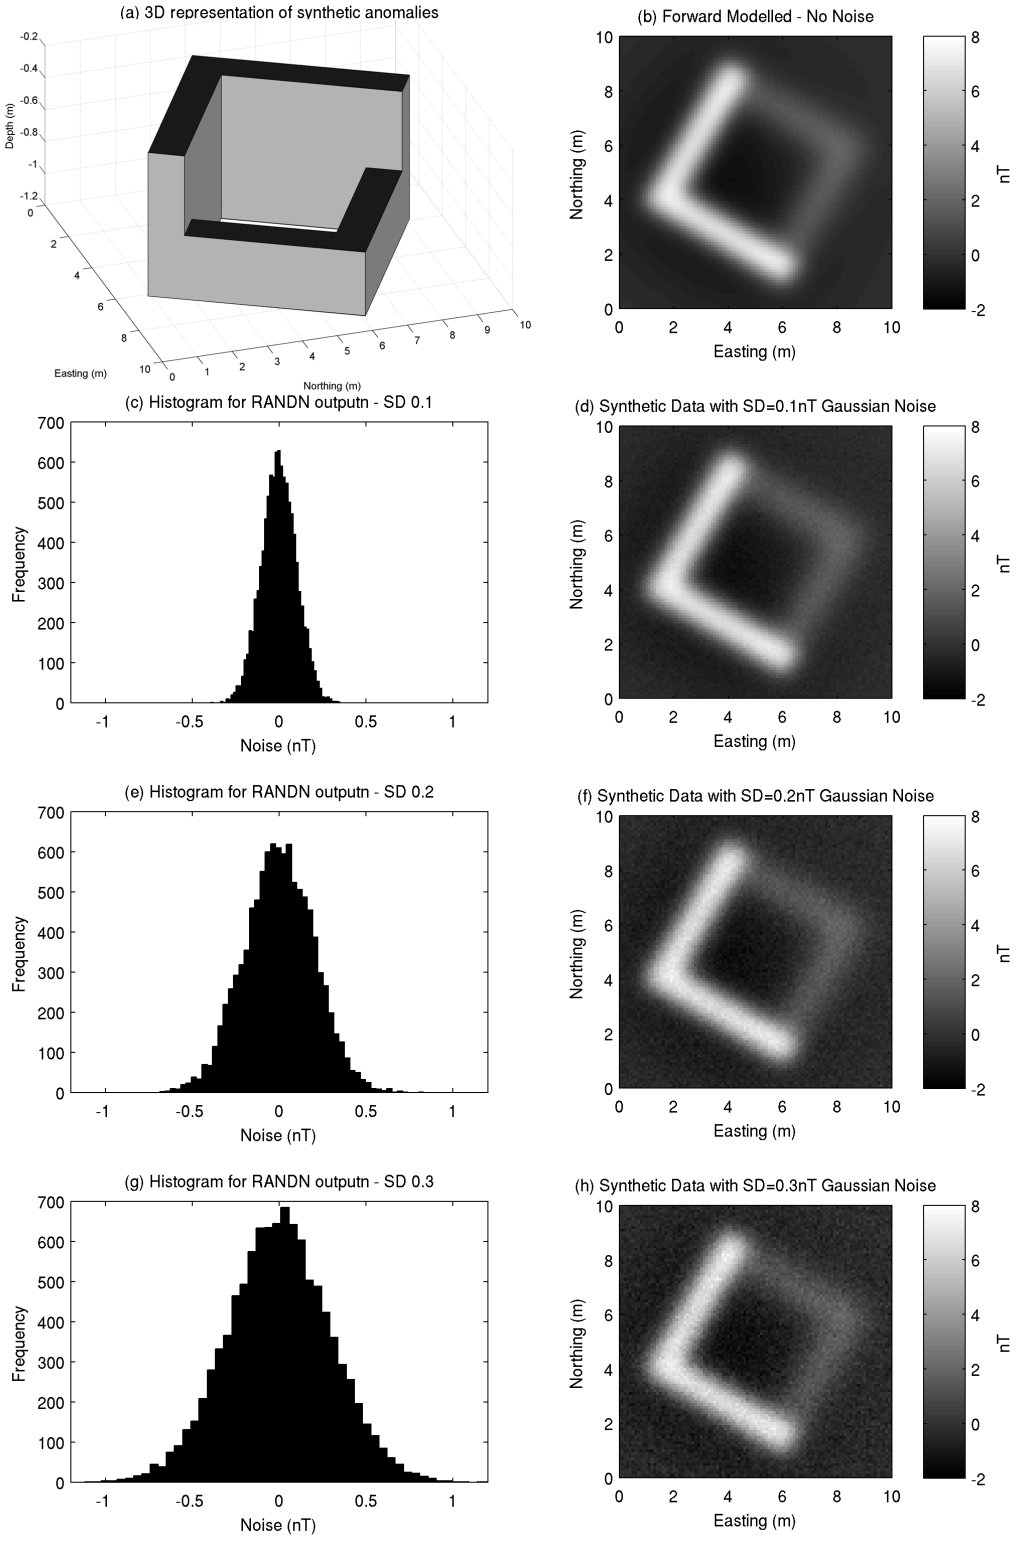

Figure 3.1: (a) 3D model of synthetic anomaly with magnetisation of 0.0375 Am^{-1}. (b) Total field data for no noise scenario. (c) 0.1 nT standard deviation, zero-mean Gaussian noise. (d) Total field data with SD=0.1 nT noise added. (e) 0.2 nT standard deviation, zero-mean Gaussian noise. (f) Total field data with SD=0.2 nT noise added. (g) 0.3 nT standard deviation, zero-mean Gaussian noise. (h) Total field data with SD=0.3 nT noise added.

3.3 Horizontal gradient magnitude

A procedure for outlining the edges of subsurface magnetic features using the horizontal gradient magnitude is described by Dole & Jordan (1978) and better defined by Cordell & Grauch (1985). Peaks in the horizontal gradient magnitude will overlie abrupt lateral changes in subsurface magnetisation, and will be directly above the edge of a magnetised body if the edge is a vertical contact and the data has been reduced to the pole. Where the body does not have a vertical contact, the horizontal gradient magnitude can be offset by an amount usually dependent on the depth below ground (Grauch & Cordell, 1987). In an archaeological context, these offsets are therefore generally quite small. An automated approach to identifying the peaks in the horizontal gradient magnitude is demonstrated by Blakely & Simpson (1986), however this approach is often limited by the signal to noise ratio found in archaeo-magnetic data.

The horizontal derivatives can be calculated in either the space or wavenumber domain. For ease of calculation, here they have been calculated in the space domain. The gradient in the *x*-direction is calculated using:

$$\frac{\partial T}{\partial x} = \frac{T_{(x+1,y)} - T_{(x-1,y)}}{2\Delta x} \tag{3.1}$$

and the *y*-direction using:

$$\frac{\partial T}{\partial y} = \frac{T_{(x,y+1)} - T_{(x,y-1)}}{2\Delta y} \tag{3.2}$$

The horizontal gradient magnitude is then calculated as:

$$\frac{\partial T}{\partial h} = \sqrt{\left(\frac{\partial T}{\partial x}\right)^2 + \left(\frac{\partial T}{\partial y}\right)^2} \tag{3.3}$$

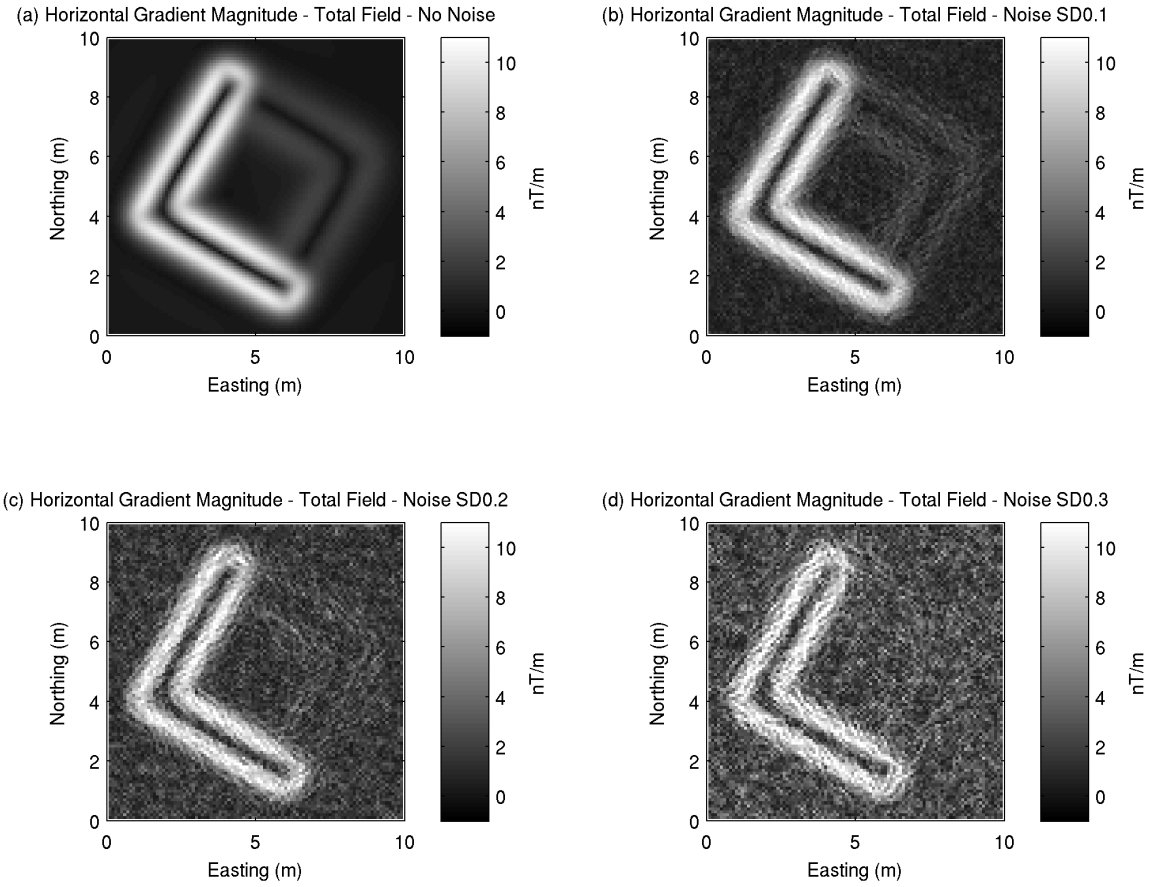

Figure 3.2: Horizontal gradient magnitude of the total field data shown in Figure 3.1. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

It can be observed in Figure 3.2a that the response to the lower part of the synthetic body has produced a weaker anomaly in the horizontal gradient magnitude ($\sim$3-4 nT/m) compared with the shallower body ($\sim$10 nT/m). This is as expected, due to a weaker response in the total field measurements and the longer-wavelength response that is typical of more deeply buried anomalies. This will make large amplitude, short wavelength anomalies easier to identify using the horizontal gradient magnitude method, which will typically show bias towards identification of shallower bodies.

Figure 3.2b shows the horizontal gradient magnitude of the synthetic model generated with zero-mean Gaussian noise with a standard deviation of 0.1 nT. Although the amplitude of the shallower anomaly means that the edges are still easily identified as maxima in the horizontal

gradient magnitude, the weaker response of the deeper body is becoming harder to identify. With increasing noise seen in Figures 3.2c and 3.2d, it becomes impossible to identify the deeper anomaly from the surrounding background noise.

3.3.1 Automated edge detection techniques

Various edge detection algorithms exist for image processing applications. They are used to identify discontinuities in digital images, which are typically changes in brightness in a greyscale image. They can therefore be applied to geophysical data to identify discontinuities in the amplitude of the measured property.

Here, an edge detection algorithm developed by Canny (1986) has been implemented. In common with the horizontal gradient magnitude, this method requires the data to be reduced to the pole. Canny (1986) identified that the optimum operator to be convolved with the signal to identify an edge, is the first derivative of a Gaussian function. Therefore, when this is convolved with the data the peaks in the result appear where edges are present.

In practice, the first step of the Canny edge detection is to create a smoothed image by convolving the input signal with a symmetric Gaussian function to reduce the effect of noise on the edge detection. The algorithm then detects zero crossings of the second horizontal derivative. These should locate the maxima and minima of the first horizontal derivative, and is therefore similar to the horizontal gradient technique.

Two thresholds are required by the edge detection algorithm. Where an amplitude is detected above the upper threshold, it is marked and declared an edge. The edge is then extended away from this point for as long as possible, until the amplitude of the edge falls below the lower threshold. This helps the algorithm detect weak edges which have comparable amplitude to noise.

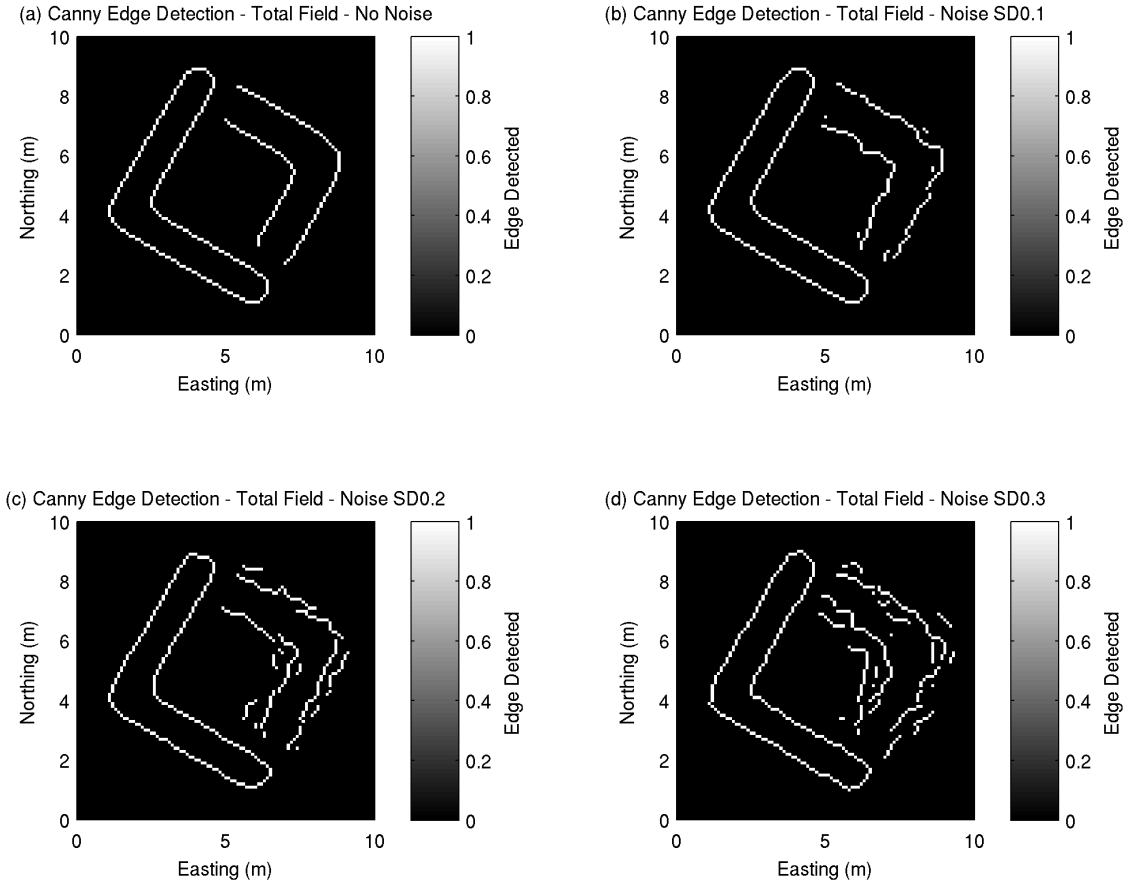

Figure 3.3: Canny edge detection of the total field data shown in Figure 3.1, using an upper threshold of 0.1563 and lower threshold of 0.0625. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

The results of the Canny edge detection routine to the synthetic anomaly are shown in Figure 3.3a for the scenario without noise, Figure 3.3b for the dataset where 0.1 nT of Gaussian noise has been applied and Figures 3.3c and 3.3d for the scenarios with 0.2 nT and 0.3 nT, respectively. An upper threshold of 0.1563 and lower threshold of 0.0625 have been applied. It can be seen that all datasets have recovered the shallower anomaly well. There is a small disconnect in the areas where the upper and lower bodies join. This is due to the overlapping gradients that exist in this area. As the noise level is increased the continuity of the deeper anomaly degrades, with an outline traceable by eye in the 0.1 and 0.2 nT scenarios (Figures 3.3b and 3.3c), yet without prior knowledge a trend cannot be identified for the deeper anomaly in the 0.3 nT scenario (Figures 3.3d), as the number of 'artificial' features produced by the presence of the noise increases.

3.4 Analytic signal amplitude

The analytic signal involves the combination of the horizontal gradient magnitude and the vertical gradient of the magnetic field. In this respect it is often referred to as the total gradient (See Equation 3.4).

$$|A| = \sqrt{\left(\frac{\partial T}{\partial x}\right)^2 + \left(\frac{\partial T}{\partial y}\right)^2 + \left(\frac{\partial T}{\partial z}\right)^2} \tag{3.4}$$

The analytic signal can also be calculated using higher order derivatives, where it is often referred to as the analytic signal of order (n).

The analytic signal method can be used as a edge detection method on its own. It was presented by Nabighian (1972) for two-dimensional profile data, and Roest *et al.* (1992) for three-dimensional gridded datasets. The amplitude of the analytic signal amplitude has maxima near the edges of magnetic sources, however is not considered to be as accurate as the horizontal

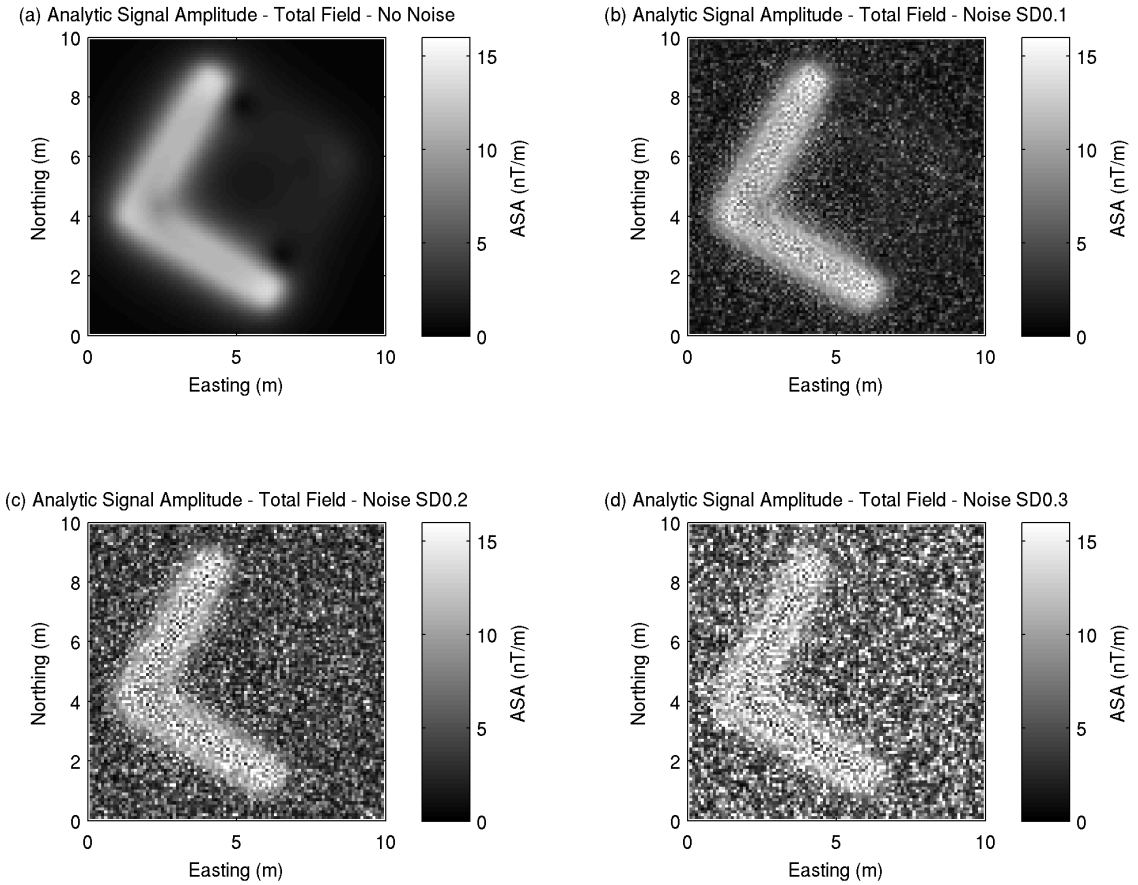

Figure 3.4: Analytic signal amplitude of the total field data shown in Figure 3.1. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

gradient magnitude of RTP data (Li, 2006). Despite this, the analytic signal amplitude transformation can be useful in areas where reduction to the pole becomes unstable, such as at low latitudes, or where significant remanent magnetisation is thought to be present.

Figure 3.4 shows the analytic signal for the synthetic scenario, with and without noise. Figure 3.4a shows a peak in amplitude over the shallow anomaly ($\sim$12-13 nT/m), with the highest peaks present over the edges of the body. The deeper anomaly is shown by much lower amplitude responses, around 2-4 nT/m. In the presence of noise (Figures 3.4b-d) the deeper anomaly quickly becomes impossible to identify as the responses are of similar amplitude to the background noise.

3.5 Theta map

The use of theta map for detecting edges in magnetic datasets is described by Wijns *et al.* (2005). Theta map combines the horizontal gradient magnitude and the analytic signal. It is defined as:

$$cos\,(\theta) = \frac{\sqrt{\left(\frac{\partial T}{\partial x}\right)^2 + \left(\frac{\partial T}{\partial y}\right)^2}}{|\mathbf{A}|} \tag{3.5}$$

where $|\mathbf{A}|$ is the analytic signal amplitude.

Theta map takes advantage of two properties of the magnetic derivatives. These are that the edges of a vertical contact will be situated not only below a peak in the horizontal gradient magnitude, but also at a zero crossing point in the first vertical derivative. Therefore, over an edge, the amplitude of the analytic signal will be largely dependent on the amplitude of the total horizontal derivative, and the value of cos(θ) will approach 1. Either side of this contact, the vertical derivative will be the significant component of the analytic signal, and the value of cos(θ) will approach 0. Thus a contact is defined by a maximum value of cos(θ) bracketed by two minimum values. Magnetically quiet areas will have a high cos(θ) value.

The advantage of the theta map method over the horizontal gradient magnitude is that it equalises the amplitude of the anomalies across the survey grid, so maxima that appear as different amplitudes in the horizontal gradient magnitude have approximately the same amplitude in the theta map. In this respect, it is often compared to an automatic gain control (AGC) filter. The advantage of this, is that subtle features in the horizontal gradient magnitude are amplified, however a disadvantage is that noise can also be amplified.

The theta map of the synthetic dataset is shown in Figure 3.5a for the zero-noise and Figures 3.5b-d, for the 0.1, 0.2 and 0.3 nT Gaussian noise examples respectively. Figure 3.5a shows the advantage of plotting the theta map, as the edges of both the shallow and deep anomalies are

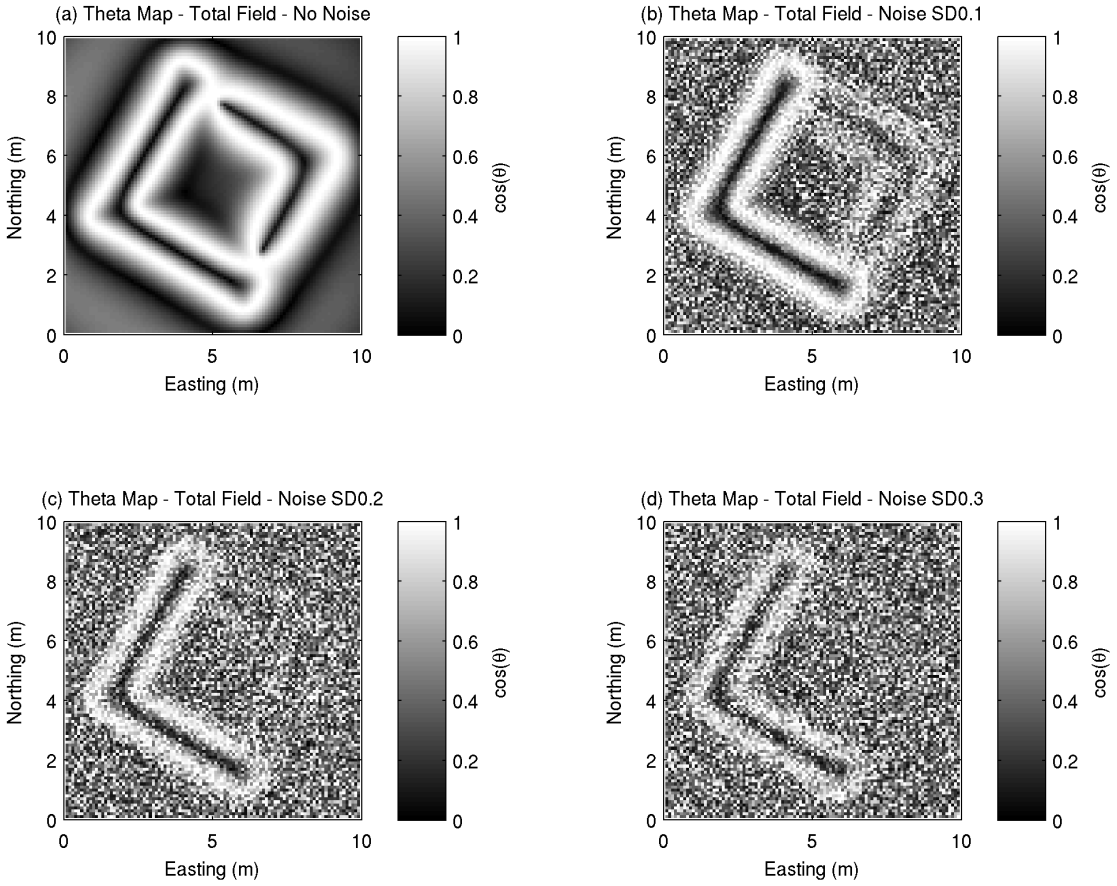

Figure 3.5: Theta map of the total field data shown in Figure 3.1. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

both detected by areas where the Theta Map is approaching 1. Over the centres of the anomalies the value of the theta map is close to zero. The theta map shows a termination of the deeper feature prior to the connection with the shallower feature demonstrating an interference effect where the closeness of the causative bodies cause the total field anomalies to overlap.

Figures 3.5b-d however, show the disadvantage of using the theta map. Although the introduction of noise into the data causes some relatively small isolated anomalies in the total field data, when the theta map is calculated the noise component can be boosted to values approaching 1, the same amplitude as the features of interest. This causes Figures 3.5b-d to be difficult to interpret, and although the high-low-high pattern across the shallower body is still evident despite the noise, the lower body is now impossible to identify in scenarios were the noise is greater than 0.1 nT.

3.6 Tilt angle

Methods using the tilt angle, and derivatives of the tilt angle have been developed to quickly identify both the horizontal location and the depth of buried anomalous bodies (Verduzco *et al.*, 2004; Salem *et al.*, 2007). The tilt angle involves the use of the horizontal and vertical gradients of total field magnetic data to produce a normalized derivative based on the ratio of the gradients. It is defined by Miller & Singh (1994) as:

$$\theta = \tan^{-1}\left[\frac{\left(\frac{\partial T}{\partial z}\right)}{\left(\frac{\partial T}{\partial h}\right)}\right] \tag{3.6}$$

Where θ is the tilt angle, and T is the magnitude of the total magnetic intensity of the anomaly. The first derivative of T in the horizontal direction is calculated as in Equation 3.3

Further work by Salem *et al.* (2007) showed that assuming the source structures have vertical contacts and the magnetic field is either vertical or reduced to the pole, the Tilt Angle can also be written as:

$$\theta = \tan^{-1}\left[\frac{h}{z_c}\right] \tag{3.7}$$

Where h is the horizontal distance and z_c is the depth to the contact. Salem *et al.* (2007) suggest that Equation 3.7 shows that the tilt angle can be used to derive the edges of the contact, as where $h=0$, $\theta=0$, and therefore the edges of vertical contacts will be present directly below the 0° contour of the tilt angle. Analysis of the depth estimates from tilt angle interpretation will be discussed in Section 4.2.

Results from the tilt angle (Equation 3.6) for the synthetic scenario can be seen in Figure 3.6 for increasing amounts of added noise. In Figure 3.6a the total field data has been normalised to values between -90° and 90°, with the shallow and deep anomalies successfully normalised to the same amplitudes. The zero contour shows a reasonable outline of the subsurface location of the anomalous body. Figures 3.6b-d show the tilt angle is highly susceptible to the influences of noise. Just as observed with the theta map, the noise is amplified to the same values as the signal. The continuity of the tilt angle response to the shallower body allows that to be identified, however identification of the deeper body is difficult, and tracing of the zero contour through the datasets with noise present is made impossible due to the short wavelength features of the noise causing multiple zero crossings.

3.6.1 Horizontal derivative of tilt angle

An alternative way to interpret the results from the tilt angle is to calculate the horizontal gradient magnitude of the tilt angle, given by Verduzco *et al.* (2004) as:

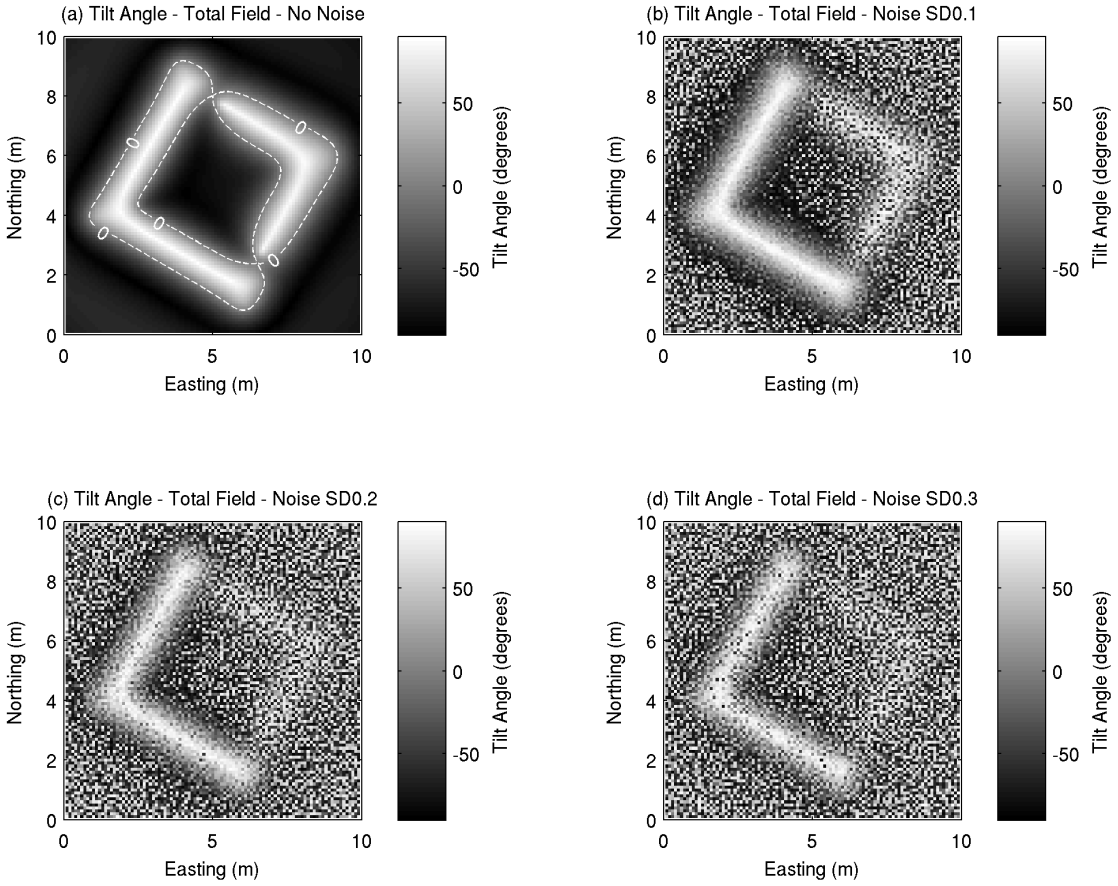

Figure 3.6: Tilt angle interpretation of the total field data shown in Figure 3.1. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

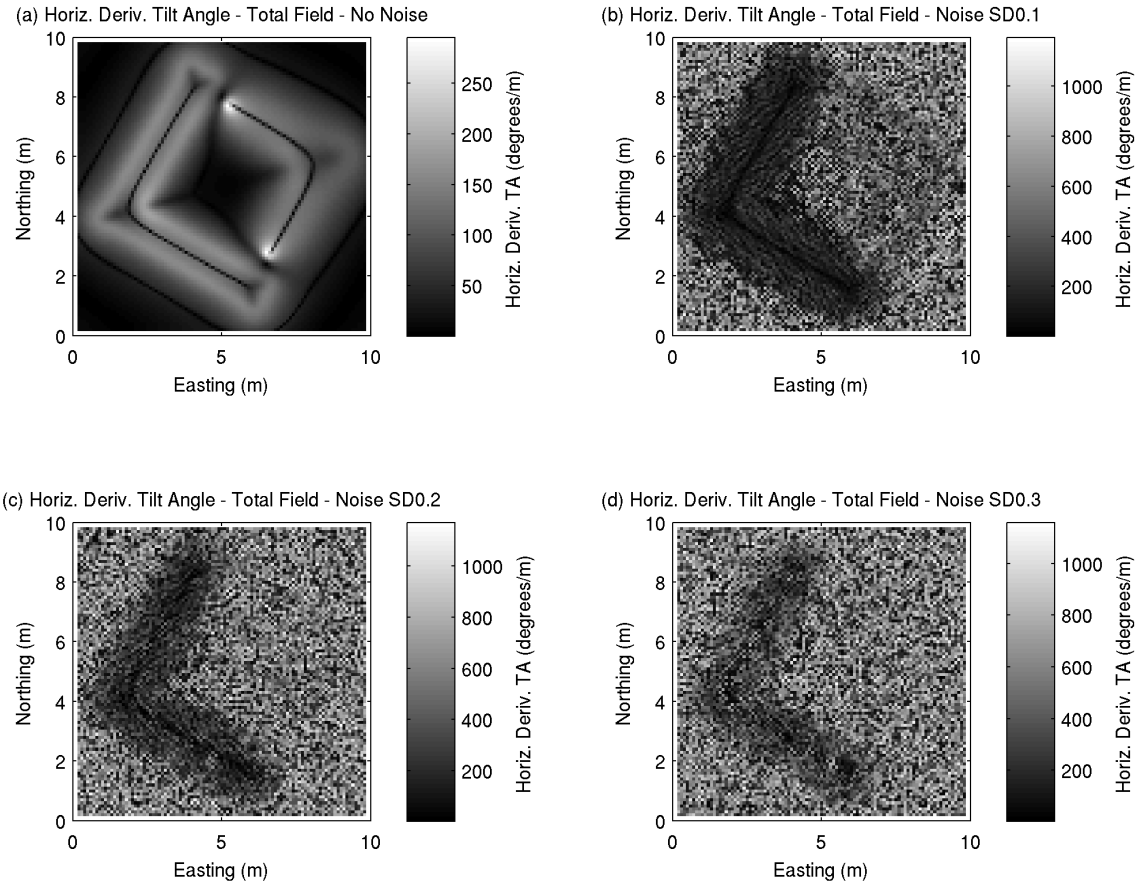

Figure 3.7: Horizontal gradient magnitude of the tilt angle interpretation of the total field data shown in Figure 3.1. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

$$\frac{\partial(\theta)}{\partial h} = \sqrt{\left(\frac{\partial(\theta)}{\partial x}\right)^2 + \left(\frac{\partial(\theta)}{\partial y}\right)^2} \tag{3.8}$$

This is similar to the local wavenumber used in methods by Thurston & Smith (1997). Subsurface contact features should be represented beneath maxima in the horizontal gradient magnitude of the tilt angle, in a way similar to the horizontal gradient magnitude of the total field, the advantage here being that the amplitude of the anomalies have been normalised prior to the gradient being analysed.

The horizontal gradient magnitude of the tilt angle is presented in Figure 3.7. Whereas the zero noise scenario shows a useful result showing peaks of $\sim$200°/m, the method is even more

susceptible to noise than the tilt angle itself, as it is a derivative of a derivative (Cooper & Cowen, 2006). If the tilt angle dataset is not interpretable due to noise, the horizontal derivative of the tilt angle will also be uninterpretable.

3.7 Power spectra analysis of derivative-based methods

Although the methods discussed above are successful in achieving their aims at making the noise-free synthetic data easier to interpret than the "measured" total field dataset, the introduction of noise causes a breakdown in the use of gradient-based methods, to the point where the total field data is easier to interpret by eye than the derivatives or combinations of derivatives. This is due to the accentuation of the high-wavenumber component of the data when derivatives are calculated.

Figure 3.8 shows the radially averaged power spectra for the synthetic dataset with the 0.2 nT standard deviation, zero-mean Gaussian noise added. The power spectrum for the total field anomaly data is represented by the solid line. Once a derivative of the total field data is calculated, high wavenumber components are accentuated relative to the low wavenumbers, as can be observed by the dashed lines in Figure 3.8. It is this effect that causes the noise to reach amplitudes which mask the signal of real anomalies as observed above, and makes gradient-based methods difficult to apply to archaeo-magnetic datasets. One way to suppress the high wavenumber component of a dataset is to apply a wavenumber filter, such as the pseudo-gravity transformation.

3.8 Pseudo-gravity Transformation

Assuming the boundaries of magnetic and gravitational sources are the same, and the intensity of magnetisation and density are proportional throughout, then magnetic potential and gravi-

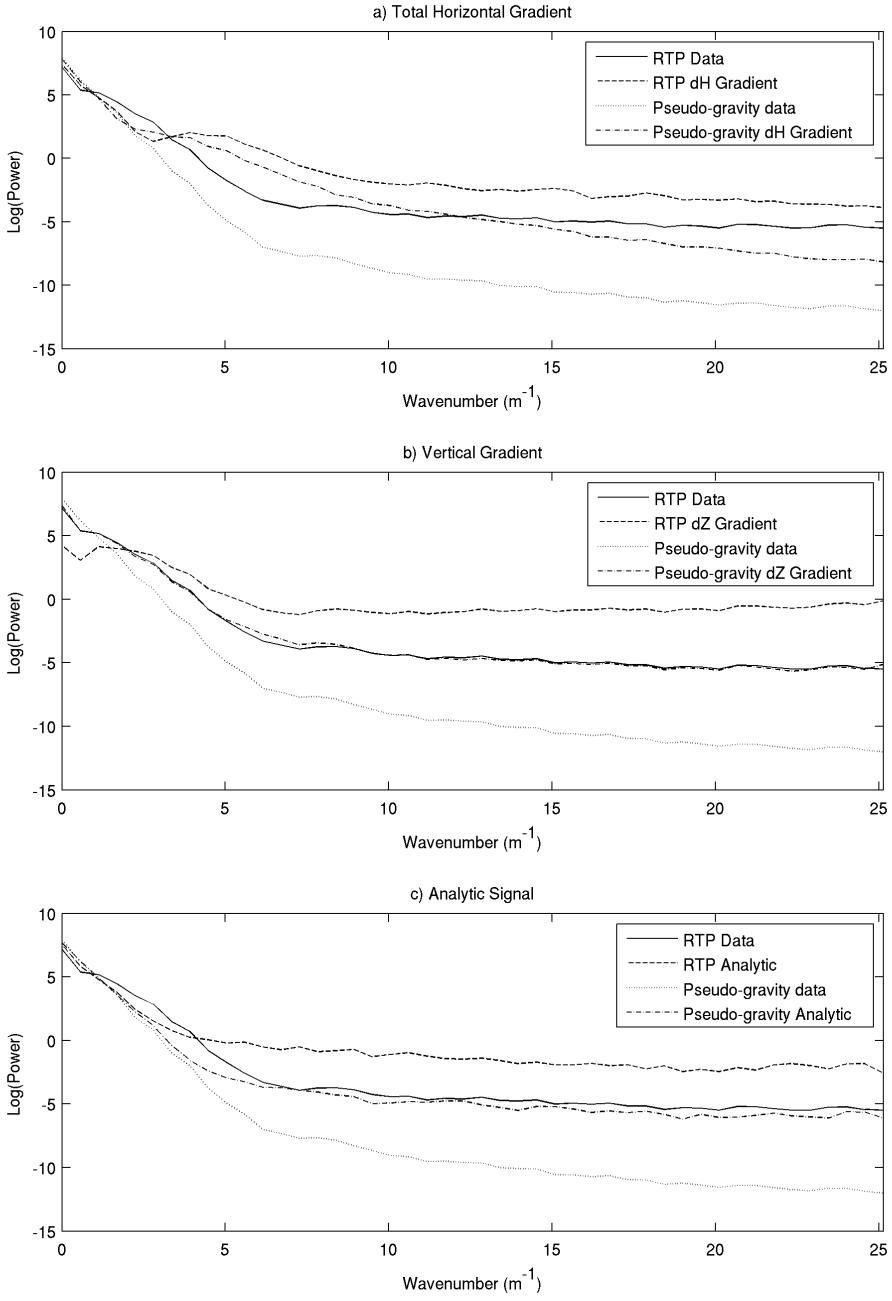

Figure 3.8: Radially averaged power spectra for total field magnetic data (solid line), generated with 0.2 nT Gaussian noise and reduced to the pole, and pseudogravity transformation (dotted line) are shown in all plots. For comparison, the power spectra for (a) horizontal gradient magnitude, (b) vertical gradient and (c) analytic signal, are shown for both RTP and pseudogravity datasets.

tational potential are linked by Poisson's relation. This relation is shown in Robinson (1971) as:

$$A_{(x,y,z)} = \frac{I}{G\rho}\frac{\partial U_{(x,y,z)}}{\partial \alpha} \tag{3.9}$$

where $A_{(x,y,z)}$ is the magnetic field potential, $U_{(x,y,z)}$ is the gravitational field potential, I is uniform magnetisation intensity in direction α, and ρ is the uniform density.

Baranov (1957) made use of Poisson's relation and introduced a linear filter which is typically referred to as the pseudogravity transformation. The output from the pseudogravity transformation is a unipolar field, and has similar properties to a gravity anomaly that would be observed should the magnetisation be accompanied with an exactly proportional density distribution. The result can simplify the interpretation of shape and location of source bodies. The transformation utility is a commonly found feature in potential field software packages, such as MagPick (Tchernychev, 2009) and Geosoft Oasis Montaj (Geosoft, 2010b). The pseudogravity datasets for the synthetic case are shown in Figure 3.9 for the scenario with no noise and 0.1, 0.2 and 0.3 nT Gaussian noise. A ratio of density to magnetisation of 1x10^2 kg m^{-3} per Am^{-1} has been used, however changing this will only act as a linear scaling factor on the output values, and therefore does not affect the wavenumber spectrum of the pseudogravity dataset.

Calculating the pseudogravity transformation of the total-field dataset has the effect of suppressing the high-wavenumber component of the data, as can be observed by comparing the dotted and solid lines on Figure 3.8. At progressively higher wavenumbers the power becomes progressively suppressed. This causes the pseudogravity transformation seen in Figure 3.9 to appear as a much smoother image than the total field data shown in Figure 3.1.

When the derivatives of the pseudogravity dataset are calculated, shown by the dot-dash lines in Figure 3.8, the resulting spectrum is much closer to the original spectrum of the total-field data (solid line), therefore the bias towards high-wavenumber features is removed, and the deriva-

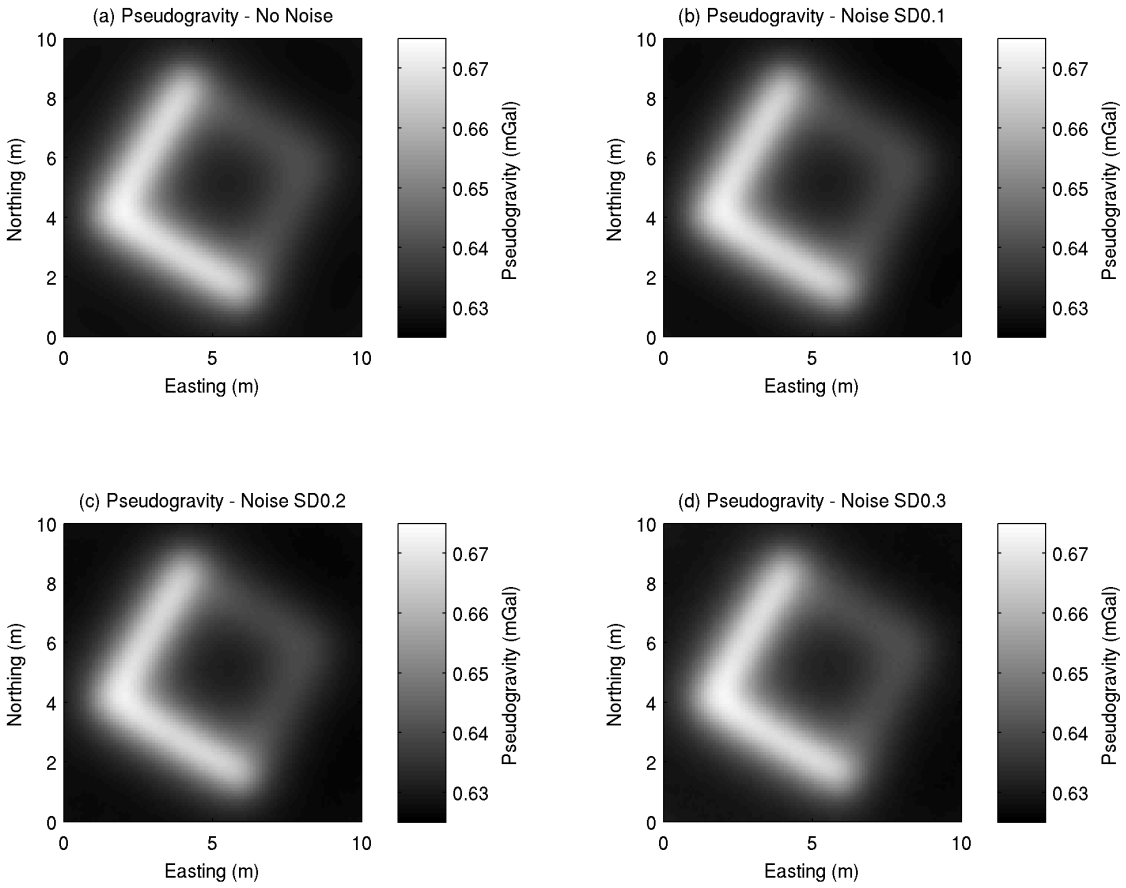

Figure 3.9: (a) Pseudogravity transformation of the synthetic data with no noise added (Figure 3.1b). (b) Pseudogravity transformation from the total field data with 0.1 nT standard deviation, zero-mean Gaussian noise added prior to transformation (Figure 3.1d). (c) 0.2 nT standarded deviation added (Figure 3.1f). (d) 0.3 nT standarded deviation added (Figure 3.1h)

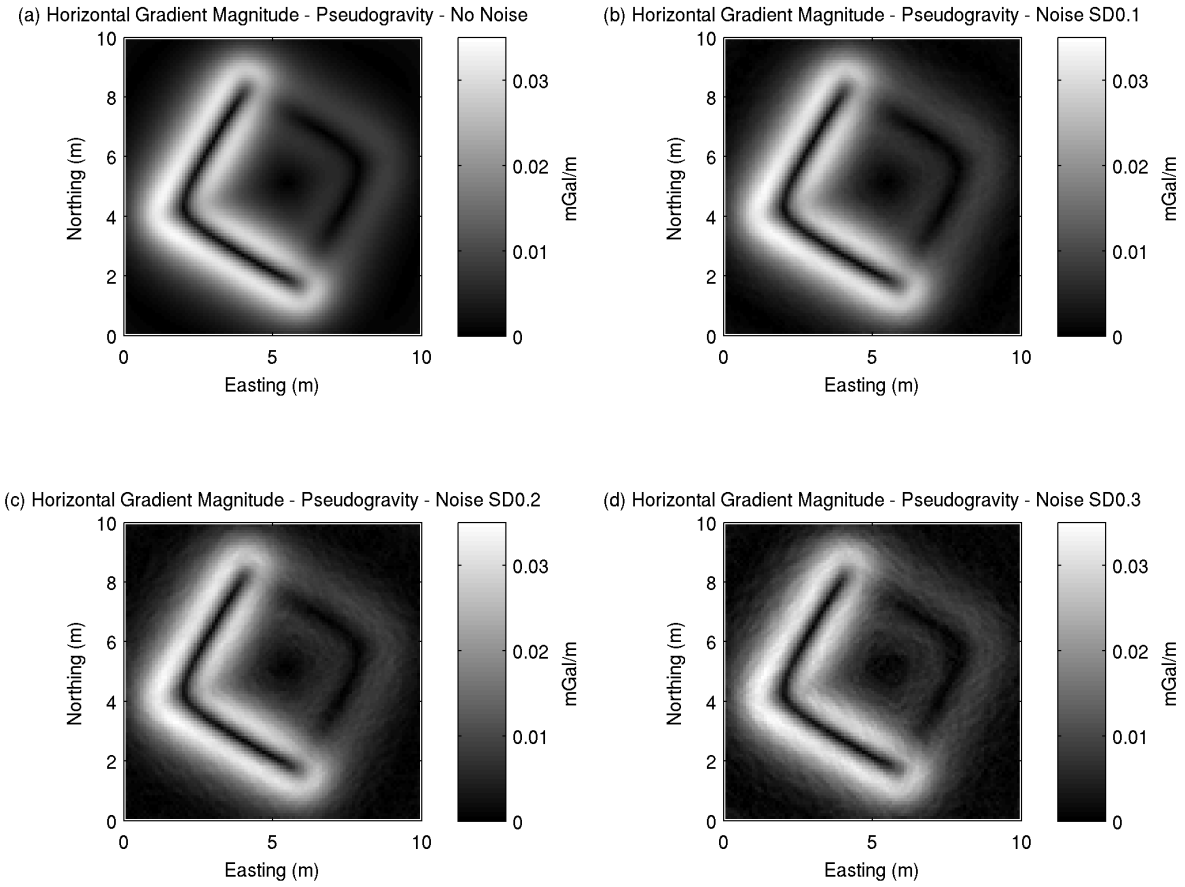

Figure 3.10: Horizontal gradient magnitude of the pseudogravity dataset shown in Figure 3.9. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied to the original total field data. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

tives are comparable to the original total field data in terms of noise. As the pseudogravity is the vertical integral of the reduced to the pole dataset, the power spectrum for the vertical derivative of the pseudogravity dataset is equivalent to that of the RTP dataset, as observed in Figure 3.8b.

3.8.1 Edge detection using pseudogravity

Derivatives calculated from the pseudo-gravity dataset have slightly altered properties from those calculated from the total field data (see Phillips *et al.* 2007), however they largely have the same desirable properties, and therefore can be used in gradient-based methods to aid edge detection without the problematic amplification of the noise identified earlier.

Figure 3.10a shows the result of the horizontal gradient magnitude method generated using the pseudo-gravity dataset. The resulting image shows similar features as seen in Figure 3.2a, with the shallower part of the body causing an anomaly of $\sim$0.035 mGal/m, and the deeper part of the body generating a lower amplitude response of $\sim$0.015 mGal/m. The higher wavenumber component present in the total field data causes Figure 3.2a to appear 'slightly sharper' than the result in Figure 3.10a, however largely the same interpretation can be made from the two datasets.

The big advantage to using the pseudo-gravity transformation can be seen in Figures 3.10b-d, where noise has been added to the original data. Whereas, in Figure 3.2, the noise meant the identification of the deeper anomaly was almost impossible, the derivative generated from the pseudo-gravity dataset is not affected to the same degree, and all results appear very similar to the zero noise scenario (Figure 3.10a). The response of the weaker, deeper body is still clearly visible above the noise threshold.

The results, without and with noise, are presented in Figure 3.11 for the Canny edge detection, Figure 3.12 for the analytic signal amplitude, Figure 3.13 for the theta map, Figure 3.14 for the tilt angle and Figure 3.15 for the horizontal gradient magnitude of the tilt angle.

When the scenarios without noise are compared for the derivatives of the total field data and the pseudogravity datasets, the results generated from the total field data are slightly preferable as those derived from the pseudogravity dataset have a slightly broader signature making the crest of anomalies slightly harder to identify. However, once noise is added, the results generated from the pseudogravity dataset are much more robust than those generated from the total field data, and show all the main features of the noise-free scenarios.

Once each technique has been applied to the pseudogravity dataset, all scenarios with additive noise have produced responses that can be used to further analyse and interpret the dataset, with the exception of the horizontal gradient magnitude of the tilt angle. As the process of calculating this is a derivative of a derivative, suppressing the high wavenumber component from the 1st

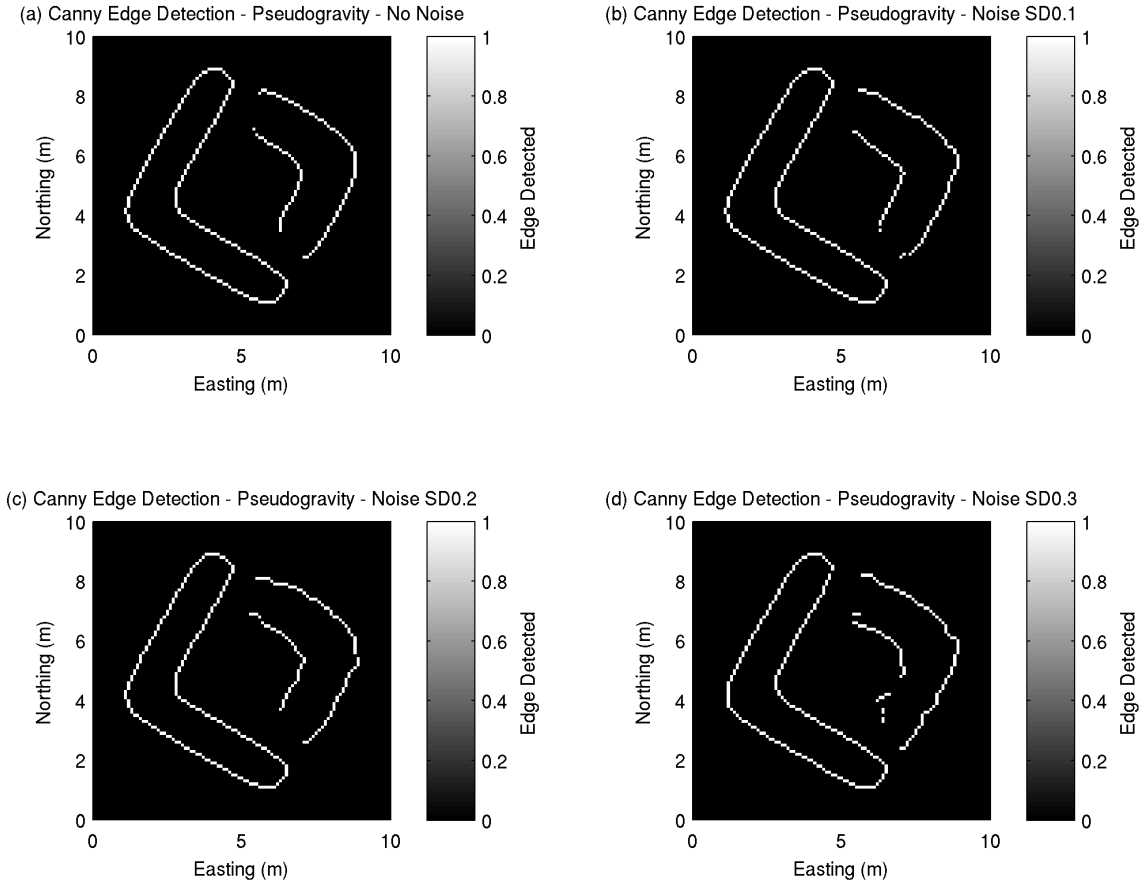

Figure 3.11: Canny edge detection of the pseudogravity dataset shown in Figure 3.9. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied to the original total field data. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

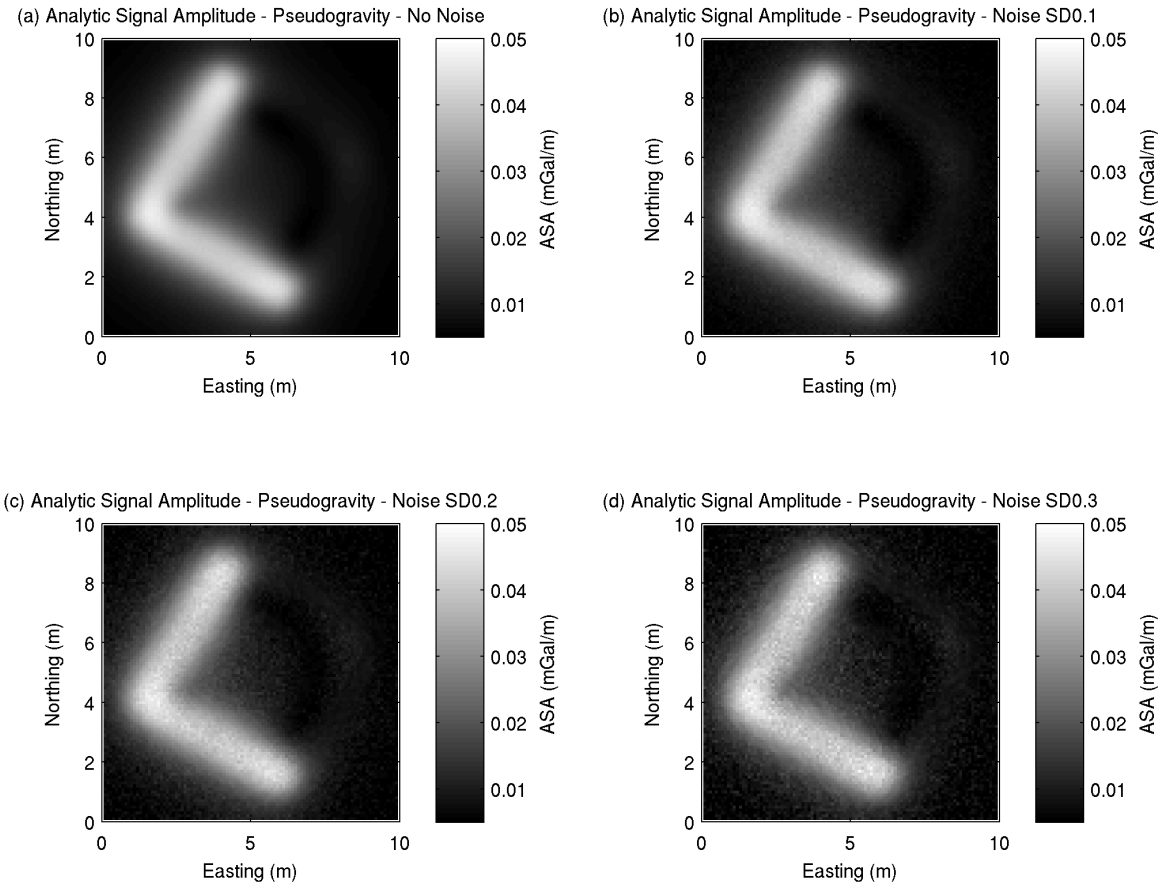

Figure 3.12: Analytic signal amplitude of the pseudogravity dataset shown in Figure 3.9. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied to the original total field data. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

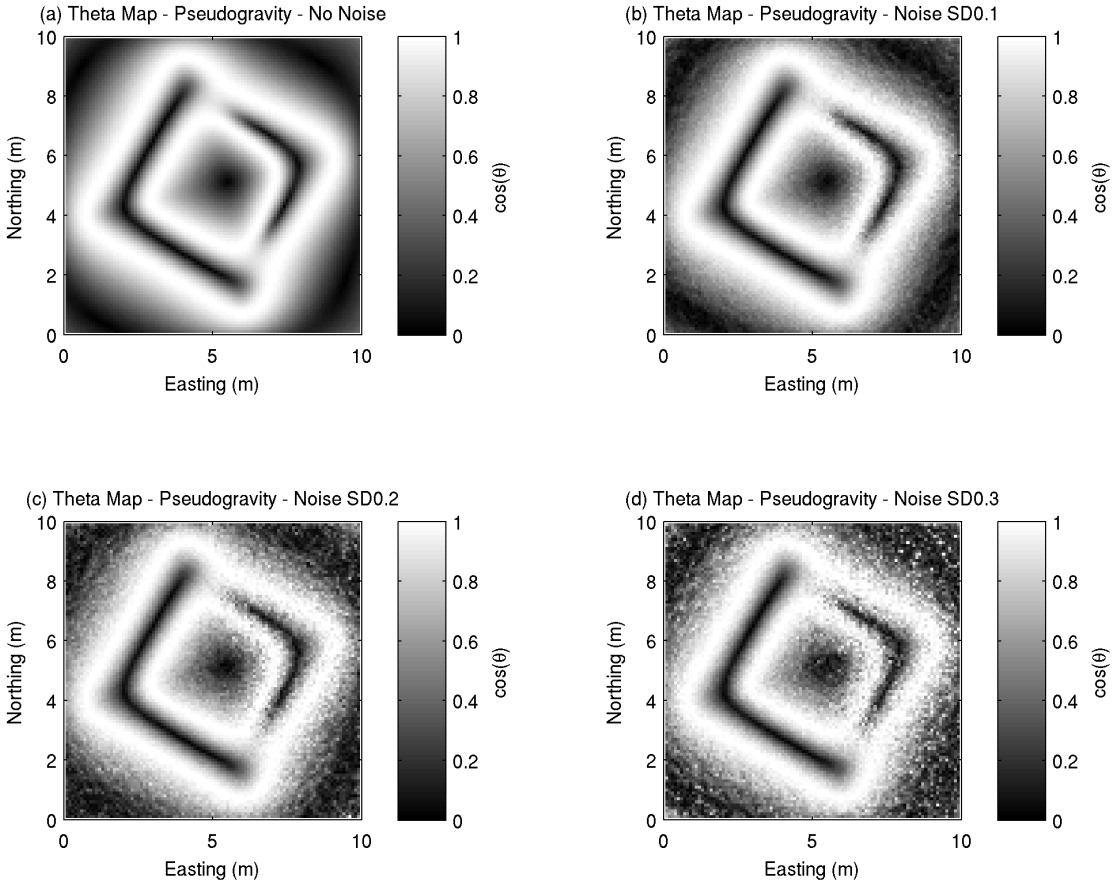

Figure 3.13: Theta map of the pseudogravity dataset shown in Figure 3.9. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied to the original total field data. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

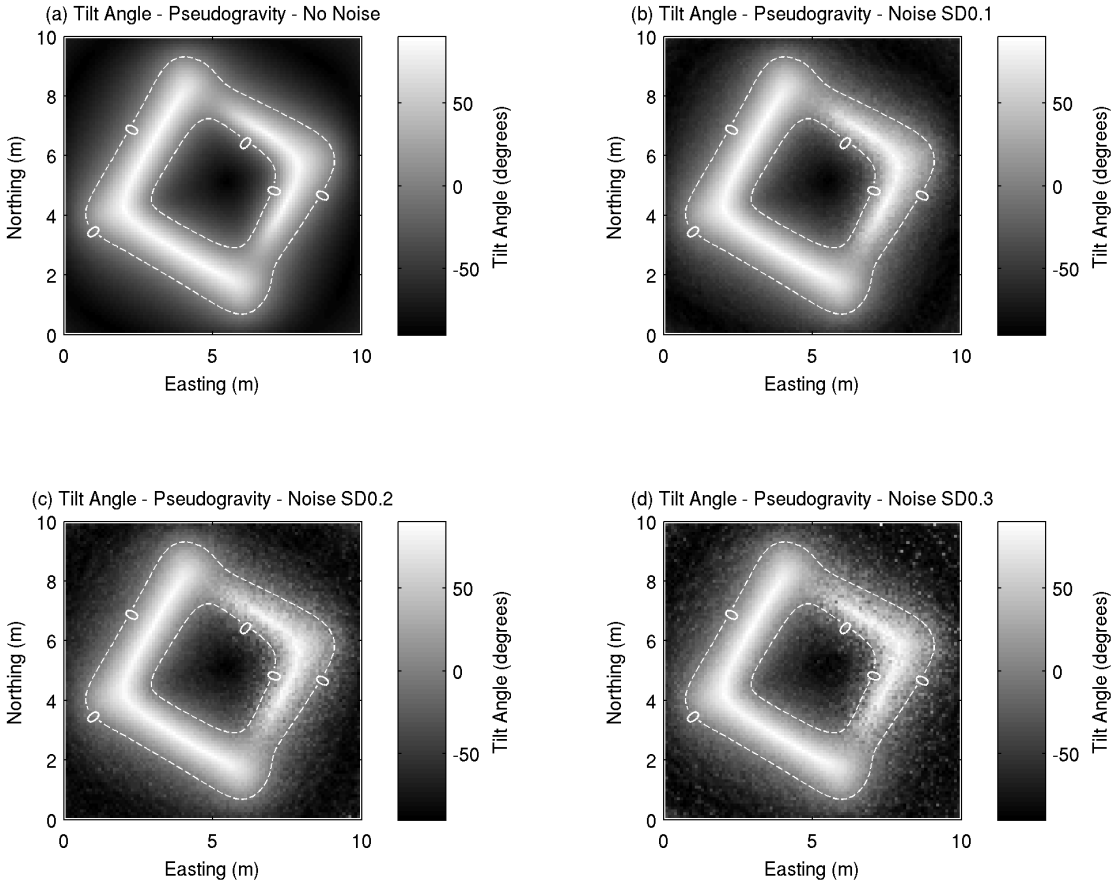

Figure 3.14: Tilt angle of the pseudogravity dataset shown in Figure 3.9. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied to the original total field data. (c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

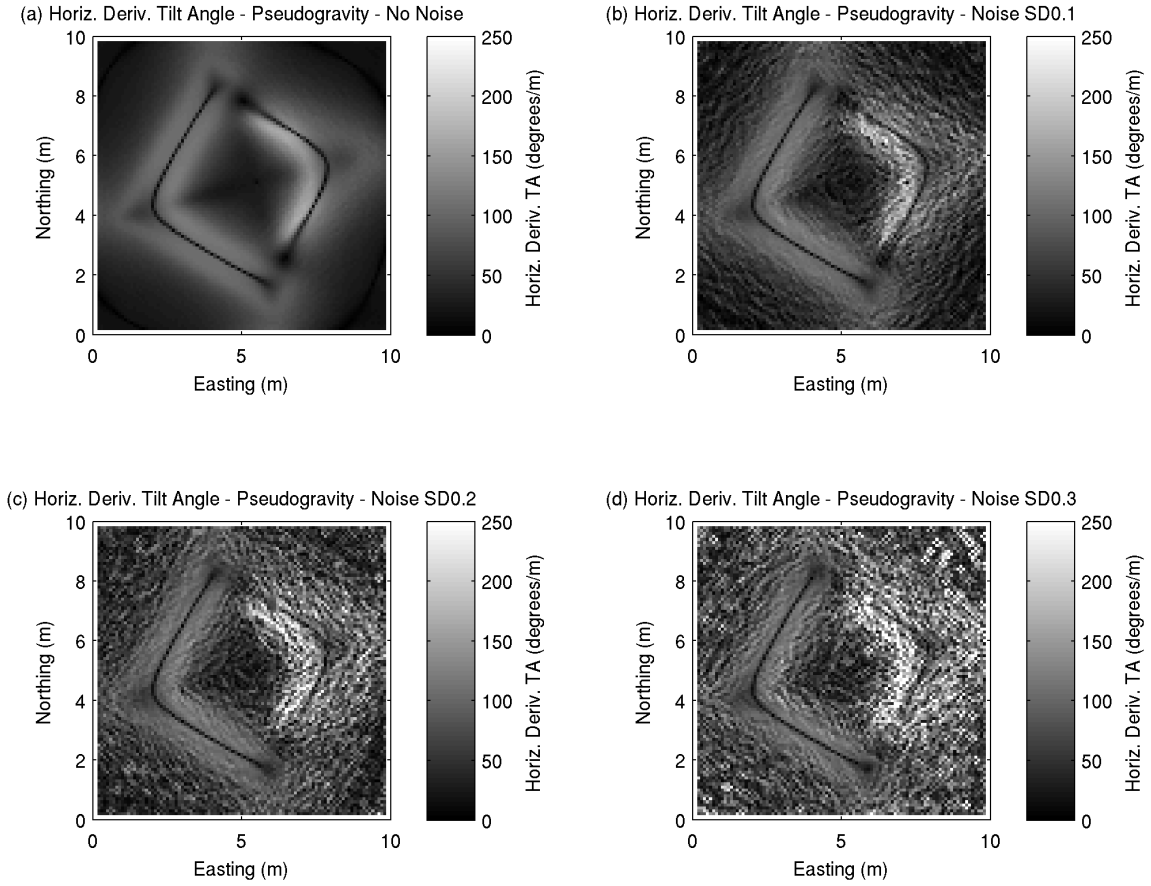

Figure 3.15: Horizontal gradient magnitude of the tilt angle for the pseudogravity dataset shown in Figure 3.9. (a) No noise has been added to the synthetic observations. (b) Noise with 0.1 nT standard deviation Gaussian noise applied to the original total field data.(c) Noise with 0.2 nT standard deviation Gaussian noise applied. (d) Noise with 0.3 nT standard deviation Gaussian noise applied.

derivative does not prevent these features being amplified once the 2nd derivative is calculated, and the information in Figures 3.15b-d is unlikely to be of use to aiding interpretation.

3.9 Case study

Brading Roman Villa is situated between the towns of Brading and Sandown on the Isle of Wight, UK. The site consists of a Roman villa, now incorporated into a visitor centre, and surrounding pasture land. The data used in this study is a subset of a much larger caesium magnetometery survey collected in 2009 by the English Heritage Geophysics Team over the surrounding pasture land (Payne, 2009). The total field data were collected with a spatial resolution of 0.5 x 0.125 m.

The subset, shown in Figure 3.16a, is a 135 x 55 m area situated $\sim$100 m NW of the villa, which shows a series of positive anomalies. The most prominent features in the data are a large circular anomaly $\sim$25 m in diameter towards the western side of the dataset, and a series of positive linear anomalies across the area. A ploughing pattern is prominent across the dataset, as well as a series of small positive peaks scattered across the area.

The pseudo-gravity transformation of the data can be seen in Figure 3.16b. Interpretation of Figure 3.16b is much harder, since the suppression of high-wavenumber components leaves only a prominent general regional trend across the area from SW to NE. The northern part of the large circular anomaly is also visible, but much harder to interpret than in the total field data.

Figure 3.17 shows the horizontal gradient magnitude of Figure 3.16. The horizontal gradient magnitude calculated from the total field data (Figure 3.17a) shows approximately half of the circular anomaly strongly, and a possible linear continuation of this to the south. More linear anomalies can be identified towards the eastern part of the study area, and the edges of several small isolated anomalies can easily be identified. Generally, features are harder to identify in

the horizontal gradient magnitude data compared to the total field data, and many anomalies appear at least partly masked by the ∼E-W ploughing pattern.

Below this, in Figure 3.17b, the horizontal gradient magnitude calculated from the pseudo-gravity dataset can be seen. The most notable feature seen in this image is the suppression of the ploughing pattern observed in Figure 3.17a. This means that many more of the observable features in the raw data can be seen in the horizontal gradient magnitude data. At least one edge of the entire circular anomaly is now identifiable, as well as two linear tangents to the circle, a much longer and continuous anomaly running from the northern part to the eastern part of the survey area, and a more obvious segregation of the isolated anomalies from the plough pattern.

The results of the theta map are shown in Figure 3.18a for the total field data, and Figure 3.18b for the pseudogravity dataset. Once again, the improved interpretability of the results from the pseudo-gravity dataset over the total field is clearly observable, with Figure 3.18b showing far more of the anomalies observable in Figure 3.16a, than Figure 3.18a, which is largely uninterpretable. The advantage of the theta map over the horizontal gradient magnitude, is that weaker anomalies are enhanced to the same amplitude as the largest anomalies, and this can be seen in the circular anomaly, where now both edges can be observed for at least part of the southern section.

Figure 3.19 shows the results of the tilt angle for the Brading data. The tilt angle calculated from the total field data (Figure 3.19a) shows the weak ploughing pattern accentuated at the expense of the other anomalies in the data, and therefore the data are hard to interpret. Figure 3.19b shows a result similar to the total field response but again with the advantage that the weaker linear anomalies have been boosted to the same amplitude as the circular anomaly. The ploughing pattern has caused significantly less disturbance to the features of interest than in Figure 3.19a.

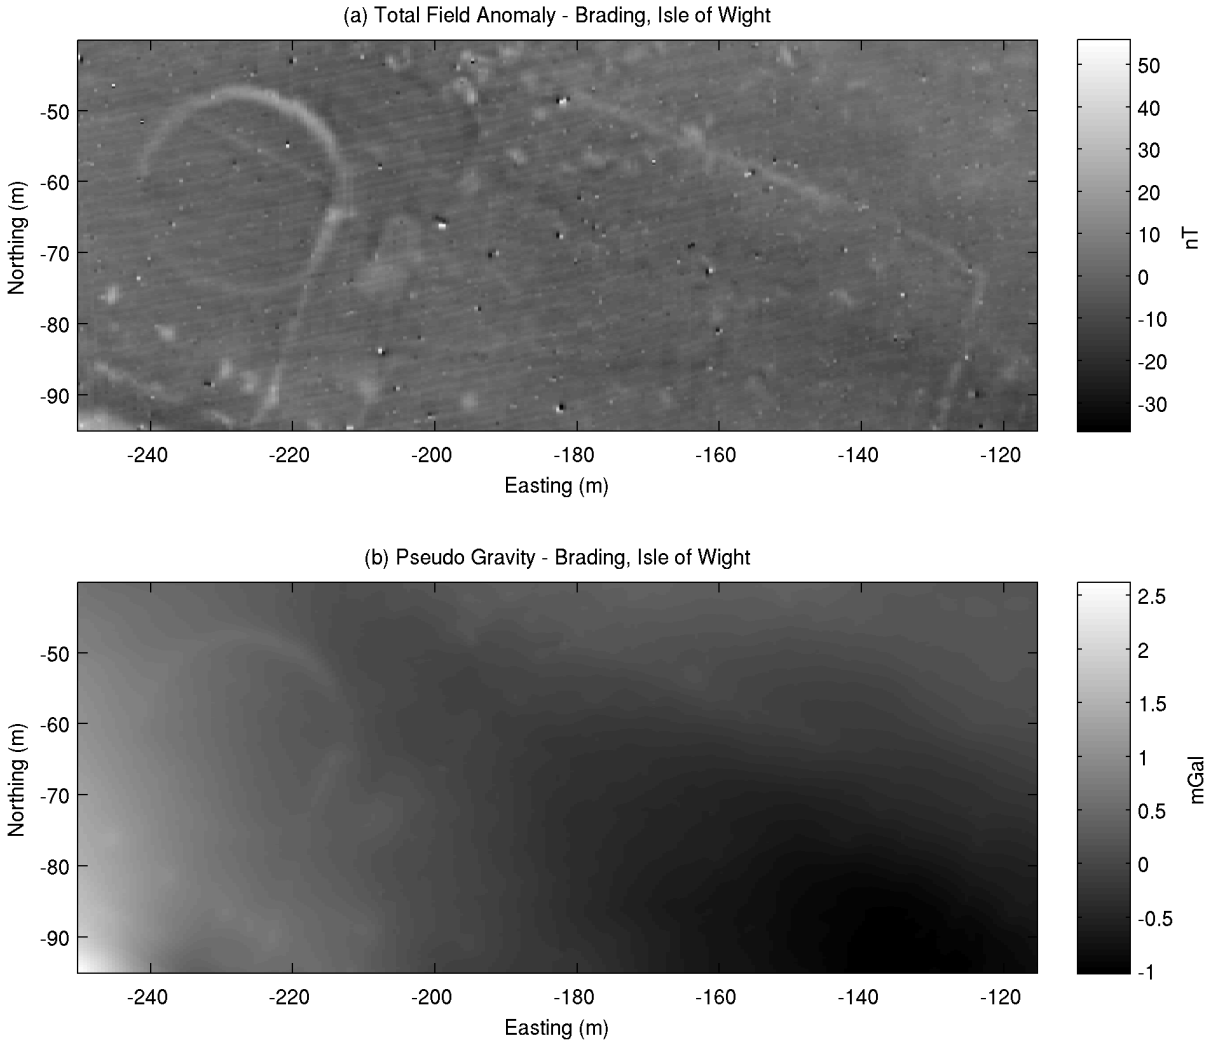

Figure 3.16: Total field and pseudogravity data from Brading Roman Villa, Isle of Wight, UK. (a) Subset of total field data from an area surrounding Brading Roman Villa, Isle of Wight, UK, showing the full dynamic range of the data collected. (b) Pseudogravity transformation of a).

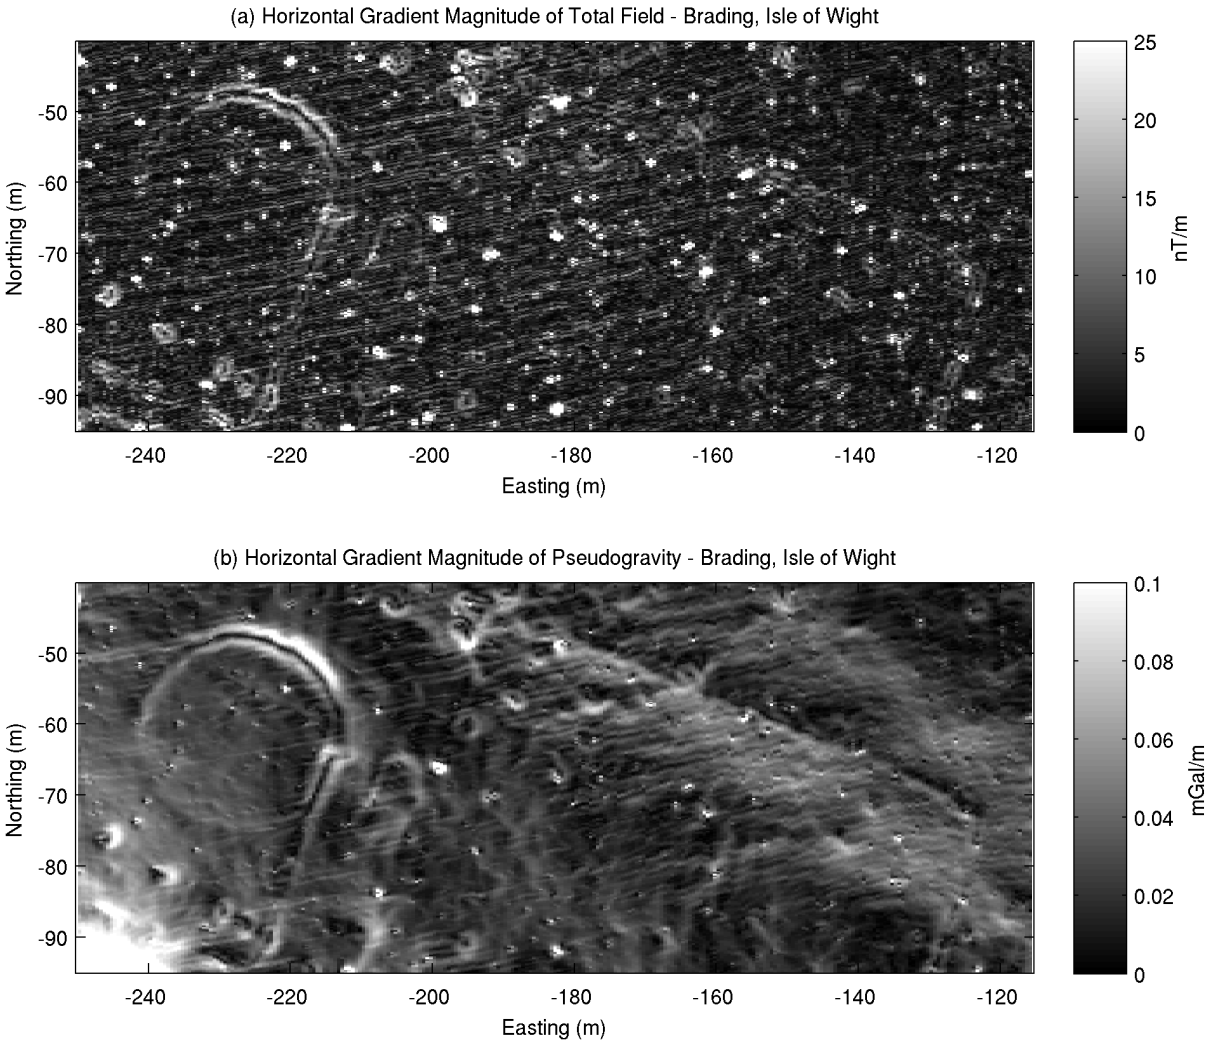

Figure 3.17: (a) Horizontal gradient magnitude of total field data from Brading. (b) Horizontal gradient magnitude of pseudo-gravity transformation

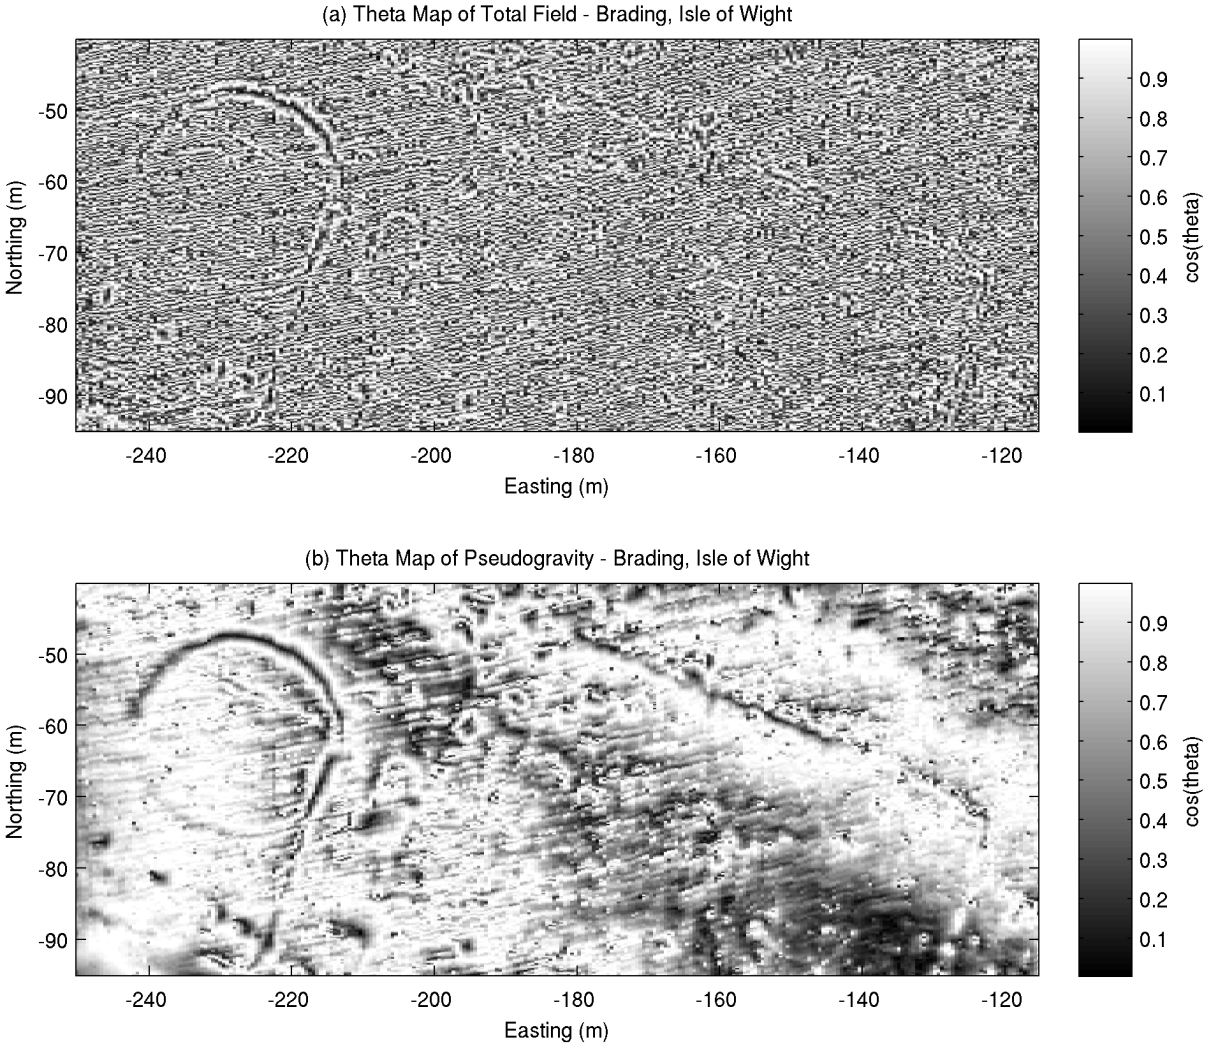

Figure 3.18: (a) Theta Map of total field data from Brading. (b) Theta Map of pseudo-gravity transformation

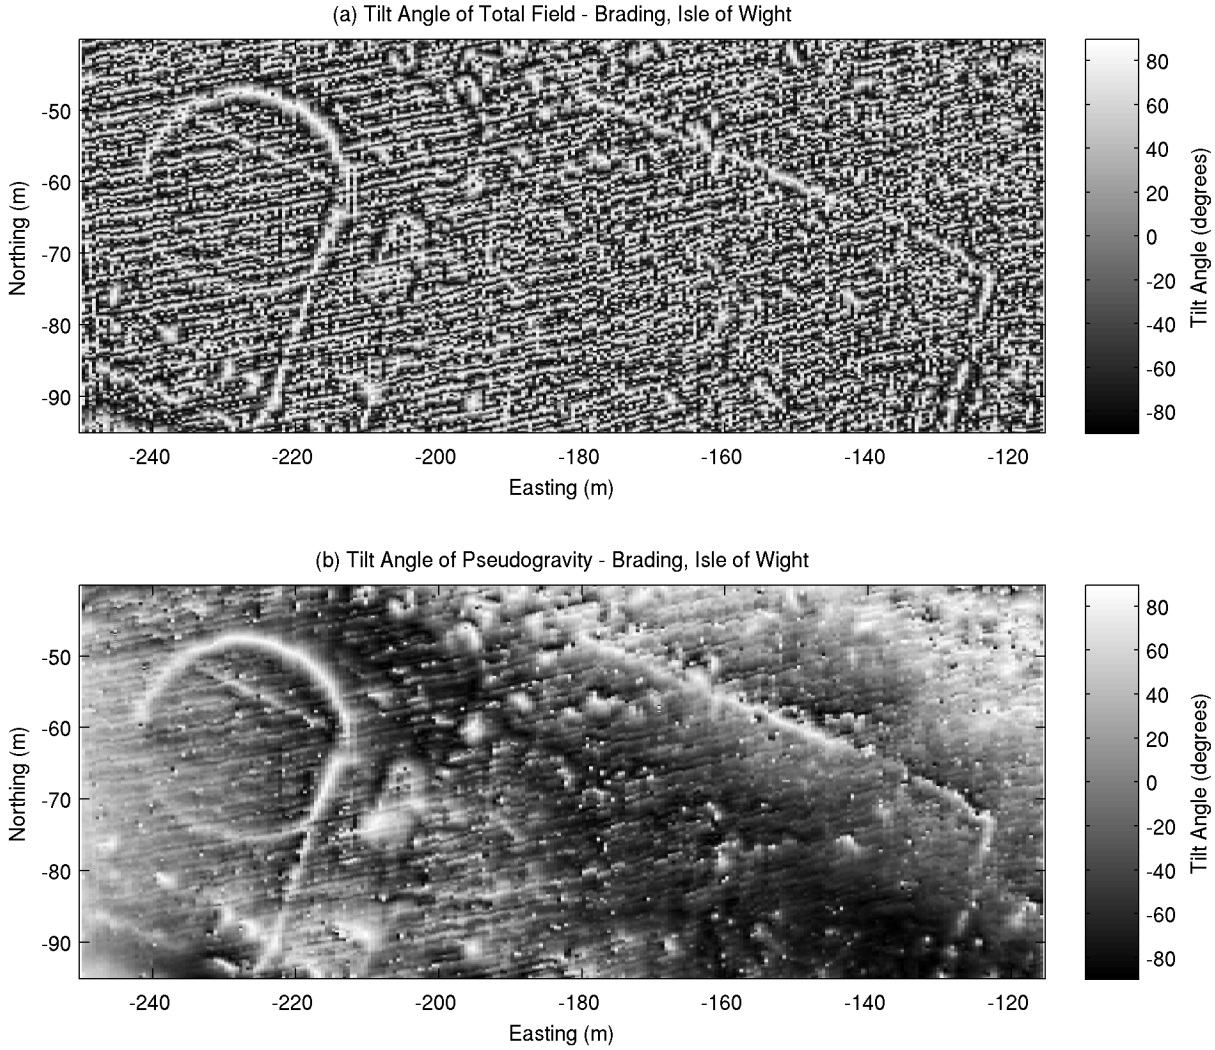

Figure 3.19: (a) Tilt derivative of Total Field data from Brading. (b) Tilt derivative of pseudo-gravity transformation

3.10 Discussion

It has been shown here that derivative-based methods have been successful at highlighting the edges of anomalous bodies and at normalising the amplitude of anomalies to improve the identification of weaker responses. With the synthetic anomaly designed to test these methods, the optimum results have been achieved from the derivatives taken from the total field measurements of the zero-noise scenario, as these results have shown sharp peaks in the vicinity of the edge locations. However, the introduction of small amounts of Gaussian noise has severely affected the users ability to interpret the responses, to the point where anomalies still identifiable in the total field data are uninterpretable in the derivative datasets.

The pseudogravity transformation is already a commonly available, and well used way to aid the interpretation of magnetic data. The particular feature of the pseudogravity transformation which makes it useful here is that the vertical gradient of the pseudogravity transformation shares the same distribution across its power spectrum as the measured total field response. This is because the pseudogravity transformation increasingly suppresses power in higher wavenumbers, which is the opposite effect of calculating derivatives, which boosts the power in the higher wavenumbers. When applied to the synthetic noise-added data, it is this effect of boosting the higher wavenumbers that causes interpretation to be difficult, an effect that is mirrored when applied to real archaeological data.

Conducting the derivative edge detection methods on the pseudogravity dataset removes this bias to high wavenumber features, and results in a dataset that has a comparable power spectra to the original data, yet has all the advantages of the various derivative-based methods at simplifying the interpretation. The pseudogravity transformation is being used here as a spectral shaping filter and the relative balance between low and high wavenumber features is restored once the derivatives are taken.

It is possible that a more bespoke filter can be designed to remove the effect of wavenumbers that do not relate to features of interest. This approach is likely to still further improve the clarity of the results of the derivative methods, however requires a prior level of interpretation, in that a wavenumber band of interest would already have been identified. The advantage of using the pseudogravity transformation is that no prior interpretation is required, and no wavenumber features are compromised.

Chapter 4

Identifying vertical extents of archaeo-magnetic targets

4.1 Introduction

Magnetic data interpretation is often limited to defining lateral extents of anomalous sources, and other geophysical techniques or excavation are then used to investigate further parameters such as depth or susceptibility. However, quantitative interpretation methods for obtaining this information do exist, and are becoming increasingly popular in aero-magnetic surveys for mineral exploration purposes, where obtaining the maximum amount of information from the data is required prior to employing a costly exploration programme. The same should apply to archaeological prospecting, however implementation continues to be rare.

There are several different types of technique that can be employed to estimate depth. Many of these are empirical methods, that use rule-of-thumb principles obtained from practical experience. They are generally designed for the analysis of individual anomalies caused by single sources, positioned in isolation. The straight-slope method (see Milsom & Eriksen (2011)) and Peter's method (Peters, 1949) are examples which use analysis of the maximum slope or gradient from a profile to make an inference to depth. Spectral methods have also been applied to derive depth information from magnetic datasets. Spector & Grant (1970) demonstrate the

effect that depth has on the Fourier transform of the signal, and calculate a depth estimate from the slope of the radially averaged power spectrum. While these methods can be quickly implemented for individual anomalies, application to large survey areas containing many anomalies can be time-intensive.

Here, methods which can be applied to the whole dataset, and give estimates for individual source bodies are investigated. The tilt-angle (Salem *et al.*, 2007), source parameter imaging (Thurston & Smith, 1997; Smith *et al.*, 1998) and Euler deconvolution (Thompson, 1982; Reid *et al.*, 1990) are all based on the analysis of the total field derivatives, while downward continuation (Desvignes *et al.*, 1999; Grant & West, 1965) utilises transformations of the signal in the Fourier domain based on the fundamental physics of potential fields.

4.2 Tilt angle

Salem *et al.* (2007) showed that depth can be calculated from the tilt angle method, where h is equal to z_c, θ will equal 45° (see Equation 3.7). Therefore, the depth to contact-like structures will be half the distance between the $\pm 45^\circ$ contours of the tilt angle. To test this, the anomaly generated from three separate prisms are analysed. All three prisms are $2\times2\times0.5$ m in dimension, and positioned between 4 to 6 mN (Figure 4.1a). The western prism is the shallowest, buried 0.3 to 0.8 meters below ground level (mbgl), and positioned 3-5 mE. The central prism is positioned 0.8 to 1.3 mbgl and situated between 9-11 mE, and the eastern prism is 1.8 to 2.3 mbgl and situated between 15-17 mE. Synthetic data has been generated at a height of 0.2 m above ground level, with an inducing field inclination of 90° and declination of 0°, to simulate a reduced to the pole dataset (Figure 4.1b). No noise has been added to this dataset.

The tilt angle interpretation of the synthetic dataset, calculated following the method described in Section 3.6, is presented in Figure 4.2. The style of presentation is the same as that used by Salem *et al.* (2007), in order to aid interpretation. The horizontal extents are shown by the

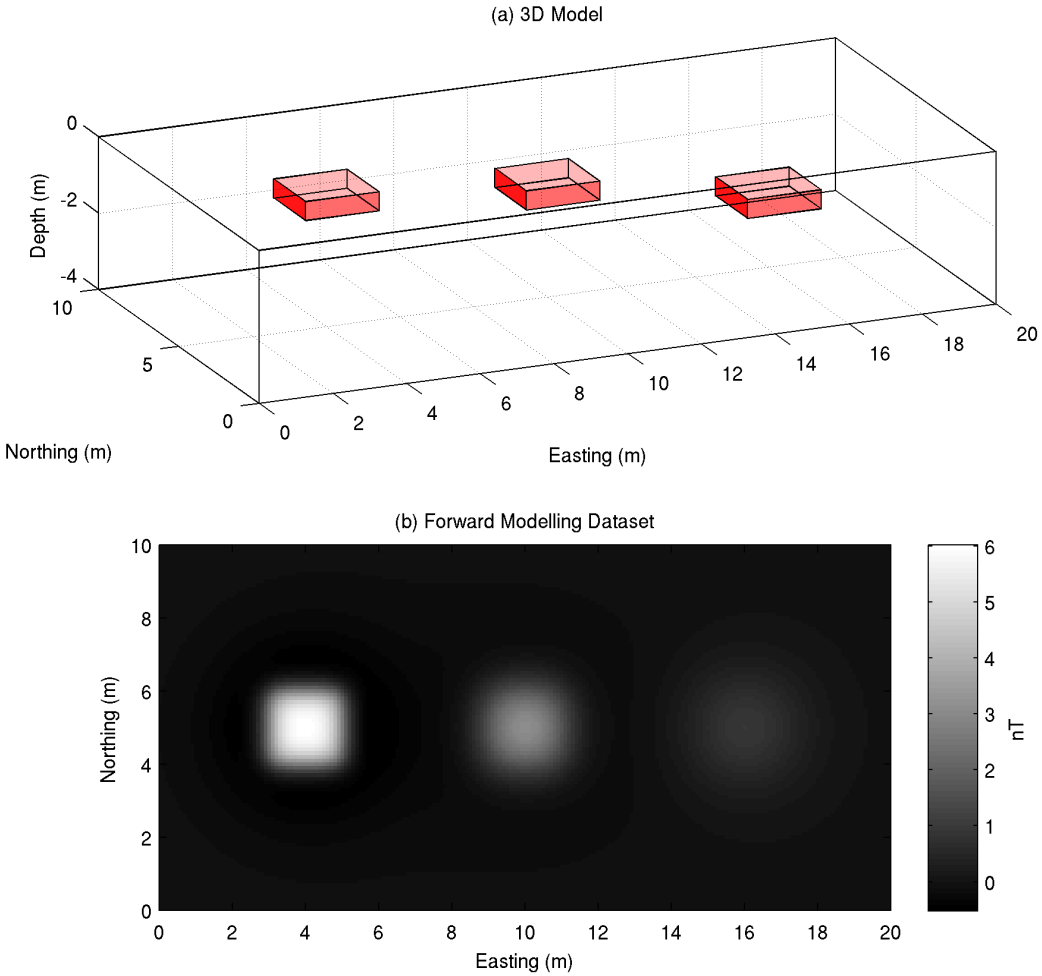

Figure 4.1: Synthetic model and dataset for a set of three prisms with different depths. (a) All three prisms are 2×2×0.5 m in dimension, with the western prism situated between 0.3 and 0.8 mbgl, the central prism situated between 0.8 and 1.3 mbgl, and the eastern prism situated between 1.8 and 2.3 mbgl. (b) Forward modelling response to a) generated 0.2 m above ground level, with 0.1 m resolution. The inducing field has been modelled at 90° to simulate a reduced to the pole dataset.

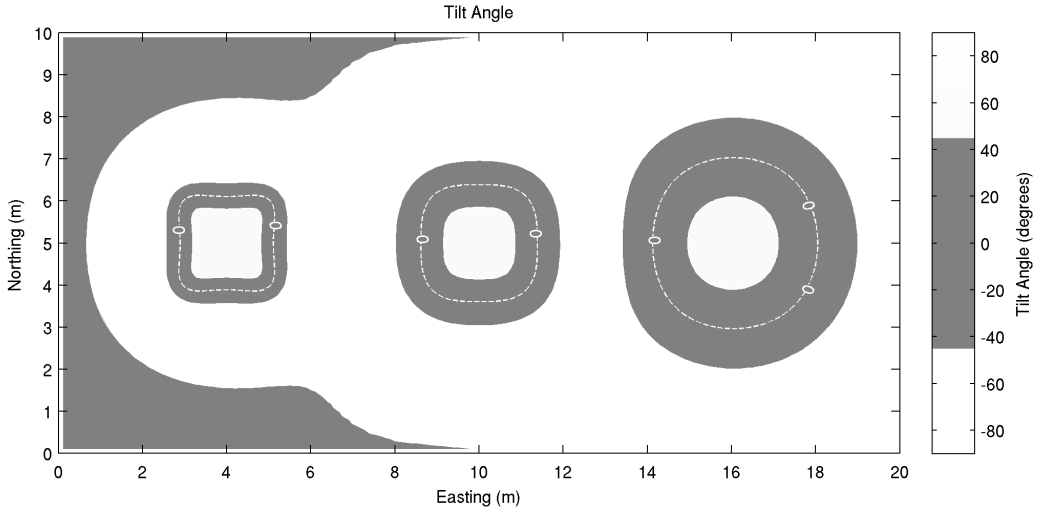

Figure 4.2: Tilt angle interpretation of the dataset shown in Figure 4.1. The data are presented as in Salem *et al.* (2007). The horizontal edges are located beneath the 0° contour, and the depth is half the distance between the 45° and -45° contours.

0° contour line. It can be seen that the outline of the 0° contour line in Figure 4.2 shows the accuracy of the edge detection degrades as the depth increases. Depth can be interpreted by analysing the distance between the 45° and -45° contours. Deeper features should be represented by increasing distance between these two contours as can be seen in Figure 4.2.

In order to study the dataset closely, a profile of the Tilt Angle dataset along 5 mN is shown in Figure 4.3. The distance between the $\pm 45^\circ$ contours for the western prism is 0.63 m, which can be measured on both sides of the anomaly in the profile shown. This suggests a depth to the top of the anomalous feature of 0.32 m below the recording level. In this case, the recording level is at 0.2 m, therefore the tilt-angle depth estimate for the western prism is 0.12 mbgl, 0.18 m shallower than in reality.

Interference from adjacent anomalies has resulted in a tilt angle that is not always symmetrical. For the central prism the distance between the $\pm 45^\circ$ contours is 1.16 m on the western side, and 1.09 m on the eastern side, resulting in depths estimates of between 0.35 and 0.38 mbgl for an actual prism depth of 0.8 mbgl. For the eastern prism, depth estimates of between 0.59

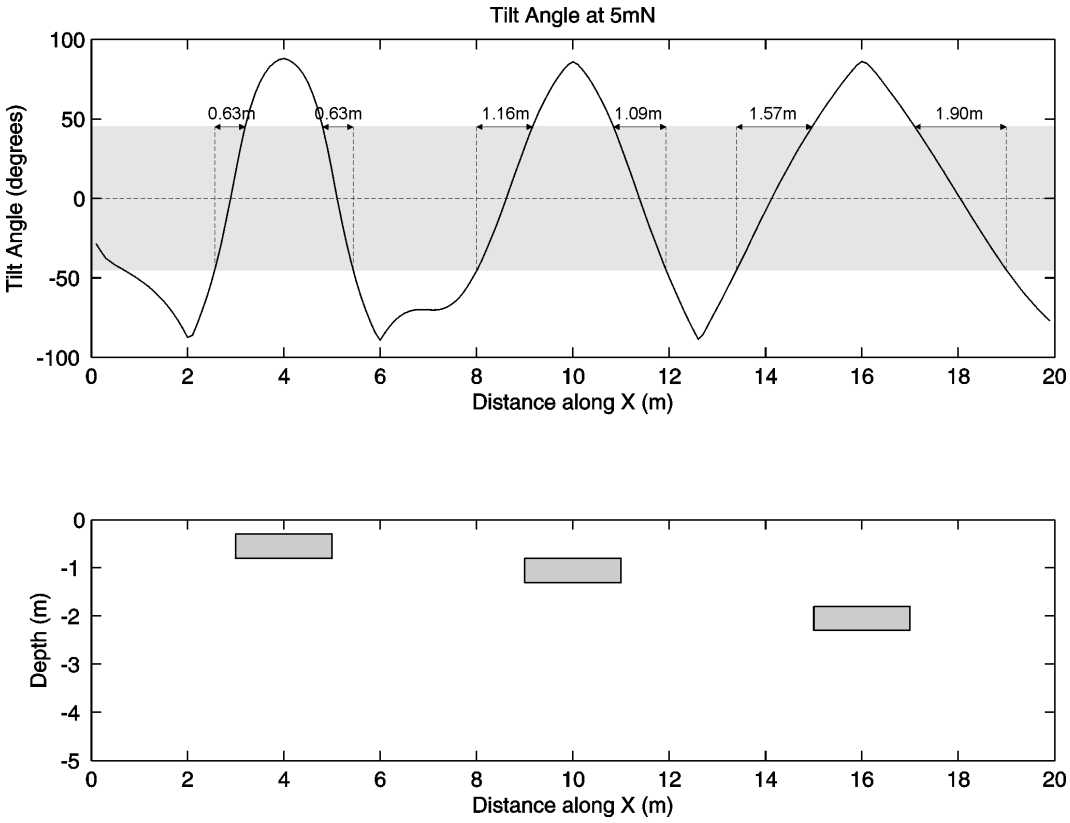

Figure 4.3: Tilt angle interpretation of the synthetic dataset presented in Figure 4.1. (top) Profile of the tilt angle along 5 mN. (bottom) Cross-section through the synthetic model.

and 0.75 mbgl have been obtained for a prism actually located 1.8 mbgl, resulting in significant underestimates for all scenarios.

4.2.1 Comparison to aero-magnetic example

The method described by Salem *et al.* (2007), was developed for mapping depth to basement features from aero-magnetic datasets, and therefore the targets investigated were of infinite vertical depth. To see the effect on the result caused by a body of finite thickness, a model similar to that used by Salem *et al.* (2007) in their Figure 1, will be analysed as an infinite thickness body, and a finite thickness body.

The model for the infinite thickness body is shown in Figure 4.4. The prism is infinitely extending in the negative and positive *x* directions, and in the negative *y* direction, but only extends north as far as 20 kmN. The depth to the top of the source body is 5000 m. Magnetisation is 0.35 Am^{-1} and a magnetic inclination of 90° has been used. In this example the recordings are taken at 250 m above ground. The calculated tilt angle profile along 20 kmE is shown in Figure 4.4. It can be seen that the edge position of the prism has been estimated accurately at 20 kmN, and that the distance between the two $\pm 45^\circ$ contours of the tilt-angle is 10516 m, giving a depth estimate of 5008 m below ground level, which is a good estimate of the actual depth, and confirms the observations made by Salem *et al.* (2007).

To investigate the usefulness of this technique on a finite thickness target, the previous example was modified so that the prism was present from 5000 to 7000 mbgl. The modelling was rerun, and the results are shown for a profile along 20 kmE in Figure 4.5. A difference in the dataset is visually obvious. The lateral position of the anomaly has again been correctly obtained from the 0° contour, but a steeper gradient of the tilt-angle has reduced the distance between the $\pm 45^\circ$ contours, and the estimate depth to the top of the prism is now underestimated, at 2279 m below ground level.

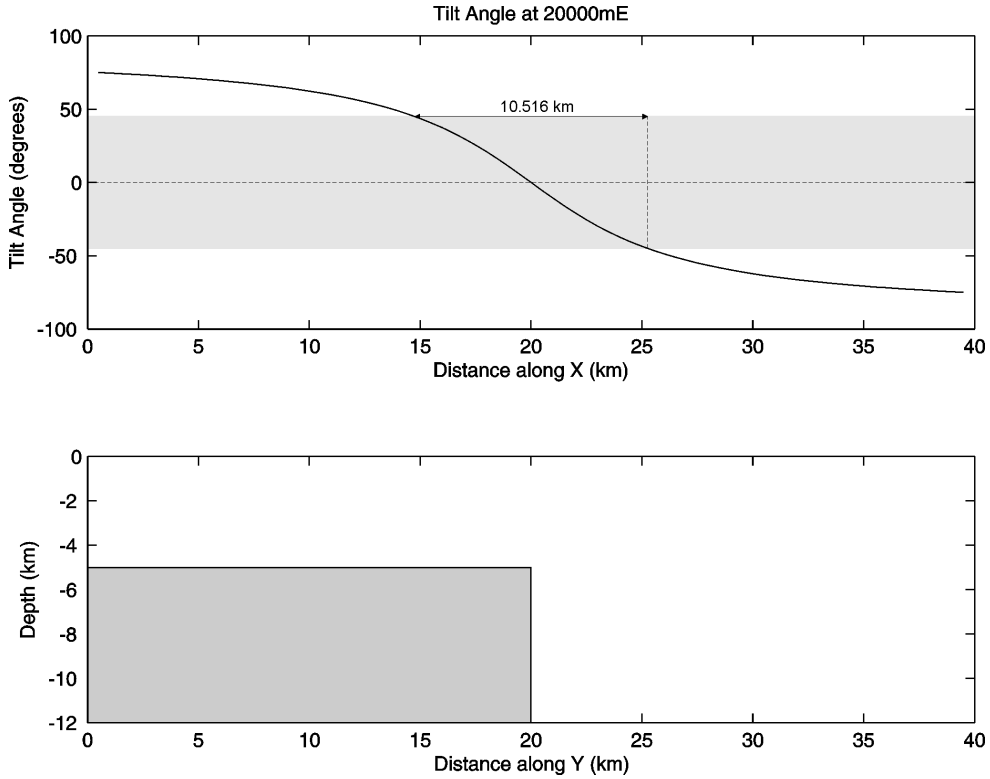

Figure 4.4: Tilt angle interpretation of the synthetic dataset presented in Salem *et al.* (2007). (top) Profile of the tilt angle along 20 kmE. (bottom) Cross-section through the synthetic model.

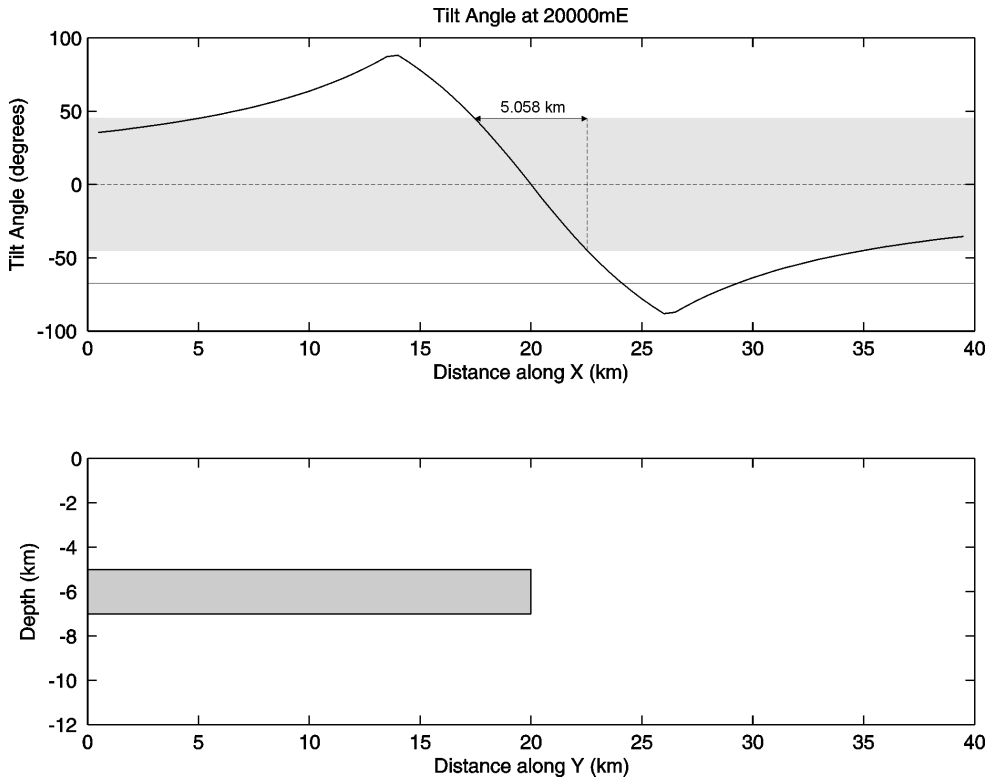

Figure 4.5: Variation of Figure 4.4, but with a finite thickness unit 5-7 kmbgl. (top) Profile of the tilt angle along 20 kmE. (bottom) Cross-section through the synthetic model.

The technique appears to break down when the anomalous feature is not of infinite vertical depth. The deeper the body is positioned, the longer the wavelength of the magnetic response produced, causing a variation in horizontal gradient magnitude, and thus tilt-angle. This is unfortunate as whereas aero-magnetic targets can often be approximated by prisms of infinite depth, archaeological targets are often confined to finite thickness, usually not exceeding $\sim$3 m. Therefore the application of the tilt angle technique for depth estimates of archaeological sources appears to be limited.

4.3 Source parameter imaging (Local wavenumber)

The source parameter imaging (SPI) method was first proposed by Thurston & Smith (1997) as an automated method to obtain location and depth information from magnetic data, as well as dip and susceptibility estimates, providing that no remanent magnetisation was present. The local wavenumber is calculated as the horizontal gradient of the tilt angle (see Equation 3.6):

$$\kappa_1 = \frac{\partial}{\partial h} tan^{-1} \left[\frac{\partial T}{\partial z} / \frac{\partial T}{\partial h} \right] \tag{4.1}$$

Smith *et al.* (1998) show that by substituting complex attributes for the horizontal and vertical gradients into Equation 4.1, it can be reduced to:

$$\kappa_1 = \frac{(n_k + 1) z_k}{z_k^2 + h^2} \tag{4.2}$$

where z is the depth to the top of a simple geometric feature, defined by the SPI structural index (n), where *n=0*, *n=1* and *n=2* for a contact, thin sheet and horizontal cylinder model respectively. If the buried body is assumed to be a contact feature (*n=0*), then positioned directly above it (*h=0*), the depth can be calculated as the inverse of the local wavenumber:

$$z = \frac{1}{\kappa_1} \tag{4.3}$$

Smith *et al.* (1998) attempted to improve this technique to remove the model-uncertainty from the depth estimates (the iSPI method). By looking at the higher-order derivatives, the second-order local wavenumber can be calculated as:

$$\kappa_2 = \frac{\partial}{\partial h} tan^{-1} \left[\frac{\partial^2 T}{\partial z^2} / \frac{\partial^2 T}{\partial z \partial h} \right] \tag{4.4}$$

which again, substituting complex attributes for the derivatives in Equation 4.4, can be reduced to:

$$\kappa_2 = \frac{(n_k + 2) z_k}{z_k^2 + h^2} \tag{4.5}$$

therefore, using Equations 4.2 and 4.5, the depth, completely independent of a chosen SPI structural index can be calculated by:

$$z = \frac{1}{\kappa_2 - \kappa_1} \tag{4.6}$$

and an estimate of the appropriate structural index derived from:

$$n_k = \frac{\kappa_1}{(\kappa_2 - \kappa_1)} \tag{4.7}$$

To analyse the suitability of the SPI technique when applied to archaeo-type datasets, the synthetic model shown in Figure 4.1 has been processed. The original total field dataset, with no noise added, is shown in Figure 4.6a. The local wavenumber, calculated using Equation 4.1, is

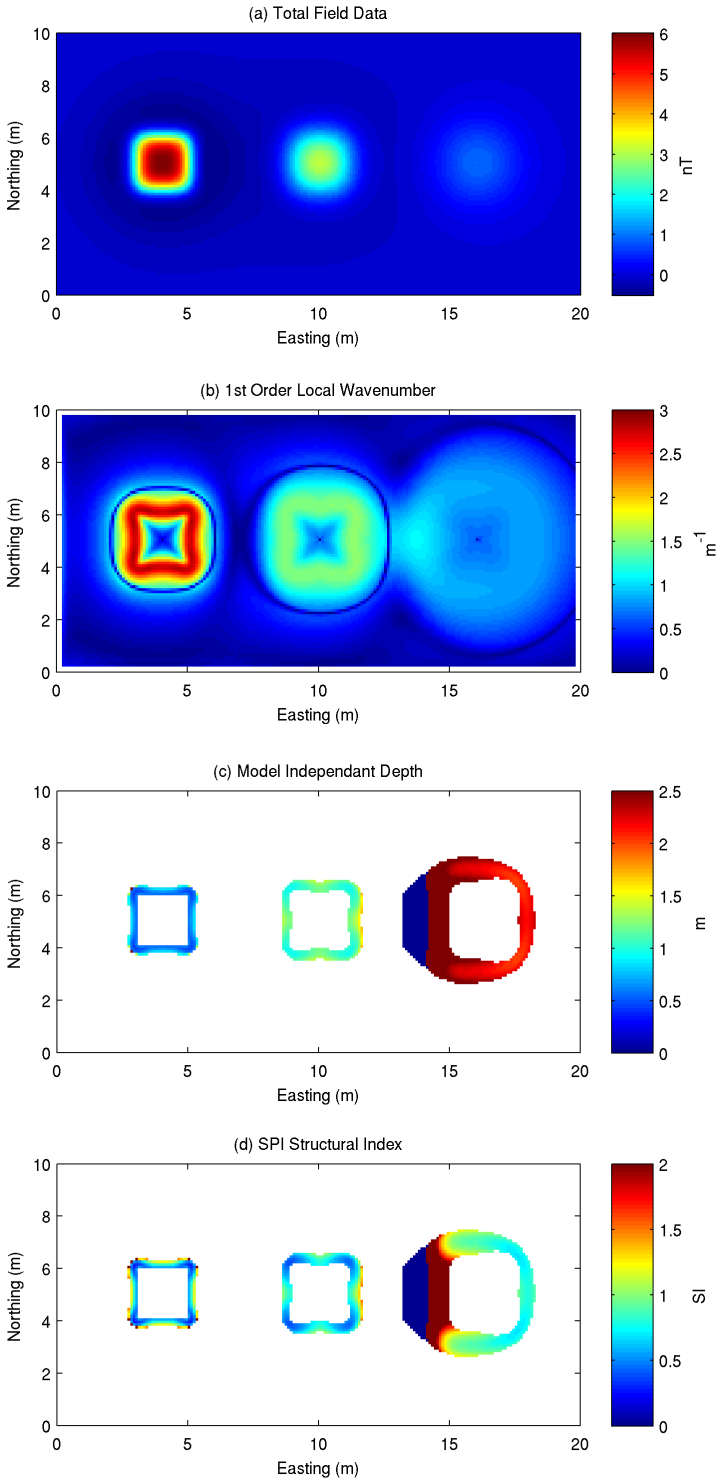

Figure 4.6: (a) Total field anomaly dataset for the compound anomaly with inclination of 90° and recording height above the ground of 0.2 m, (b) 1st-order local wavenumber (c) Model-independent depth estimate (only shown for peaks in b)) (d) SPI structural index (only shown for peaks in b)).

presented in Figure 4.6b. As discussed in Section 3.6.1, the local wavenumber shows maxima above the edges of contact-like features, and can therefore be used for edge detection purposes. Using Equation 4.2, the depth beneath these edges can be calculated. The local wavenumber of the maxima around the westernmost prism is 2.83 m^{-1}, which, assuming a structural index of a contact, results in a depth of 0.35 m below the recording level, or 0.15 mbgl. Assuming a structural index of a thin sheet results in a depth estimate of 0.71 m below the recording level, or 0.51 mbgl, and for a horizontal cylinder, a depth estimate of 1.06 m below the recording level, or 0.86 mbgl.

The maximum local wavenumber surrounding the central prism is 1.52 m^{-1}, which yields a depth for a contact feature of 0.46 mbgl, and for the eastern prism, the maximum wavenumber for the area unaffected by the interference from the central prism is 0.84 m^{-1}, which results in a depth estimate of 0.99 mbgl. In order to obtain correct results for the synthetic prisms, a structural index in between that of a contact and thin sheet must be applied.

The improvements made to the method by Smith *et al.* (1998), allow automatic calculation of the structural index as well as model independent depths (iSPI). The model independent depths are calculated using Equation 4.6, and are shown in Figure 4.6c. The results are very encouraging. Taking the values of the model independent depths from the areas where the highest local wavenumbers were recorded, the depths below recording level for the prisms from west to east are 0.46 m (should be 0.5 m, 0.3 mbgl+0.2 m recording height), 0.94 m (should be 1.0 m) and 2.02 m (should be 2.0 m). The calculated structural index (Figure 4.6d) for each feature as identified from points towards the corners of the anomalies are, moving west to east, 0.28, 0.39 and 0.69, confirming the observations made above, that the prisms represent features between the contact and thin sheet approximations.

While this technique has produced some successful results, the usability with regards to real archaeo-magnetic datasets is limited. The results with 0.2 nT of synthetic noise added are shown in Figure 4.7. The high wavenumber content of the signal means that no useful interpretation

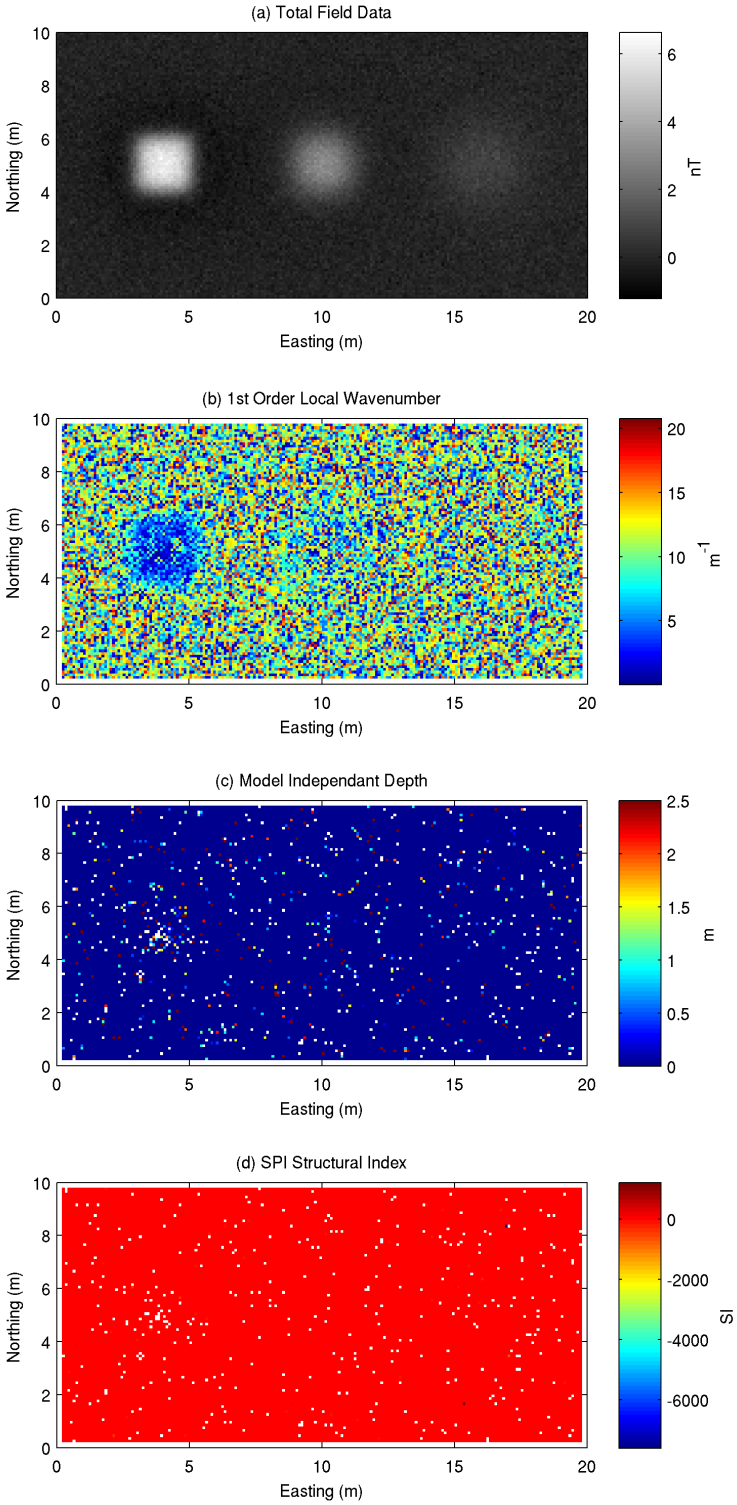

Figure 4.7: (a) Total field anomaly dataset for the compound anomaly with inclination of 90° and recording height above the ground of 0.2 m, with 0.2 nT standard deviation Gaussian noise added. (b) 1st-order local wavenumber (c) Model-independent depth estimate (d) SPI structural index.

can be made. The original method based on the work by Thurston & Smith (1997) relies on the local-wavenumber, or horizontal gradient of the tilt angle, which was shown in Figure 3.7 to be severely affected by noise, due to being a gradient of a gradient. The second order local-wavenumber required for the model independent depth, and structural index calculations, is yet another order of derivative compared to the local-wavenumber. This is further demonstrated in Figure 4.8, where the radially averaged power spectra for the total field data, without noise, and the derivations of the tilt angle, 1st order wavenumber and 2nd order wavenumber are shown. Each derivative increases the power in the higher wavenumbers relative to the lower wavenumbers as shown in Chapter 3. Using the pseudogravity dataset to calculate the tilt angle would result in a spectra closely matching the original total field data, however in the SPI method a further derivative is required, and the iSPI method, requires a further derivative again, therefore in no real-case archaeological scenario is the noise level likely to be small enough for these methods to be applicable.

4.4 Downward continuation

A useful property of potential field data is that once the data are recorded at a certain recording level, the data that would be recorded had the instrument been placed higher, or lower to the ground, can be calculated. The process of calculating the field that would be observed at intervals below the recording level is called downward continuation. Downward continuation makes two important assumptions. Firstly, no sources can be present between the original data level, and the continuation level, and secondly, all sources are positioned below the magnetometer, i.e. in the ground.

Downward continuation is of interest to geophysicists conducting magnetic surveys, as it has the effect of accentuating the amplitude and high wavenumbers of the anomalies recovered. It is described by Blakely (1995) as

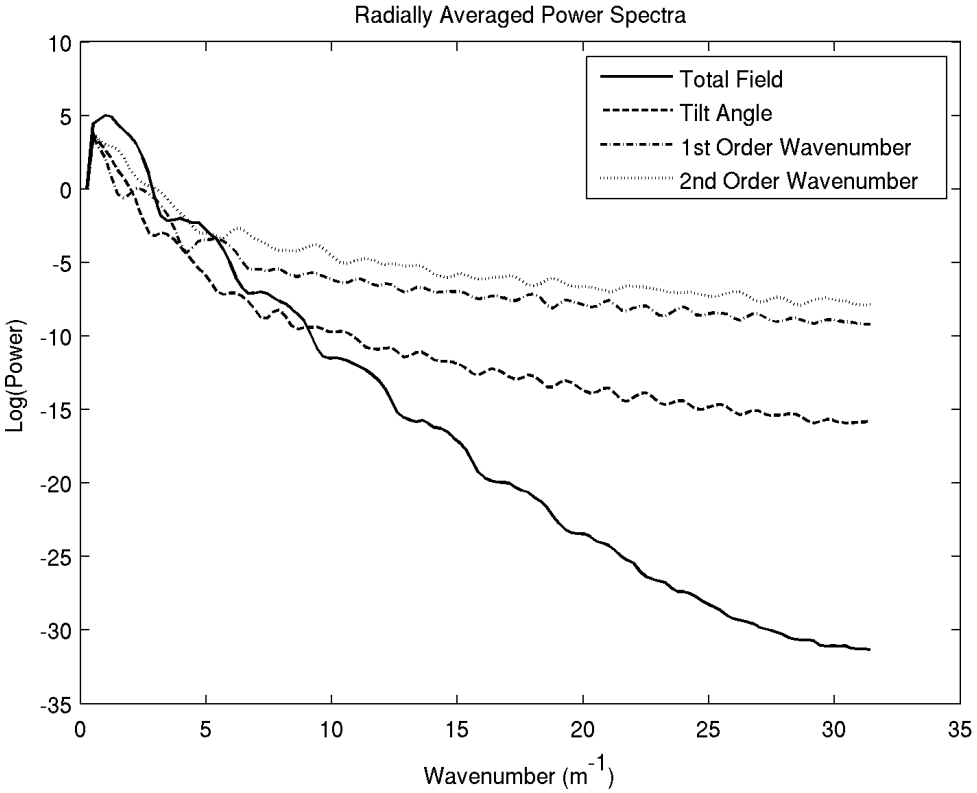

Figure 4.8: Radially averaged power spectra for the iSPI method. The calculation of the tilt angle, 1st order wavenumber and 2nd order wavenumber show increasing energy in the higher wavenumbers, resulting in a dataset from which it is difficult to derive meaningful information.

$$F(k_x, k_y, \Delta z) = F(k_x, k_y, z)e^{+|k|\Delta z}. \tag{4.8}$$

where Δz is the distance to be downward continued, $F(k_x, k_y, z)$ is the original data in the Fourier domain, and $|k|$ is the absolute value for the wavenumber. The transformation amplifies the highest wavenumbers of the data depending upon the distance to downward continue. Therefore, it is possible that small variations in the original dataset, including noise, may become large and unrealistic anomalies.

As well as a method of enhancing the appearance of anomalies, downward continuation has been used to yield estimates of the susceptibility of an equivalent strata (Grant & West, 1965; Desvignes *et al.*, 1999). The susceptibility of an equivalent strata is a laterally varying susceptibility model, with infinite depth, that can account for the observations. A different model is calculated at each continuation level. Desvignes *et al.* (1999) set this process out as a three stage method, involving the calculation of the magnetic potential, downward continuation to the level required, followed by conversion to magnetisation, and susceptibility.

To test the effects of downward continuation on the dataset presented in Figure 4.1, the data were transformed from the recording level, at 0.2 m above ground (Figure 4.9a), to ground level (Figure 4.9b), to the top of the upper prism at 0.3 mbgl (Figure 4.9c), and finally through the first prism to the base at 0.8 mbgl (Figure 4.9d). Downward continuation for short distances below the recording level has shown to be a useful tool in this scenario, as Figure 4.9b show all three anomalies with sharper edges and larger amplitudes. Figures 4.9c and 4.9d, however, highlight a commonly observed problem with downward continuation. Blakely (1995) states that the procedure is an unstable calculation, and that small errors become accentuated with increasing depth. As a method for counteracting this, the downward continuation procedure is often treated as an inverse problem, with regularisation imposed to stabilise the calculation. The procedure has been repeated, this time using the inverse approach of Tchernychev (2009), and the results are presented in Figure 4.10. This method has produced improved results, and

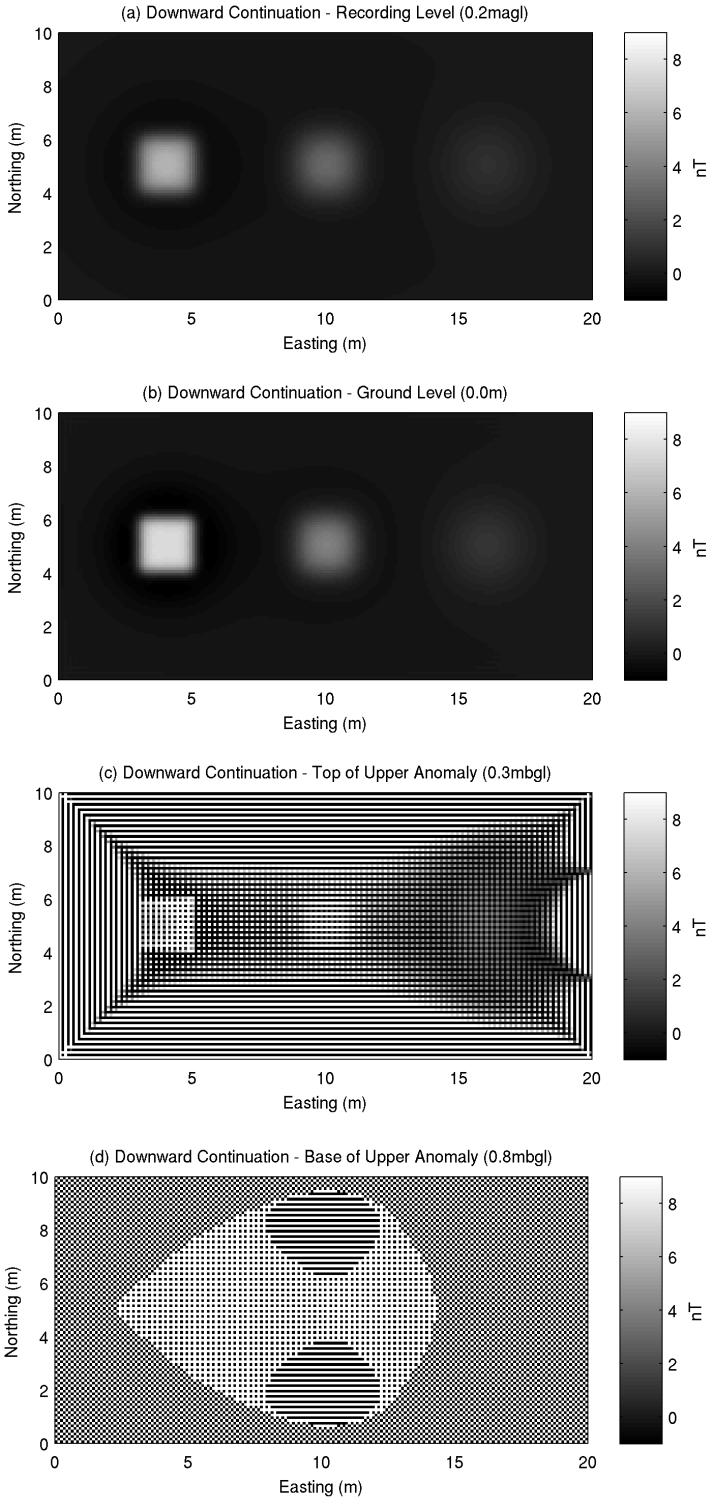

Figure 4.9: Downward continuation of the dataset shown in Figure 4.1. (a) Original recording level. (b) Downward continued 0.2m, to ground level. (c) Downward continued to 0.5m, the top of the shallowest prism. (d) Downward continued to 1.0m, the base of the shallowest prism.

stable downward continued datasets can be seen at the 0.3 and 0.8 mbgl depths in Figure 4.10c and 4.10d, respectively. However, this approach has allowed continuation to depths beyond the upper surface of the first prism, thus invalidating one of the major assumptions that no sources are present between the recording level and the continuation level. Figure 4.10d shows heavily alternating positive/negative features, however away from this westernmost anomaly the dataset shows no evidence that the results are invalidated.

Desvignes *et al.* (1999) show that the potential (V) and its downward continuation can be calculated by:

$$V(k_x, k_y, \Delta z) = F(k_x, k_y, z)\frac{\sqrt{k_x^2 + k_y^2}}{(ik_y cosI + \sqrt{k_x^2 + k_y^2} sinI)^2} e^{\Delta z\sqrt{k_x^2+k_y^2}}. \tag{4.9}$$

Using this method, the downward continuation of the potential is performed and shown in Figure 4.11, for the same depths of 0.2 m above ground (Figure 4.11a), ground level (Figure 4.11b), 0.3 mbgl (Figure 4.11c) and 0.8 mbgl (Figure 4.11d). Figures 4.11a and 4.11b, show a much smoother image than seen with the downward continuation of the total field. This is to be expected for the reasons discussed in Chapter 3, since the magnetic potential is proportional to the pseudogravity dataset. As the continuation level increases, the results once again become unstable, with the datasets generated at 0.3 mbgl and 0.8 mbgl becoming heavily contaminated with artefacts of the calculation.

The problems associated with downward continuation are particularly highlighted once Gaussian noise is added. Figure 4.12a shows the same dataset from Figure 4.1, but this time with 0.2 nT standard deviation Gaussian noise added, representing a light to moderate noise scenario. The presence of the noise causes unstable results even at small continuation distances. Figure 4.12b shows the calculation at 0.3 mbgl using the method shown in Equation 4.8. The results show rapid variations from positive to negative values across the grid with no correlation to the original data. The only stable results to downward continuation in the presence of noise, have been attained once regularisation is applied using the method by Tchernychev (2009). Fig-

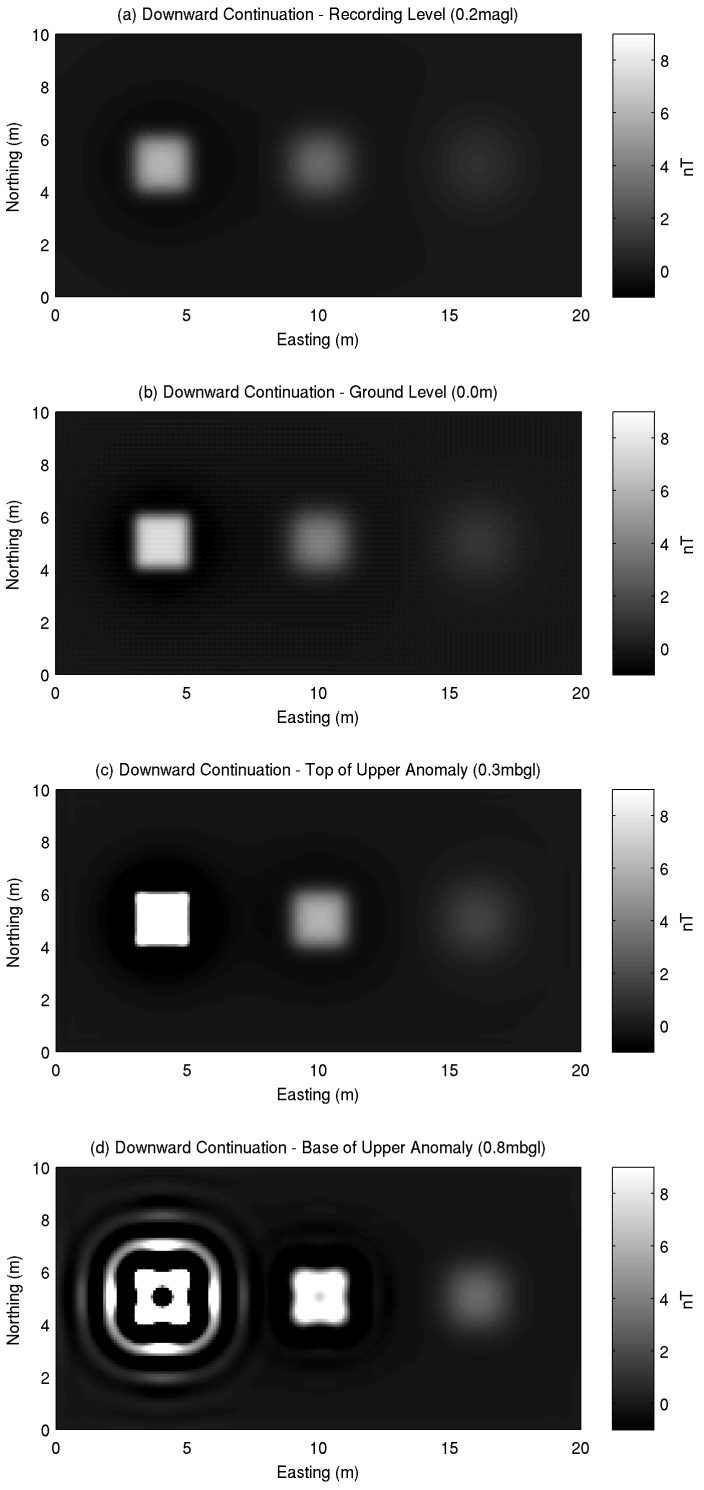

Figure 4.10: Downward continuation of the dataset shown in Figure 4.1, with regularisation imposed for stability. (a) Original recording level. (b) Downward continued 0.2m, to ground level. (c) Downward continued to 0.5m, the top of the shallowest prism. (d) Downward continued to 1.0m, the base of the shallowest prism.

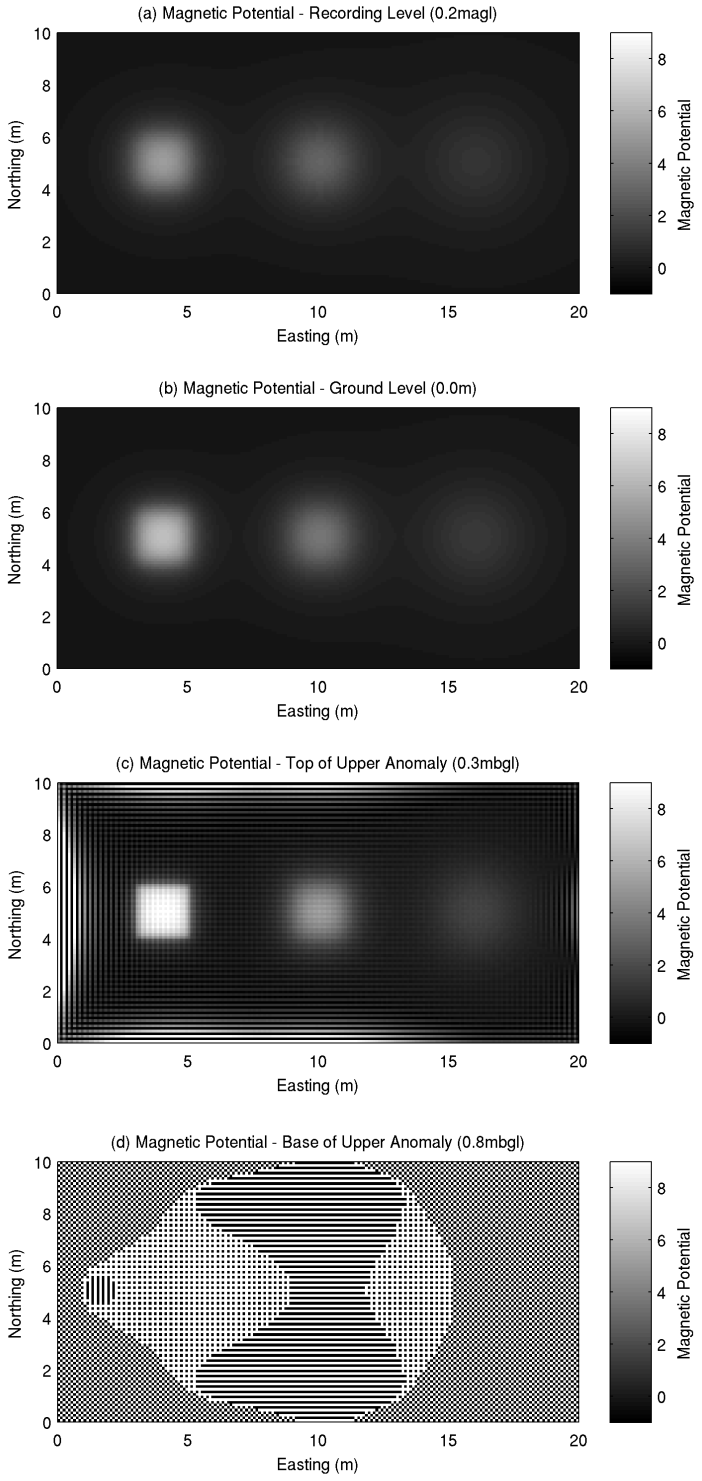

Figure 4.11: Downward continuation of the magnetic potential of Figure 4.1 (a) Original recording level. (b) Downward continued 0.2m, to ground level. (c) Downward continued to 0.5m, the top of the shallowest prism. (d) Downward continued to 1.0m, the base of the shallowest prism.

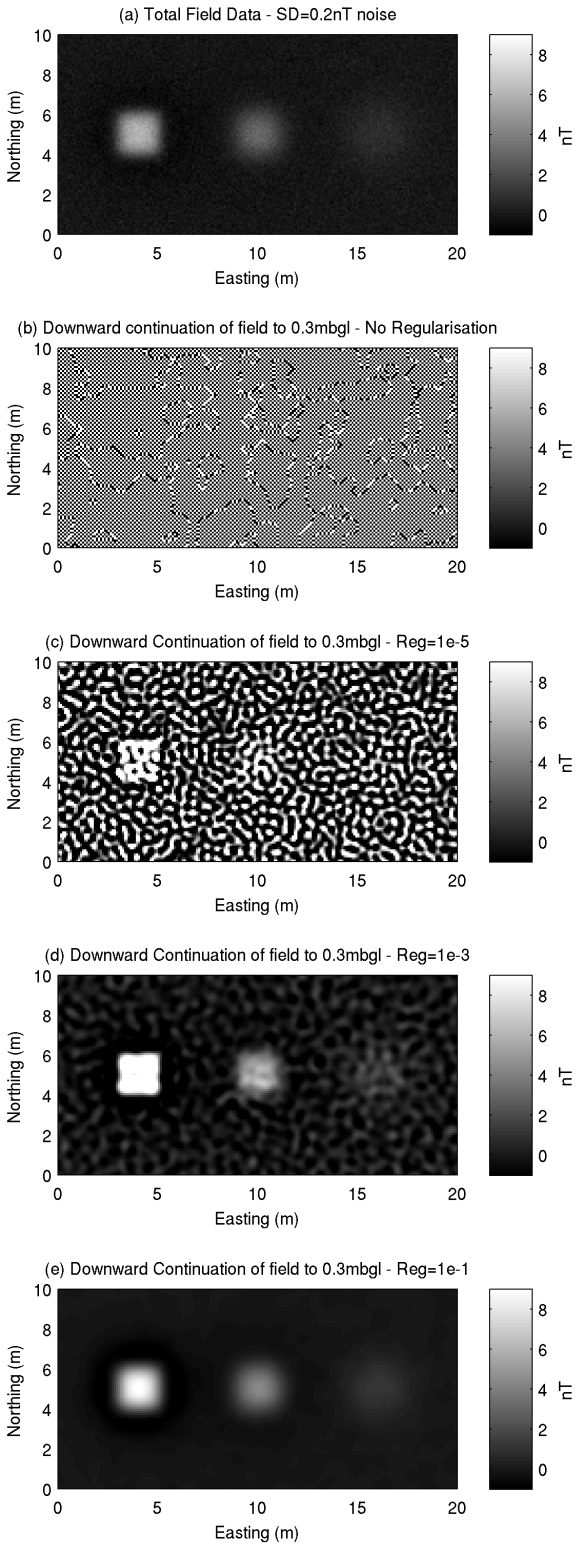

Figure 4.12: Downward continuation of the synthetic dataset with 0.2 nT standard deviation Gaussian noise added. (a) Synthetic dataset, generated by adding noise to Figure 4.1. (b) Downward continuation to 0.3 mbgl with no regularisation. (c) Downward continuation to 0.3 mbgl using a regularisation parameter of 1×10^{-5}. (d) Downward continuation to 0.3 mbgl using a regularisation parameter of 1×10^{-3}. (e) Downward continuation to 0.3 mbgl using a regularisation parameter of 1×10^{-1}.

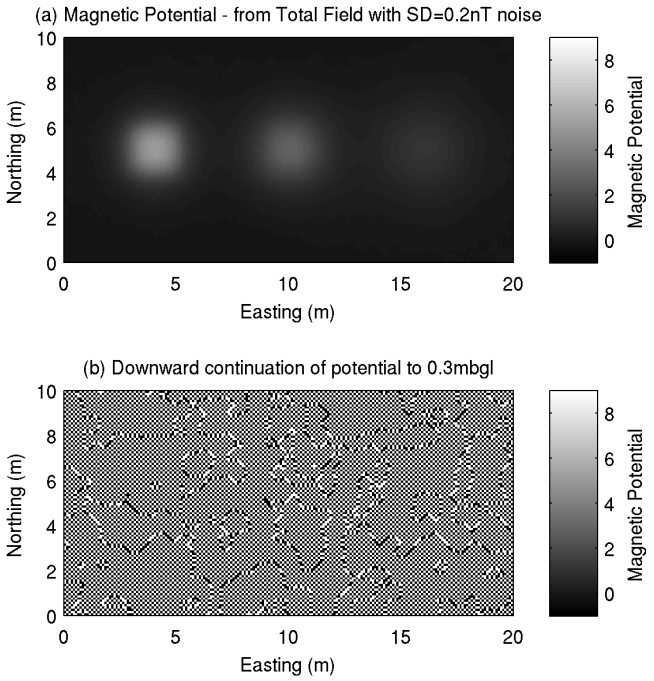

Figure 4.13: Downward continuation of the magnetic potential of the synthetic dataset with 0.2 nT standard deviation Gaussian noise added. (a) Magnetic potential of the total field response to Figure 4.1 with zero-mean Gaussian noise with a standard deviation of 0.2 nT added. (b) Downward continuation to of the magnetic potential to 0.3 mbgl, the top of the shallowest prism.

ures 4.12c, 4.12d and 4.12e show regularisation parameters of 1×10^{-5}, 1×10^{-3} and 1×10^{-1} respectively. As the regularisation parameter is increased the influence of the noise is reduced yet the sharpness of the recovered features is compromised. A full discussion of the effects of regularisation for inversion purposes will be discussed in Section 5.5.

A similar result is observed when downward continuing the magnetic potential using Equation 4.9. The magnetic potential of Figure 4.12a is presented in Figure 4.13a, yet once downward continued to the depth of the shallowest body the result has become uninterpretable (Figure 4.13b).

N	Type of Body
0	Contact (Edge feature)
1	Line of mass (sheet-like feature)
2	Line of dipoles (cylinder/pipe)
3	Point dipole (sphere of uniform magnetisation)

Table 4.1: Structural Index for use in Euler deconvolution of total-field magnetic data.

4.5 Euler deconvolution

Euler deconvolution was developed by Thompson (1982) for profile datasets and expanded to work with 2D gridded datasets by Reid *et al.* (1990), as a method for obtaining depth estimates to particular types of bodies from any potential field dataset. Examples applied to archaeological situations can be seen in Desvignes *et al.* (1999), Murdie *et al.* (1999) and Piro *et al.* (2007). For magnetic data, the method works by calculating estimates for the depth and lateral position of sources from analysis of the total field data, and its derivatives. These functions are linked by Euler's homogeneity equation, which is defined as:

$$(x - x_0)\frac{\delta T}{\delta x} + (y - y_0)\frac{\delta T}{\delta y} + (z - z_0)\frac{\delta T}{\delta z} = N(B - T) \tag{4.10}$$

where (x_0, y_0, z_0) is the location of the magnetic source, T is the total magnetic field at (x,y,z), B is the regional value of the total field, and N is the structural index. The structural index selected is very important to the final output parameters. The value is related to the attenuation with distance of the potential field, and varies depending on the shape of the source. The structural index for certain idealised shapes is shown in Table 4.1.

A dipole is represented by a structural index of 3, as the magnetic field from a dipole source is attenuated approximately at the rate of $1/r^3$. A source that can be represented by a line of dipole attenuates at a lower rate, at approximately $1/r^2$, and sheet-like features at a rate of $1/r^1$, hence the structural index of 1.

Euler deconvolution works by moving a square window of a specified size across the grid of data. Usually a pre-defined structural index is used, however there have been attempts to solve for the structural index as well (Gerovska & Arauzo-Bravo, 2003). Within each data window, there are a number of data points, and at each data point Euler's equation is solved using a least squares method. When applying Euler deconvolution in Geosoft Oasis Montaj (Geosoft, 2010a), the output for each position of the data window is a solution for the location in the x, y and z-directions, a background value for the field at this point, B, and uncertainties for the solution in both a lateral position and for depth, based on the standard deviation. These uncertainties are presented as a percentage of the depth calculated.

Many solutions can be calculated as the window is moved across the survey area, and some sort of selection criteria is often used to judge which solutions to keep. A common approach works on the basis that a single target will be detected multiple times by a moving window. Therefore, confidence about solutions can be obtained by locating the solutions that cluster at a point (Ugalde & Morris, 2010). An alternative to identifying clustering is to filter the results based on the uncertainty values, and on the distance, or offset, of the solution from the location of the data window from which it was calculated.

Figure 4.14 shows results for the synthetic dataset shown in Figure 4.1, without added noise. Only the default filtering for depth uncertainties less than 15% are applied here. A structural index of 1 has been chosen, and a window size of 20 samples. Figure 4.14a shows a plan view of the data area, with the colour representing the depth of the solutions. Figure 4.14b shows an X-Z profile. The colourbar in this case highlights the position of the solutions in the Y-direction. Figures 4.14c, 4.14d and 4.14e show the error in depth, offset and background, respectively. These are the main parameters for filtering. Looking at Figure 4.14c, it is apparent that generally the error in depth is low for most of the solutions ($< 4\%$). Trends in the solutions can be identified, with the error increasing as the solutions are closer to the surface. The error in depth can therefore be used to filter these near surface solutions, as well as poorly constrained solutions located between the central and eastern prisms. Figure 4.14d shows the offset between

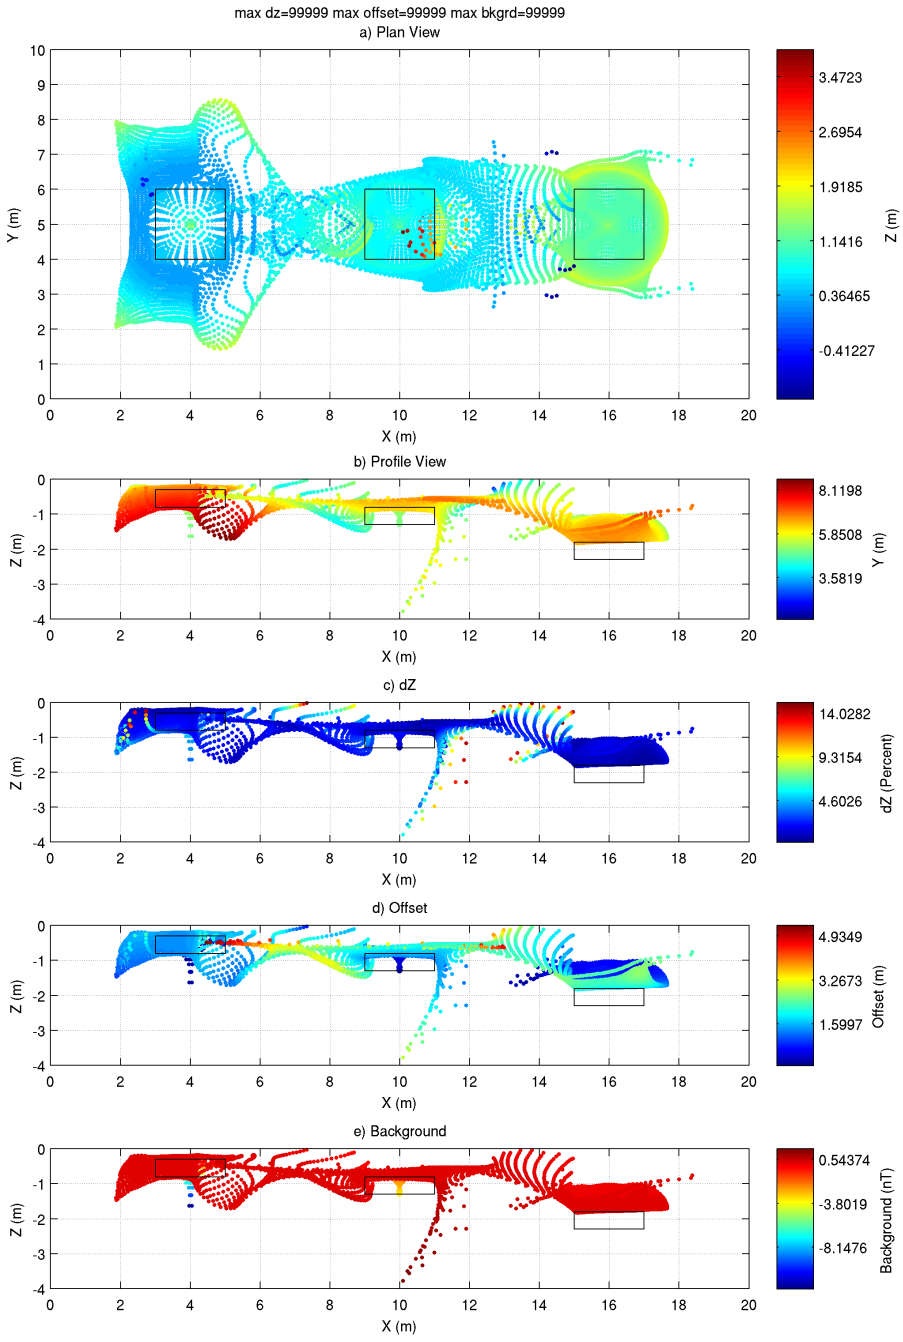

Figure 4.14: Euler deconvolution solutions for $N = 1$, with no synthetic noise added. (a) Plan view of all solutions with $dZ \leq 15\%$. (b) Profile view of all depth solutions. The colour scale shows the Y-axis. (c) Colour scale showing the error in depth of each solution. (d) Colour scale showing horizontal offset of each solution. (e) Colour scale shows the Background removed from each solution.

the windowed dataset and the solution locality. Due to the low amplitude and shallow nature of archaeological anomalies, they are unlikely to extend great distances across the survey area. It is therefore unlikely that solutions with a large offset between the data window and solution are caused by archaeological sources, and can therefore be discounted. This can be seen in Figure 4.14, where solutions with high offsets are located in the areas between the prism outlines, therefore this is a useful parameter to use as a filter. Figure 4.14e shows the background that is removed from each window prior to calculating the solution. Generally this value is quite low, as is expected due to the input being an anomaly dataset, without the absolute value of the Earth's field added. However, in areas located directly above the prism locations, the low gradients in the field lead to values for the background being calculated up to -10 nT. This has the effect of causing the solution to be positioned deeper than when a near zero-background is removed, and can be seen by the deeper than desired solutions below the western and central prisms.

Figure 4.15 shows the Euler deconvolution once the filtering has been applied. Only solutions with a depth error $\leq$2 %, an offset of $\leq$1 m, and background $\leq$0.2 nT are shown. This has resulted in solutions that plot towards the four corners of the western prism, at a depth just above the centre. These therefore are good results. For the central prism, solutions have also been obtained in the corner areas, however they plot towards the top of the prism's actual depths, and for the eastern prism, a relatively large scatter of solutions have been obtained surrounding the horizontal extents of the prism, yet all depth estimates are shallower than the actual depth.

The process has been repeated, this time using a structural index of 2. The results are shown in Figure 4.16. For the western prism the solutions once again are located towards the corners, slightly more towards the centre of the prism than the previous scenario, with depth estimates around the depth of the base of the actual prism position. A similar result is obtained for the central prism, while the solutions for the eastern prism are located towards the centre and top of the actual prism location. The positioning of the solutions towards the centre of this prism is likely to be reflecting its anomaly appearance more as a point source, rather than a feature

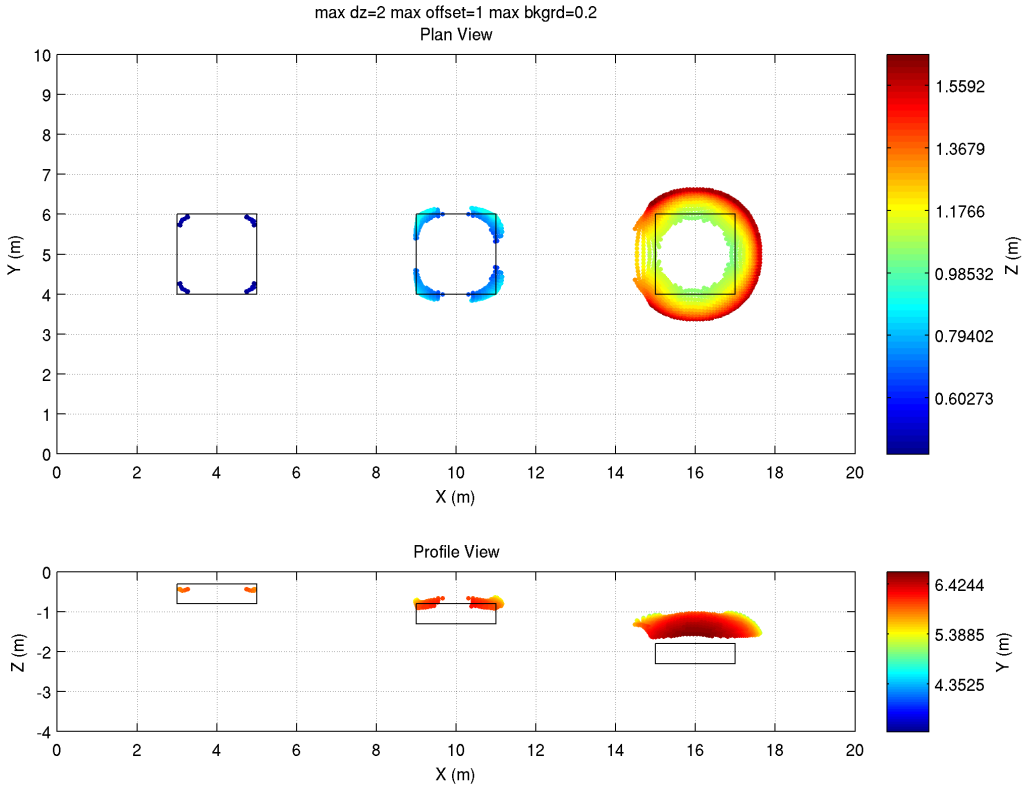

Figure 4.15: Euler deconvolution solutions for N=1, with no synthetic noise added. (a) Filtered dataset in plan view, colour representing depth (b) Profile (X-Z) View, colour representing Y.

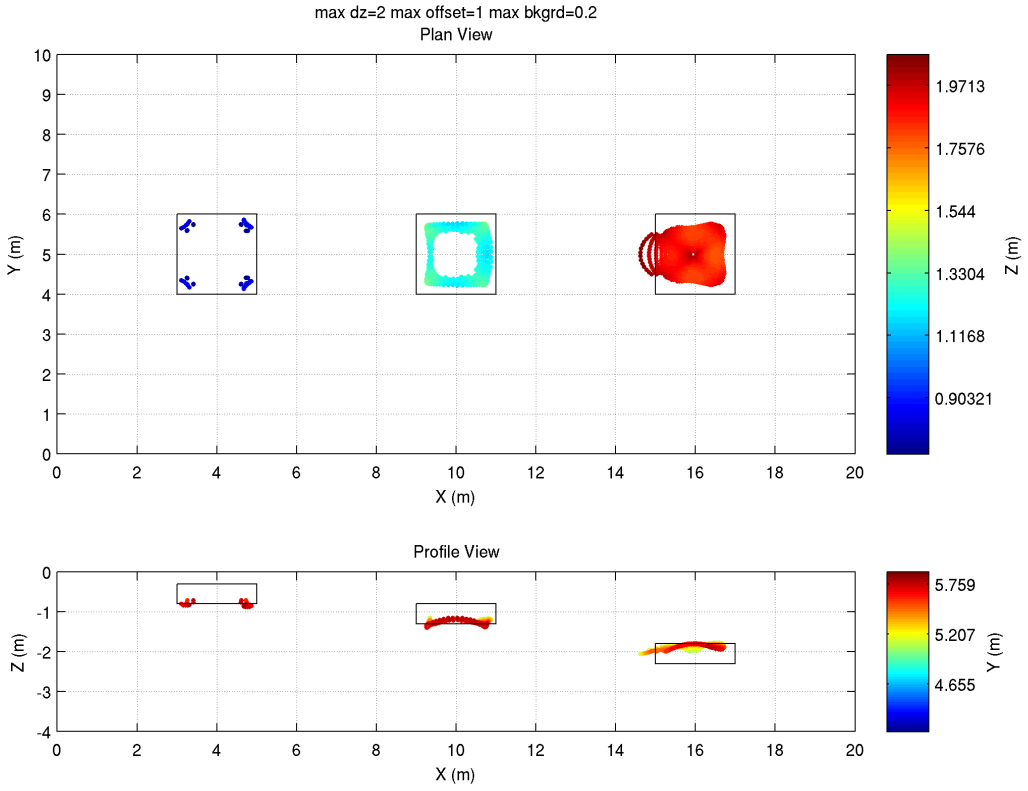

Figure 4.16: Euler deconvolution solutions for N=2, with no synthetic noise added. (a) Filtered dataset in plan view, colour representing depth (b) Profile (X-Z) View, colour representing Y.

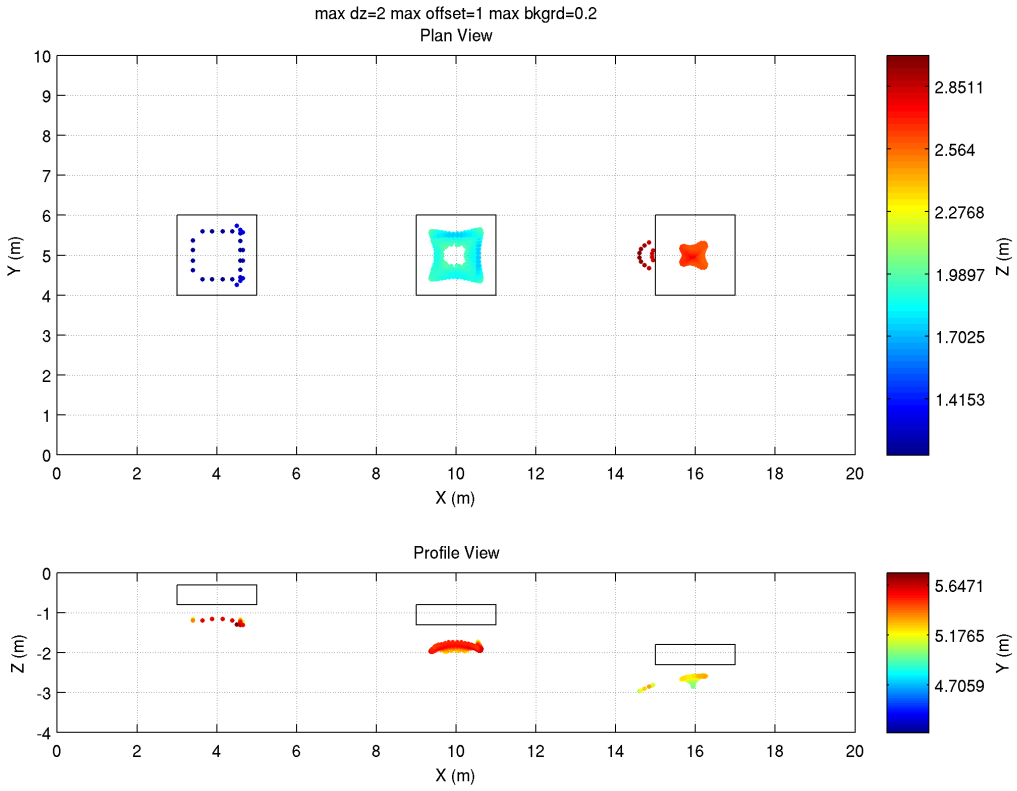

Figure 4.17: Euler deconvolution solutions for N=3, with no synthetic noise added. (a) Filtered dataset in plan view, colour representing depth (b) Profile (X-Z) View, colour representing Y.

that can be represented by infinite extents. The increased depth of the prism results in a lower amplitude, low-wavenumber response in the total field data.

Finally, the results using a structural index of 3 are shown in Figure 4.17. This has produced solutions for all three prisms which plot within the horizontal extents of the prism, yet at a depth $\sim$0.5 m deeper than the base of each prism. The results highlight how, if the source body does not conform to a specific idealistic shape then the appropriate structural index will change, depending on the depth to the source, even though the shape of each prism is identical. This is a phenomena that was identified by Ravat (1996), who noticed that steel drums buried shallowly are located accurately by a structural index of $\sim$2, with ones buried below a depth of 1 m, requiring a structural index of $\sim$3.

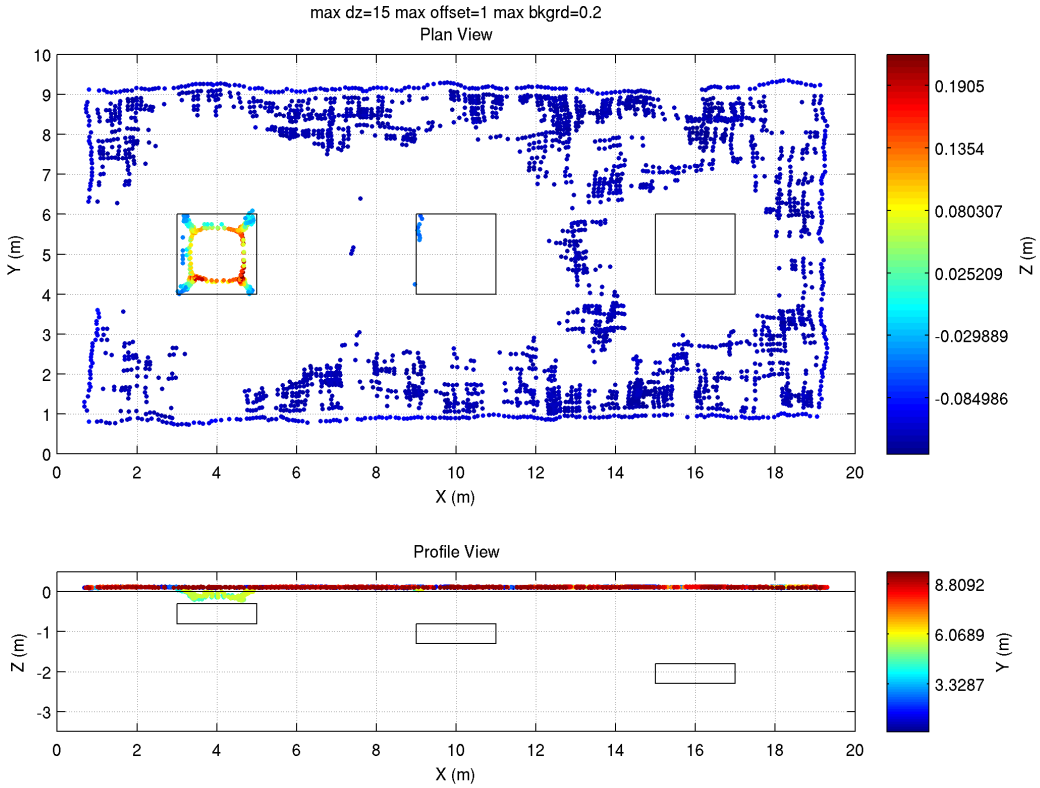

Figure 4.18: Euler deconvolution solutions for SI=2, with 0.2 nT standard deviation Gaussian noise added. (a) Filtered dataset in plan view, colour representing depth (b) Profile (X-Z) View, colour representing Y.

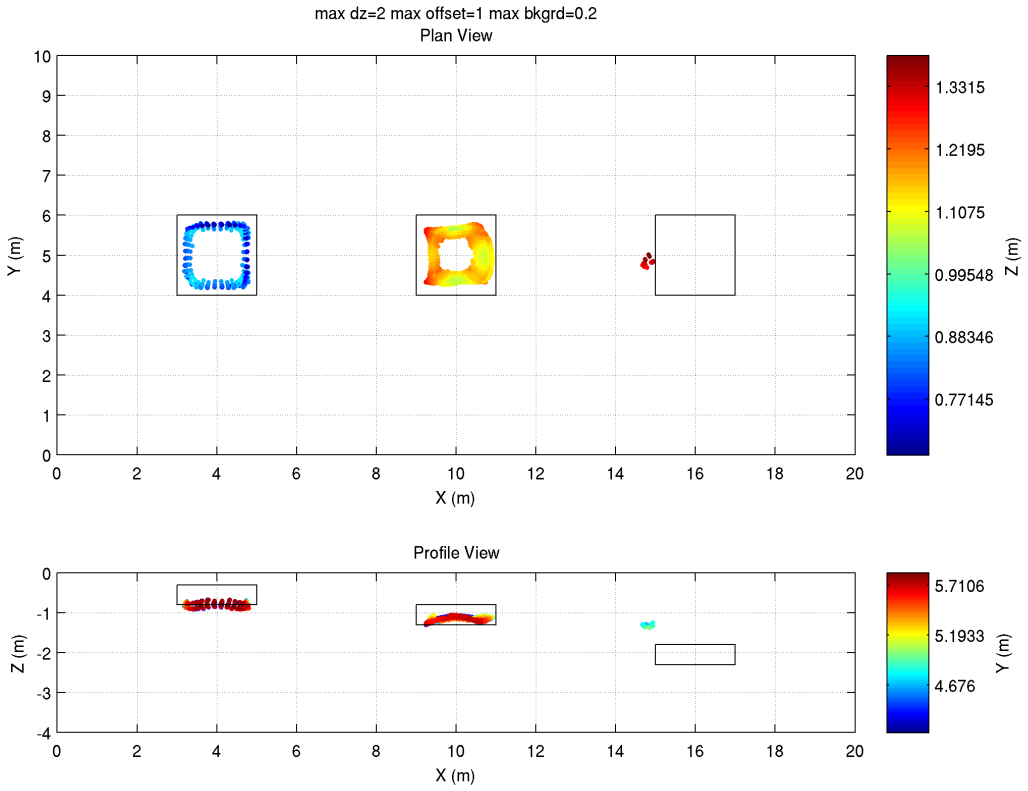

Figure 4.19: Euler deconvolution solutions for N=2, with 0.2 nT standard deviation Gaussian noise added, upward continued 0.2 m. (a) Filtered dataset in plan view, colour representing depth (b) Profile (X-Z) View, colour representing Y.

The effect of the presence of noise in the dataset is shown in Figure 4.18. Here Gaussian noise of 0.2 nT standard deviation has been applied to the synthetic dataset. The results show the solutions for a structural index of $n = 2$. As identified in Chapter 3, the noise has a large effect on the calculated gradients, which in turn have affected the Euler solutions. The filtering has been relaxed to allow uncertainties of depth up to 15% in Figure 4.18, and the majority of the solutions have been calculated as being above ground. A few solutions plot in the area around the top of the shallowest, western prism, however no solutions have been obtained for the central or eastern prisms.

Upward continuation

Upward continuation is a commonly applied approach to overcoming the problems associated with noise, prior to performing Euler deconvolution. Salem & Ravat (2003) suggested that in the presence of noise the results of the Euler deconvolution are unstable and incorrect, and small levels of upward continuation are required to produce stable and accurate results. Upward continuation is the opposite of downward continuation, as the field is calculated as it would be further away from the sources. This process has the effect of suppressing high-wavenumber features and producing a smoother dataset. The solutions obtained by upward continuation of the synthetic case with noise, by 0.2 m, prior to Euler deconvolution, are shown in Figure 4.19. An improvement is observed as the high-wavenumber noise is largely filtered out, resulting in smoother gradients, and more accurate solutions. However, Salem & Ravat (2003) also noted that upward continuation is not applicable above a certain height, as the higher continuation levels cause the anomalies to become smaller in amplitude, and broader, to the extent that nearby anomalies begin to overlap. Once the gradients in the data window are recording the effects of more than one body, the results will be invalidated. Salem & Ravat (2003) therefore recommend that the ideal level of upward continuation is the smallest which produces consistent estimates of causative body localities.

While upward continuation is shown to be a useful tool in filtering the data prior to Euler deconvolution, it further highlights the problems associated with non-idealistic shaped bodies, where the optimum structural index will vary as distance between source and receiver increases. At great enough heights above the surface, it is possible that every anomaly that has finite dimensions will appear to be represented as a dipole, or $n = 3$. To study this, the dataset shown in Figure 4.1, without synthetic noise added, has been upward continued to a height of 3.0 m above the original instrument level. A structural index of $n = 3$ is chosen, and the average depth from remaining solutions after filtering, representing each of the three prism, has been calculated at each continuation level. Should the source body be an idealistic dipole feature,

upward continuation should have no effect on the final depth solution obtained. Figure 4.20 shows the results for each prism, with the western prism shown on the left, central prism in the middle, and the results from eastern prism shown on the right. Initially, for all prisms, upward continuation for small distances results in a decrease of the depth estimate to each prism. The initial estimates are all deeper than the actual base of each prism, as seen in Figure 4.17. After continuing the field up to $\sim$1 m for the western prism and $\sim$2 m for the central prism, the structural index of 3 begins to be appropriate and the depth estimates are positioned within the body's actual extents. Continuing the field beyond these heights results in an unexpected increase in the depth estimate. This is the point where the overlapping effect of the magnetic responses begins for these particular anomalies. At these heights the anomalies also represent minimal amplitudes for the total field. At an upward continued height of 1.6 m, the amplitude of the response for the eastern prism is just 0.22 nT, and falls to 0.09 nT when it is continued to 3.0 m above the measurement level.

The results show that overlapping anomalies are not the only problem associated with upward continuing to ever increasing levels. The assumptions on the optimum structural index that will apply at each continuation level also has to be taken into account.

4.6 Discussion

As discussed in Chapter 3, derivative based methods are being increasingly adopted in the processing and interpretation of aero-magnetic datasets. In these situations, often the distance to the source bodies are large, and the amplitude of the susceptibility contrast high, resulting in assumptions and simplifications of the targets to simple geometrical features. The tilt angle depth estimation is proposed as a quick and easy process to allow the depth to a contact-like feature to be calculated, however, due to the very near-surface nature of archaeo-magnetic surveys, it is unlikely that an archaeological feature can be simplified to such a source shape. As the target moves away from this ideal-case scenario, the method breaks down, and underestimates the

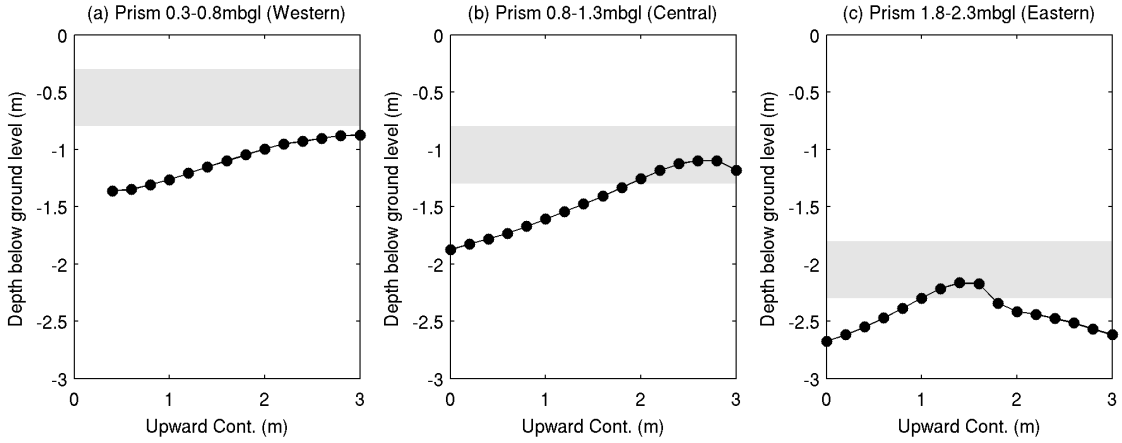

Figure 4.20: Effect of upward continuation on the Euler deconvolution results for synthetic data with no noise added. The actual depths of each anomaly is shown in grey. The depth is an average of all the filtered solutions relating to that anomaly. The structural index of $N = 3$, and filter parameters are *dz<1*, *offset<1* and *background<0.2* have been applied. As the dataset is continued upwards the Euler solutions appear shallower. It is likely that at a certain height all anomalies will appear as dipolar features and can be represented by a structural index of $N = 3$. However, as the field is continued upwards, the amplitude of the total field data is significantly reduced, and a point will be reached where the anomalies can no longer be detected.

actual depth. Even if the source could be assumed to be a contact feature, the problems associated with noise levels in calculating the tilt angle are likely to be a limiting factor without some prior filtering (see Chapter 3). Typical noise content of the data is also likely to be the limiting factor for the implementation of the SPI method to archaeology, which once again shows highly promising results for a zero-noise scenario.

Euler deconvolution is a method that has been applied to archaeological datasets in the past. This approach is also affected by the noise content of the data, however, as demonstrated in previous studies, a good method of filtering out the effects of this noise is to upward continue the signal prior to solving Euler's homogeneity equations. This has been shown to be a good procedure for producing results with increased confidence, however an essential assumption to calculating accurate readings is choosing the correct structural index, once again from a list of idealistic geometrical shapes. Although it is possible for the structural index for a body to lie between one of these members, it has been shown here that when upward continued, the optimum structural index of a non-idealistic bodies will vary with continuation distance, or

depth below the surface. Without prior knowledge it is unlikely that an optimum would be known, and by examining various choices of N the user can often obtain uncertainties larger than the derived depth estimate.

The methods presented here can be used to produce good results for interpretation of depth, however, very careful implementation is required, often using knowledge that is not likely to be available. The increasing uptake in aero-magnetic interpretation is unlikely to be matched by archaeo-geophysicists, as noise content and assumptions regarding the source body shape are likely to hinder implementation. Euler deconvolution has proven to be the most practical for application to archaeo-magnetic datasets. As long as the assumptions are noted and the solutions filtered appropriately, estimates for the depth to simple shaped sources can be quickly derived.

Chapter 5

Inversion of synthetic archaeological-type magnetic datasets

5.1 Introduction

Although interpretation of the magnetic field through either visual or derivative-based means can tell us about the horizontal location, and estimate the depths of simple structures, ideally we want to know about the distribution of magnetic susceptibility that produces the anomalies, as these can then be interpreted in terms of buried archaeological features. This is where magnetic inversion is required. The computing power required to invert magnetic data is quite high, and this has led to several examples where a parameter-based approach has been applied, simulating the archaeological structures as simple bodies or layers, and adjusting parameters to fit to the observed data (Powell, 1967; Herwanger *et al.*, 2000). An alternative approach is a generalised inversion scheme, which involves representing the subsurface with a mesh of small cells, which are iteratively inverted to produce a model where the physical distribution of magnetic susceptibility is calculated within this mesh (Li & Oldenburg, 1996; Pilkington, 1997).

5.2 The inverse problem

When the subsurface distribution of magnetic susceptibility is known, it is possible to calculate what the observed data would be if you carried out a survey over the area by forward modelling. The forward modelling operator will include the relevant physical equations and the spatial distribution of the causative bodies and observation points. It is a well-posed problem and leads to a unique solution.

Inverse modelling is more complicated than forward modelling. Typically an archaeo-magnetic survey results in a grid of data collected just above the surface. What is required for interpretation is information about the subsurface properties, which caused the readings that were observed (see Figure 5.1). Inverse modelling is an iterative procedure that calculates a model, which, if forward modelled would match the observed data to within pre-defined limits.

The problem with inverse modelling is the non-uniqueness of the final model. There will be a number of different models that fit the observed data equally well as only a finite number of data are collected. Also, the data will include a noise-component, due to instrument inaccuracy, positional error or another external source. In order to identify a single model to select, the user needs to provide *a priori* information, or apply constraints to the inversion parameters to ensure they are within realistic values. For the MAG3D code that will be investigated here, this is applied to the inversion through the use of the model norm.

5.3 Objective function

A simple and commonly used approach to constraining the model output from the inversion routine is to attempt to recover a smooth variation in the susceptibility throughout the model. This is often used where no other *a priori* information exists, and is based on the theory of Occam's Principle, where the simplest explanation of the observations is the most probable. This

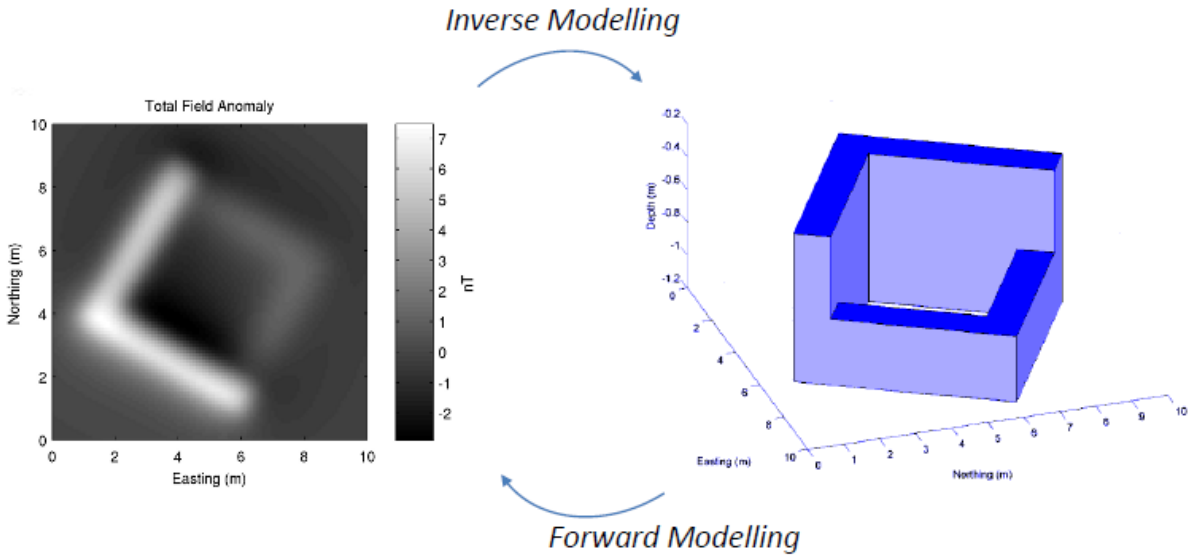

Figure 5.1: Forward and inverse modelling. Typically, total field magnetic data will be collected just above the survey, yet the information desired from the survey is what is below the survey causing the observations. Inverse modelling uses the observation dataset, to produce a model of the physical properties of the subsurface. Forward modelling takes this model and produces a synthetic dataset that would be collected over this model. An inversion works by iteratively changing the model until the calculated forward model matched the observations within a pre-defined misfit.

approach, however, may not be appropriate for archaeological purposes. Anomalies present in archaeo-magnetic data are usually caused by abrupt changes in the subsurface properties caused by human activity placing material with differing properties to the surrounding soil, such as the case with foundation materials, or infills where ditches have been dug and over time soil with contrasting magnetisation fills the ditch. Either scenario is unlikely to result in a gradual change in magnetisation, and the smooth-model constraint may not be appropriate (Piro *et al.*, 2007).

5.3.1 Model norm

Should knowledge of the subsurface be available, it may be desirable to manipulate the model norm to calculate a model as close as possible to this. MAG3D does this by including a reference model in the model norm. The model norm is defined by Li & Oldenburg (1996) as:

$$\phi_m(m) = \alpha_s \int_v w_s\{w(z)[m - m_{ref}]\}^2 dv + \alpha_x \int_v w_x \left\{ \frac{\partial w(z)[m - m_{ref}]}{\partial x} \right\}^2 dv$$
$$+\alpha_y \int_v w_y \left\{ \frac{\partial w(z)[m - m_{ref}]}{\partial y} \right\}^2 dv + \alpha_z \int_v w_z \left\{ \frac{\partial w(z)[m - m_{ref}]}{\partial z} \right\}^2 dv \quad (5.1)$$

The constants α_s, α_x, α_y, and α_z control the relative importance of each term. The first term of Equation 5.1 controls the smallness of the model, which is how closely the model matches the reference model, m_{ref}, or a zero-halfspace if no reference model is provided. If α_x, α_y, and α_z are set to equal zero, the inversion will produce a 'small' model that stays as close as possible to the provided reference model. The latter three terms control the smoothness in the x, y and $z-$ directions. Setting high values for α_x, α_y and α_z, will force the inversion to penalise rapid changes of susceptibility in each direction. Should there be a preference for constraining smallness or smoothness between individual cells in the mesh, this can be controlled by w_s, w_x, w_y, and w_z. These are spatially dependent sensitivity weighting functions, and unlike the α-parameters which are single values, w-parameters are declared for each boundary within the mesh. The function $w(z)$ is a depth weighting function that will be discussed in more detail below in Section 5.4.

5.3.2 Data norm

To solve the inverse problem it is important that the model norm is designed to produce the type of model desired, but also that the forward modelled response from the final model matches the observations measured at the surface. Each datum includes information on the Earth's magnetic field at that point, plus a quantity of noise. During the inversion it is important to find a model which satisfies the data, without trying to fit the noise by placing anomalous susceptibilities in the subsurface, thus producing an unrealistic solution. Quantifying the noise level of the data is therefore important. Noise, in terms of an inversion routine, is not just the traditional types of

noise associated with collecting the observations, but also anything to do with the discrepancies between representing the subsurface as a series of cells and the real Earth structure (Oldenburg & Li, 2005).

To define the measure of the fit between the calculated forward modelled response of the model to the set of observations measured during the survey, a data norm is declared. This is stated by Oldenburg & Li (2005) as:

$$\phi_d = \sum_{j=1}^{N} \left(\frac{d_j^{obs} - F[m]}{\epsilon_j} \right)^2 \tag{5.2}$$

where $F[m]$ is the calculated forward modelled data, d_j^{obs} are the observed field measurements, and ϵ_j is the standard deviation of the $j'th$ datum. The noise is assumed to be Gaussian and independent of the datum amplitude.

5.3.3 Regularisation

The inversion seeks to produce a final model by minimising an objective function. The objective function incorporates both the model norm and the data norm, therefore the inversion will satisfy the constraints placed on the model, be it smallness or smoothness, as well as fitting the data. The objective function is defined as:

$$\phi(m) = \phi_d + \beta\phi_m \tag{5.3}$$

where β is known as the trade-off, or regularisation parameter. Changing this parameter allows a trade-off between fitting the data, and fitting the model requirements. A low value of β will produce a model that preferentially fits the data, while a high value of β will place a higher emphasis on meeting the requirements of the model-norm rather than a good fit to the observed

data. If the level of noise in the data is over-estimated or under-estimated in the data-norm, it is possible to compensate for this in the choice of regularisation parameter. Selecting the optimum value for β will be discussed in Section 5.5.

5.4 Depth weighting

Should the inversion be performed with no depth weighting function added to the objective function, the anomalous values of susceptibility that allow the model to fit the observed data will always be placed towards the top of the mesh. Positioning the anomalous susceptibilities here requires the smallest values, and therefore the model will be the closest to matching the halfspace should no reference model be present. Placing anomalous susceptibilities deeper would require a larger spatial area, or higher susceptibility to generate the same anomaly at the recording station.

To overcome this Li & Oldenburg (1996) devised a depth weighting function that matches the decay of the sensitivity kernel of the mesh being used. A column of cells is taken from the centre of the mesh, and the decay in the response of an anomalous body placed at every cell in this column is measured. By devising a depth weighting function that is the inverse of this, it becomes equally likely that the inversion will place an anomalous susceptibility at any position throughout this column.

Li & Oldenburg (1996) identified the function that matches the decay kernel to be of the form $(z + z_0)^3$. The exponent is set to 3, as the decay of the magnetic anomaly of a small cell is approximately related to that of a dipole where the decay of the magnetic field is a function of distance cubed. The value of z_0 is a function of the cell size and height of the recording instrument. The depth weighting function used in MAG3D is therefore defined by Oldenburg & Li (2005) as:

$$w^2(z) = \frac{1}{(z + z_0)^\gamma} \tag{5.4}$$

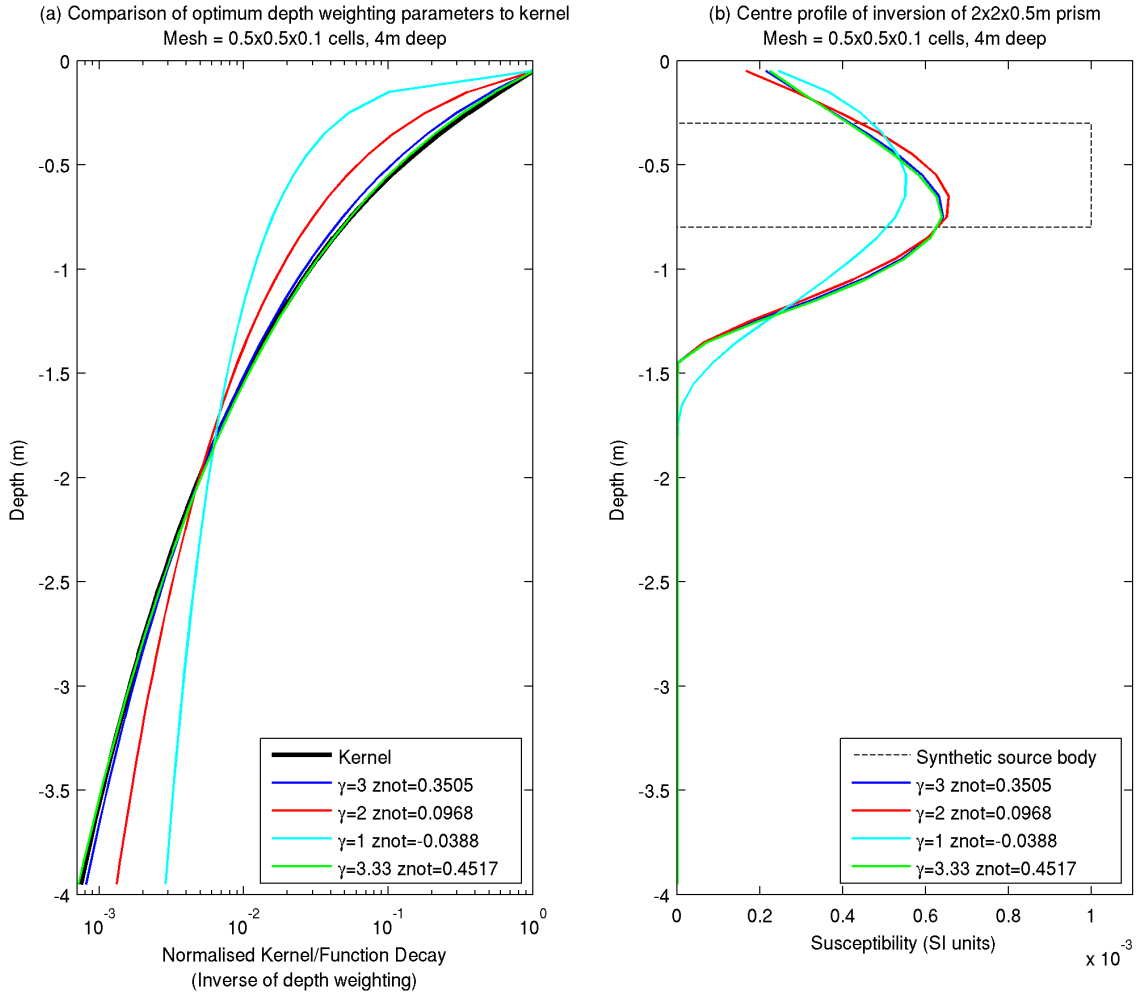

Figure 5.2: Optimum depth weighting parameters to the kernel response. (a) Depth weighting parameters for various values of γ, for a 4m deep mesh of cells $0.5\times0.5\times0.1$ m. The optimum parameters for values of $\gamma = 1, 2$ and 3 are shown as well as an optimum where γ and z_0 were allowed to vary. Solutions derived by Nelder-Mead simplex direct search. (b) Profiles through the centre of a synthetic $2\times2\times0.5$ m prism, inverted using the depth weighting parameters shown in a).

where γ in this case is the exponent.

The decay kernel for a typical archeo-type mesh, of cells $0.5\times0.5\times0.1$ m, to a depth of 4 m, is shown in Figure 5.2a as the solid black line. In order to design the depth weighting function to be used by MAG3D, four functions of form $(z+z_0)^\gamma$ are also shown. For the blue, red and cyan lines, the value of γ has been set at 3, 2 and 1 respectively, and the value of z_0 optimised using a Nelder-Mead simplex direct search (Lagarias *et al.*, 1998) within MATLAB. It is shown in Figure 5.2a that a γ value of 3 closely matches the decay kernel, with $z_0 = 0.3505$. Optimised

functions of $\gamma = 2$ (line of dipoles) and $\gamma = 1$ (thin sheet) do not create as good a match to the kernel, and if used as the weighting function would force the inversion to preferably place non-zero susceptibility values elsewhere in the column.

The optimisation was repeated, this time allowing both the value of γ and z_0 to be optimised by the search routine. The output of this was a function of $(z + 0.4517)^{3.33}$. However, this function offers only a slight improvement over the $(z + 0.3505)^3$ function.

Inversions of a 2×2×0.5 m prism placed 0.3 to 0.8 meters below ground level (mbgl), with positivity and smoothness constraints applied, were performed for all the depth weighting functions identified in Figure 5.2a. Profiles through the centre of the body are presented in Figure 5.2b. If the depth weighting is satisfactory, giving equal chance of the column to be a non-zero value, then the depth of the body is controlled by the lateral extents of the anomaly and the constraints applied to the inversion in the model norm. In this case, Figure 5.2b shows that the peak susceptibility is positioned at approximately the same depth (0.5-0.7 mbgl) for all scenarios. It is noticeable that for the depth weighting of $\gamma = 1$, the inversion has placed higher susceptibilities at shallower depths. Overall however, all depth weighting functions have produced satisfactory solutions for this particular scenario, and changing the value of γ has not significantly affected the final result.

5.5 Choosing a regularisation parameter

The unconstrained optimisation problem is solved in MAG3D by minimising Equation 5.3, where β is the regularisation parameter. In default mode, MAG3D calculates the model which minimises ϕ_m subject to a pre-defined misfit of ϕ_d. This is appropriate if the errors in the data are known, however this is nearly always not the case when dealing with real field data. The trade-off between the smoothness of the model, and the fit of the data can be controlled by adjusting β, therefore estimating β is equivalent to estimating the level of noise in the data.

Should β be too high, the inversion will not attempt to place the fine structure into the model that will allow it to fit the data, and the result will be a model that does not reflect all the information contained in the data. Should the value of β be too low, the inversion will place excessive structure into the model in order to fit data which are merely errors.

An example of underfitting and overfitting the data is demonstrated in Figure 5.3. Figure 5.3a shows a scenario where the data have been underfit. The forward response from the final model does not match the observed data well, and therefore the subsurface model is not a suitable representation of the earth structure. Further adjustments to the model in Figure 5.3b, particularly in the near surface have resulted in a much improved fit to the observed data. Overfitting the data is shown in Figure 5.3c. The difference between the observed and calculated datasets is the smallest of the three models, however, if noise is present in the observed data, attempts to fit this with the calculated data will result in excess and unrealistic structures being placed into the model. Here it can be seen that attempts to fit the noise have resulted in separating both anomalies into less homogeneous features, with a particular increase in structure very close to the surface where high-wavenumber signals can be produced.

There are a number of methods that can be used for identifying an optimum value for the regularisation parameter (β). These include the chi-squared method, analysis of the L-curve, and generalised cross-validation.

5.5.1 Chi-squared method

If assumptions are made that the noise contained within the data has a known standard deviation ϵ_j, has Gaussian distribution with a zero-mean, and that the noise affecting each datum is uncorrelated across the survey area, then the value of β to be used in the inversion can be calculated.

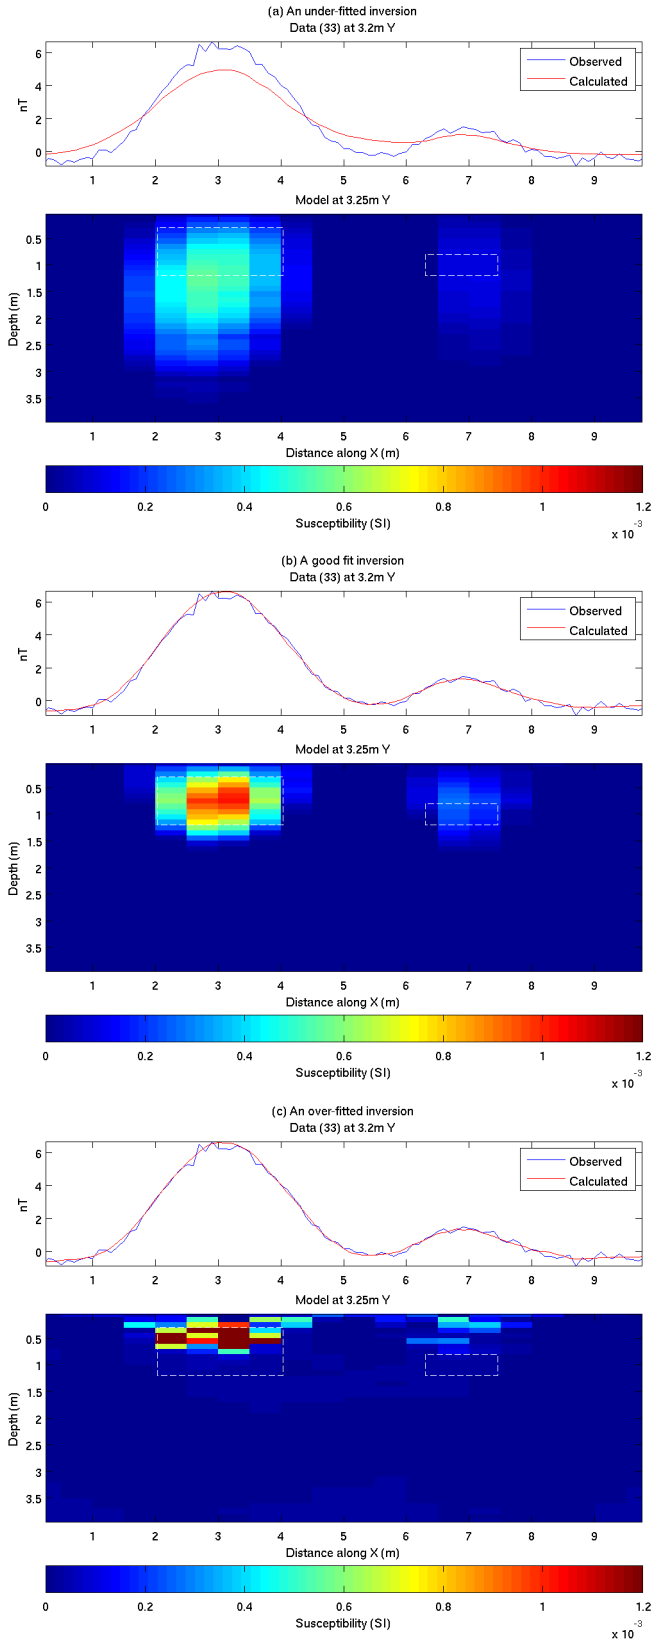

Figure 5.3: Profiles through inversion results for (a) an under-fit example ($\beta = 1 \times 10^{12}$), (b) a good fit example ($\beta = 1 \times 10^{10}$), (c) an over-fit example ($\beta = 1 \times 10^{8}$). Note the under-fit model does not produce data that fit the observed values, while the over-fit example has included excessive structure in an attempt to fit the noise in the observed dataset.

When the assumptions above are true, the data norm (Equation 5.2) should equal the number of data points *N* (Oldenburg *et al.*, 1993). This is because ϕ_d becomes the chi-squared variable χ^2_N with N degrees of freedom (Li & Oldenburg, 1996). The value of β is then selected that minimises ϕ_m subject to ϕ_d equalling *N*.

While this method is relatively quick to run, it requires a lot of information concerning the noise to be known or assumed. Should an incorrect value for ϵ_j be prescribed to the inversion, the result will not be the optimum model, and will either be under-modelled (if the prescribed standard deviation is too high) or over-modelled (if the standard deviation is too low). The method is therefore particularly useful for synthetic modelling where the noise can be controlled, but may not produce the optimum model for real data scenarios, where parameters are likely to be unknown.

5.5.2 L-curve

A common approach to identifying the value of β to be used in the inversion is to create an L-curve (Lawson & Hanson, 1974; Hansen, 1992). The L-curve is created to visualise the trade-off between the data misfit and the model norm.

For a range of β, the data norm (ϕ_d) is plotted against the model norm (ϕ_d). The plot typically consists of two parts, forming an L-shape, where no regularised solutions can exist in the portion of the graph below the curve. For high values of β, ϕ_d can be reduced greatly with only small increases in ϕ_m. This is where small amounts of structure can be added to the model and generate large improvements to the fit to the data as the low-wavenumber features in the data are being matched. Where β is small, large increases in ϕ_m are required to generate only small decreases in ϕ_d. This is the region where excessive structure is being added to the model to fit the high-wavenumber noise in the data. This additional structure is unlikely to be realistic.

In between these zones, there will ideally be a sharp corner, or 'knee'. This is the point where the coherent signal in the data has been modelled, yet the noise has still been misfit. Using the L-curve method, highly correlated errors will still lead to an increase in ϕ_m, as there is no way of distinguishing correlated errors from signal produced by real source bodies (Hansen & O'Leary, 1993). Hansen (1992) proposed the β at the point of maximum curvature of the $log(\phi_d)$,$log(\phi_m)$ data as the chosen regularisation parameter.

Figure 5.4 shows multiple L-curves for the compound synthetic anomaly presented in Figure 3.1 with zero-mean Gaussian errors of standard deviation of 0.2 nT applied. Where the correct 0.2 nT error has been applied to each datum in the datafile, the generated L-curve is represented by the solid line. As the correct error has been assigned in the datafile using the chi-squared approach attempts to minimise ϕ subject to $\phi_d = N$, in this case $\phi_d = 10201$. It can be seen in Figure 5.4 that the intersection of the solid line and $\phi_d = 10201$ is towards the 'knee' of the L-curve, and therefore the correct regularisation parameter has been obtained to achieve the solution. However, should the error applied to each datum in the datafile not be correctly assigned, this will not be the case. The dotted line in Figure 5.4 shows where an overestimated noise level of 0.5 nT has been assigned to the datafile, where only SD=0.2 nT has actually been applied to the data. This has resulted in shifting the L-curve down the y-axis. Now, the intersection between the L-curve and $\phi_d = 10201$ results in a model that has underfit the data, and the chosen value of β is not appropriate. Plotting the full L-curve allows this observation to be made, and also for a more appropriate value of β to be selected from the solution closest to the corner. The opposite effect can be observed where the noise level has been underestimated at 0.1 nT in the datafile. This has the effect of shifting the curve up the y-axis, and now a solution of $\phi_d = 10201$ is not achievable, yet a suitable β can be chosen by analysis of the L-curve.

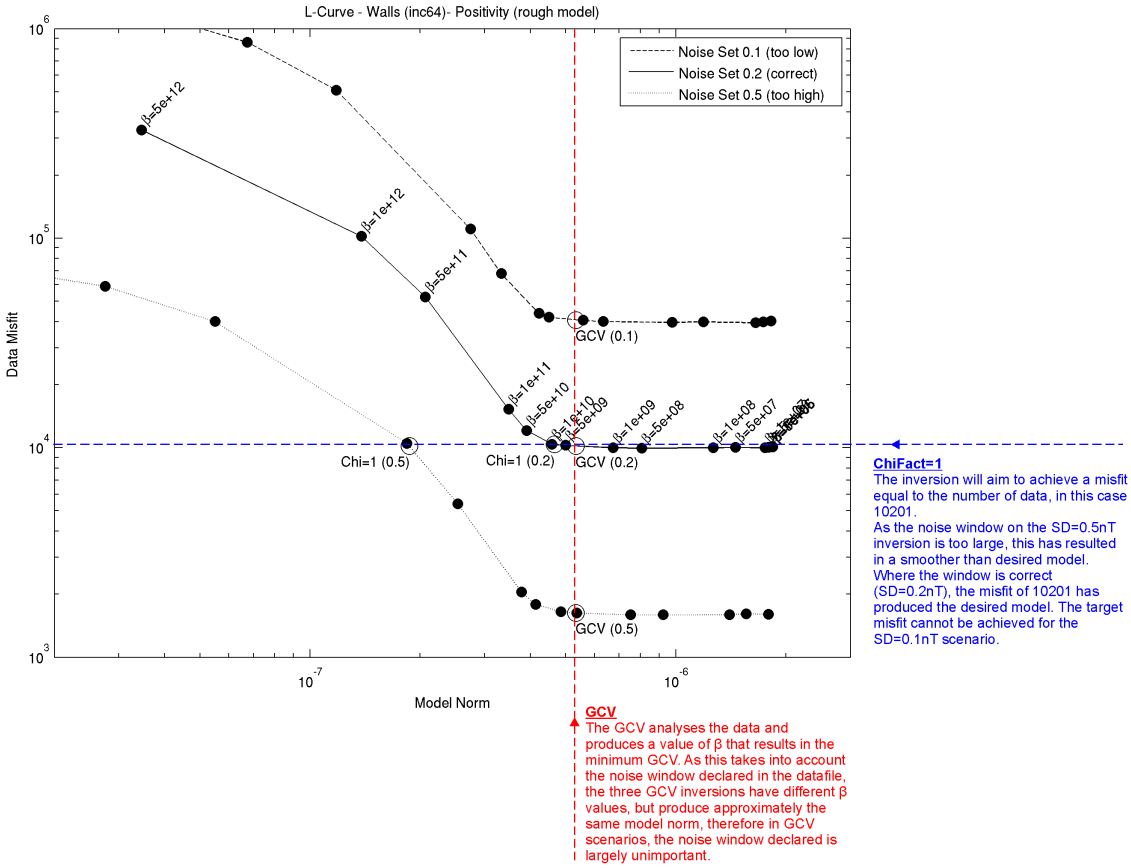

Figure 5.4: L-curves for the correctly stated noise (solid line), overestimated noise (dotted line) and underestimated noise (dashed line). Solutions found where *ChiFact=1*, will result in a data misfit equal to *N*, which is shown by the blue line. Solutions found using the GCV technique all recover the same model, despite the difference in noise window. This is due to the GCV outputting a β that is dependant on the noise window declared. All solutions will lie on the red line, despite the noise level is declared by the user.

5.5.3 Generalised cross validation

A third method of choosing an appropriate value for β is generalised cross validation (Golub *et al.*, 1979). Generalised cross validation (GCV) tests that the β chosen does not make the inversion too sensitive to any individual data point, and therefore should be able to predict missing data points. Cross validation is done by removing a datum, and calculating a solution to the now $(m-1), n$ problem, using a range of β values. The predicted data value for the missing datum is calculated by forward modelling and compared to the original. GCV differs from cross validation as it is independent of the particular ordering of the data (Hansen & O'Leary, 1993). The β that produces the smallest GCV is chosen as the most appropriate for use in the final inversion.

The GCV technique is known however, to not always be robust. Hansen (1992) shows that the GCV function can often be very flat, making a numerical estimate of the minimum error prone, and Hansen & O'Leary (1993) discuss the problems associated with different type of noise while calculating the GCV. While the GCV will provide a good estimate of β in the presence of random noise, once again, it cannot distinguish between correlated noise and signal. Therefore in the presence of correlated noise the estimate of β will be underestimated, and the final model will overfit the data.

GCV analysis of the data is shown in Figure 5.5, for the three scenarios. The selected value of β is the one which produces the minimum GCV. Each scenario has resulted in a different value of β, however when these values are plotted on the L-curve, they all plot towards the 'knee' (see Figure 5.4).

In this case, and in situations where accurate values for the noise level is not available, parameter estimation using either the L-curve or GCV analysis can successfully identify the regularisation parameter that achieves the best fit to the signal without over-fitting the noise in the data. These methods are applicable even where incorrect data errors are assigned in the datafile.

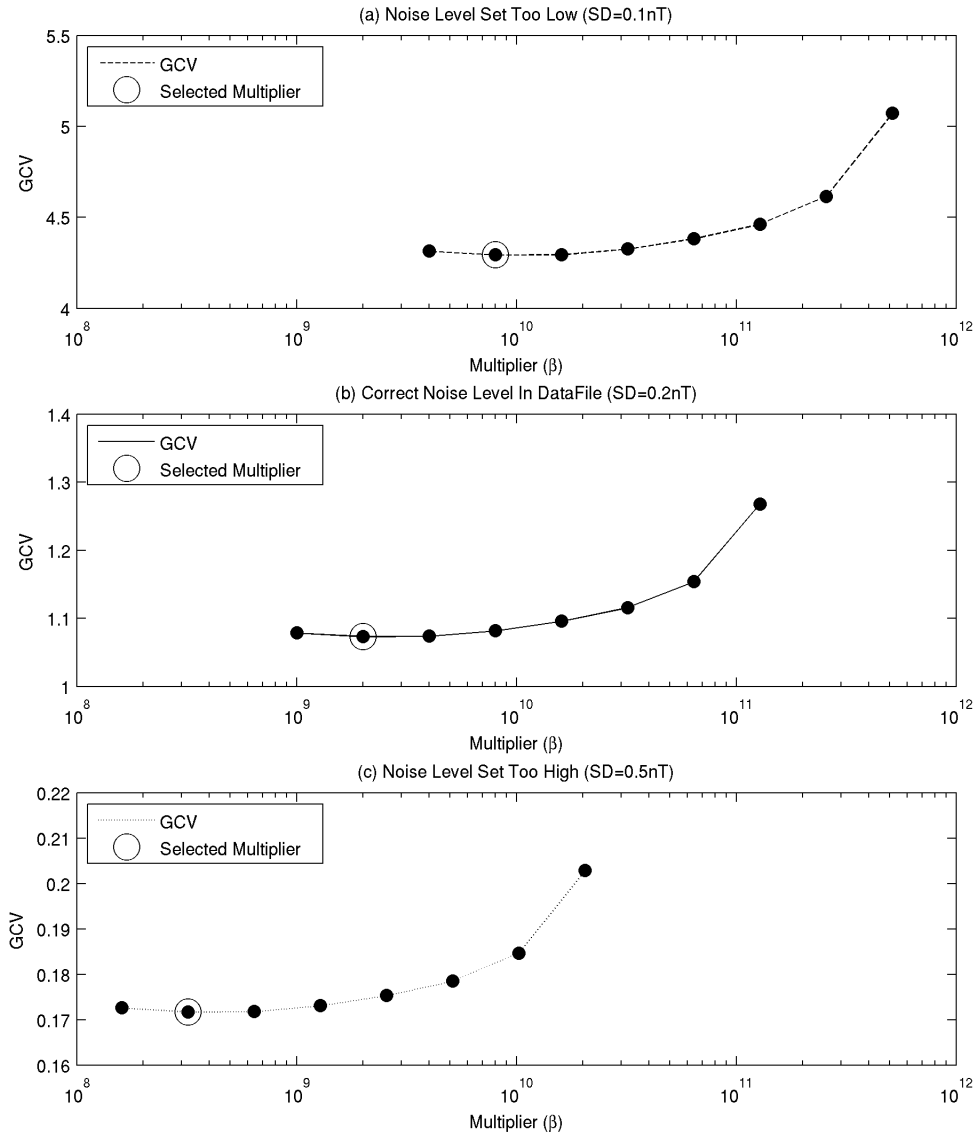

Figure 5.5: Generalised Cross Validation results for a range of β, where a noise level of SD=0.2nT has been added, yet differing values included in the datafile. (a) Noise level has been underestimated (SD=0.1nT). (b) Correct level of noise in datafile. (c) Noise level has been overestimated (SD=0.5nT). Despite the uncertainty in the noise level being declared in the datafile the GCV analysis has calculated a β value that compensates for this, and produces approximately the same model norm for each scenario (see Figure 5.4 where all GCV scenarios plot towards the 'knee' of the L-curves.)

5.6 Inversion of a synthetic anomaly

5.6.1 Mesh design

The subsurface distribution of anomalous magnetic susceptibility is unlikely to be represented accurately by dividing the subsurface into a series of 3D cells. The subsurface is likely to have a continuously varying structure, and therefore if the mesh is to represent the Earth without influencing the result, the cells should be small enough to account for this variation. In reality, computational demands also come into consideration when choosing the parameters used to construct the mesh.

Mesh cell size

The size of the cells within the mesh control the resolution of the final model that is produced. It is therefore important that the cells are small enough to recover the subsurface structure, yet smaller cells lead to a greater number of calculations for the same volume, and hence have a large impact on the amount of processing time required to produce the model, as well as making the problem more under-determined and hence increasing uncertainty. Ideally therefore, the size of the mesh cells should be the largest possible that does not have an adverse effect on the model recovered for the data misfit.

The importance of cell size will be illustrated by considering the simple model introduced in Chapter 3. Synthetic data generated from the model introduced in Figure 3.1, with an inclination of 64°, declination of 0°, Gaussian noise with 0.2 nT standard deviation, and a receiver height of 0.2 m above the ground surface, has been inverted using MAG3D. As sharp discontinuities within the model were expected α_s was set equal to 1, with α_x,α_y and α_z set to zero to eliminate these terms from the objective function. As only positive sources were desired, positivity constraints on the modelled susceptibilities were imposed.

Three models have been created using a mesh of the same volume, but with cell sizes of small $0.1\times0.1\times0.1$ m cells, $0.5\times0.5\times0.1$ m cells and large $1.0\times1.0\times0.1$ m cells. The models produced are largely comparable in terms of the structure that has been recovered, with the coarser resolution particularly observable in the $1\times1\times0.1$ m mesh. The calculated data and data misfit from each inversion, are shown in Figure 5.6. The smallest cell size has produced calculated data that replicates the observations well, and the data misfit highlights areas where the inversion has not attempted to fit the Gaussian noise (Figures 5.6a and 5.6b). These plots are very similar to those produced with the $0.5\times0.5\times0.1$ m mesh (Figures 5.6c and 5.6d), suggesting that the smaller cell size may not have been required to produce a model that more accurately fits the observations, and that the extra computing time was unnecessary. The larger cell size (Figures 5.6e and 5.6f), shows a significant difference. The calculated data appears pixelated, suggesting that a higher model resolution is required to accurately match the observations. This is highlighted by significant data misfits particularly in the area of the larger anomaly.

Padding cells for regional trends

Li & Oldenburg (1996) suggest designing a mesh where cells extend horizontally, outside of the area covered by the data. The modelled values in these cells are unlikely to be well constrained by the data as they are distant from the nearest datum. Due to the nature of the attenuation with distance observed in magnetic data, these cells are larger than those immediately under the data area, and the same principle can be used to justify having larger cells at the base of the mesh.

The purpose of having the mesh extend outside the area of the observations is that there may be large scale magnetic features outside of the survey area that cause anomalies to be registered within the data area. Should the mesh only cover the area immediately beneath the data area, then these anomalies would have to be modelled by placing incorrect susceptibilities in parts of the model that the real signal does not derive from.

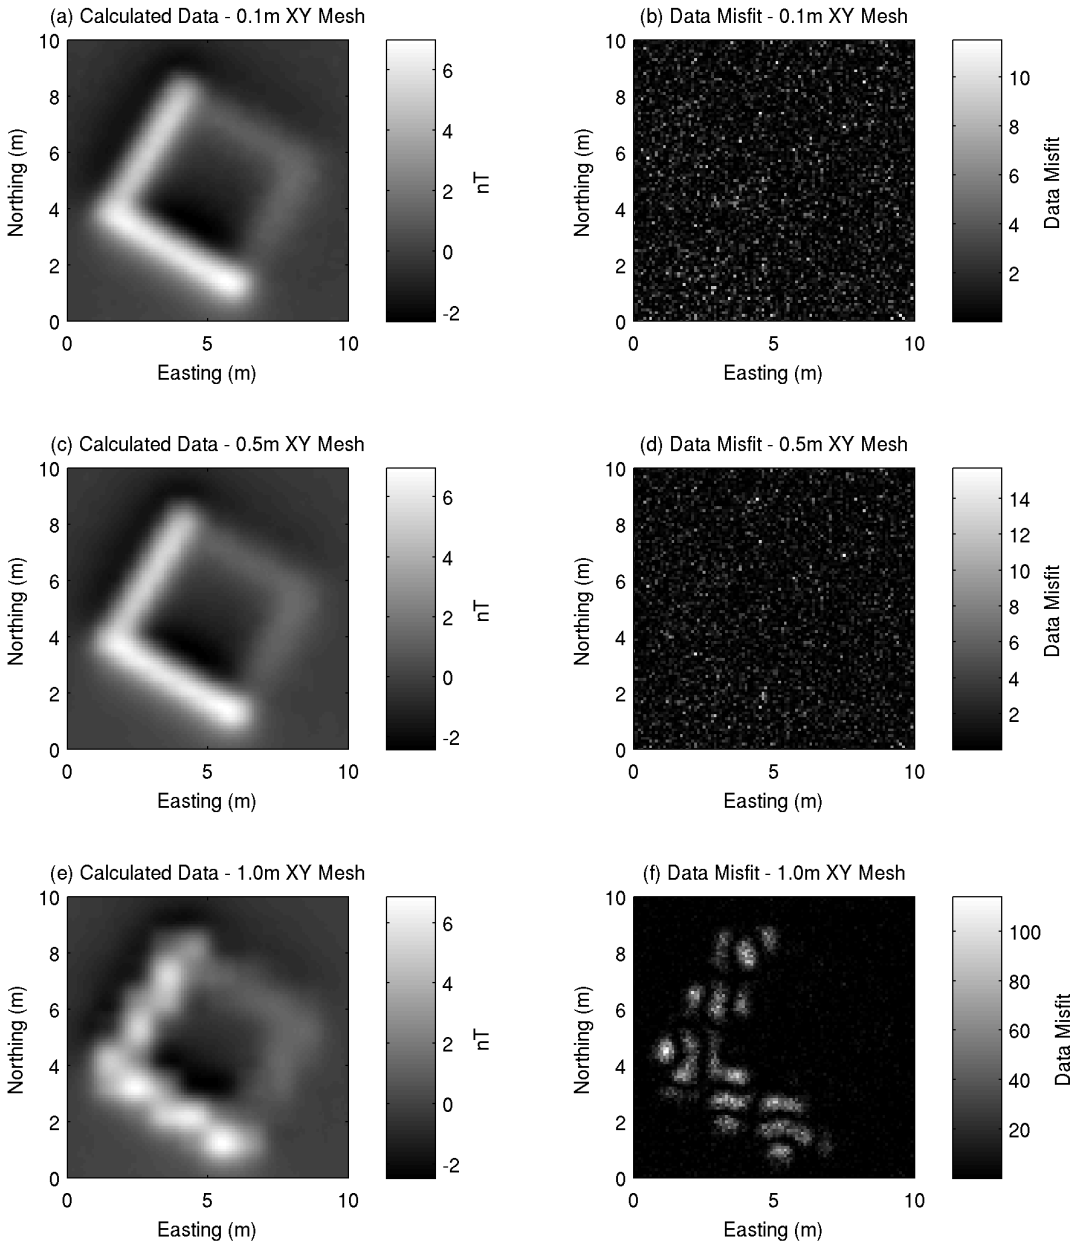

Figure 5.6: Effect of varying the X-Y size of mesh cells. (a) The calculated dataset from the final model produced with a mesh consisting of cells $0.1\times0.1\times0.1$ m. (b) Misfit associated with a). (c) Calculated dataset from the final model produced with a mesh consisting of cells $0.5\times0.5\times0.1$ m. (d) Misfit associated with c) (e) Calculated dataset from the final model produced with a mesh consisting of cells $1.0\times1.0\times0.1$ m. (f) Misfit associated with e)

The effect of these cells on the final result will be illustrated by inverting the synthetic scenario described in Section 5.6.1, using the same inversion parameters, and this time with a mesh with central 0.5×0.5×0.1 m cells beneath the data area, extending to a depth of 4 m, and with additional 1×1×0.1 m, and 2×2×0.1 m cell around the horizontal extents.

A value of β of 1×10^{10} was identified from analysis of the L-curve (Figure 5.7a). It can be seen in Figures 5.7b and 5.7c that the observed and calculated data show all the same features, and the plot of the data misfit (Figure 5.7d) does not show any coherent signal suggesting that the difference between the observed and calculated datasets is just the Gaussian noise added to the observations.

The inverse model is shown in Figure 5.8, where the lateral positions of the anomalous bodies has been reproduced successfully. The vertical extents recovered by the inversion differ from those used to generate the observations. The shallower body is shown to appear at approximately the correct depth (0.3 mbgl) and reaches a peak amplitude of ∼0.0008-0.0009 SI, at depths of around 0.8 to 1.4 mbgl. In reality the base of this feature should be at 1.2 mbgl, however evidence of this feature is shown in slices down to approximately 3.5 mbgl. Also present in this model are anomalously high susceptibility values in the 3 m wide padding cells that surround the small cells of the core of the mesh. High values are particularly noticeable towards the southern side of the mesh, and have peak amplitudes of 0.0004 SI increasing towards the base of the mesh.

To identify the effect these features had on the inversion, the susceptibilities from the core of the mesh were set to zero, leaving only the susceptibilities in the padding cells. These were then forward modelling to see the effect these cells have on the calculated dataset. The results of this forward modelling are shown in Figure 5.9.

It is demonstrated that the padding cells have been used by the inversion to create a regional trend across the area. In this case, due to higher susceptibilities being placed south of the data area, a negative magnetic field has been calculated for the southern-most observation points.

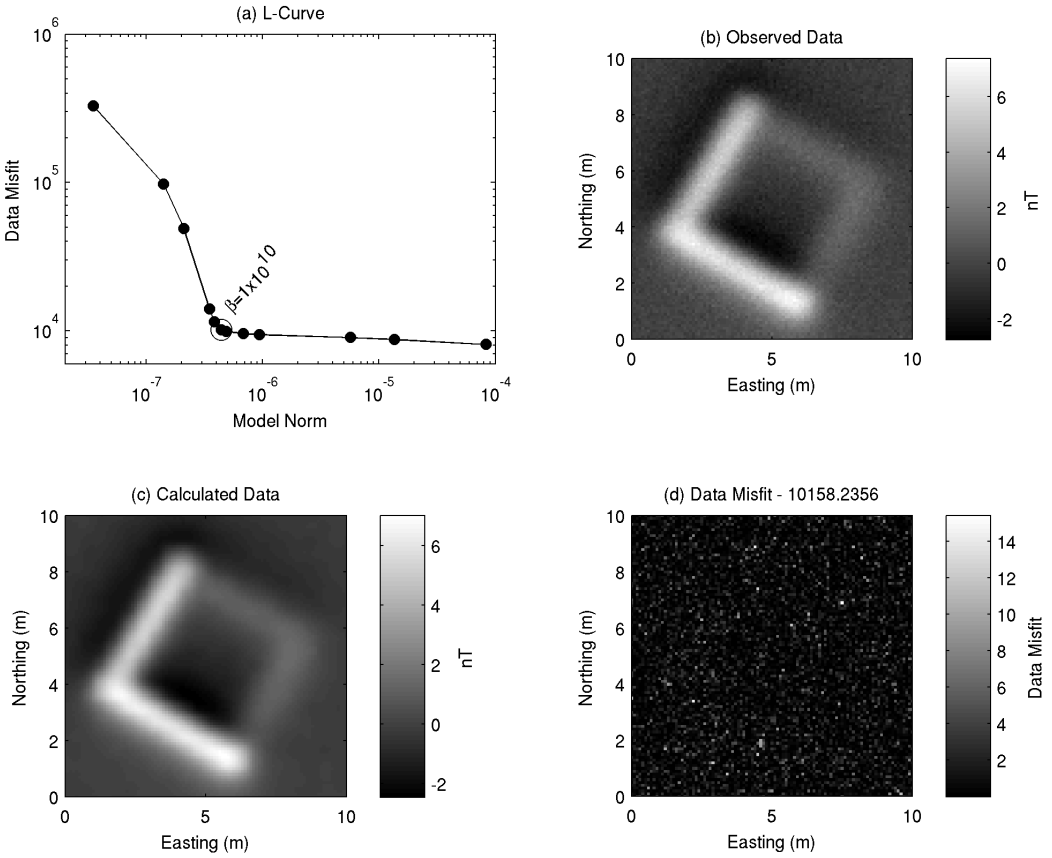

Figure 5.7: Analysis of data misfit for the inversion of the synthetic body in Figure 3.1, using a mesh with padding cells around the horizontal extents. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2. As there is no coherent signal of the archaeological signal in the data misfit plot, it is encouraging that all the signal has been incorporated in the modelling, and the differences highlighted represent the unwanted Gaussian noise.

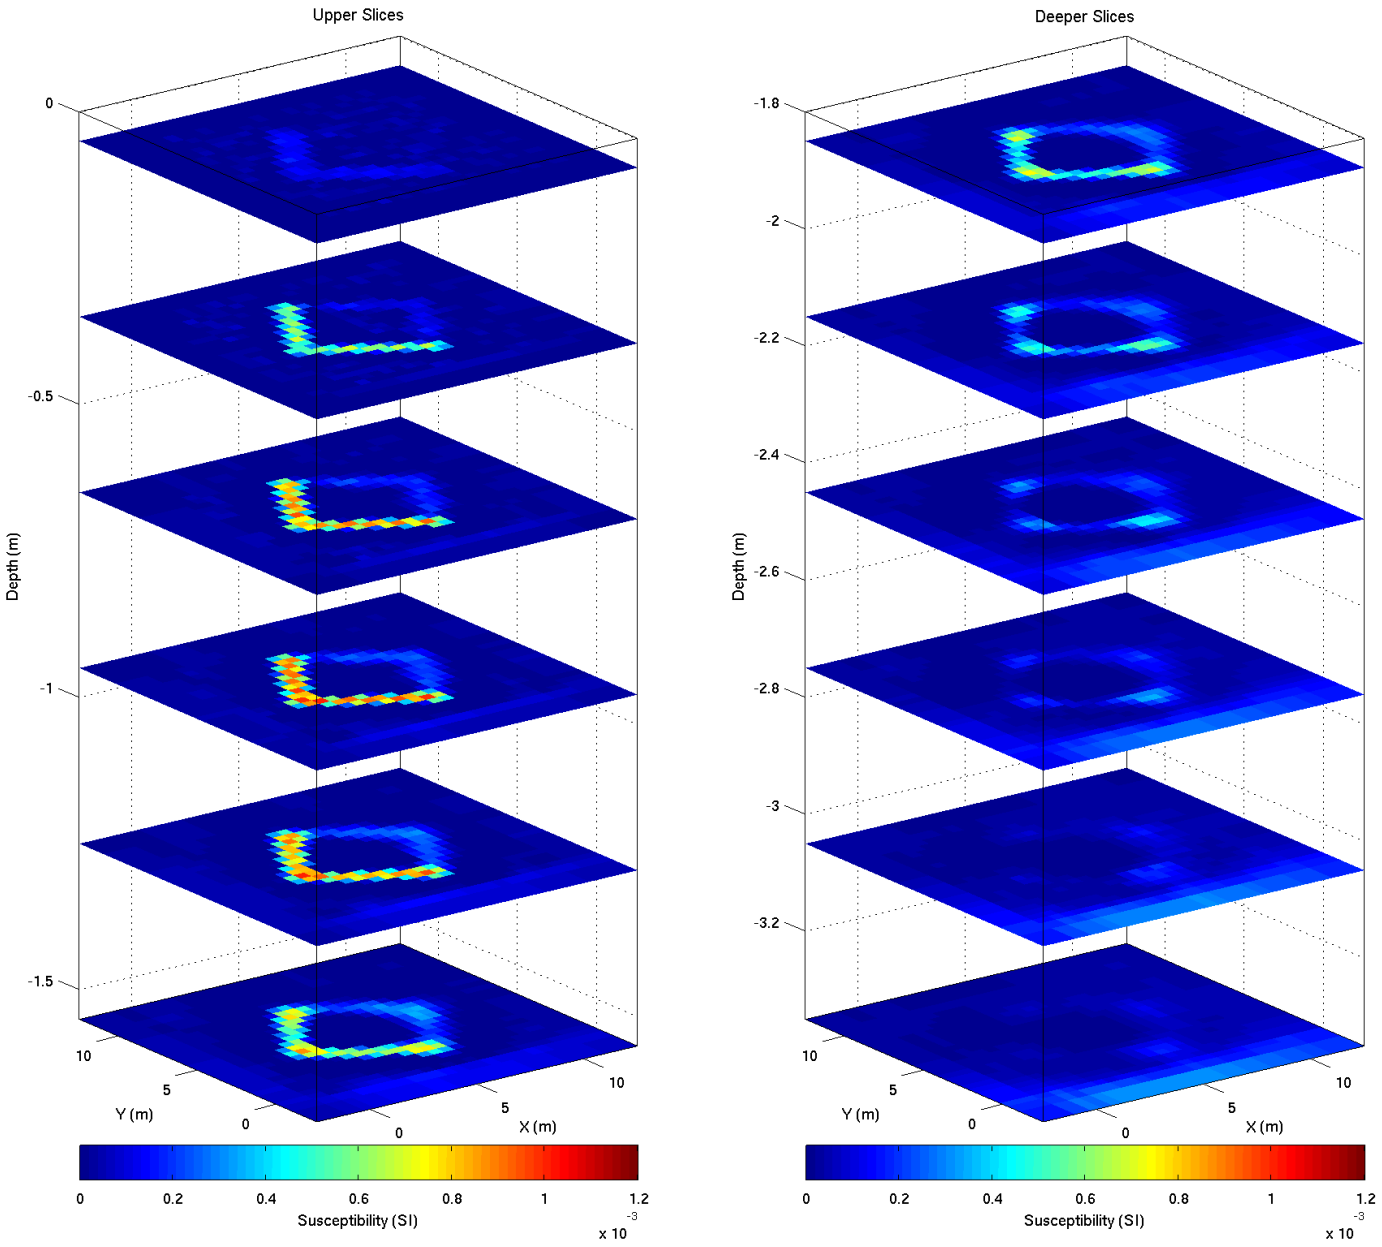

Figure 5.8: Final inversion model for the synthetic body presented in Figure 3.1, with padding cells surrounding the centre mesh. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl. The archaeological features are shown to have much larger vertical extents here than in reality, and there is a high susceptibility placed into the poorly constrained padding cells.

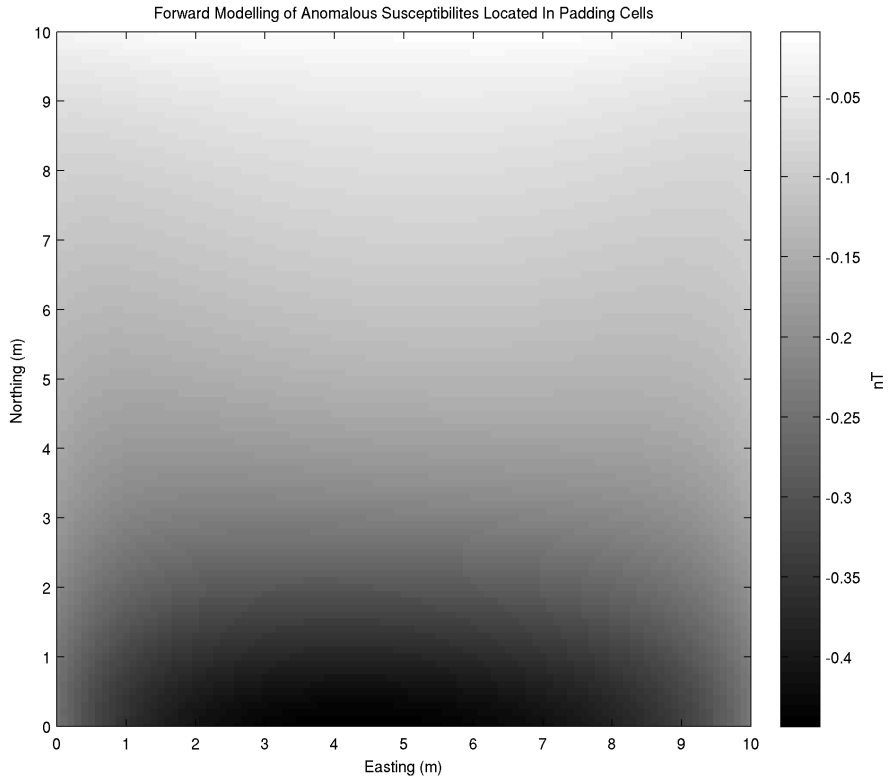

Figure 5.9: Forward modelling of the anomalous susceptibilities placed within the padding cells of the inversion shown in Figure 5.8. The padding cells have been used to produce a gradient across the survey area from -0.45 nT in the south to 0 nT in the north.

This is as expected within an inclined field. A gradient has been placed across the survey area from -0.45 nT in the south, to $\sim$0 nT in the north. Li & Oldenburg (1996) recommended using padding cells in areas where they will be poorly constrained by the data, as they are supposed to account for all the regional trends that should not be accounted for within the main part of the mesh. However, in this case, there is no regional trend across the area, and all anomalies should be accounted for by susceptibilities within the main part of the mesh. The effect of this negative trend across the survey area has resulted in higher susceptibility being required from the main bodies, and hence could explain the deeper than expected extent of the simulated archaeological bodies.

The inversion was repeated, this time using a smaller mesh volume without the padding cells around the horizontal extents. The value for β of 1×10^{10} was selected by analysing the L-curve shown in Figure 5.10a. Despite the absence of padding cells around the horizontal extents of the data area, the observed and calculated datasets (Figures 5.10b and 5.10c) show a good match with no coherent features being missed by the modelling (Figure 5.10d). The final inversion model is shown in Figure 5.11.

The model shows slight increases in susceptibility in the vicinity of the bodies at very shallow depths, with a much larger increase in susceptibility seen in the slice at 0.3-0.4 m. This is the depth at which the shallower body should appear, and this feature is seen to have a strong presence in the next two slices before weakening in the slice at 1.2-1.3 m and considerably in the final slice shown at 1.5-1.6 mbgl. The deeper of the two features also shows a maximum susceptibility in the 0.9-1.0 m and the 1.2-1.3 m slices as would be expected, however the maximum susceptibility of this deeper feature does not reach similar amplitudes to the shallower feature, and higher than expected susceptibilities exist in the slices above where this body should be present. Despite this, removing the padding cells from the mesh prior to inversion has dramatically improved the spatial recovery of the bodies.

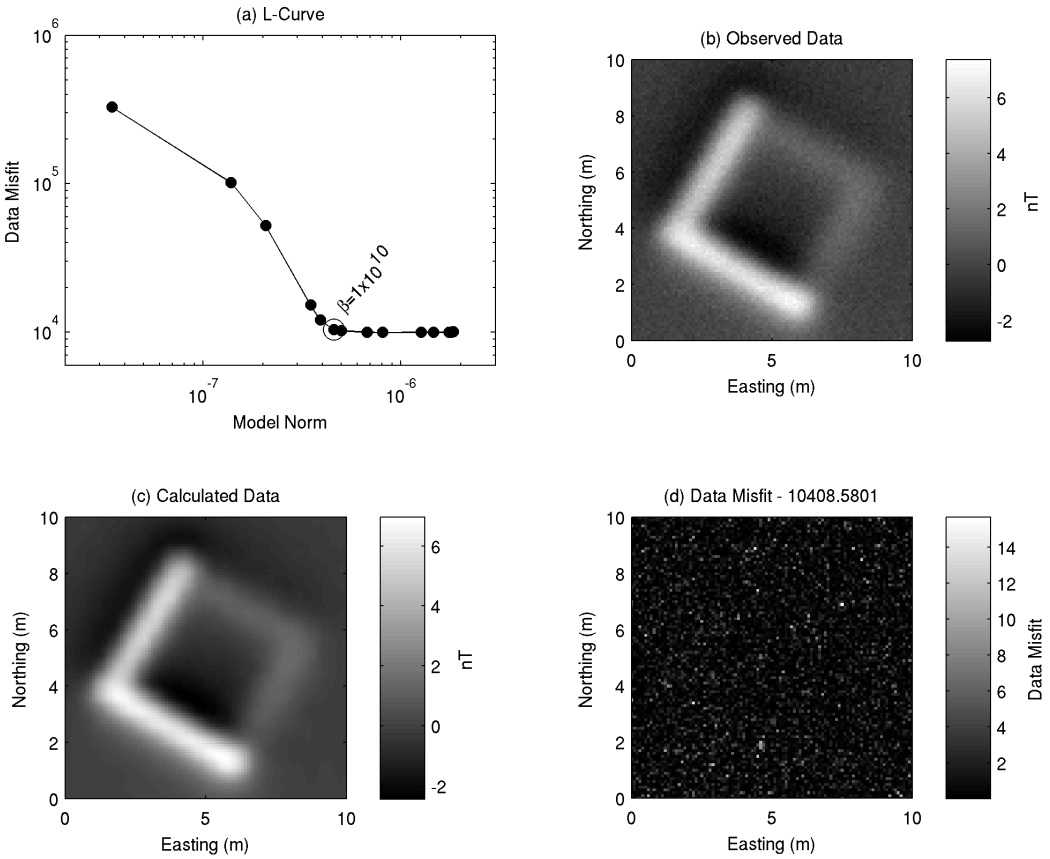

Figure 5.10: Analysis of data misfit for the inversion of the synthetic body in Figure 3.1, using a mesh without padding cells. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

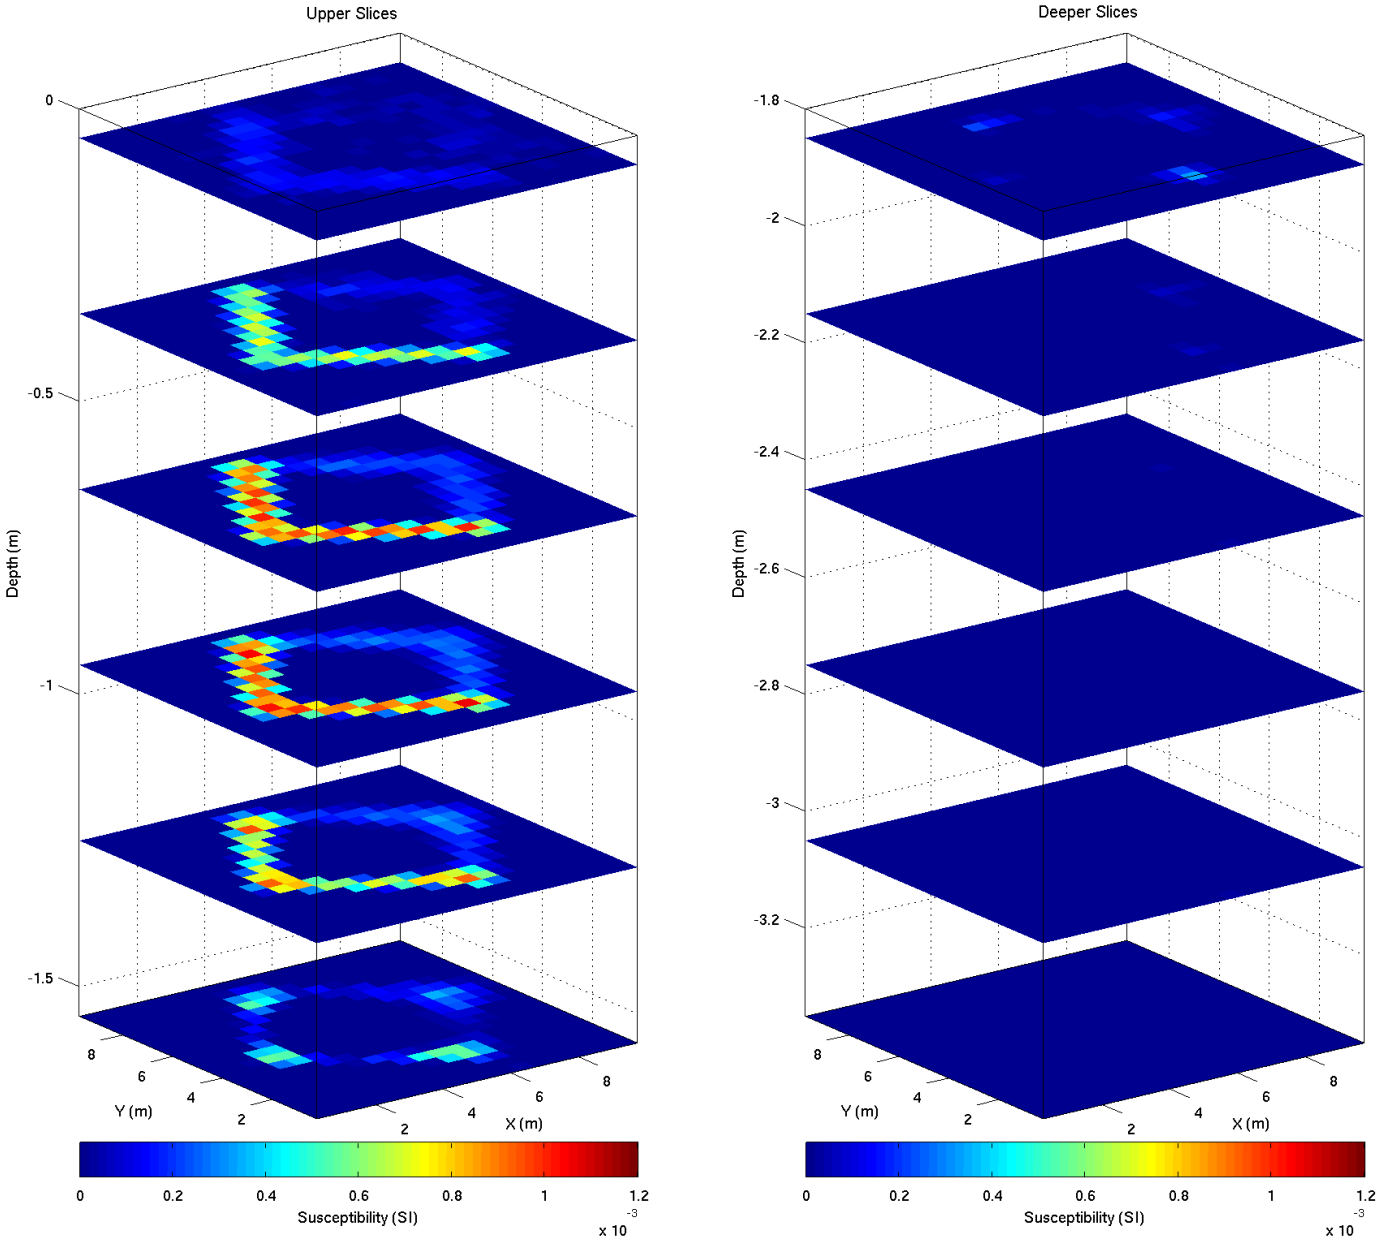

Figure 5.11: Final inversion model for the synthetic body presented in Figure 3.1, without padding cells. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl. The shallower feature, shown here at the bottom and left is positioned between 0.3 and 1.2 mbgl and has been recovered well. The deeper feature, shown at the top and right should be positioned between 0.8 and 1.2 mbgl. Both bodies were generated with a susceptibility of 0.001 SI, so the recovered susceptibility for the deeper feature is lower than expected.

To look at the responses to the two bodies in more detail, two profiles at 2.75 mE, 6.25 mN and 7.25 mE, 3.75 mN are shown in Figure 5.12. The first of these represents a profile directly through the centre of the shallow, thicker body, and the second, vertically through the centre of the deeper, thinner body. The first profile (black line) shows a good, if somewhat, smoothed relation to the actual outline of the body (dashed line). The model shows a gradually increasing susceptibility from the shallow subsurface, with a peak susceptibility of 0.0011 SI recovered by the inversion. There is no evidence of this feature below 1.5 mbgl.

The second profile, for the deeper feature has not recovered the synthetic body as accurately. The recovered susceptibility does not rise as quickly as for the shallower anomaly, however there is an increase from below 0.35 m until a peak susceptibility of $\sim$0.0003 SI is recovered at 0.95 mbgl. Although the peak susceptibilities are somewhat lower than the values used to generate the synthetic data (0.001 SI), the depths of the peak susceptibilities match well with the depths of the deeper body (0.8-1.2 mbgl).

Problems of non-uniqueness

The model recovered by the inversion, for the deeper body, highlights what is known as the inverse problem. As the number of unknowns in this scenario is the same as the number of cells in the mesh ($50\times50\times40$ = 100000), and the number of data is 10201, the inversion is a largely under-determined problem. This means, with far more unknowns than knowns, there are a large number of solutions to the model that will produce the same fit to the observations.

In order to demonstrate the inverse problem a 2×2 m prism has been modelled in Figure 5.13. First a dataset with the top of the prism at 0.8 m and base at 1.2 m is represented by the black cross-section in Figure 5.13b. The forward response to this prism is shown for an E-W profile across the centre in Figure 5.13a, by the black line, and a vertical profile through the centre of the prism shows the susceptibility as 0.001 SI (Figure 5.13c). The depths and susceptibility of this prism are the same as the deeper body in the previous inversion.

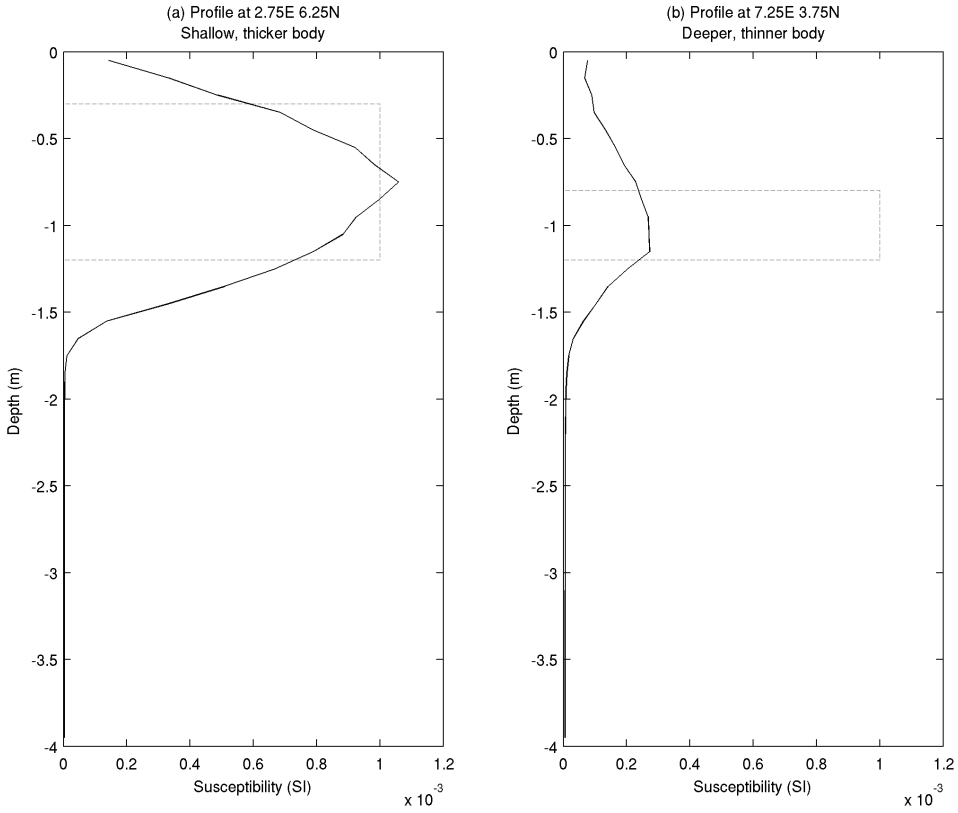

Figure 5.12: Vertical profiles through the susceptibility model shown in Figure 5.11, at (a) 2.75mE, 6.25mN, representing a profiles through the centre of the shallower body, and (b) 7.25mE, 3.75mN, representing a profile through the deeper feature. The actual profile used to generate the observations is shown by the dashed grey line, and the profiles from the inversion is shown by the black line.

Following this, the the X-Y dimensions of the prism remained the same, yet the top of the prism was fixed to 0.6 mbgl, and the Nelder-Mead simplex direct search algorithm within Matlab was used to calculate a base to the prism, and a susceptibility that best matched the forward data from the 0.8-1.2 m example. This resulted in a cross-section that is shown in blue in Figure 5.13b, with the forward data and susceptibility profile shown by the blue lines in Figure 5.13a and 5.13c, respectively. Note, that in the cross-section, the blue prism also encompasses the area shown by the black prism. The modelling has shown a good, comparable match to the forward data of the original prism, and resulted in a thicker prism with a base at 1.58 mbgl, and a lower susceptibility of 0.00043 SI. This process was repeated for a prism with the top fixed at 0.4 mbgl (red) and 0.2 mbgl (cyan). Although some differences can be observed between the calculated datasets, they are largely similar, and in the presence of noise, these differences could be accounted for by the noise level of the data.

This highlights the inverse problem, as any one of these bodies could be responsible for the data recorded at the surface. Looking at the effects observed in Figure 5.13, it can be seen that the cause of the lower than expected susceptibility recovered for the deeper body in Figure 5.12 is likely to be due to presence of higher susceptibilities shallower than the actual body depths. This has resulted in a larger anomalous zone of lower amplitude.

5.6.2 Effect of regional trends

Whereas the scenarios presented in Figures 5.8 and 5.11 suggest that the presence of padding cells around the centre of the mesh hinder the final inversion result, a regional trend was not actually present within the observed dataset. It is quite possible that regional trends will exist in archaeological surveys, not just due to geological reasons, but also due to uncontrollable sources such as fences, buildings and cars close to survey areas.

To test the effectiveness of the padding cells in the presence of a regional trend, a new dataset was generated from the anomaly shown in Figure 3.1, with an inclination of 64° and a declina-

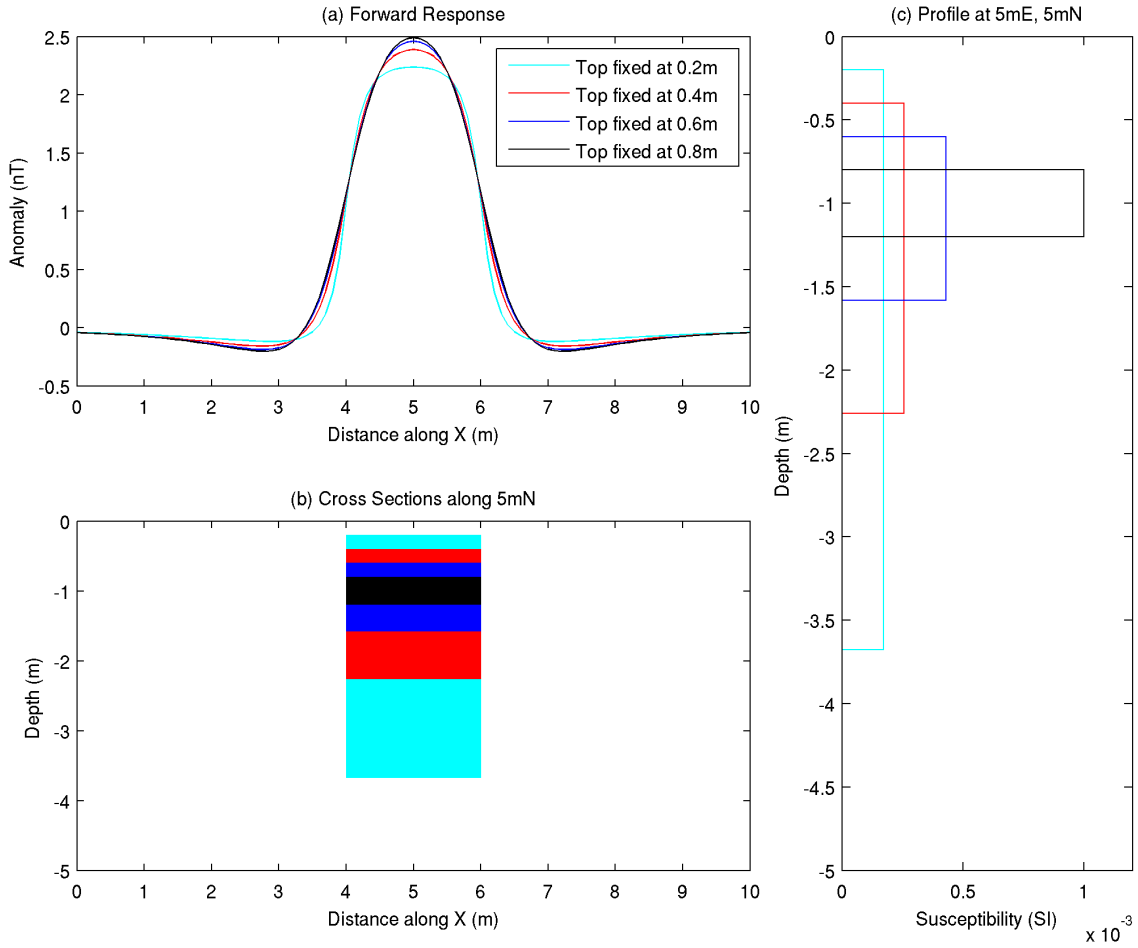

Figure 5.13: Demonstration of the inverse problem. (a) Forward modelled data along an E-W profile across the centre of a 2 x 2m body of varying depths and susceptibility. The profiles are generated from the cross-sections shown in (b) with susceptibilities shown in (c). All models, despite being very different generate a set of data at the surface that are comparable at the surface, highlighting the problem with under-determined inverse problems.

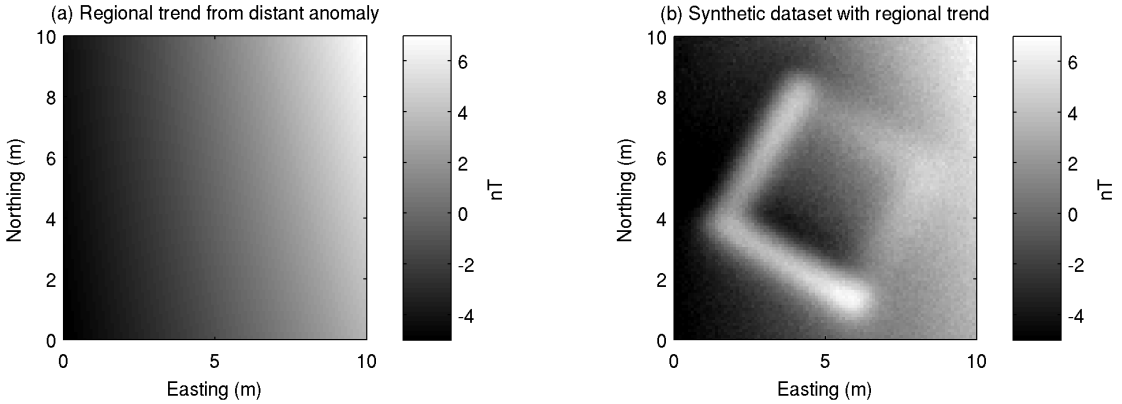

Figure 5.14: Synthetic dataset with regional trend. (a) The regional trend caused by a source at 50 to 100 mE and 50 to 100 mN at a depth of 20 to 70 mbgl. (b) The synthetic dataset added to the regional trend, with 0.2 nT Gaussian noise added.

tion of 0°, to represent a typical mid-latitude location. This time the additional signal from a body positioned off-grid at 50 mE to 100 mE, and 50 mN to 100 mN at a depth of 20 to 70 m has been added to the dataset. The susceptibility of this body was set to 0.133 SI, which within a field of 37.5 Am^{-1} resulted in a magnetisation intensity of 5 Am^{-1}. The effect of the regional trend is shown in Figure 5.14a, and with zero-mean Gaussian noise of 0.2 nT standard deviation added, along with the anomaly from Figure 3.1, the new synthetic dataset with regional trend can be seen in Figure 5.14b

The inversion has been repeated using the same parameters described above, and using the mesh with $1\times1\times0.1$ m, and $2\times2\times0.1$ m padding cells around the horizontal extents. The final model has again been selected by analysis of the L-curve (Figure 5.15a), in this case selecting an optimum β of 1×10^9. Depth slices from the model can be seen in Figure 5.16. It is shown that the inversion has placed very high susceptibilities into the padding cells of the mesh. At shallow depths this is only in the inner set of $1\times1\times0.1$ m cells. This is particularly noticeable in the depth slice at 0.35 mbgl. Deeper than this however, both the $1\times1\times0.1$ m and $2\times2\times0.1$ m cells contain high susceptibilities to account for the regional trend, and the amplitude of these susceptibilities generally increases with increasing depth. Below a depth of $\sim$2 mbgl, very high susceptibilities are placed within the central part of the mesh as well.

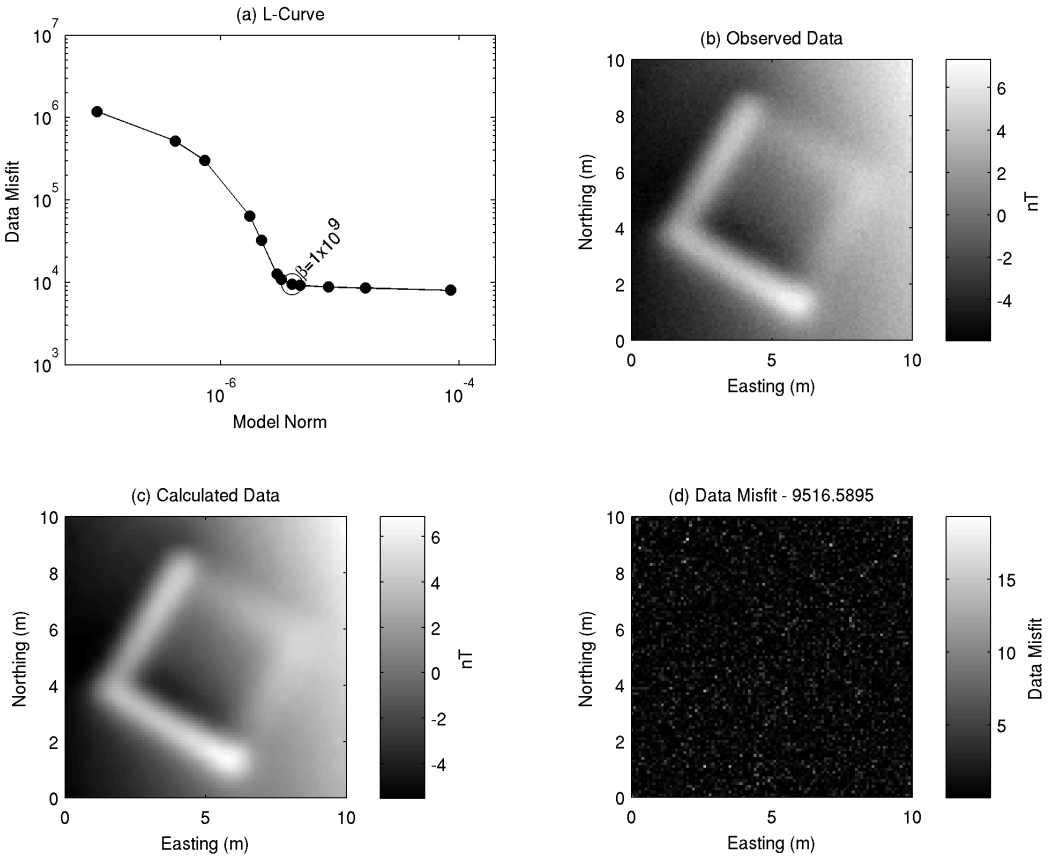

Figure 5.15: Analysis of data misfit for the inversion of the synthetic body with regional trend shown in Figure 5.14, using a mesh with padding cells around the horizontal extents. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

A strong increase in susceptibility can be observed by the arrival of the shallower archaeological body at 0.35 mbgl, and an increase in amplitude can also be observed for the deeper body in the 0.95 mbgl slice. The base of both of these bodies becomes hard to interpret below the 1.85 m slice due to the high susceptibilities introduced in the central area discussed above.

Where no regional trend was present within the dataset, it has been shown that the presence of the padding cells was a hindrance to accurately recovering the synthetic model. Therefore, to test the influence that a regional trend would have on the model where no poorly-constrained padding cells are present, the inversion was repeated, this time using just the core mesh of 0.5×0.5×0.1 m cells. Selecting an optimum value for β was difficult for this inversion, as the L-curve (Figure 5.17a) does not show a distinctive "knee", instead a smooth curve is derived. A value of $\beta = 1 \times 10^{11}$ was eventually chosen, however the Chi-squared and GCV solutions are also highlighted on Figure 5.17a to demonstrate the uncertainty. There are also noticeable differences between the observed (Figure 5.17b) and calculated (Figure 5.17c) datasets, with the calculated dataset showing a distinctly pixelated image, a possible indication that the mesh cell sizes are too large to accurately recreate the observations. These differences are highlighted in the data misfit (Figure 5.17d) which shows that a poor match has been made to the amplitudes of the magnetic field towards the western side of the survey area in particular.

The final inversion model is shown in Figure 5.18. The model is largely made up of very low susceptibilities, with large values being placed at the extreme edges of the mesh, towards the southern base of the model. Large susceptibilities are placed at shallow depths, which causes the pixelisation as identified in Figure 5.17c. Little evidence of the archaeological bodies is present in the model, as the archaeological features have been modelled in the same anomalous susceptibility distributions used to replicate the regional trend.

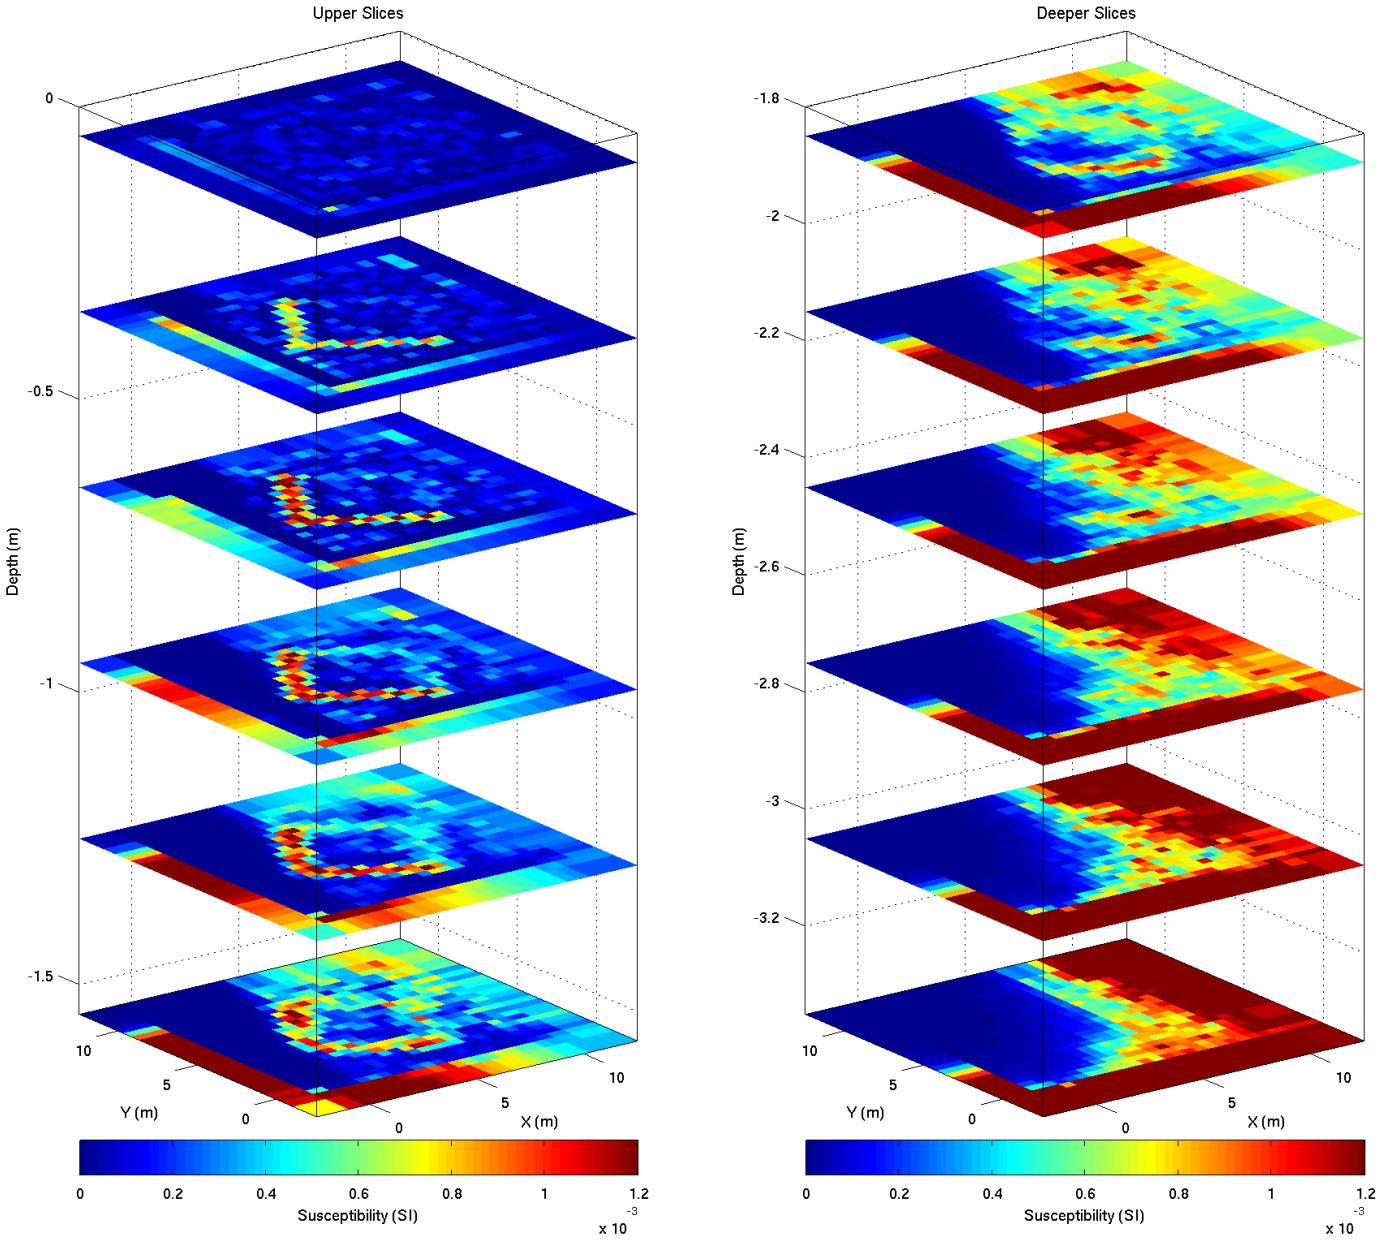

Figure 5.16: Final inversion model for the synthetic dataset shown in Figure 5.14, with padding cells around the horizontal extents of the mesh. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl. The inversion has placed very high susceptibilities into the padding cells, and towards the base of the mesh, to account for the regional variation across the survey area.

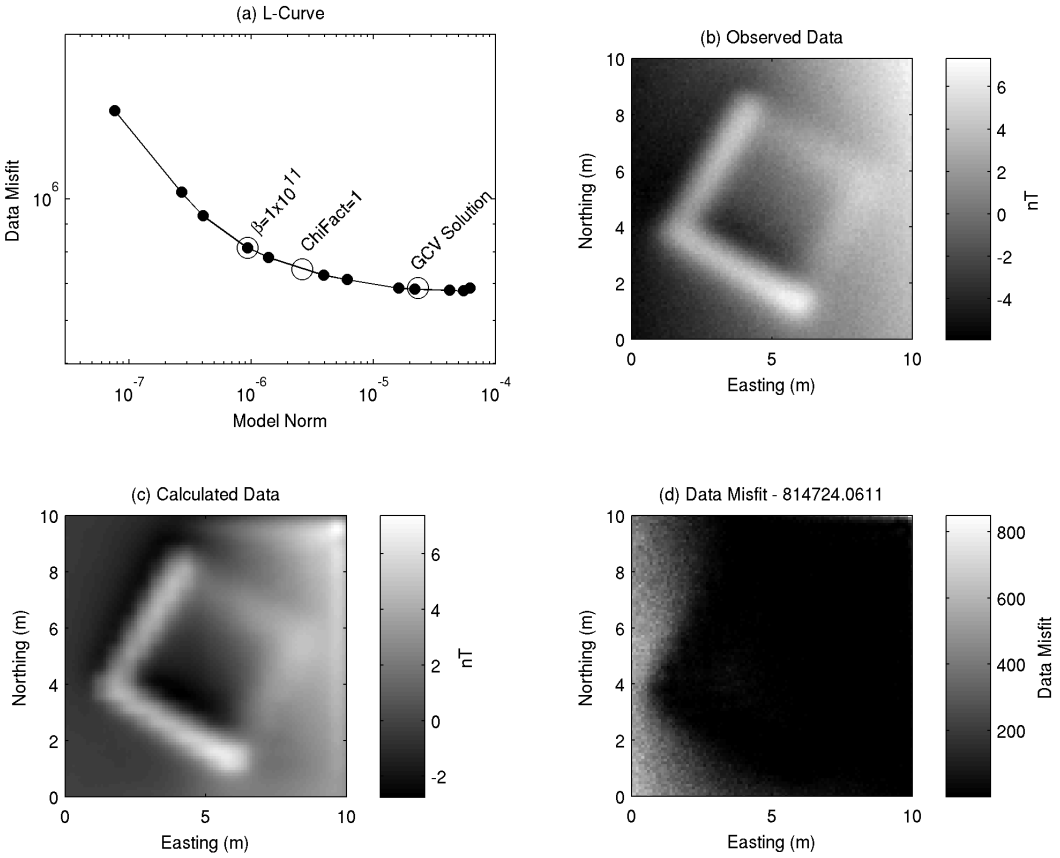

Figure 5.17: Analysis of data misfit for the inversion of the synthetic body with regional trend shown in Figure 5.14, using a mesh without padding cells around the horizontal extents. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2. Note the scale is much larger here than for Figure 5.15d.

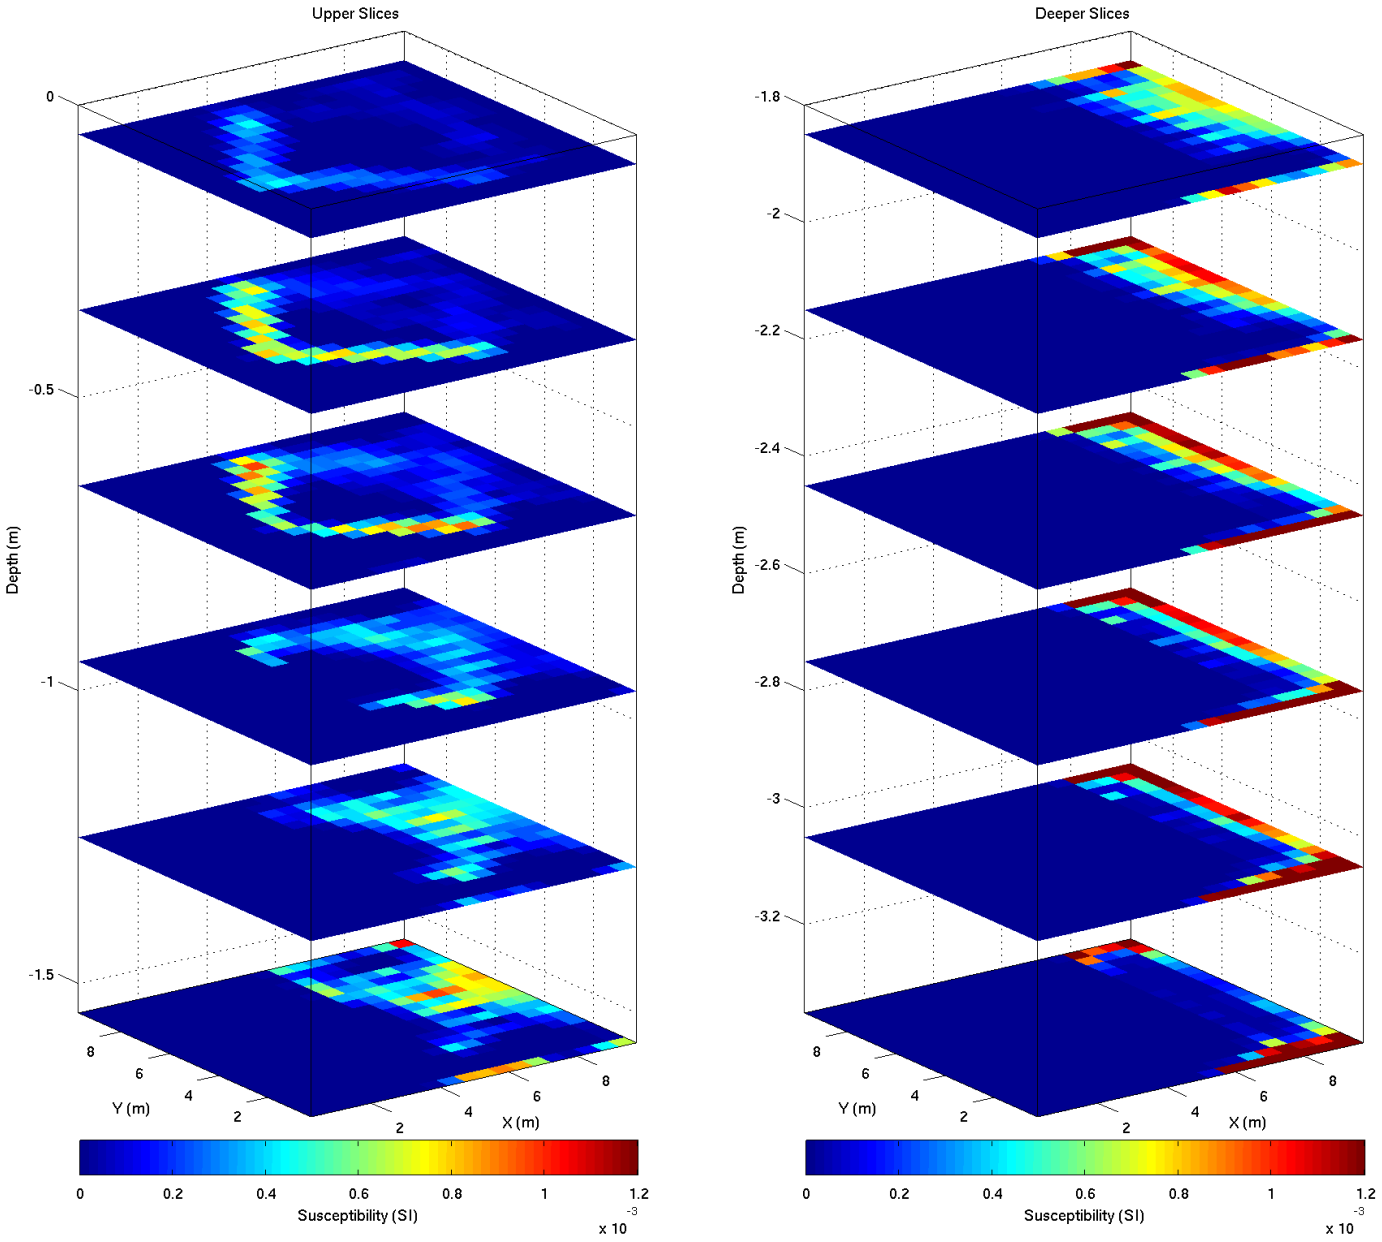

Figure 5.18: Final inversion model for the synthetic dataset shown in Figure 5.14, without padding cells around the horizontal extents of the mesh. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl. The absence of the padding cells around the exterior of the survey area forces the inversion to recreate all the signal in the observations by placing anomalies within the mesh, even though in this scenario they derived from outside.

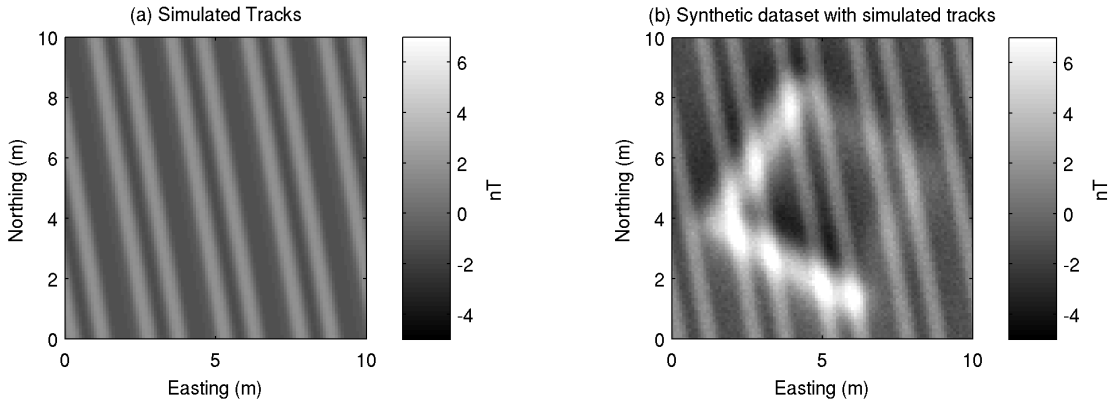

Figure 5.19: Synthetic dataset with simulated tractor tracks. (a) Pairs of 0.5 m wide tracks along a bearing of 170°, 1 m apart (b) The synthetic dataset added to the tracks, with 0.2 nT Gaussian noise added.

Recent agricultural noise in datasets

Topographical anomalies or compressed areas of soil caused by ploughing, or more significant relic features such as ridge-and-furrow produce a response in near-surface magnetic surveys that are usually identified by their repetitive, linear pattern. While they are accurate recordings of the subsurface magnetic properties and micro-topographical features of the site, it is often the anomalies due to deeper features which are the primary focus of the survey. These target anomalies can be masked by the near-surface pattern.

To simulate this type of noise, the magnetic anomaly generated by a repeating 0.5 m wide, by 0.1 m deep prisms of infinite length have been produced. A 1 m gap is included between each pair of prisms, which are separated by 0.5 m. This is designed to replicate a scenario similar to track marks caused by a tractor passing repeatedly over the survey area. The agricultural signal is generated along a bearing of 170°. The susceptibility of the agricultural signal is also set to 0.001 SI, resulting in responses between -1.5 and 2 nT (Figure 5.19a). This signal has then been added to the forward modelled data from Figure 3.1 to produce a dataset for inversion (Figure 5.19b). Although it is not actually "noise", the agricultural signal is referred to here as the agricultural noise as it is the undesired part of the dataset.

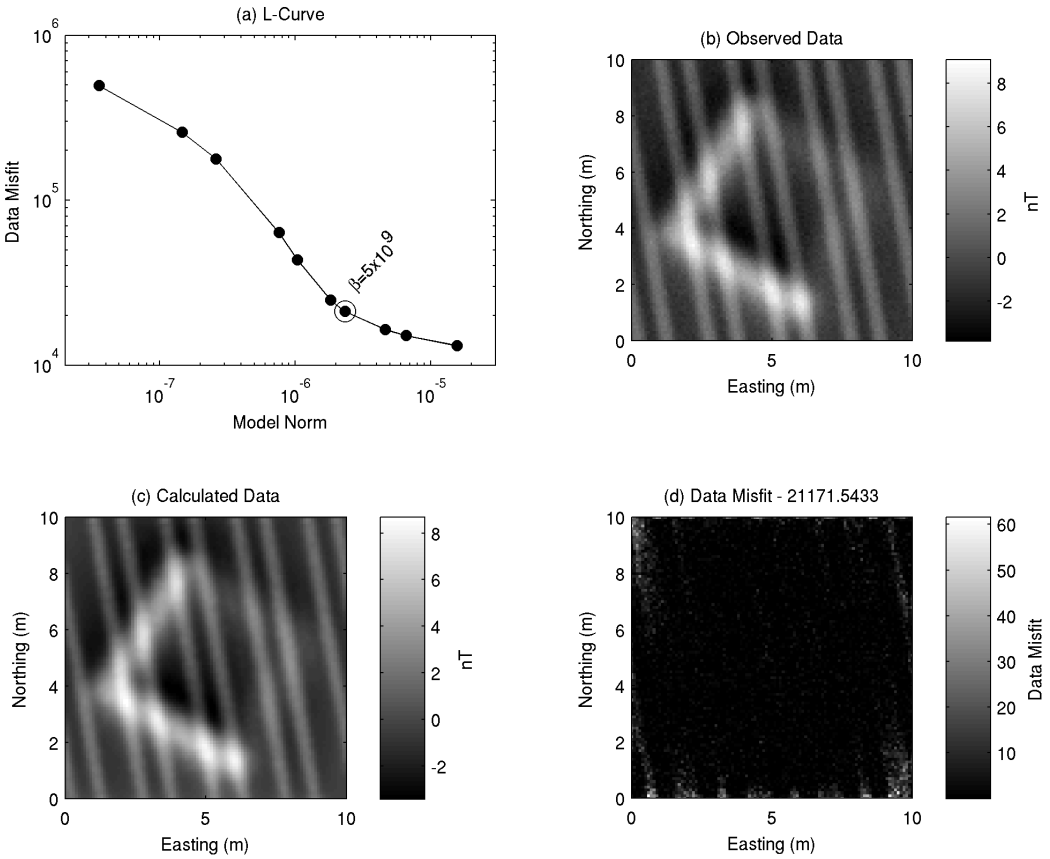

Figure 5.20: Analysis of data misfit for the inversion of the synthetic body in Figure 3.1 with simulated tractor tracks superimposed. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

The L-curve shows a broad curve, without a distinct "knee" (Figure 5.20a). A value of $\beta = 5 \times 10^9$ has been chosen and the plot of data misfit shows a good match to the observations has been made in the central part of the survey area, with higher misfits present around the edges (Figure 5.20d).

The final model is shown in Figure 5.21. A good recovery of the linear agricultural signal is present in upper slices with an amplitude of $\sim$0.0006-0.0008 SI. The linear features are present at greater depths in the model than they should be, still visible in slices down to 0.35 mbgl whereas they should terminate at 0.1 mbgl. The slice at 0.35 mbgl shows a transition within the model below which evidence of the linear pattern is no longer prominent, and the archaeological bodies are recovered with high susceptibilities $>$0.001 SI. The archaeological features have a relatively sharp base at 0.85 mbgl, and there are no anomalous susceptibilities associated with these bodies below $\sim$1 mbgl, which is shallower than the desired 1.2 mbgl.

The presence of the agricultural anomalies superimposed over the archaeological anomalies has caused a significant difference in the success of recovering the features of interest (Figures 5.11 and 5.21). As it is not desired to recover the agricultural signal in the model, it would be advantageous to remove these features from the dataset prior to the inversion, and only invert the anomalies that are of interest in the final model. This is likely to improve the chances of successfully recovering the archaeological sources, and methods to achieve this will be investigated in Chapter 6.

5.6.3 Remanent magnetisation

Any remanent magnetisation that may be preserved, through the processes described in Section 2.5, will be in addition to the magnetisation induced when the body is placed within the Earth's magnetic field. MAG3D assumes that all remanent magnetisation is negligible and that all the magnetic responses can be accounted for by the induced magnetisation. If this is not the case, then the inverse problem becomes more complicated as, instead of solving the inverse

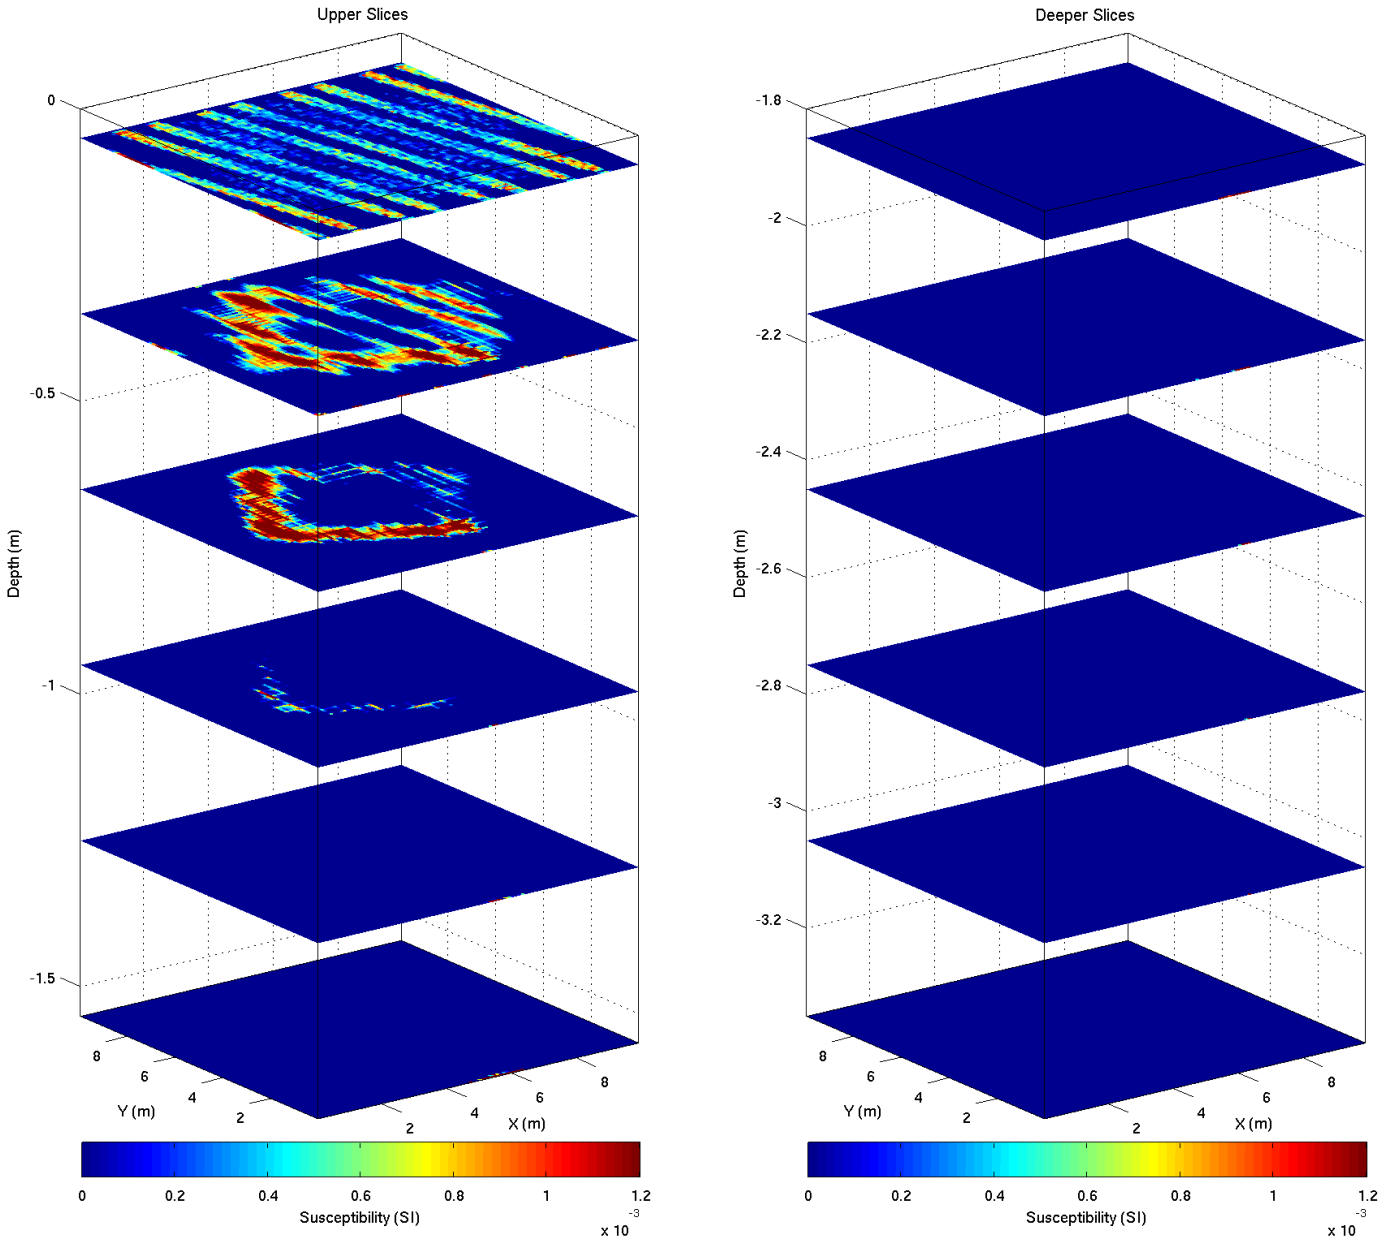

Figure 5.21: Final inversion model for the synthetic dataset shown in Figure 5.19, without padding cells around the horizontal extents of the mesh. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl.

problem for a set of susceptibilities, the inversion also has to produce solutions to the inclination, declination and magnetisation, of the remanence. This is coupled to the problem of ambiguity of not being able to distinguish between induced magnetisation and remanent magnetisation (Williams, 2008).

In theory, should all the magnetisation not be in the direction of the inducing field, the resulting model will be compromised and additional artefacts will be introduced to the model. In order to identify the scale of these artefacts another synthetic dataset has been generated consisting of three $2\times2\times0.5$ m prisms, all buried at 0.3-0.8 mbgl (Figure 5.22a). Each one of these has been given a different magnetisation vector. The first, seen on the left, has been forward modelled with an inclination of 64° and a declination of 0°, which is the same as the inducing field for the whole area. The central prism has a magnetisation vector of inclination 90° and declination of 0°, and the prism on the right of Figure 5.22a has an inclination of 64° and a declination of 40°. The model has been forward modelled with the addition of 0.2 nT standard deviation Gaussian noise and the synthetic set of observations at a height of 0.2 m above ground level is shown in Figure 5.22b.

To see what artefacts would be introduced owing to these incorrect magnetisation vectors, the dataset shown in Figure 5.22b has been inverted. Cell sizes of $0.5\times0.5\times0.1$ m have been used, and the total mesh size was $20\times10\times4$ m. An inclination of 64° and declination of 0° has been used to represent the inducing field, therefore matching the most westerly prism. The final model was selected by analysis of the L-curve (Figure 5.23a), which produced a good fit to the data (Figure 5.23d). No coherent pattern of misfit can be identified, however increases in the data misfit can be seen around the southern part of the central prism and the eastern side of the easternmost prism. This is not a surprise as these are the areas where the varying magnetic vectors cause variations in the total field response compared to the western prism, which is inline with the inducing field (see Figure 5.22b).

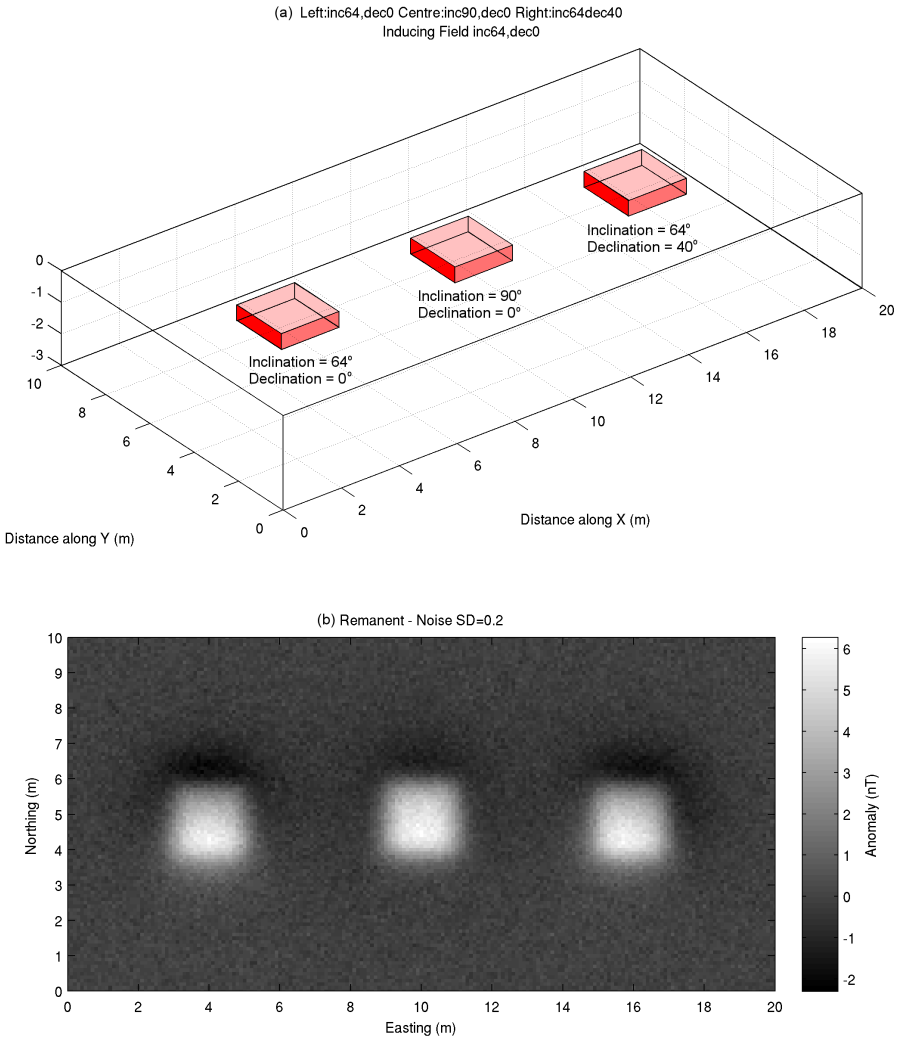

Figure 5.22: Synthetic dataset of three prisms each with different magnetisation vectors, all within an inducing field of 64° inclination, 0° declination. (a) 3D model showing all prisms to be placed between 0.3 and 0.8 mbgl. (b) Dataset achieved from forward modelling a) with 0.2 nT standard deviation Gaussian noise added.

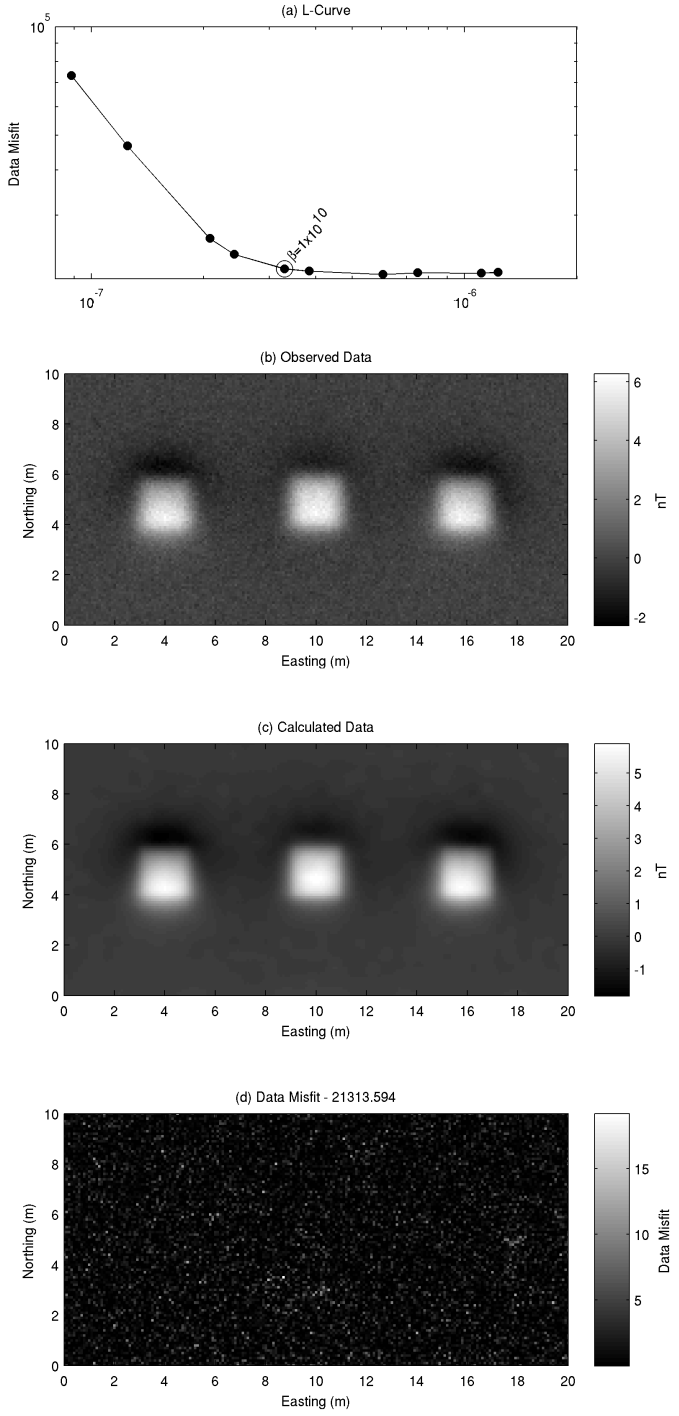

Figure 5.23: Analysis of data misfit for the inversion of the synthetic body in Figure 5.22a. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

The final inversion model is shown in Figure 5.24, which shows every third depth slice from the surface down to a depth of 2.2 mbgl. The control test here is the western prism which shows a strong increase in amplitude at 0.35 mbgl, reaching a peak susceptibility of ∼0.001 SI around 0.75 mbgl, reducing in amplitude at 1.25 mbgl and finally completely disappearing by 1.4 mbgl.

Key differences are present between this prism and the central and eastern prisms, where the negligible remanence assumption is violated. The central prism is positioned between 9-11 mE and 4-5 mN in Figure 5.22, yet in the recovered model the body has been recovered moving progressively northwards as depth increases. To show this N-S profiles over the centre of each prism is shown in Figure 5.25. An obvious difference can be seen between the control prism (Figure 5.25a) and the central prism (Figure 5.25b). This difference has been introduced into the model to match the higher total field response that is generated over the northern part of the prism due to the 90° inclination. Similar observations can be seen with the eastern prism, which has been recovered dipping towards the NW in order to match the lower total field response towards the east of the prism, due to the 40° declination.

The results to the varying magnetisation vectors are surprising, as although the recovered bodies do differ from the control prism, the majority of the susceptibility anomaly is modelled within the prism boundaries, even where the magnetisation vector differs to the inducing field. This is particularly the case at very shallow depths, with the differences increasing at greater depths. For archaeological anomalies that are typically buried shallowly in the subsurface, remanent magnetisation may be less of an issue than with larger scale anomalies buried deeply with high susceptibilities.

5.6.4 Negative apparent susceptibilities

As discussed in Sections 2.6 and 2.7, archaeological anomalies can appear as having negative susceptibilities if they consist of lower susceptibility material than the surrounding soil. This is typical of features such as foundations, constructed by low susceptibility materials such as

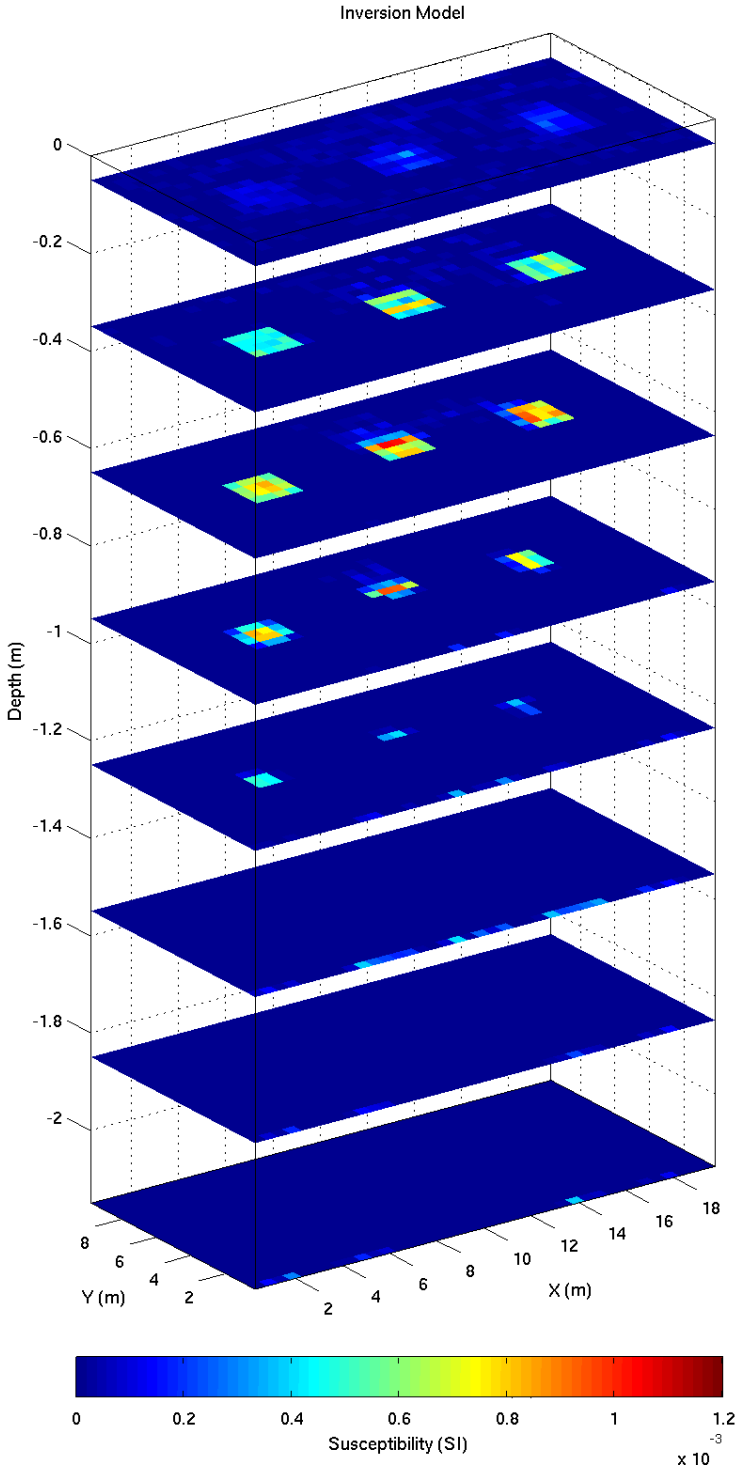

Figure 5.24: Final inversion model for the synthetic dataset shown in Figure 5.22b. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85 and 2.15 mbgl. The actual depths of each prism is between 0.3-0.8 mbgl.

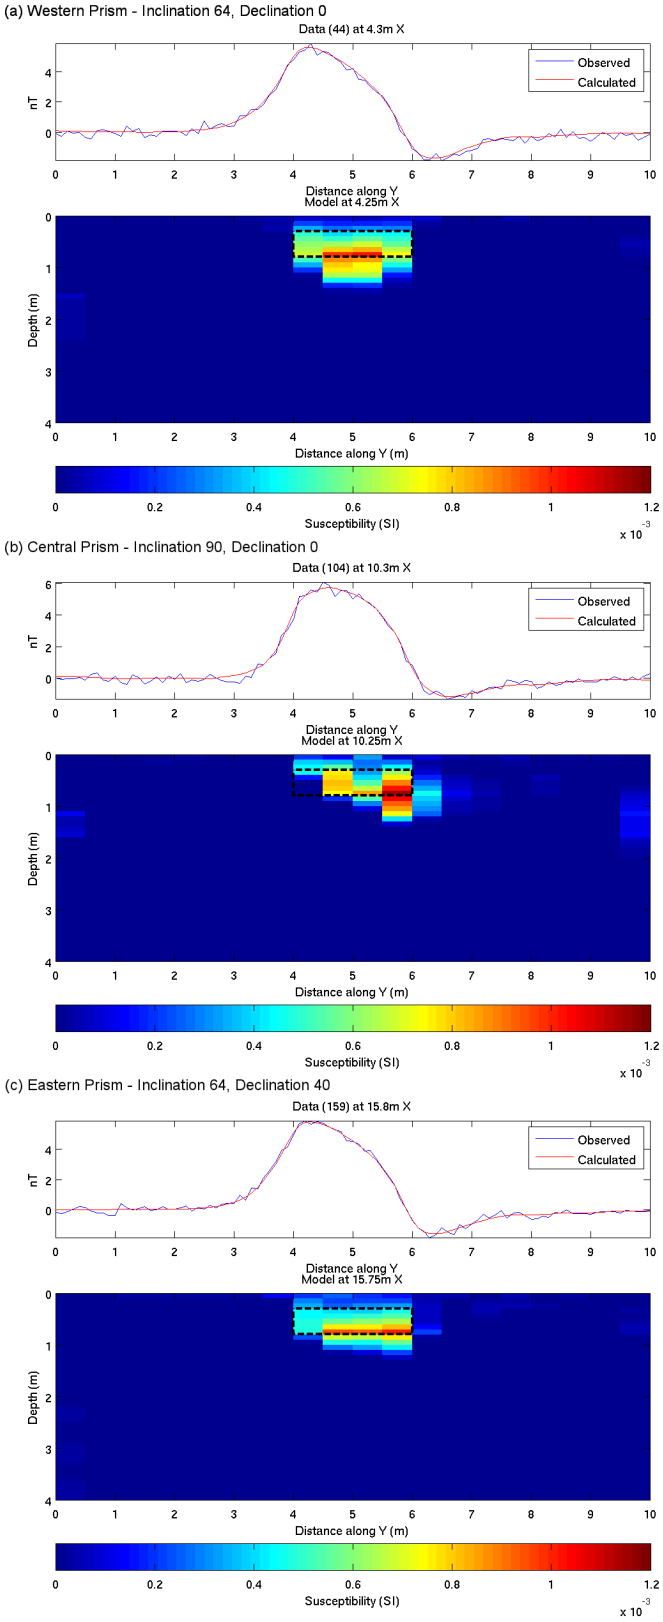

Figure 5.25: N-S Profiles through inversion model shown in Figure 5.24, showing the observed and calculated data. (a) Profile through the control prism whether the inclination and declination match the inducing field. (b) Profile through the central prism where inclination is 90° and declination 0°. Note the prism moves northward with increasing depth. (c) Profile through the eastern prism where inclination is 64° and declination 40°.

limestone and placed within a high susceptibility soil. Once the processing of the data has been conducted to reduce the total field measurement to a total field anomaly dataset, these features are likely to show themselves as negative anomalies, and when modelled, should produce negative susceptibilities.

In order to produce a model with both positive and negative susceptibilities, bound constraints need to be applied. The inverse problem is thus solved by minimising the objective function (Equation 5.3) subject to:

$$m^- < \mathbf{m} < m^+ \tag{5.5}$$

where **m** is the susceptibility, and m^- and m^+ are a minimum and maximum susceptibility assigned by the user. In all the above examples, positivity has been constrained as all the recovered bodies were desired to consist of positive susceptibilities. This is achieved by setting m^- to zero, and m^+ to a large value. Should negative susceptibilities be desired in the model, m^- must be chosen as a negative value.

To produce a model to test MAG3D's ability to model negative susceptibilities, two identical prisms of 2×2×0.5 m have been placed between 0.3 and 0.8 mbgl (See Figure 5.26). The western prism has a susceptibility of 0.001 SI, and the eastern prism -0.001 SI. These have been forward modelling within the inducing field of 37.5 Am^{-1}, with inclination 64° and declination set to 0°, to produce the synthetic observations shown in Figure 5.26b. The positive prism has resulted in a maximum total field response of ∼6 nT and the negative body has produced a response of the same amplitude, but of opposite polarity, around -6 nT.

An inversion has been performed keeping the parameters the same as those used above, but this time specifying maximum and minimum susceptibility limits of -0.05 and 0.05 SI respectively. A mesh of 0.5×0.5×0.1 m cells covering the data area of 16×10 m and extending down to a depth of 4 mbgl. The L-curve produced (Figure 5.27a) has highlighted a very strong "knee"

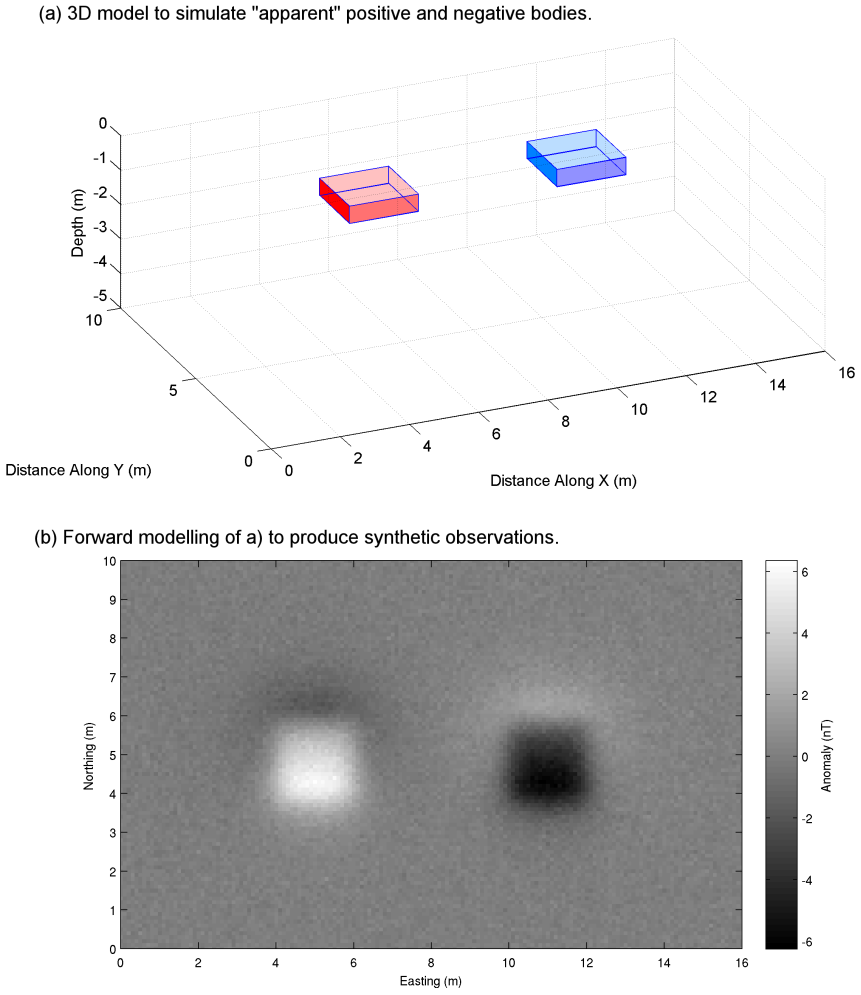

Figure 5.26: Synthetic dataset of two prisms with different polarity magnetisation. Positive on the left, negative on the right. The inducing field is 64° inclination and 0° declination. (a) 3D model showing both prisms to be placed between 0.3 and 0.8 mbgl. (b) Dataset achieved from forward modelling with 0.2 nT standard deviation Gaussian noise added.

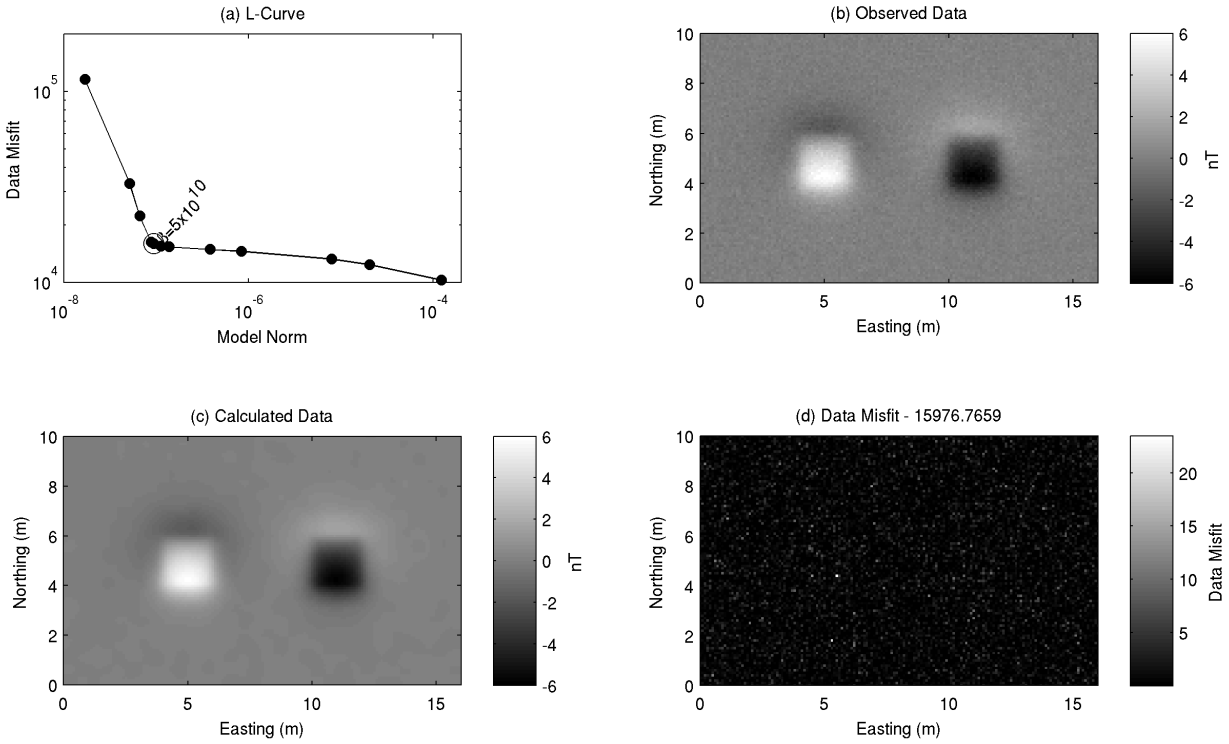

Figure 5.27: Analysis of data misfit for the inversion of the synthetic body in Figure 5.26. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

allowing the selection of an optimum value for β of 5×10^{10}. The data misfit is shown in Figure 5.27d and shows that all the main features have been successfully modelled, with the misfit showing only Gaussian noise.

The final inverse model is presented in Figure 5.28, and shows a disappointing result. Both prisms have been modelled by anomalous susceptibilities which slowly increase in amplitude from the surface to a peak of $\sim$0.0006 SI for the positive body, and -0.0006 SI for the negative body, at around 2 mbgl. Below this depth there is still evidence of both bodies, with the susceptibility decreasing, yet still anomalous to the base of the mesh at 4 mbgl. This has obviously resulted in modelled bodies which are much larger than the 0.5 m thick prisms used to generate the observations.

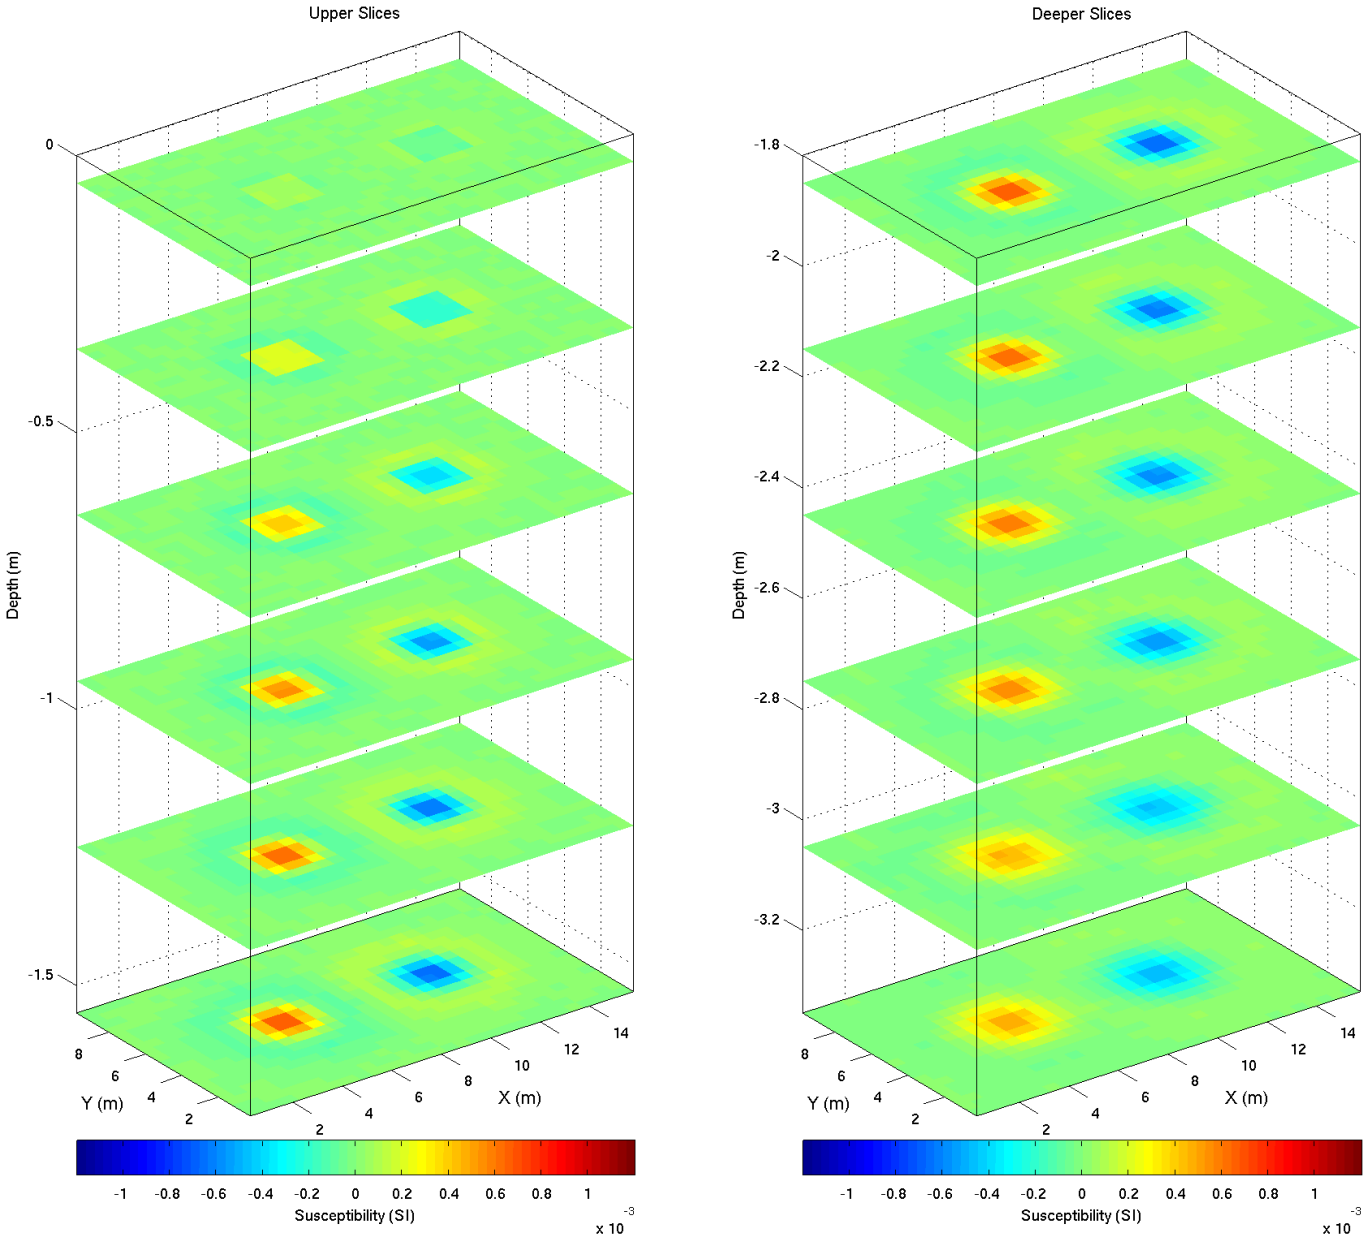

Figure 5.28: Final inversion model for the synthetic dataset shown in Figure 5.26, with bound constraints of -0.05 and 0.05 SI imposed. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl.

The reasons this model creates such a good fit to the data is highlighted in an E-W profile through the model at 4.25 mN, and showing the nearest survey line of calculated and observed data (Figure 5.29). This profile passes through the centre of both prisms, which are clearly observable as the high and low susceptibility anomalies. Focussing on the positive prism, it can be seen that the large positive anomaly is adjacent to lower amplitude negative anomalies on either side. These negative susceptibilities are about -0.0001 to -0.0002 SI. To illustrate the effect these negative features are having on the data misfit, all negative susceptibilities have been removed from the model shown in Figure 5.28. The remaining positive susceptibilities in the model have been forward modelled and the E-W profile shown in Figure 5.29 is presented again in Figure 5.30. It is shown that without the negative susceptibilities on either side of the positive anomaly the calculated data shows a much broader signal and does not produce a good fit to the data either side of the prism location. The negative susceptibilities thus are placed into the model to create a shorter wavelength feature in the calculated dataset, rather than placing the anomaly shallower, which in this case would result in a better fit to the actual prism's location. This shows a disadvantage of using bounds rather than constraining positivity or negativity, as the prisms have been modelled using a combination of features with opposite polarities.

Ideally, a positive anomaly would be modelled by a positive feature, and a negative anomaly modelled by a negative feature, as this would often be the most likely and simplest scenario, and will reduce the number of possible solutions to the inverse problem. Two separate methods have been designed to achieve this.

Placing a positivity constraint on the inversion will obviously force the inversion to fit the anomalies identified in the observations with positive susceptibility features in the model. This has been done in Figure 5.31a. A much improved fit to the actual prism location can be seen for the positive prism as the inversion has been forced to move the body shallower in order to create a good fit to the negative measurements seen on either side of the large positive anomaly. Without allowing regional trends or negative susceptibilities only a poor fit to the negative prism can be made.

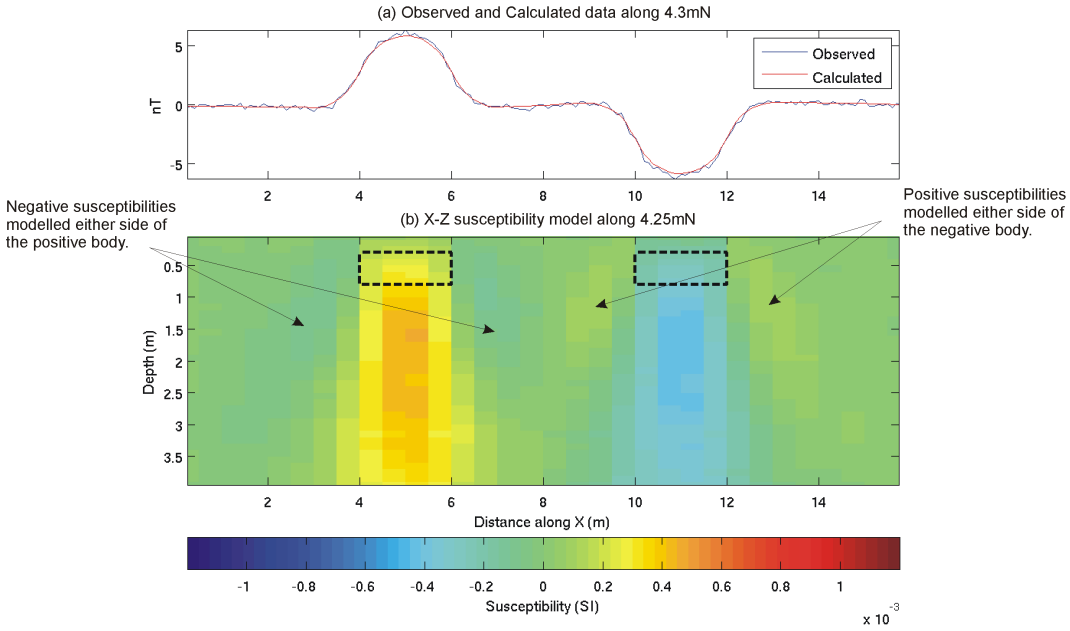

Figure 5.29: East-west profile along 4.25 mN through the inversion model shown in Figure 5.28. The model has been created to fit the observations by creating a negative-positive-negative pattern for the positive anomaly, and vice-versa for the negative anomaly.

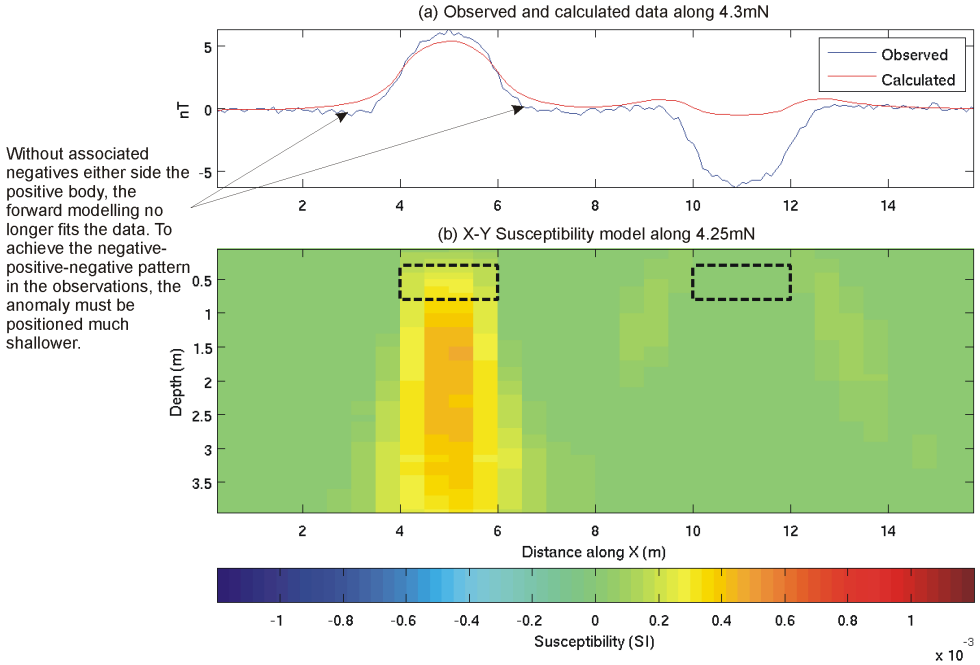

Figure 5.30: Forward modelling of the positive susceptibilities only, in the models shown in Figures 5.28 and 5.29. It is demonstrated that the positive susceptibilities are not sufficient to fit the observed dataset, and the negative susceptibilities that make up the negative-positive-negative pattern observed in Figure 5.29 are essential for providing a good fit.

The same procedure can be conducted for a negativity constrained inversion. A limitation of the MAG3D code, is that m^+ must be greater than zero. In order to overcome this problem, to constrain negativity, the total field observations have had their polarities switched and the modelling is constrained for positivity. In order to correct the final model, the polarity of the recovered susceptibility model is reversed again, resulting in a model which fits the original data, with negativity constraints applied. Forward modelling has shown this process to be an acceptable way for constraining negativity. The negativity constrained inversion is shown in Figure 5.31b, and shows an excellent fit to the negative prism, with a large data misfit in the area of the positive prism.

Making a reference model - Method 1

The model norm, shown in Equation 5.1, penalises structure in the model dependant upon how far it deviates from the reference model. Until now, all inverse models presented have been modelled using a zero halfspace model, in order to recover models with the lowest susceptibilities required to fit the data, and therefore recovering the smallest model. It is possible to use prior knowledge to generate a reference model and thus bias the inversion result, however in the absence of such information the positivity and negativity constrained inversion results can be used.

Figure 5.32 shows a very quick and simple way of generating a reference model. The positivity model has provided a good fit to the positive prism, and failed to fit the negative prism, where the model largely consists of zero susceptibilities. The negativity constrained inversion has produced a good fit to the negative prism, and has zero susceptibilities in the vicinity of the positive prism. Therefore for each cell within the mesh, it is possible to add the susceptibility recovered in the positivity model to the value in the corresponding cell of the negativity model to produce a reference model. This is shown in Figure 5.32. The reference model has been forward modelled to identify how good the data fit is, where it can be seen that the amplitudes

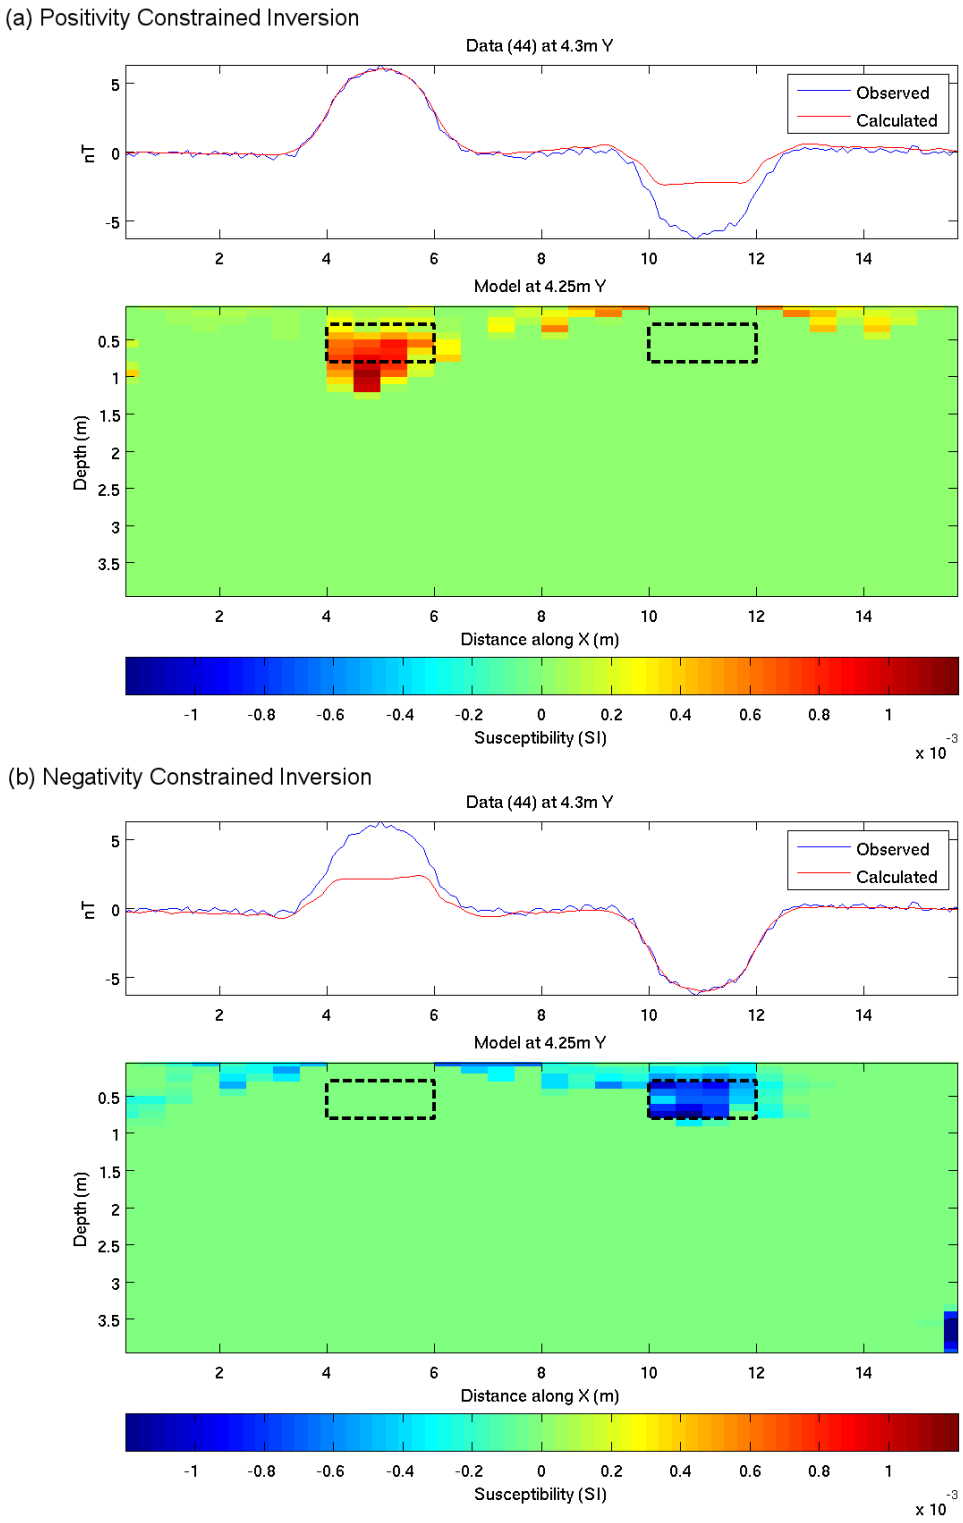

Figure 5.31: E-W profiles through constrained inversion models, with observed and calculated datasets. (a) Positivity constrained by setting m^-=0 and m^+=1. (b) Negativity constrained with m^-=-1 and m^+=0.

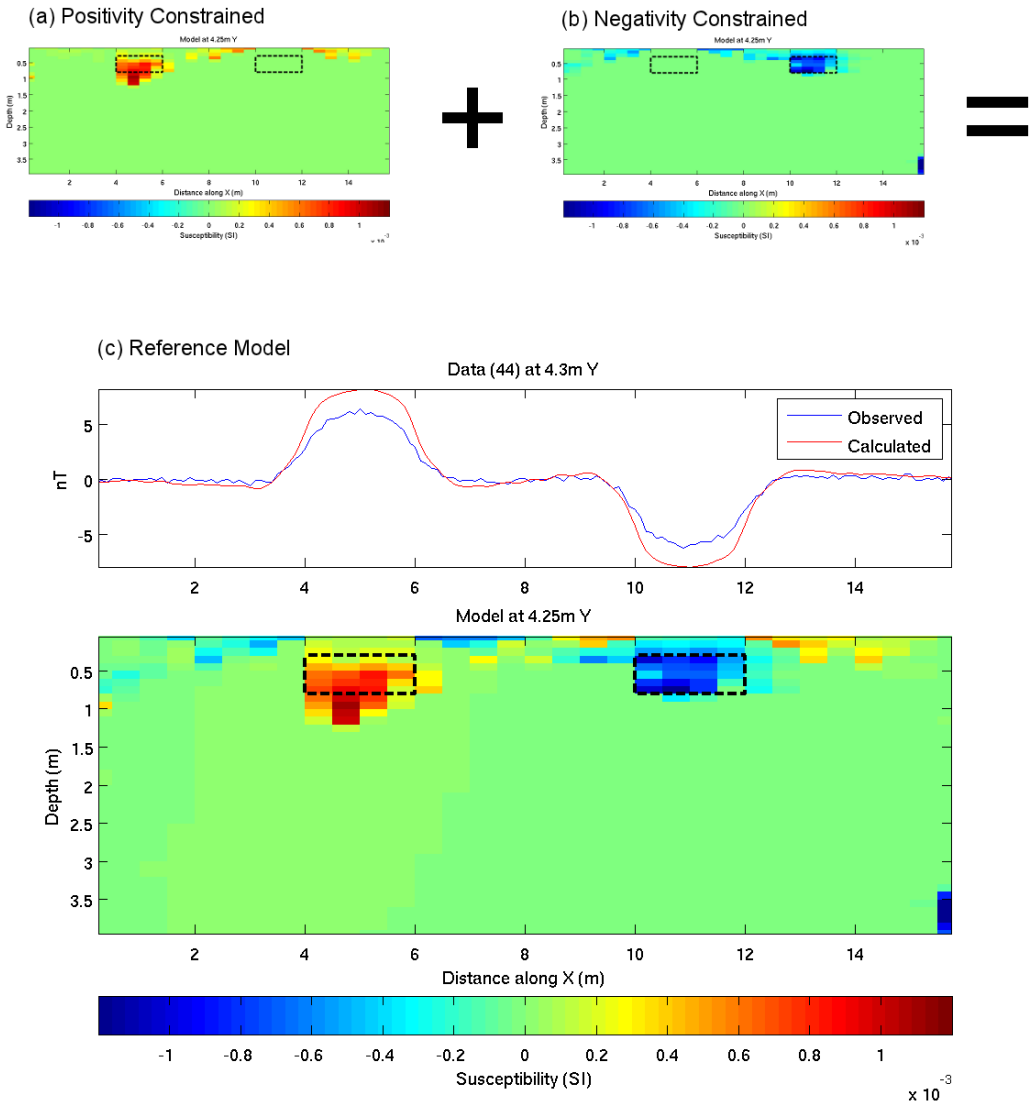

Figure 5.32: Reference model created by adding the value of each cell from the positivity model (a), to the corresponding cell in the negativity model (b). The final model (c), has been forward modelled and shows amplitudes larger than the observations. This is due to the attempts to fit to opposite polarity prism in each model, and a final inversion, using the reference model in (c) is required to produce a model which fits the observations.

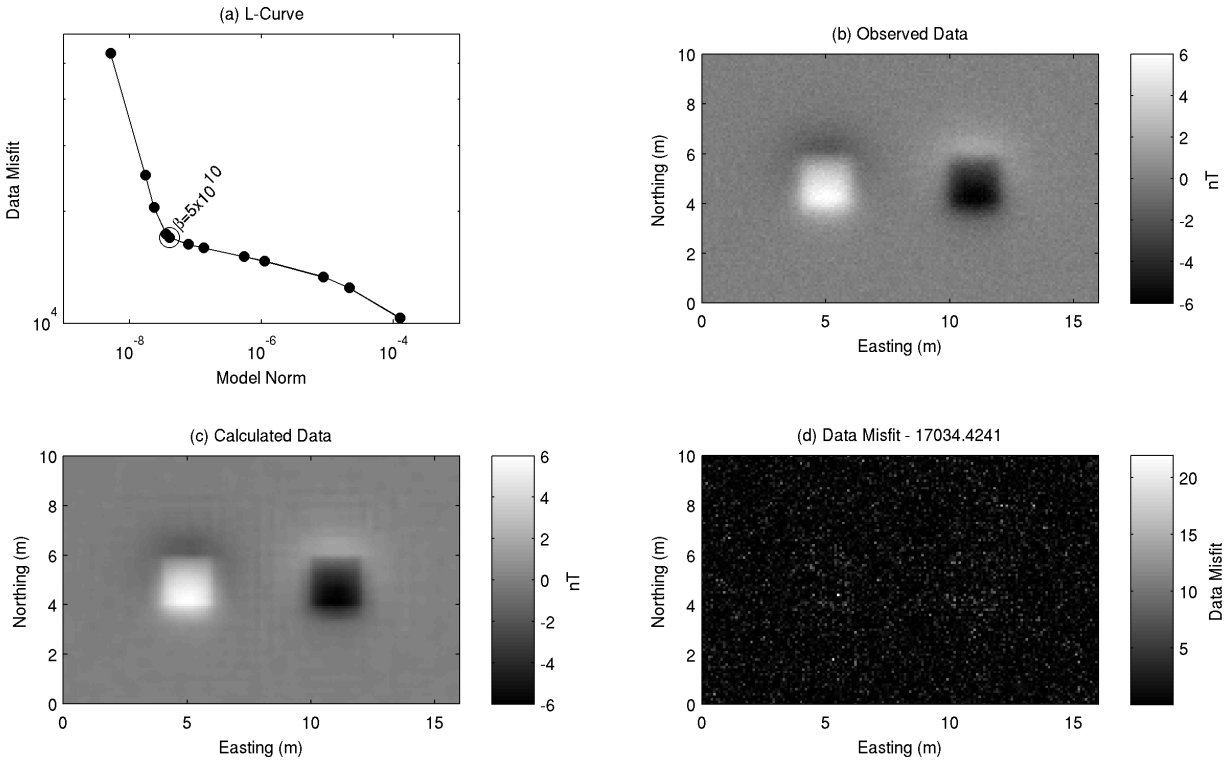

Figure 5.33: Analysis of data misfit for the inversion of the synthetic body in Figure 5.26a. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

of the positive and negative anomalies are over-estimated. This is due to attempts to fit the opposite polarity prism by placing anomalies in the immediate sub-surface. As this model does not produce a good fit to the data, another inversion is required with the bounds set to allow both positive and negative susceptibilities, yet using this model as the reference model should ensure that the structure of the final model produces a similar model to the reference and does not revert towards the model seen in Figures 5.28 and 5.29.

The inversion has been re-run using the model shown in Figure 5.32c as the reference model, and bounds set as $-0.05 < m < 0.05$. The L-curve for this inversion is shown in Figure 5.33a, and an optimum value of β of 5×10^{10} is highlighted. The inversion has produced a model close to the reference model, but that also fits the data. Slices through the final model can be seen in

Figure 5.34, and the E-W profile through the model showing the corresponding observed and calculated data is shown in Figure 5.35.

Several unwanted features that were present in the reference model have been carried over into the final model. Anomalous susceptibility features can be seen in the very top slice of Figure 5.34. These are present due to the attempts of the positivity and negativity constrained inversions to fit the prism of opposite polarity. They can be seen in the profiles of Figure 5.31, and contribute to the over-modelled amplitudes seen in the forward modelling of the reference model (Figure 5.32). High amplitude features in the final set of cells around the edges of the mesh can also be seen. These have not been removed in the final inversion as they are present in the reference model. Instead, to lower the forward modelled amplitude of the responses to the reference model, opposite polarity features have been placed into the model in the areas below the prisms. This is particularly visible in Figure 5.35. The prisms themselves have been modelled well, with the anomalous susceptibilities being placed largely within the volume of the actual prism locations, and the susceptibility approaching 0.001 and -0.001 SI for the positive and negative prisms respectively.

Making a reference model - Method 2

The method described above for creating a reference model could be improved by reducing the amplitude in the outer layer of cells in the mesh to zero. This would remove the attempts to fit the opposite polarity prism, by modelling very near-surface and in the extremes of the mesh. However, as seen in Chapter 3, edge detection allows the horizontal extents of anomalies to be accurately identified, and therefore these techniques can be used to make informed choices about where to pick and choose from the corresponding positivity and negativity constrained models when creating a reference model.

Figure 5.36 shows the edge detection process for the dataset shown in Figure 5.26. Firstly the data are reduced to the pole (Figure 5.36a), and the pseudo-gravity transformation is applied

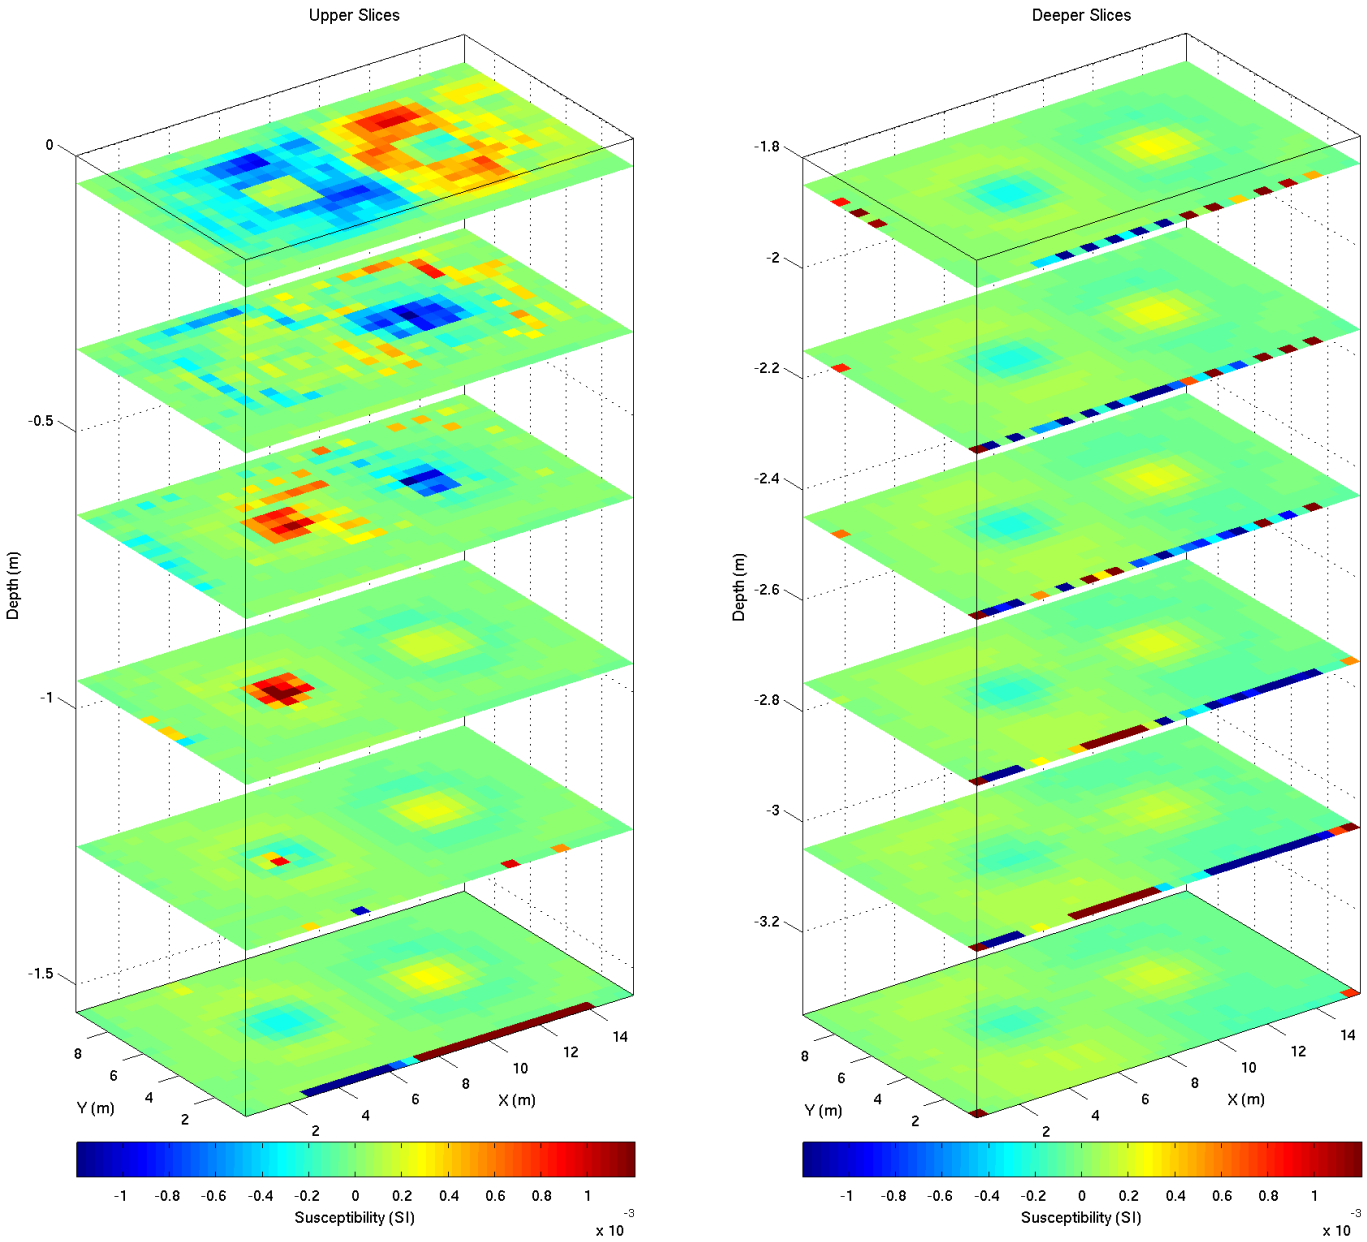

Figure 5.34: Final inversion model for the synthetic dataset shown in Figure 5.26, with bound constraints of -0.05 and 0.05 SI imposed, and a reference model applied using the technique shown in Figure 5.32. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl.

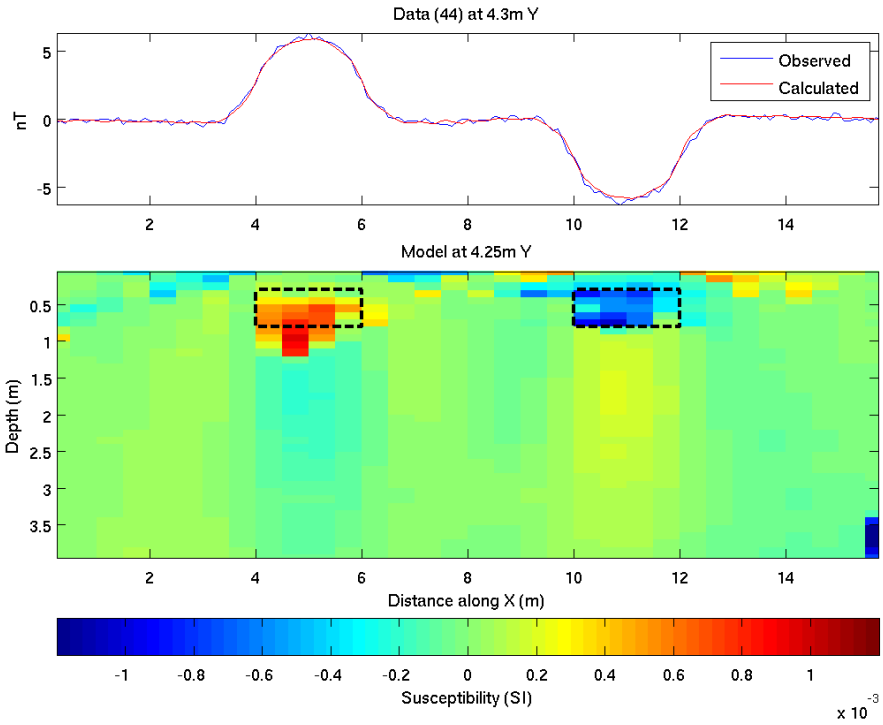

Figure 5.35: E-W profile through final inversion model, produced by using Figure 5.32 as a reference model, and relaxing the susceptibility bounds to allow both positive and negative susceptibilities.

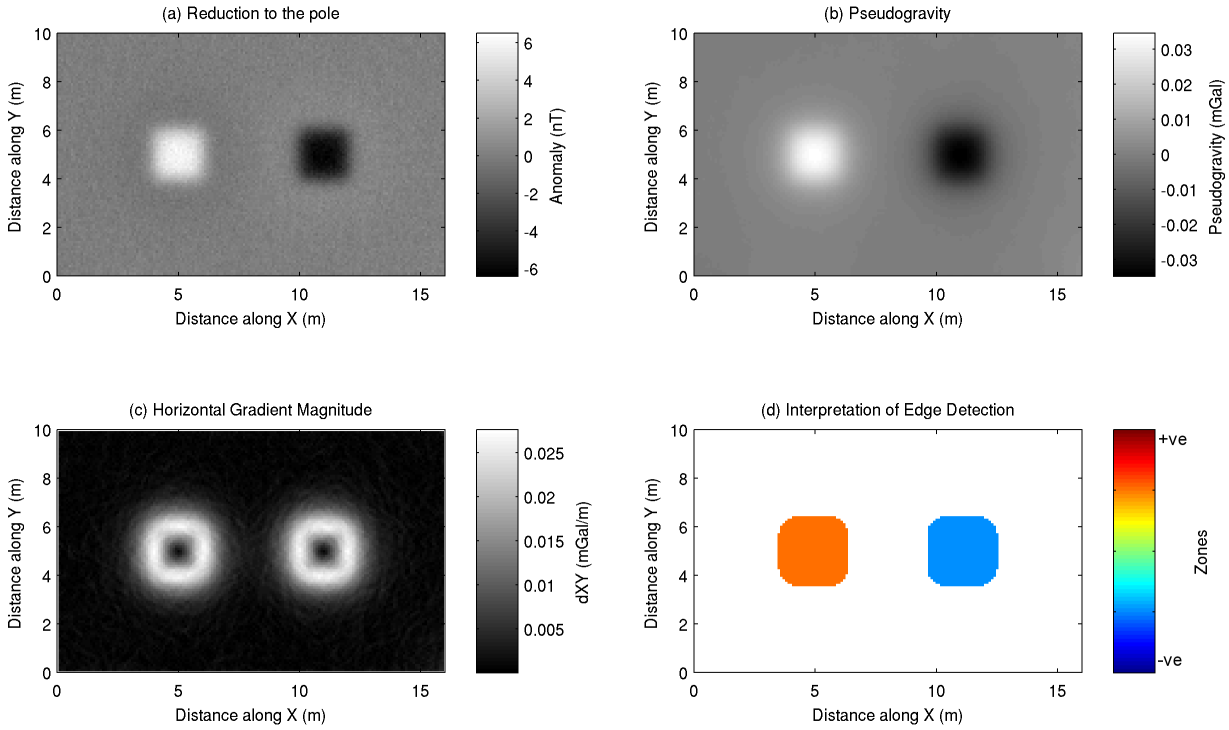

Figure 5.36: Edge detection of dataset presented in Figure 5.26b. (a) Firstly the dataset it reduced to the pole. (b) Pseudo-gravity dataset (c) Horizontal gradient magnitude of pseudo-gravity. (d) Interpretation of c) with indication of polarity of amplitude made from a).

(Figure 5.36b). Then the horizontal gradient magnitude of the pseudo-gravity dataset is calculated (Figure 5.36c) and edges are defined as being located at the peak amplitudes of the horizontal gradient magnitude. This allows the two bodies to be defined by their lateral locations, and an indication can be made as to whether they relate to positive or negative anomalies. This interpretation is shown in Figure 5.36d.

Next, for the surface area that relates to an indicated positive anomaly in Figure 5.36d, the corresponding volume immediately below this point is taken from the positivity model, and vice-versa from the negativity model. This results in a model that is populated immediately below the anomaly areas, and is empty in all other cells. This model can then be used as a reference model for a subsequent inversion, and is shown in Figure 5.37, and the E-W profile along 4.25 mN in Figure 5.38.

The dataset has been re-inverted, this time using Figure 5.37 as the reference model, and the L-curve for the inversion is shown in Figure 5.39a, with an optimum value for β of 1×10^{11}. The final model had deviated away from the reference model sufficiently enough to produce a good comparison between the observed (Figure 5.39b) and calculated (Figure 5.39c) datasets. The data misfit as shown in Figure 5.39d shows no obvious correlated areas of misfit across the survey area.

Slices from the final inversion model are shown in Figure 5.40, and the E-W profile along 4.25 mN in Figure 5.41. The model is very similar to the reference model, and does not show the excess structure in the outer cells as seen with the previous method for creating a reference model in Figure 5.34. The positive and negative prisms have both been modelled well, with the negative prism almost entirely contained within the actual prism volume, whereas the positive prism has been modelled $\sim$0.1-0.2 m too deep, as was the case with the original inversion where the positivity constraint was applied.

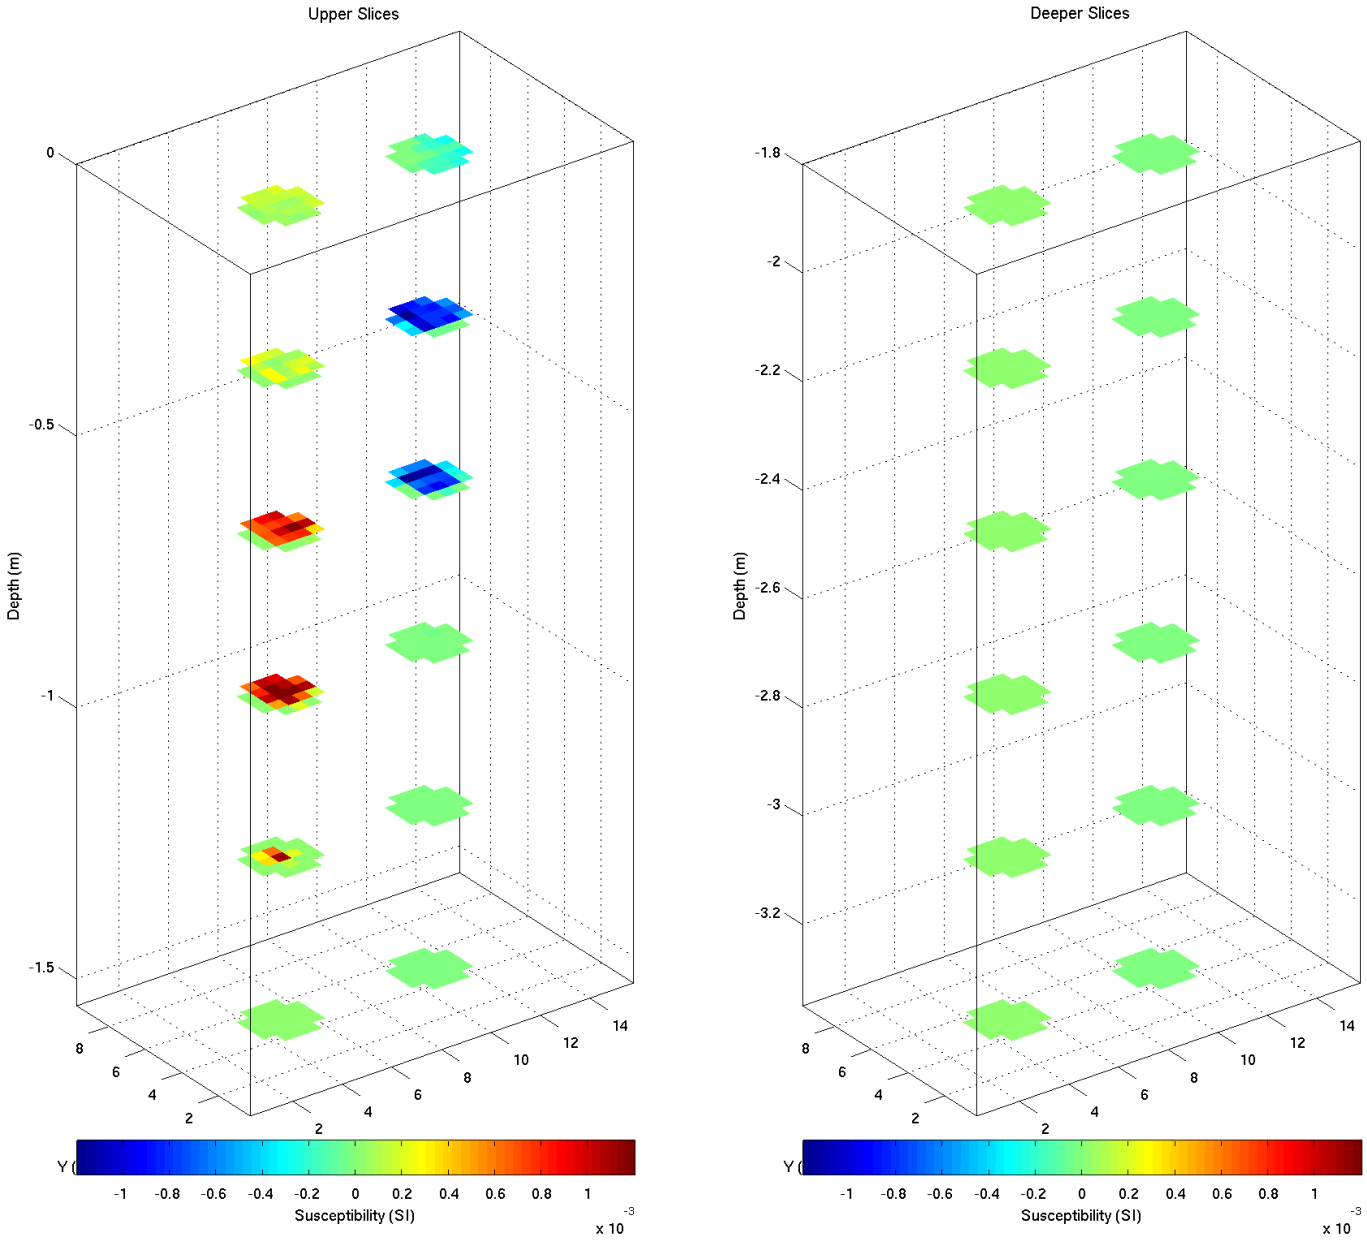

Figure 5.37: Reference model for inversion created by using theta map interpretation to aid picking from positivity and negativity constrained models. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl.

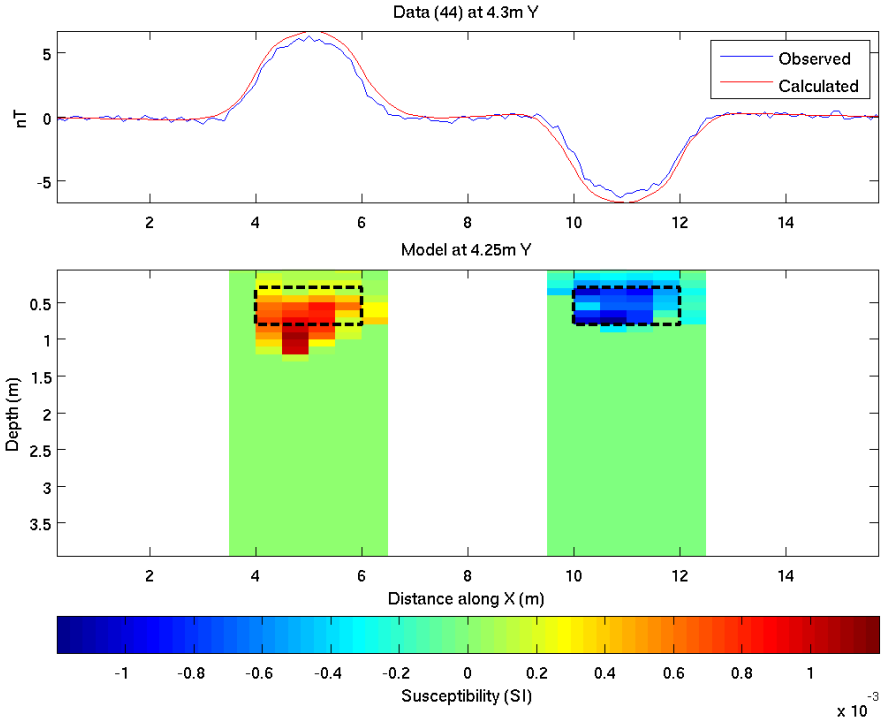

Figure 5.38: E-W profile along 4.25 mN through reference model shown in Figure 5.37. The model has been forward modelled to produce the calculated dataset shown in red. The calculated data produced by this reference model produces a much improved fit to the observations than observed in Figure 5.32.

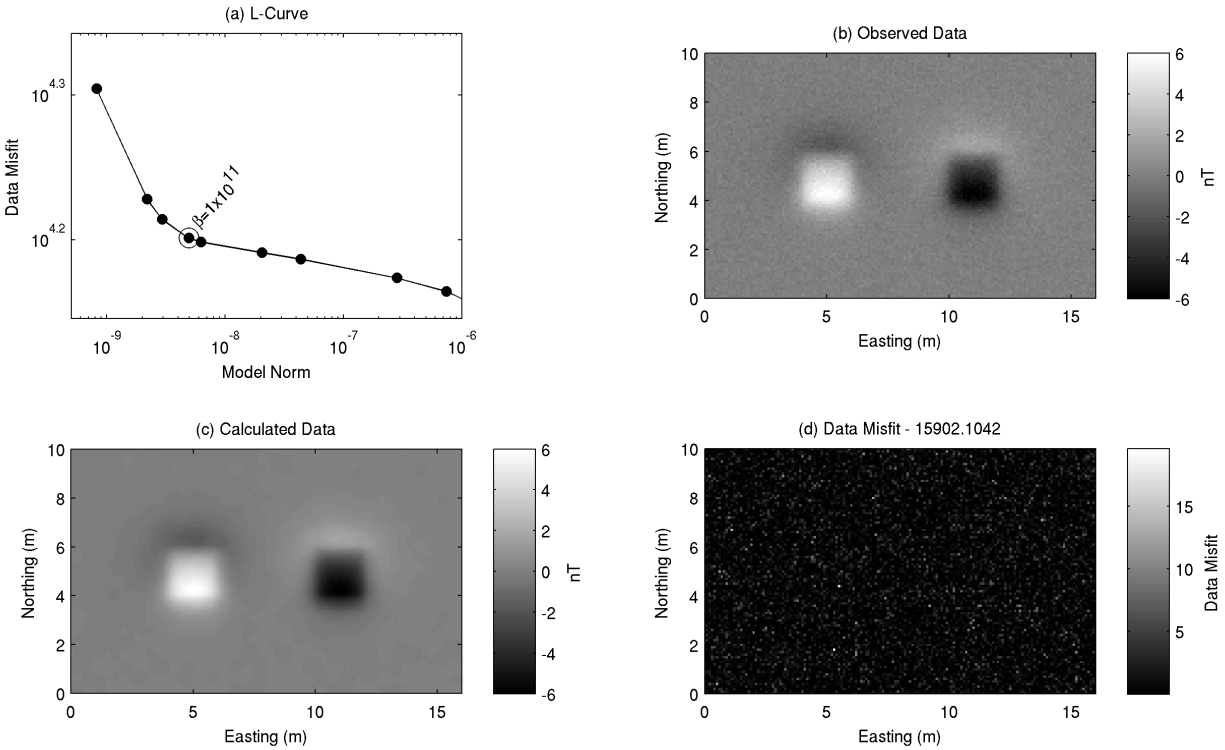

Figure 5.39: Analysis of data misfit for the inversion of the synthetic body in Figure 5.26a, using a reference model shown in Figure 5.37. (a) L-curve analysis, with the chosen value of β shown. (b) Synthetic total-field data generated at a height of 0.2 m above ground level, with Gaussian noise of SD=0.2nT added. This is used as the "observed" data in the inversion. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

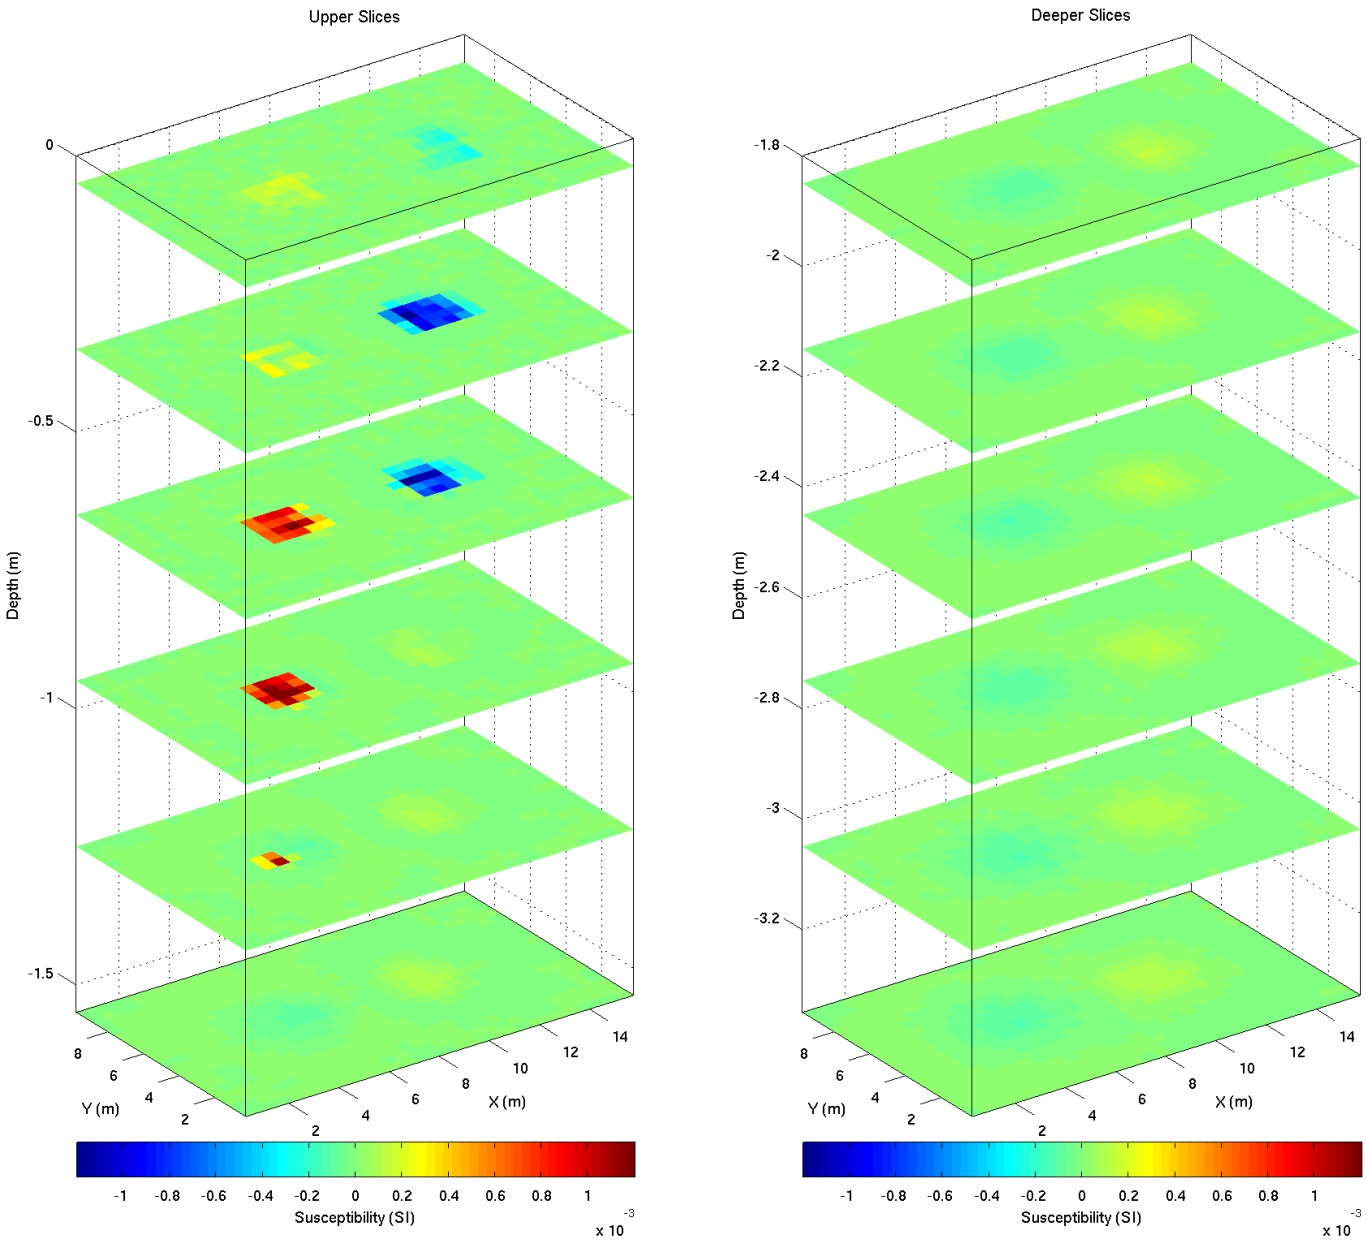

Figure 5.40: Final inversion model for the synthetic dataset shown in Figure 5.26, with bound constraints of -0.05 and 0.05 SI imposed, and a reference model applied using the technique shown in Figure 5.37. Slices shown have their centres at 0.05, 0.35, 0.65, 0.95, 1.25, 1.55, 1.85, 2.15, 2.45, 2.75, 3.05 and 3.35 mbgl.

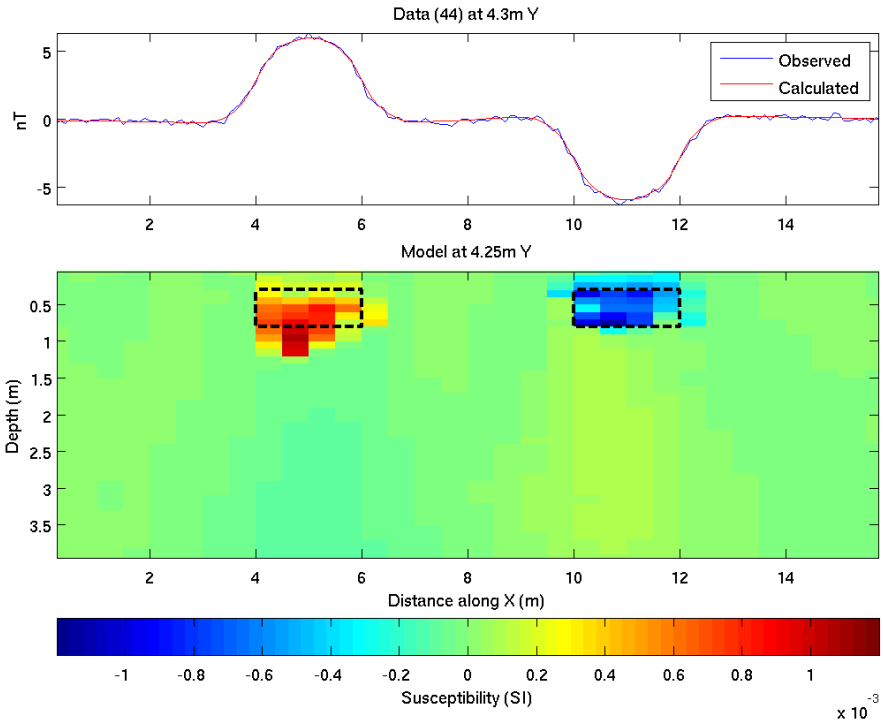

Figure 5.41: E-W profile through final inversion model, produced by using Figure 5.37 as a reference model, and relaxing the susceptibility bounds to allow both positive and negative susceptibilities.

5.7 Conclusions

It has been demonstrated here that it is possible to invert synthetic archaeological datasets using the MAG3D code. Due to the code being designed for larger scale aero-magnetic surveys, some of the default methods for the inversion are no longer appropriate, and care must be taken to recover a model that is the most likely to represent the subsurface archaeology.

MAG3D attempts to minimise an objective function that is a trade-off between a model norm and a data misfit. There are three choices for selecting the value of the regularisation parameter (β):

- The Chi-squared method works by assuming the value for the standard deviation of the noise in the datafile is known to be correct by the user. Making this assumption the inversion can produce a result rapidly, however it is shown here that if the user has marginally under or over-estimated the noise level, this will have a dramatic effect on the final model recovered. In real-world archaeo-magnetic surveys it is unlikely that accurate measurements of the noise level will be available, and therefore this technique of estimating β is unlikely to be useful.
- Generalised Cross Validation (GCV) of the dataset is a statistical approach to calculating a value for β that has been shown here to produce successful results, even when noise levels have been incorrectly assigned to the datafile.
- To give the user maximum control over the selection of β, calculating the model norm and data misfit for several values of β and identifying the "knee" in the subsequent L-curve is demonstrated to be a favourable way of analysing the trade-off, despite the extended run times required for multiple inversions.

Typically for inversions run using MAG3D, a mesh design featuring a central core consisting of small cells in the volume of interest immediately below the survey area, with larger

cells surrounding this, is favoured. This approach has not been very successful at recovering archaeological-type anomalies. The regional cells, poorly constrained by nearby data, have been shown to account for significant regional trends across the survey area, even for datasets where no regional trend is present. This has a dramatic effect on the ability to accurately recover the archaeological features.

- It is therefore essential that any regional trends that exist in the dataset are removed prior to inversion, and that only the core central cells of the mesh are used in the inversion to ensure that all features in the dataset are accounted for in the mesh immediately below the observation points.

- Similarly, where anomalies are present in the data that are thought to derive from within the mesh-extents, but are not the primary focus of interest, it is also desirable that these are filtered prior to inversion, as it has been shown that any mis-modelling of these features will have a significant knock-on effect in the modelling of the features of real interest.

Due to the nature of archaeological features buried within susceptible soils, archaeo-magnetic anomalies are often both positive and negative with respect to the background level of the magnetic field. It is therefore required that the inversion can recover a model that includes both positive and negative apparent susceptibilities.

- This can be achieved by setting constraints on the magnitude of the susceptibility that covers both positive and negative susceptibilities. A problem with this, is the design of the model norm, which requires the model to stay as close as possible to a reference model. In the absence of a reference model a zero halfspace will be used. It is demonstrated here that the inversion will model dipolar anomalies with both positive and negative susceptibility contrasts. Using opposite polarity susceptibilities to control the wavelength of the forward modelled signal allows depth uncertainty to occur, as short wavelength features

can be modelled deeply that otherwise can only be modelled by moving the body to much shallower depths in the model.

- In order that anomalous bodies are only modelled by single polarity features, a positivity or negativity constraint can be applied to the inversion. This will result in some of the anomalies in the dataset being modelled in one inversion, and some in the opposite polarity inversion, and it is obviously desirable to combine these within a single model. This combined model can then be placed within the model norm of a subsequent inversion as a reference model, to ensure that the forward modelled dataset from the final model fits the observed dataset.

- An easily accomplished reference model is formed by simple addition of the respective cells from the positivity and negativity models. The final model using this technique has been shown to recover the synthetic bodies well, however, as a consequence of attempting to fit negative anomalies with positive susceptibilities and vice-versa, many unwanted features remain in the final model. It is therefore useful to be more selective in which parts of each model are taken to be included within the reference model.

- Using the techniques introduced in Chapter 3, lateral edge detection has been used to identify areas of the surface below which anomalous bodies are likely to be present. An indication of the polarity of each body can be taken by analysis of the polarity of the observations, and a much more refined reference model can be formed. This has been shown to be successful in recovering the synthetic bodies spatial location, extents and susceptibility, without introducing large unrealistic anomalies across the remainder of the model's volume.

Chapter 6

Preparation of archaeo-magnetic data prior to inverse modelling

6.1 Introduction

The previous chapter highlighted that not all magnetic datasets are suitable for inverse modelling. Common problems are associated with datasets containing information relating to sources distal to the survey area, or the presence of repetitive near-surface linear anomalies that are not the primary target. Also, where sources contain remanent magnetisation in a direction that differs from the direction of the inducing field, inversion was shown to calculate incorrect and unrealistic models of the subsurface.

Prior to inverting datasets which include these characteristics, it is therefore necessary to filter the data, so that the inversion only attempts to recover the features of interest. Here, high-pass filtering will be investigated as a method for separating local and regional anomalies, and directional filtering will be used to remove unwanted linear features which may represent modern or historic land uses, or surveying artefacts.

It is not possible to remove the influences of remanent magnetisation from a magnetic dataset unless the location and direction of the causative body is known. Therefore, here, attempts to

identify the presence of remanence are made in order to assist the identification of datasets that are suitable for modelling with a single magnetisation vector.

6.2 Regional fields

It was shown in Section 2.11.2 that the measured magnetic signal can be converted into the Fourier domain to allow processing or filtering based on the wavenumber or spatial frequency of the signal. This is generally done by applying weights to the power at particular wavenumbers, in order to boost or suppress the contributions these wavenumbers have in the filtered dataset.

Filtering in the Fourier domain can be represented as:

$$G(k_x, k_y) = A(k_x, k_y).F(k_x, k_y) \tag{6.1}$$

Where *G* is the Fourier representation of the filtered dataset, *A* is the filter operator applying the weighting, and *F* is the Fourier representation of the signal. k_x and k_y are the wavenumbers in the *x*- and *y*-directions, and related to wavelength (λ) by $k = \frac{2\pi}{\lambda}$. The final filtered dataset is obtained by calculating the inverse Fourier transform.

The filter operator (*A*) is usually a two dimensional function. If the same weighting was applied to all wavenumbers the function will be space invarient, and simply result in multiplying the signal across the survey area. A filter that applies only at *F(0,0)* would result in a DC-shift of the dataset.

6.2.1 High-pass filtering

High-pass filtering can be applied to remove the regional anomalies and separate the features of interest. The wavenumber-dependant attenuation of the magnetic field with distance results in anomalies caused by features located at a large distance away from the measurement point, represented by long wavelength (low wavenumber) features in the data. Therefore, high-pass filters can be applied to suppress the energy in the low wavenumbers and remove the effect of the fields generated by distal features.

Here, three types of filter are compared to identify their suitability in removing the regional trend from the compound synthetic dataset introduced in Figure 5.14. These are a simple high-pass filter, the Butterworth filter and the Cosine roll-off filter.

The simple high-pass filter can be shown as:

$$\begin{aligned} A(k_x, k_y) &= 0, \text{ for } k < k_0 \\ A(k_x, k_y) &= 1, \text{ for } k > k_0 \end{aligned} \tag{6.2}$$

where k_0 is chosen by the user, below which all wavenumbers will be discounted.

Gibbs's phenomena

The filter shown in Equation 6.2 will result in an abrupt change in the Fourier domain around k_0. This has the effect of creating a sharp discontinuity in the spectrum, and increasing the likelihood of Gibbs's phenomena. Gibbs's phenomena occurs when a signal is represented by a Fourier series. Increasing the number of Fourier coefficients improves the fit to the function, except in the area around the discontinuity. Here, an 'overshoot', and 'undershoot' exist either side of the discontinuity. Increasing the number of coefficients results in a narrowing of the region of poor fit, but does not reduce the amplitude of the 'overshoot' or the 'undershoot'.

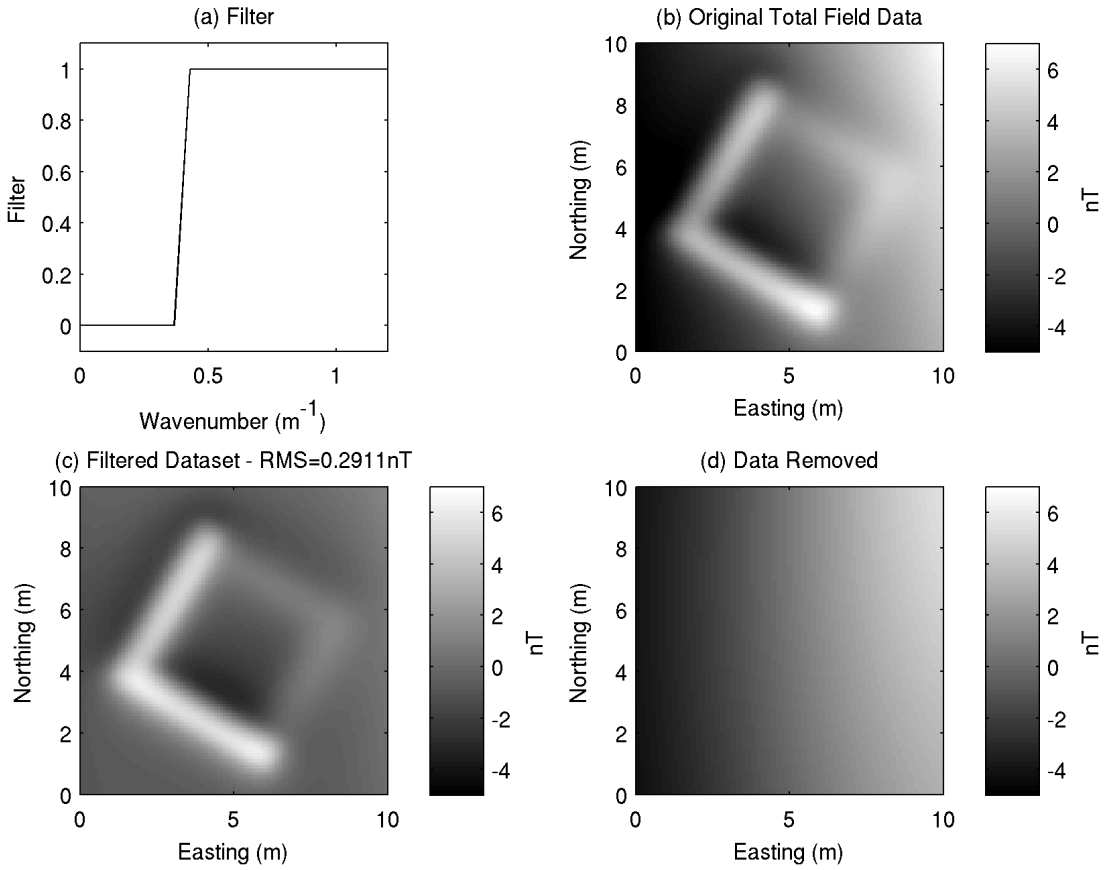

Figure 6.1: High pass filtering of regional trends from compound synthetic dataset. (a) Simple filter removing energy from wavenumbers less than 0.4 radian/m. (only shown up to 1.2 m^{-1}) (b) Synthetic dataset with regional trend, as shown in Figure 5.14. (c) Filtered dataset. (d) Data removed by filtering.

When the inverse Fourier transform is applied, this can result in ringing being observed in the output.

The simple high-pass filter has been applied to the synthetic dataset as shown in Figure 5.14, with the aim of removing the regional trend and producing a dataset that only represents the response to the compound body. A value for $k_0 = 0.4$ has been applied, and the filter (A) can be seen in Figure 6.1a. The filter has been applied to separate the observed total field dataset (Figure 6.1b), into the local (Figure 6.1c) and regional (Figure 6.1d) datasets. In this case, it is the local dataset that is desired for further modelling, and the regional trends are desired to be removed.

It can be seen that the filter has satisfactorily removed the trend which is represented by high amplitudes to the east, and low amplitudes to the west. In this scenario, we know the 'ideal' outcome, as it is the response to the compound anomaly shown in Figure 3.1b. Therefore to calculate the difference between the filtered dataset and the optimum, a simple root mean squared (RMS) approach has been used to quantify the difference between the 'ideal' and filtered output. The RMS error is defined as:

$$\text{RMS error} = \sqrt{\frac{\sum_{i=1}^{n}(T_i - \bar{T}_i)^2}{n}} \tag{6.3}$$

where T is the data from Figure 3.1b and $\bar{T}$ is the filtered output. The RMS error between the filtered dataset using the simple high-pass filter, and the ideal scenario is 0.29 nT. Although the result does not appear to be affected by Gibb's phenomena, the Butterworth and Cosine roll-off filters have been developed to offer ways to taper the transition between the 'pass' and 'reject' areas of the filter.

6.2.2 Butterworth filter

The Butterworth filter is shown by Geosoft (2010b) to be:

$$A(k_x, k_y) = 1 - \frac{1}{\left[1 + \left(\frac{|k|}{k_0}\right)^n\right]} \tag{6.4}$$

where k_0 is the central wavenumber, defining the pass/reject point of the filter. With the Butterworth filter a smooth transition between the 'reject' and 'pass' sections of the filter is applied, dependant on the degree of the filter, n. Higher values of *n* result in a more abrupt transition, whereas lower values of *n* will produce a smoother transition that will suppress wavenumbers further away from the central wavenumber.

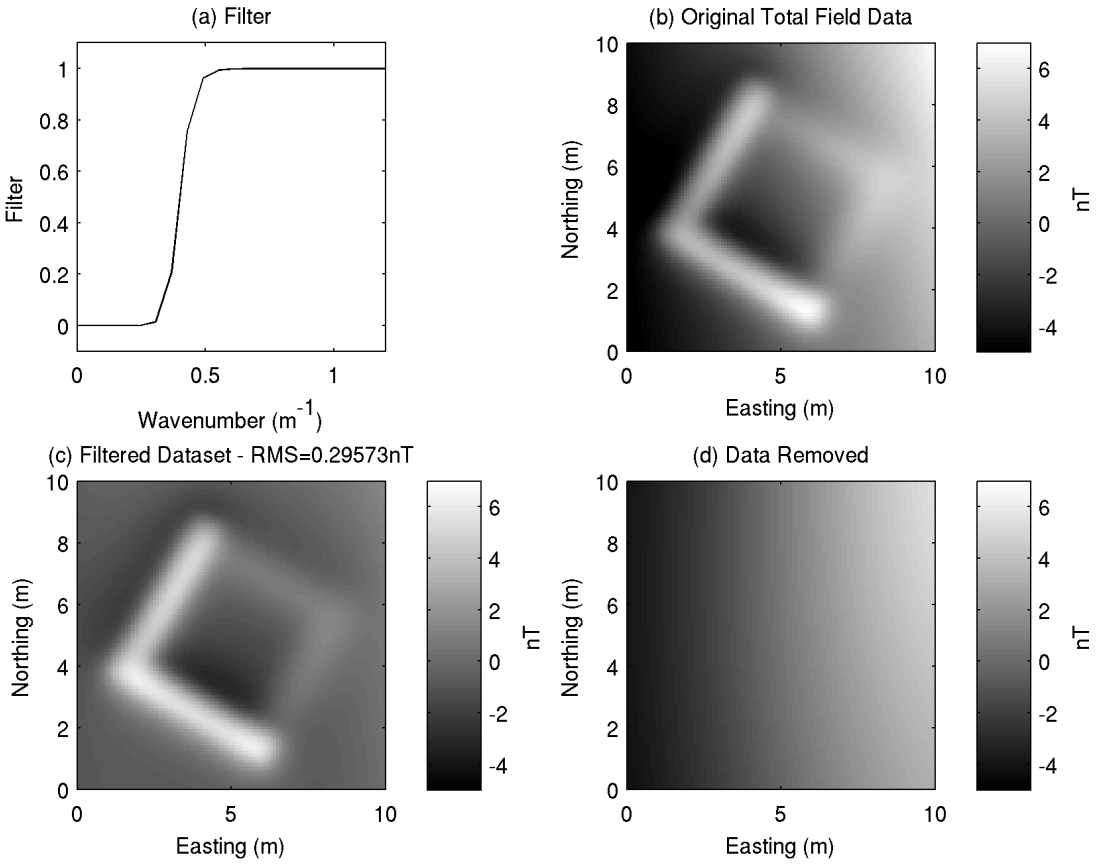

Figure 6.2: Butterworth filtering of regional trends from compound synthetic dataset. (a) Butterworth filter using $k_0 = 0.4$ and $n = 16$ (only shown up to 1.2 m^{-1}). (b) Synthetic dataset with regional trend, as shown in Figure 5.14. (c) Filtered dataset. (d) Data removed by filtering.

The results of filtering the regional trend from the synthetic dataset using a Butterworth filter of $k_0 = 0.4$ and $n = 16$ are shown in Figure 6.2. Again the filter has performed well at separating the regional and local datasets, and the output shows no sign of the trend from west to east. The RMS error between the filtered dataset and the ideal scenario in this case is 0.30 nT.

6.2.3 Cosine roll-off filter

The advantage of the Butterworth filter is that the transition is tapered around a central wavenumber, however the amount of filtering either side of this central wavenumber is dependant on the degree (*n*) applied, and is not strictly defined by the user. The Cosine roll-off filter allows the user to ensure which parts of the spectrum are filtered. It is defined as (Geosoft, 2010b):

$$\begin{aligned} A(k_x, k_y) &= 0, && \text{for } k < k_0 \\ A(k_x, k_y) &= 1 - cos^n \left[\frac{\pi}{2} \left(\frac{|k| - k_0}{k_1 - k_0} \right) \right], && \text{for } k_0 \leq k \leq k_1 \\ A(k_x, k_y) &= 1, && \text{for } k > k_1 \end{aligned} \tag{6.5}$$

The user is required to input two wavenumbers, k_0 and k_1. The transition between 'reject' and 'pass' is conducted solely between these two parameters, at the degree specified by *n*

Figure 6.3 shows the application of the Cosine roll-off filter to the synthetic scenario, using the parameters of $k_0 = 0.3, k_1 = 0.5$ and $n = 1$. Again, the filter has performed well at separating the local (Figure 6.3c) and regional (Figure 6.3d) fields from the original data. Comparison to the ideal scenario shows an RMS error of 0.31 nT.

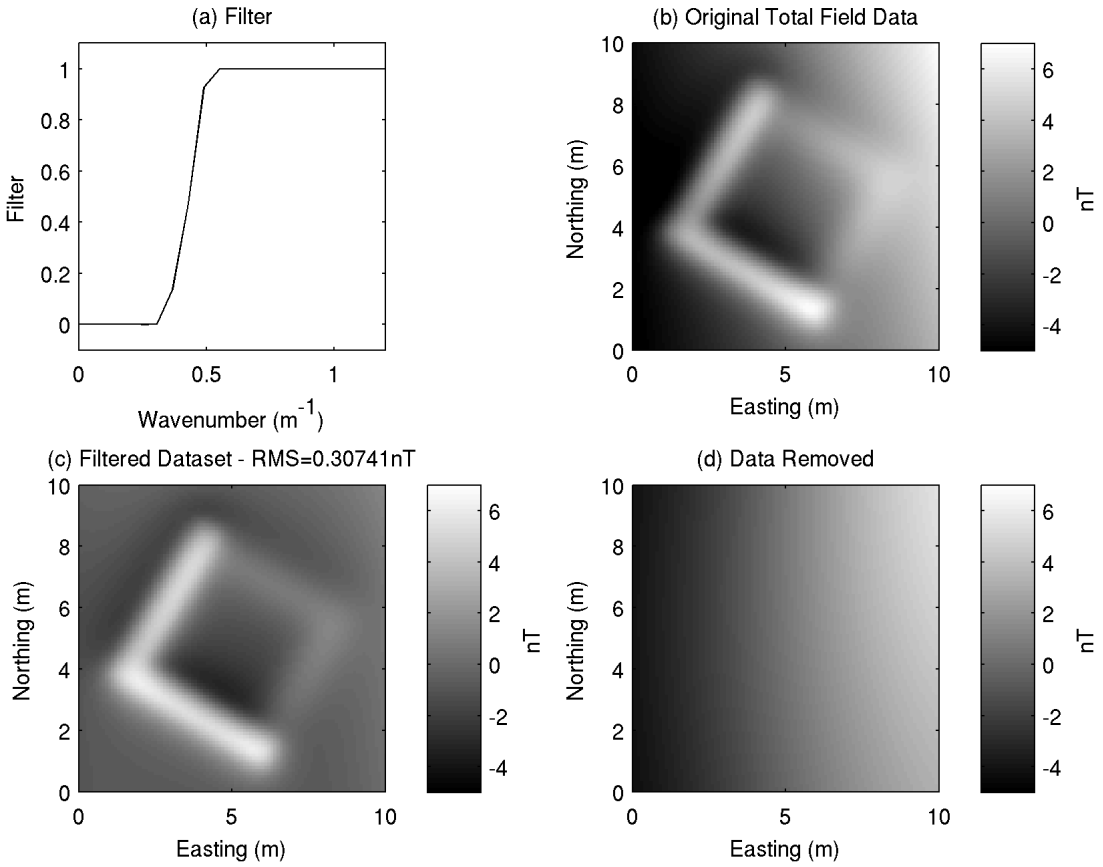

Figure 6.3: Cosine rolloff filtering of regional trends from compound synthetic dataset. (a) Cosine roll-off filter using $k_0 = 0.3, k_1 = 0.5$ and $n = 1$ (only shown up to 1.2 m^{-1}). (b) Synthetic dataset with regional trend, as shown in Figure 5.14. (c) Filtered dataset. (d) Data removed by filtering.

Filter	k_0	k_1	n	**RMS Error**
High Pass	0.2	-	-	1.24 nT
High Pass	0.3	-	-	0.61 nT
High Pass	0.4	-	-	0.29 nT
High Pass	0.5	-	-	0.38 nT
High Pass	0.6	-	-	0.52 nT
Butterworth	0.4	-	2.0	0.75 nT
Butterworth	0.4	-	4.0	0.43 nT
Butterworth	0.4	-	8.0	0.33 nT
Butterworth	0.4	-	16.0	0.30 nT
Cosine roll-off	0.1	0.7	1.0	0.41 nT
Cosine roll-off	0.2	0.6	1.0	0.35 nT
Cosine roll-off	0.3	0.5	1.0	0.31 nT

Table 6.1: Comparison of high-pass filtering techniques for removal of regional trends from a synthetic scenario

6.2.4 Comparison of high-pass filters

The three filters have been tested using a range of constants in order to obtain the optimum result. The RMS error for each combination can be seen in Table 6.1. In reality the optimum result will not be quantifiable, and decisions on the parametrisation are likely to be determined by the user based on visual inspection of the filtered dataset.

For the simple high-pass filter, an optimum central wavenumber of 0.4 m^{-1} is identified from Table 6.1. When a central wavenumber is chosen lower than this, the regional trend is not completely removed from the signal, and a central wavenumber choice higher than 0.4 m^{-1} begins to shown evidence of the removal of the synthetic archaeological body.

The optimal parametrisation of each filter has led to similar results, as can be observed in Figures 6.1, 6.2 and 6.3. Very little difference can be seen in the RMS errors derived from the results, so it could be said that each filter may be used to successfully separate the local and regional fields for this particular scenario. The problems associated with Gibb's phenomena are not obviously seen here. This is not an entirely surprising result however, as the sharp discontinuities introduced into the spectrum by the application of the simple high-pass filter will be

positioned around the central wavenumber. In this scenario, a filter with a central wavenumber of 0.4 m^{-1} will result in unrealistic energy with a wavelength of 15.7 m being introduced into the filtered dataset. Any excess energy introduced into this part of the spectrum would therefore represent itself as a gentle trend in the output local dataset, and is unlikely to significantly affect the RMS error, or appear as an obvious ringing pattern in the final dataset of this size.

In order to easily visualise the effects of choosing a filter with an abrupt transition between the pass and reject wavenumbers, the same regional trend was modelled on a 50×50 m area (Figure 6.4a). This time the trend is shown without the response to the synthetic archaeological body. Figures 6.4b, 6.4c, 6.4d and 6.4e, demonstrate filtering of the dataset by applying a Butterworth filter of $k_0 = 0.4\ m^{-1}$, and ever increasing degrees of n, to show the output from a very smooth transition in the filter to a more abrupt, and finally, completely abrupt transition with the simple high-pass filter (Figure 6.4f). As the degree of the filter is increased, the signal with a wavelength of 15.7 m becomes more pronounced, producing a Gibb's phenomena effect of ∼-0.5 to 0.5 nT. Therefore, even though the effects of Gibb's phenomena are difficult to visualise in the filtering of the 10×10 m area, care should be taken to reduce the likelihood of introducing unwanted signal.

6.3 Linear patterns

Relic agricultural features such as ridge and furrow, or more recent ploughing furrows and seed beds, often have a magnetic signature which is recorded in magnetic surveys over archaeological sites. These types of anomalies appear as regular linear patterns across entire sites. Gaffney & Gater (2003) suggest aligning the survey direction parallel to topographic expressions of the agricultural features, using pattern identification from aerial photographs if upstanding topographic features cannot be identified on site. If this approach is adopted, it should be possible to remove the effect of the agricultural features using de-striping algorithms typically used to correct directional errors and sensor offsets from zigzag traverses (see Section 2.11.1). However,

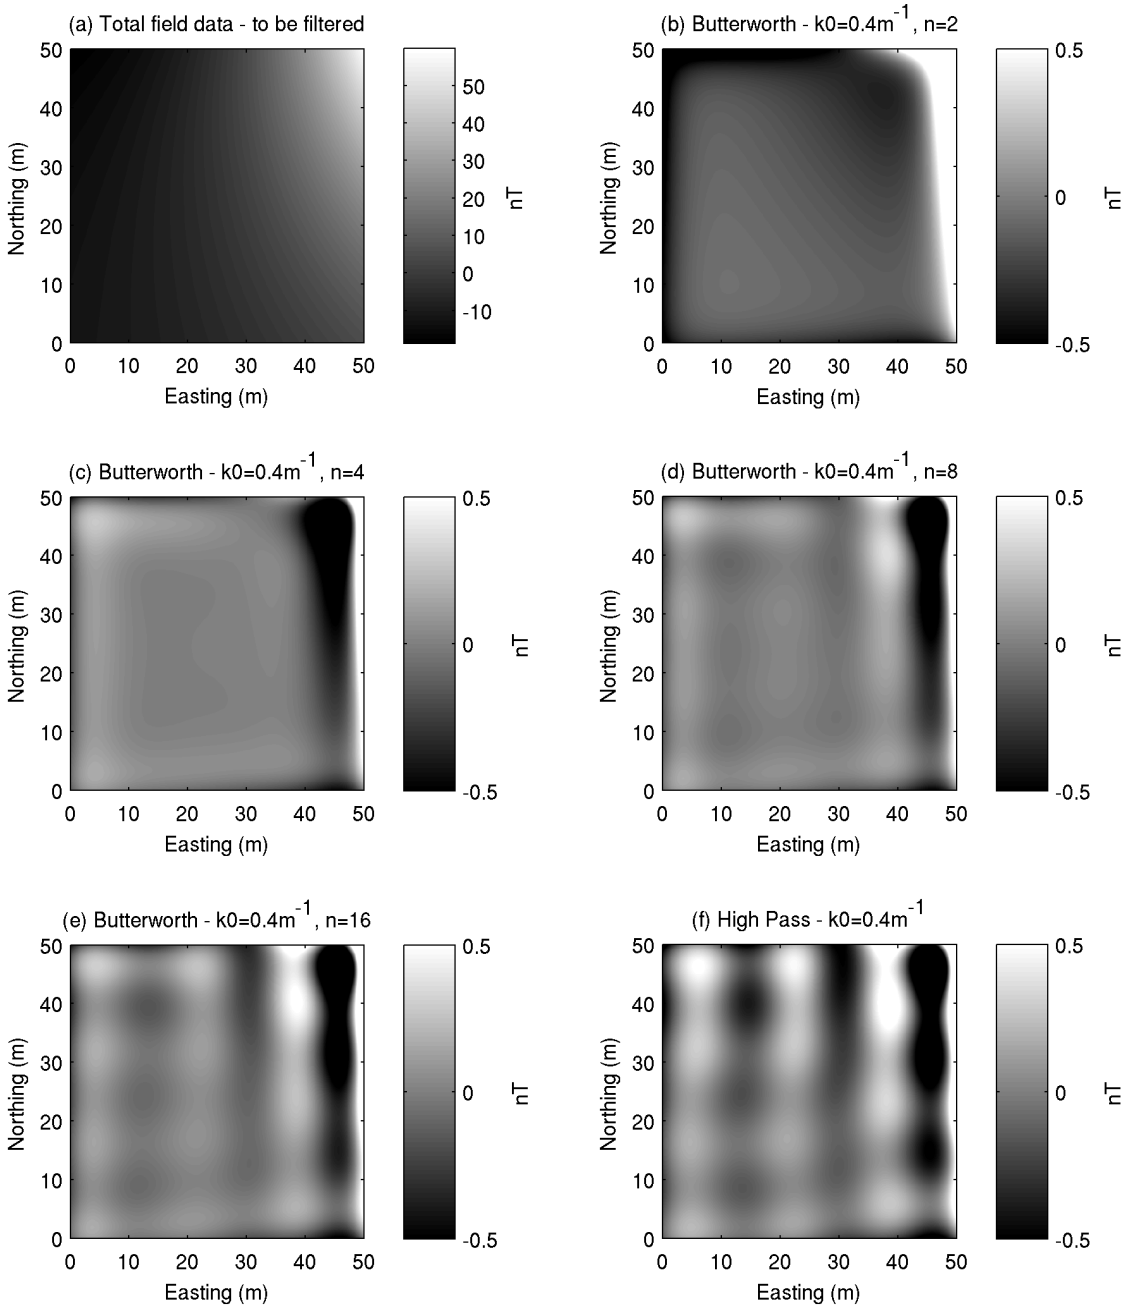

Figure 6.4: Analysis of the risk of introducing Gibbs phenomena during high-pass filtering. (a) Regional trend across a 50×50 m data area. (b) Filtering of a) with a Butterworth filter of $k_0 = 0.4\ m^{-1}$ and $n = 2$ (very smooth transition between pass and reject parts of the filter). (c) Filtering of a) with a Butterworth filter of $k_0 = 0.4\ m^{-1}$ and $n = 4$. (d) Filtering of a) with a Butterworth filter of $k_0 = 0.4\ m^{-1}$ and $n = 8$. (e) Filtering of a) with a Butterworth filter of $k_0 = 0.4\ m^{-1}$ and $n = 16$. (f) Filtering of a) with a high-pass filter with central wavenumber at $k_0 = 0.4\ m^{-1}$ (very sharp transition between pass and reject parts of the filter).

when it is not possible, or convenient, to incorporate the direction of unwanted features into the survey design, it becomes necessary to identify and remove them using directional filters during post-processing.

6.3.1 Spectra of repeating near-surface anomalies

The Fourier transform of a continually repeating sine wave will have a peak at a particular wavenumber which is defined by the wavelength of the signal. Figure 6.5a shows a sine wave with a wavelength of 1 m. The Fourier domain representation of this signal is shown in Figure 6.5b. As expected the energy of the function peaks at 2π divided by the wavelength, therefore here is 6.28 m^{-1}.

Assuming that the input signal is composed of the combination of the magnetic field generated by the underlying archaeology of interest, and the sine wave (simulating a plough pattern) superimposed, the perfect filter to remove the effect of the sine wave in the Fourier domain would be the inverse of Figure 6.5b. However, this information is unlikely to be available, yet it could be possible to design a filter to suppress the energy around a peak wavenumber. In reality, the responses from ground disturbed by agricultural processes are unlikely to produce a perfect sine wave anomaly signal, and of course, the real archaeological signal is also likely to have significant energy at similar wavenumbers to the plough pattern also.

Figure 6.5c shows the magnetic field anomaly generated by a repeating 0.5 m wide by 0.1 m deep 2D prism oriented perpendicular to the profile. The signal appears similar to the sine wave anomaly (Figure 6.5a), however the energy spectrum is more complicated and consists of several peaks and troughs throughout the Fourier domain (Figure 6.5d). Figure 6.5e shows prisms similar to Figure 6.5c, however this time a 1 m gap is between each pair of prisms. This is designed to replicate a scenario similar to track marks caused by a tractor. The Fourier transform of the signal is shown in Figure 6.5f. As in the previous example, the energy spectrum is made up of a number of peaks and troughs at a range of wavenumbers.

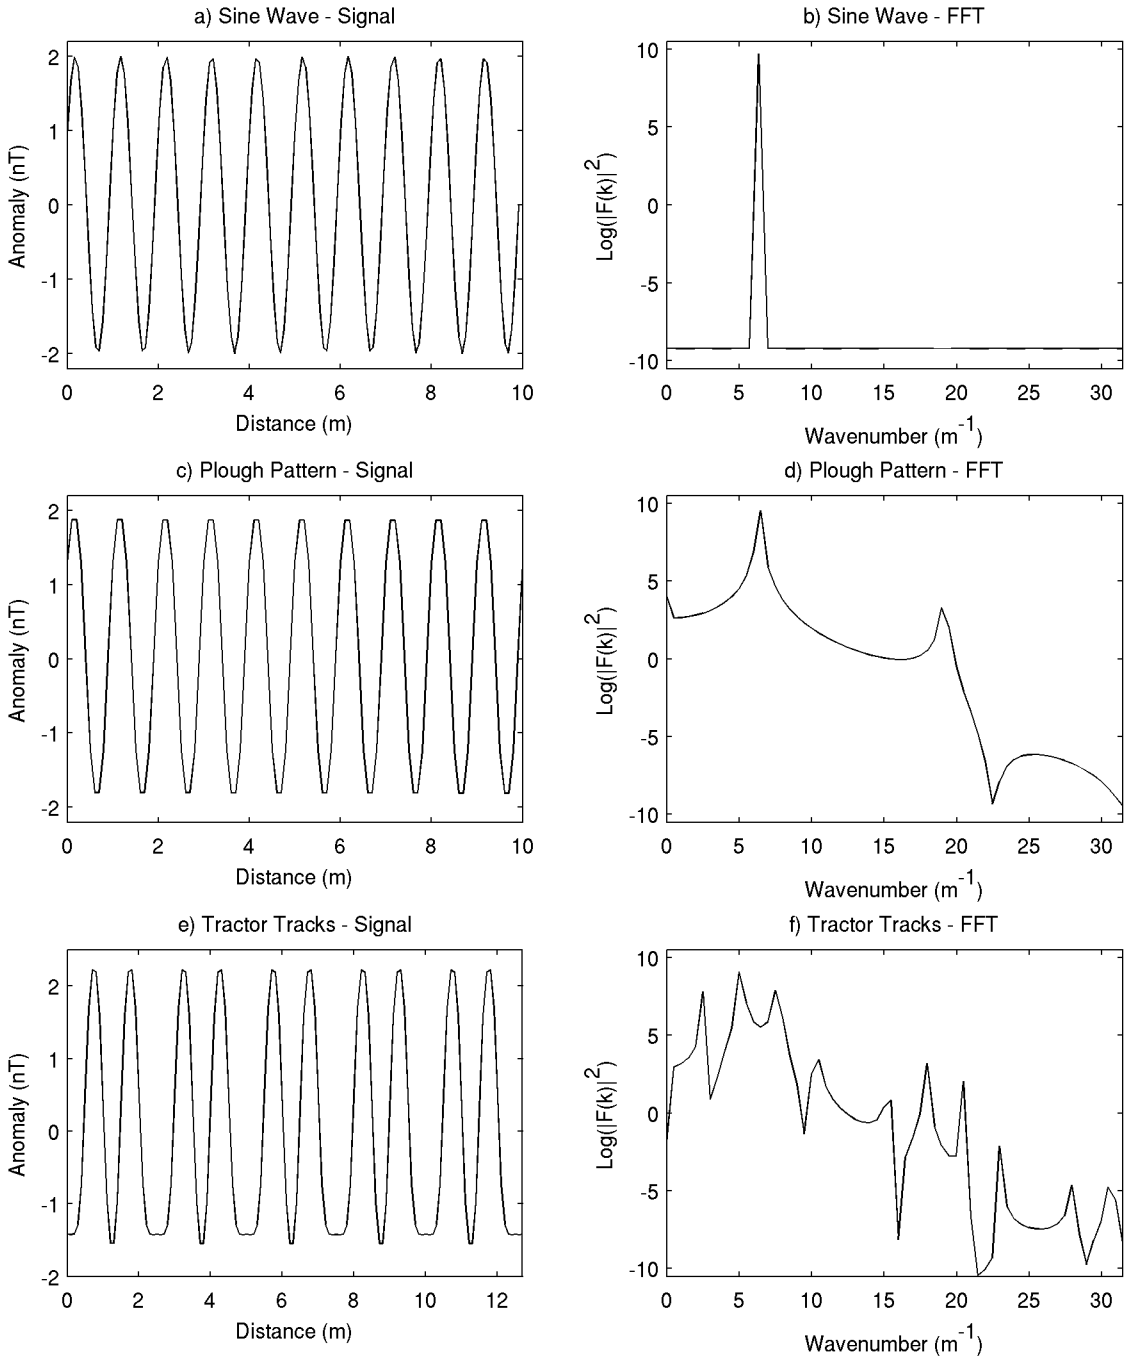

Figure 6.5: (a) Sine wave of wavelength=1m. (b) FFT of a). (c) Repeated "plough" anomaly of prisms 0.5m wide, 0.1m deep and infinite length positioned 0.5m apart. (d) FFT of c) (e) Repeating "tracks" anomaly pattern of pairs of ∞ x 0.5 x 0.1m prisms positioned 0.5m apart, separated by 1m. (f) FFT of e)

It can be seen in these 2D examples that filtering of realistic signals from agricultural sources is much more complicated than for the sine wave example, as the energy of the anomaly is spread out across the entire spectrum.

6.3.2 Synthetic dataset

In order to test the effectiveness of various directional filtering algorithms, they have been applied to the synthetic dataset shown in Figure 5.19, which is shown again in Figure 6.6. The dataset consists of a combination of the archaeological (Figure 6.6a) and agricultural signal (Figure 6.6b), and is shown in the space domain in Figure 6.6c and the Fourier domain in Figure 6.6d. Although it is not actually "noise" the agricultural signal is referred to here as the agricultural noise, as it is the undesired part of the dataset. The agricultural signal is generated as described in Figure 6.5e, and has a bearing of 170°.

The aim of the directional filtering routines presented here, is to remove the agricultural noise in order to produce a dataset that most closely represents the archaeological signal. The effect of the agricultural noise in the Fourier domain can be seen as a linear peak in Figure 6.6d running from the origin to $\sim$30 Kx, 5 Ky and mirrored in the negative Kx domain.

6.3.3 Directional pass/reject filter

Once the bearing of the features to be removed is identified, it is possible to isolate and remove all power from the wavenumbers representing this bearing. This method is presented by Fuller (1967) and Bezvoda *et al.* (1990), and typically is done by selecting a minimum and maximum bearing. Any wavenumbers within this zone are set to zero in the Fourier domain, all other wavenumbers remaining the same. The inverse Fast Fourier Transform (FFT) is then performed and the power of lineations along the selected bearings will be reduced.

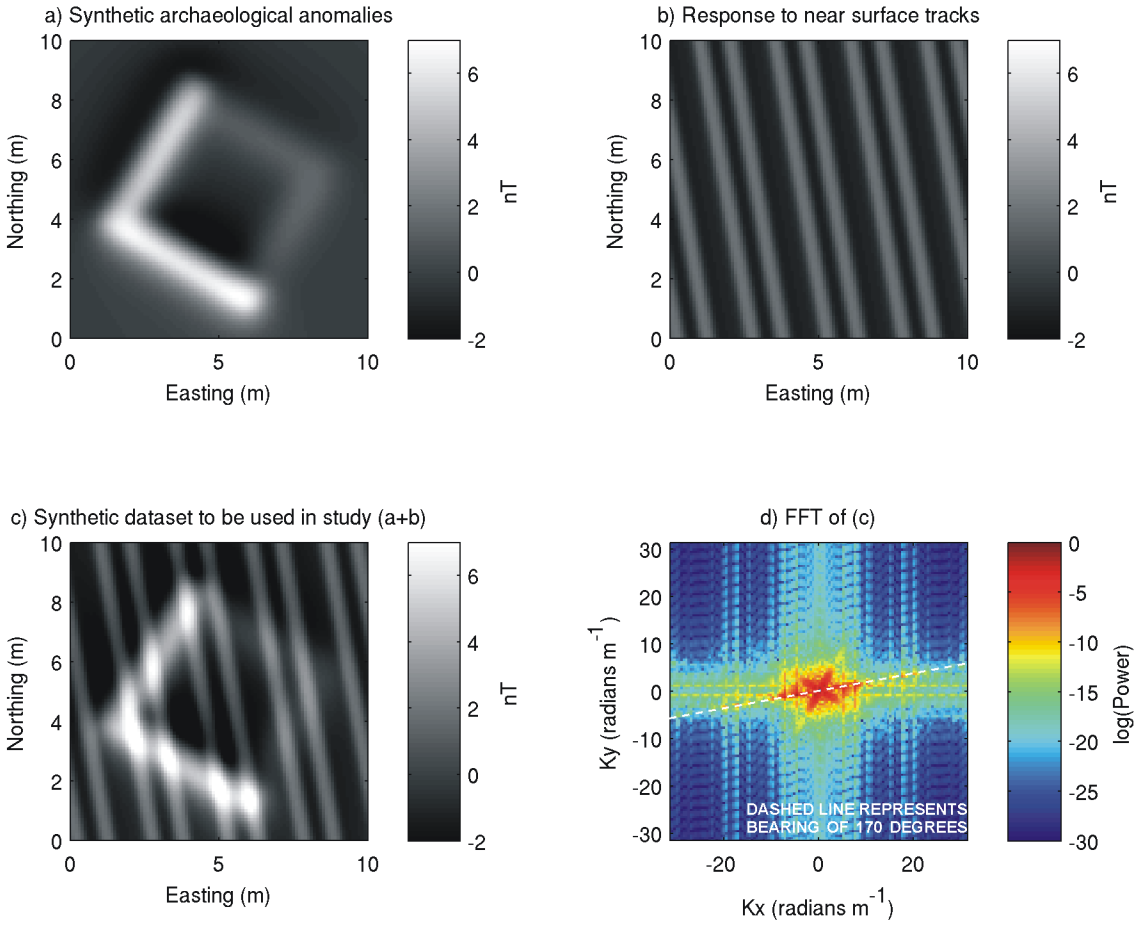

Figure 6.6: (a) Synthetic archaeological signal required from the data (b) Near-surface agricultural "tracks" required to be filtered out by the directional filtering. (c) Synthetic dataset to be used, combining a) and b). (d) 2D FFT of c), showing the effect of the agricultural noise as a linear peak running roughly left to right.

In reality, editing the spectra to zero within the filter zone would create a sharp discontinuity within the spectra, and increases the likelihood of Gibbs phenomena, where ringing is often observed once the spectra is transformed back into the space domain. In order to reduce the effects of Gibbs phenomena, the filter function should be allowed to change from 1 to 0 in a smooth way. Scollar (1970) discusses several methods for this, one of which is to cosine taper the filter function. This is the method that has been used here.

Linear anomalies present in the space domain can be identified in the Fourier domain along an azimuth rotated 90° from the azimuth observed in the space domain. For example, a N-S trending pattern, will appear as peaks in the left to right direction in the Fourier domain, and an E-W trending pattern will result in peaks in the vertical direction of the Fourier domain. Figure 6.7a shows the filter operator to be multiplied by the Fourier domain data. The minimum azimuth of the filter is set as 160° and the maximum is 180°. In between these two angles the filter operator equals zero, and no power will be retained in this area. In order to reduce Gibbs phenomenon, tapering has been conducted between 155° and 160° and 180° and 185°. The filter is multiplied by the Fourier domain data and the resulting filtered spectrum is shown in Figure 6.7b. The inverse FFT is then performed and the filtered dataset in the space domain is presented in Figure 6.7c. The filter has performed a good job at removing the majority of the agricultural noise, as little of the 170° pattern remains in the filtered dataset. In order to identify the data that has been removed, the filtered dataset is subtracted from the original dataset, and the data removed by the filtering is shown in Figure 6.7d. A small amount of the amplitude of the larger archaeological anomaly has been removed from the data, and is shown as a higher amplitude in the removed data towards the SW corner.

6.3.4 Directional cosine tapering of filtered zone

An alternative to completely removing the power from wavenumbers relating to the filtering direction, is to increasingly suppress the power depending on how close the data are to the

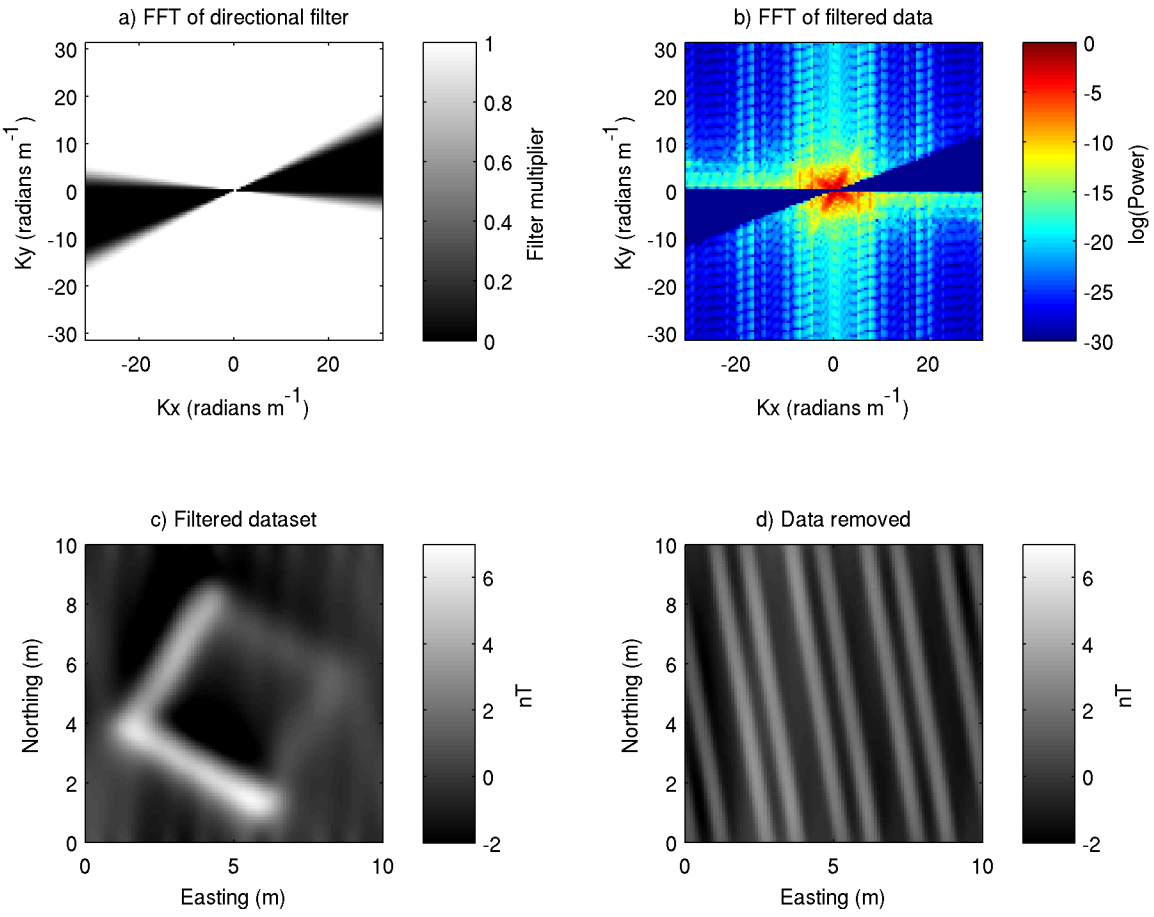

Figure 6.7: Directional pass/reject filter. (a) Filter to be applied to the Fourier domain. (b) FFT of filtered data. Note the reduction in power between 160° and 170°. (c) Post-filtering dataset (d) Data removed from the original data by filtering

filtering direction. This is the approach of directional cosine tapering. The algorithm is stated by Geosoft (2010b) as:

$$A(\theta) = |cos^n \left(\alpha - \theta + \frac{\pi}{2}\right)| \tag{6.6}$$

where α is the filtering direction, and the value of *n* is the degree of the cosine function, which can be varied to alter the shape of the function. Directional cosine filtering is commonly used and can be seen included in decorrugation techniques by Ferraccioli *et al.* (1998) and Pilkington & Roest (1998), as well as archaeological examples such as Linford *et al.* (2007).

Three scenarios have been tested here, with varying values for *n*. Figures 6.8a-c, show the filter to be applied to the Fourier domain for *n=0.5*, *n=1.0* and *n=2.0* respectively. Figures 6.8d-f, show the filtered Fourier domain, Figures 6.8g-i show the final filtered dataset, and Figures 6.8j-l show the data that has been removed by the filtering process. It can be seen in the Fourier domain plots that the obvious difference between the Cosine tapering technique, and the previous technique is that very little of the spectrum is actually removed, instead it is increasingly suppressed depending upon how close the data is to the filter azimuth.

Figure 6.8g shows that the *n=0.5* value has not been sufficient to remove all the agricultural signal, as striping can still be observed across the dataset. It is notable however, that part of the archaeological anomaly has been filtered out, as Figure 6.8j shows evidence of this in the removed data. This will result in a decrease in amplitude of the signal of interest in the filtered dataset. This is a problem that increases as the value of *n* is increased. Whereas Figures 6.8h and 6.8i show little evidence of the agricultural noise, the amplitude of the archaeological signal is increasingly reduced as more of the signal of interest is removed from the dataset. In Figure 6.8i, the response of the upper archaeological anomaly is $\sim$4 nT, instead of the $\sim$7 nT seen in Figure 6.6a. Whilst in this scenario both the archaeological bodies are still evident in the results, smaller amplitude anomalies would also have been reduced and may go undetected.

This is particularly a problem for archaeological surveys, where signals are typically only a few nT in amplitude. Furthermore any subsequent quantitative interpretation techniques, which rely on the amplitude of the total field signal, will be compromised.

6.3.5 Power/azimuth threshold (PATH) filter

An alternative method to the directional pass/reject filter and cosine taper filter, would be to add an additional amplitude constraint to the filter. A minimum and maximum azimuth is defined by the operator, as well as an amplitude threshold. Anything in the Fourier domain between the two angles and above the threshold is removed. All other data remains. The procedure is demonstrated in Figure 6.9. This techniques solves two of the drawbacks identified in the previous techniques. 1) The amount of the spectrum that is removed by filtering will be small, and not all the power at certain wavenumbers will be removed/reduced. 2) All data outside of the azimuths declared will remain the same, causing no loss of amplitude in features along different azimuths to the filter.

This filtering can seen in Figure 6.10a. A minimum azimuth of 160°, maximum of 180° and a normalised energy density threshold of -14 has been selected. The choice of the threshold value is important to the final result, and will be dependent upon the spectrum of the specific problem. Selecting a threshold value that is too low will result in a larger proportion of the spectrum becoming filtered than required, and the result will tend towards that obtained by the directional pass/reject filter. Choosing a threshold value that is too high will filter too little of the spectrum, and the undesired linear patterns will remain in the filtered dataset. It is therefore ideal to select the highest possible threshold that removes the undesired features.

Any data within these constraints are removed and shown by the areas which are not white in Figure 6.10a. The Fourier domain is subsequently interpolated across the removed data to avoid sharp discontinuities and Gibb's phenomena. Comparatively, much less data has been filtered than Figures 6.7a and 6.8a-c. The filtered Fourier domain can be seen in Figure 6.10b. The fil-

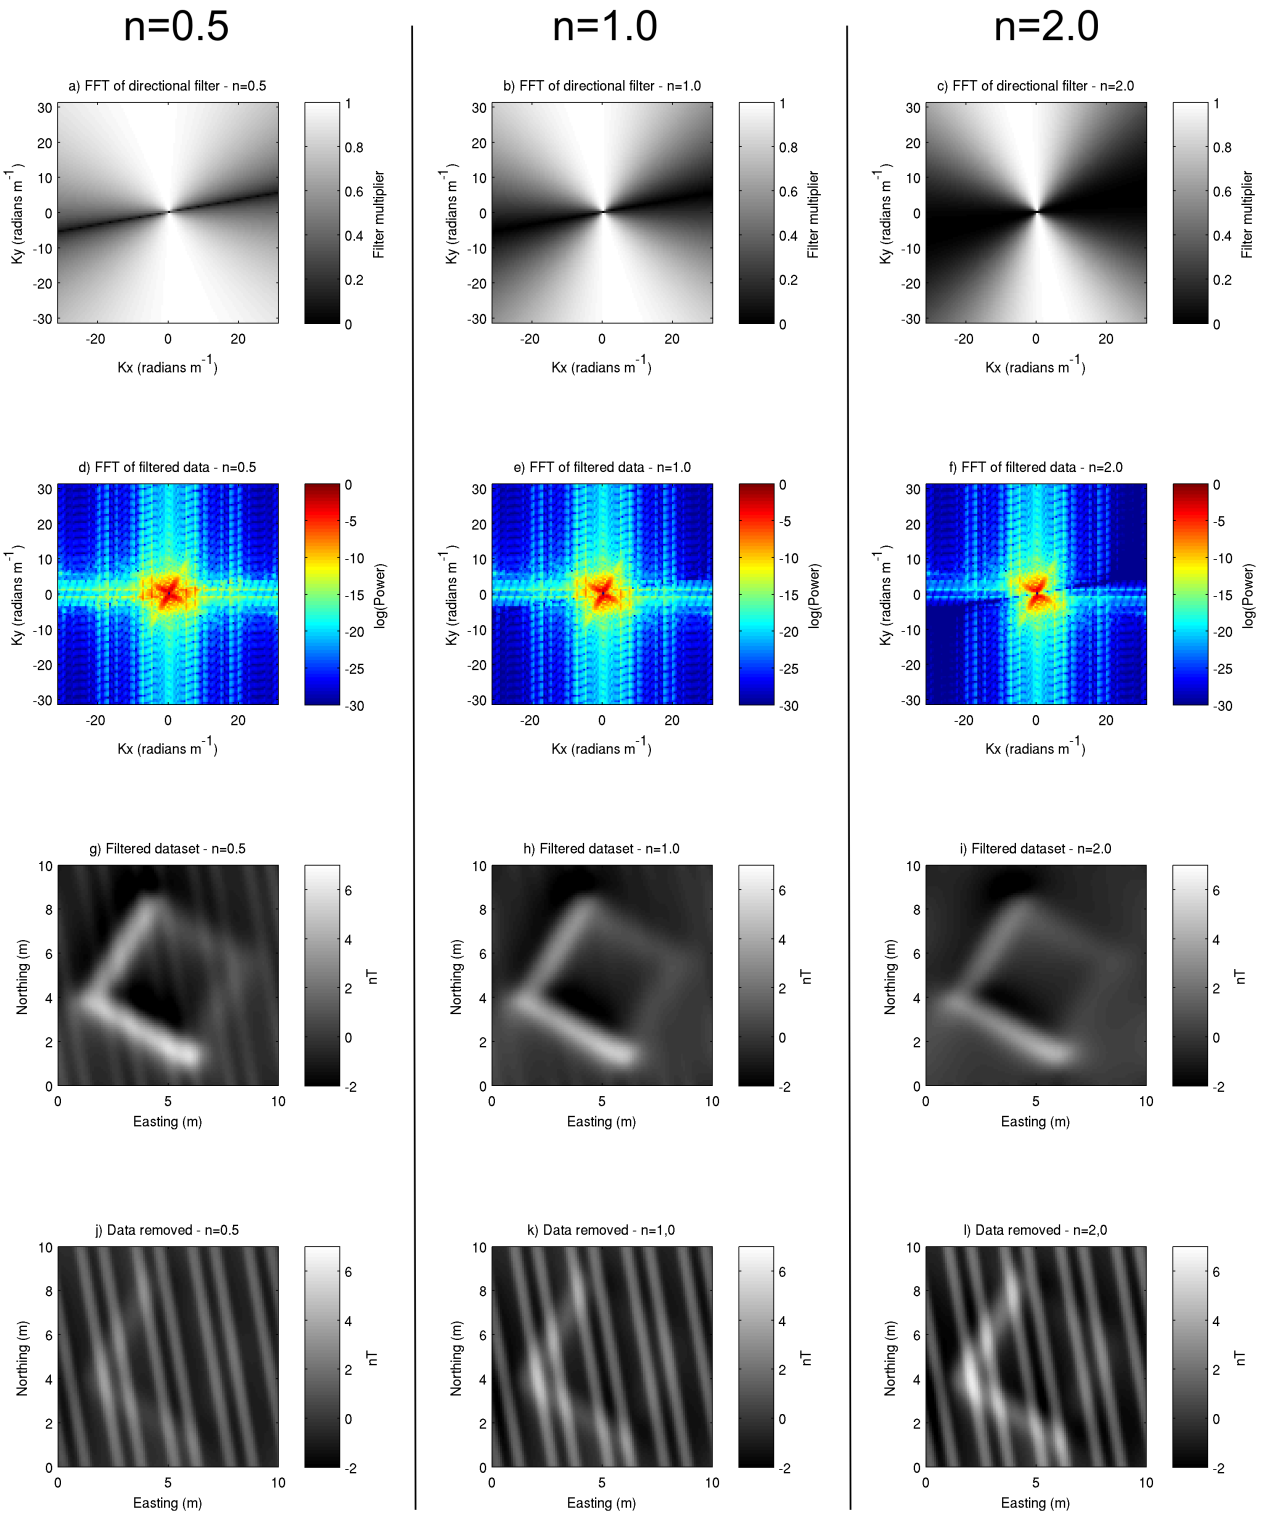

Figure 6.8: Illustration of the effect of the cosine taper filter, using α=*170°*. For *n=0.5*, the filter to be applied to the Fourier domain (a), Filtered Fourier domain (d), filtered dataset (g) and data removed by filtering (j) are shown. The respective plots for *n=1* are shown in figures (b), (e), (h) and (k), and for *n=2* in (c), (f), (i) and (l).

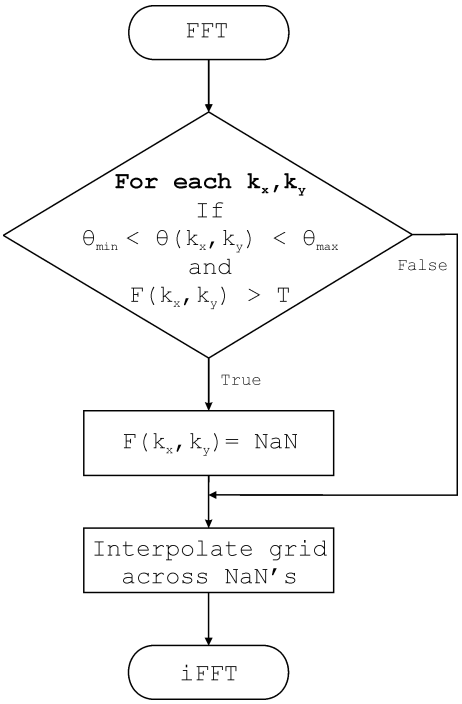

Figure 6.9: Flow chart outlining the procedure of the PATH filter, where $F(k_x, k_y)$ is the function of the magnetic data represented in the Fourier domain, and the user inputs, θ_{min}, θ_{max} and T are the minimum and maximum angles to be filtered, and the minimum normalised power threshold, respectively.

tered data (Figure 6.10c) shows good removal of the 170° agricultural noise. The removed data (Figure 6.10d) clearly shows the agricultural noise, with little evidence of the archaeological anomaly.

Comparing the directional filtering techniques

It can be seen in Table 6.2 that the power/azimuth threshold filtering has produced a much lower RMS error than any other technique. The second best technique based on the value of the sum of squares is directional cosine tapering with an n-value of 0.5, however, as demonstrated in Figure 6.8 this has left evidence of the linear anomalies present in the image. As *n* is increased, the value of the sum of squares is dramatically increased, showing that although the images look the clearest, a large proportion of the desired archaeological signal has been removed leading to a dataset that is the furthest away from the desired output.

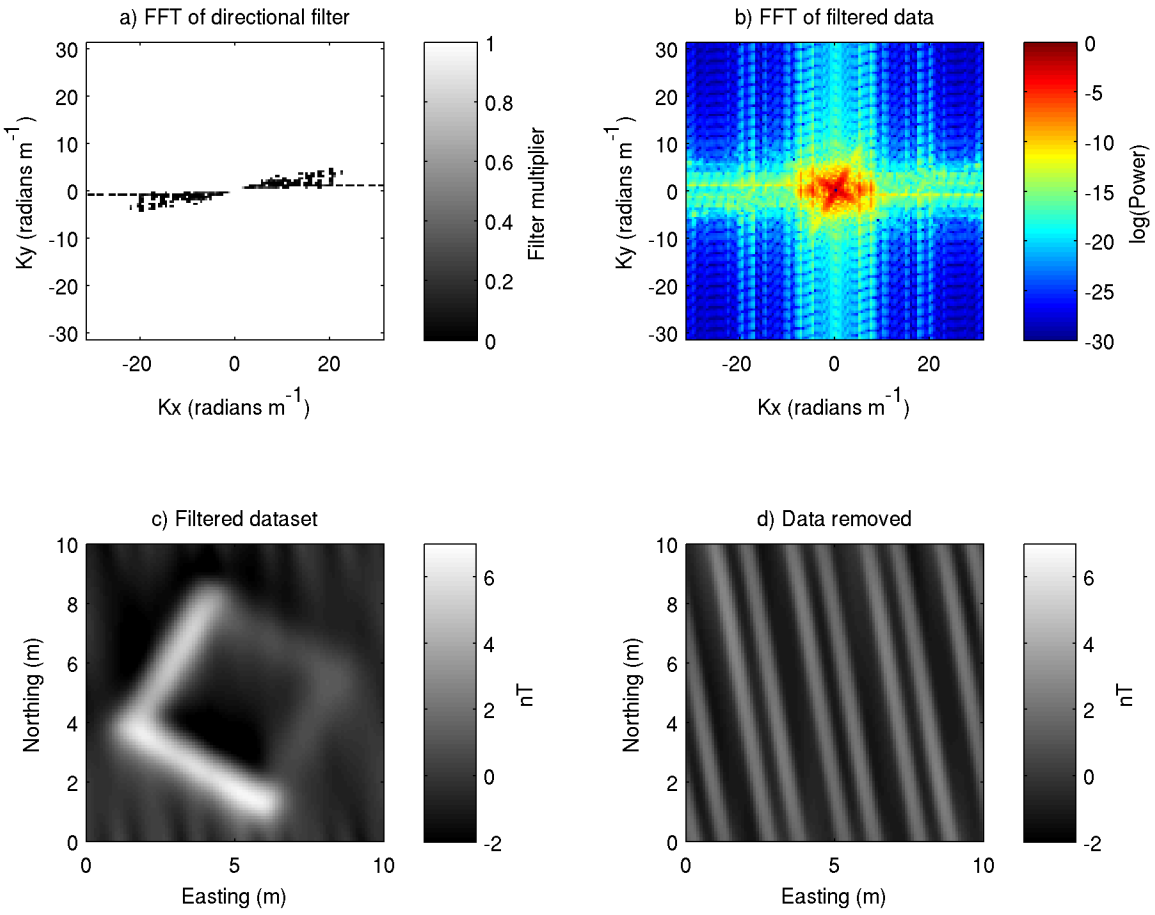

Figure 6.10: Power/azimuth threshold (PATH) filter. (a) Filter to be applied to data 160> α >180 and normalised energy density >-14. (b) Filtered data in the Fourier domain. (c) Filtered dataset. (d) Data removed by filtering

Filtering method	**Additional information**	**RMS Error**
Directional pass/reject		0.63 nT
Cosine tapering	n=0.5	0.59 nT
Cosine tapering	n=1.0	0.79 nT
Cosine tapering	n=2.0	1.07 nT
Power/azimuth threshold	Linear interpolation	0.37 nT
Power/azimuth threshold	Cubic interpolation	0.38 nT
Power/azimuth threshold	Nearest neighbour interpolation	0.46 nT
Power/azimuth threshold	No interpolation	0.38 nT

Table 6.2: Root mean square of the differences between the desired data and the filtering output.

The power/azimuth threshold filtering requires the filtered parts of the spectrum to be interpolated to remove the problem of sharp discontinuities in the spectrum. To test this, three interpolation techniques, triangle-based linear, triangle-based cubic and nearest neighbour interpolation have been applied to the data, as well as a scenario where no interpolation has been applied. In this case, Table 6.2 shows the linear interpolation to be the most effective, while the un-interpolated scenario is the second best. In addition to the increased likelihood of Gibbs phenomenon, when the spectrum is not interpolated prior to the inverse Fourier transform, previously high-powered wavenumbers are altered to low-power, and a striping pattern can again be present, this time with negative polarity. This is illustrated in Figure 6.11.

6.3.6 Case study

Brading Roman Villa is situated between the towns of Brading and Sandown on the Isle of Wight, UK. The site consists of a Roman villa, now incorporated into a visitor centre, and surrounding pasture land. The data used in this study is a subset of a much larger Caesium magnetometery survey collected in 2009 by the English Heritage Geophysics Team over the surrounding pasture land (Payne, 2009). The data were collected with a spatial resolution of 0.125 x 0.5 m.

The subset, shown in Figure 6.12b, is a 135 x 55 m area situated $\sim$100 m NW of the villa, showing a series of positive anomalies. The most prominent features in the data are a large circular anomaly $\sim$25 m in diameter towards the western side of the dataset, and a series of positive linear anomalies across the area. An agricultural pattern is prominent across the dataset, as well as a series of small positive peaks scattered across the area.

The Fourier domain representation of the data is shown in Figure 6.12a, highlighting the areas where the filtering has been applied in white. The reject criteria for the filter were an angle $>$75 and $<$85°, with a normalised energy density amplitude of $>$-12. The post-processed inverse FFT is shown in Figure 6.12c, and the removed data in Figure 6.12d. It can be seen that the filtered

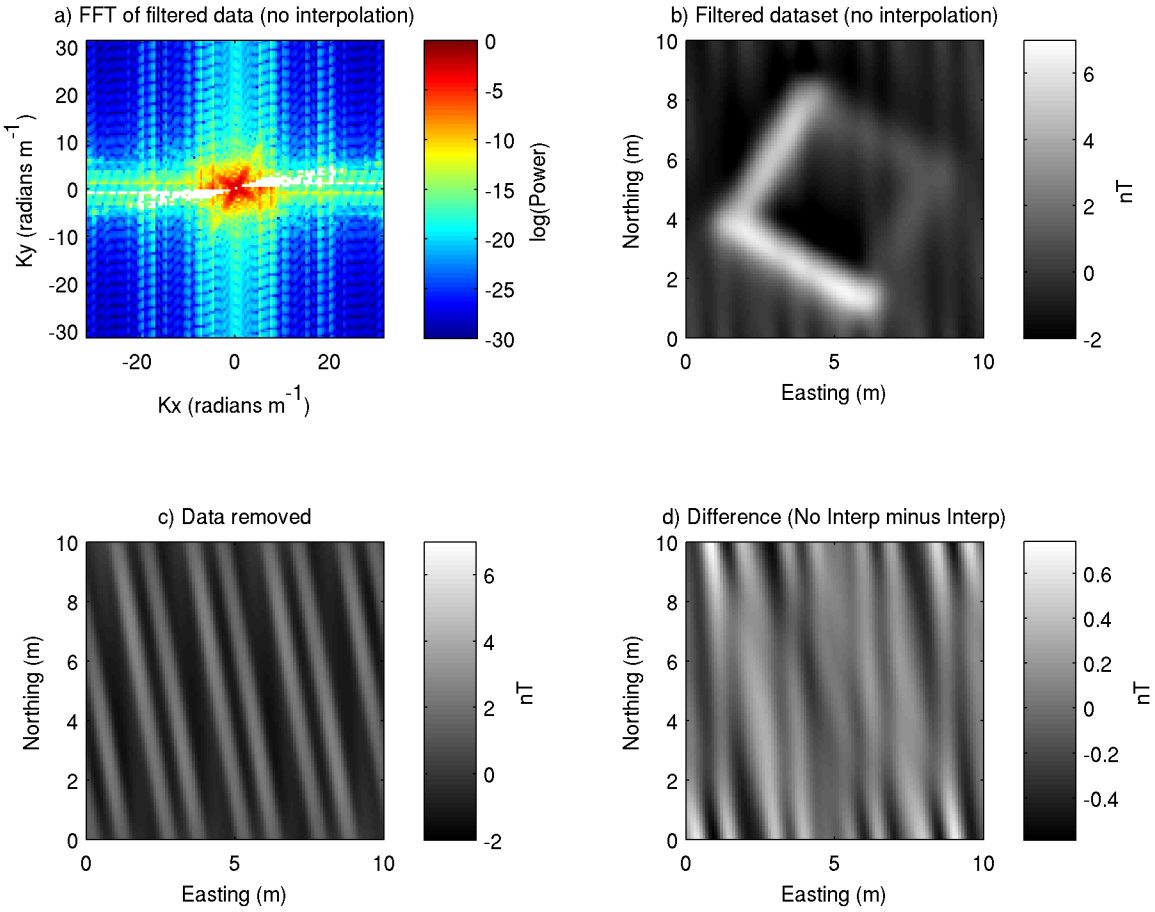

Figure 6.11: Illustration of the effect of interpolating across the filtered spectrum. (a) Filtered data in the Fourier domain (no interpolation). (b) Filtered dataset. (c) Data removed by filtering. (d) Non-interpolated data minus interpolated data (Figure 6.11b minus Figure 6.10c). The effect of not interpolating the Fourier domain prior to iFFT has created a set of linear anomalies closer to a N-S direction. Note the scale axis is much smaller on this plot, as these differences are much smaller amplitude.

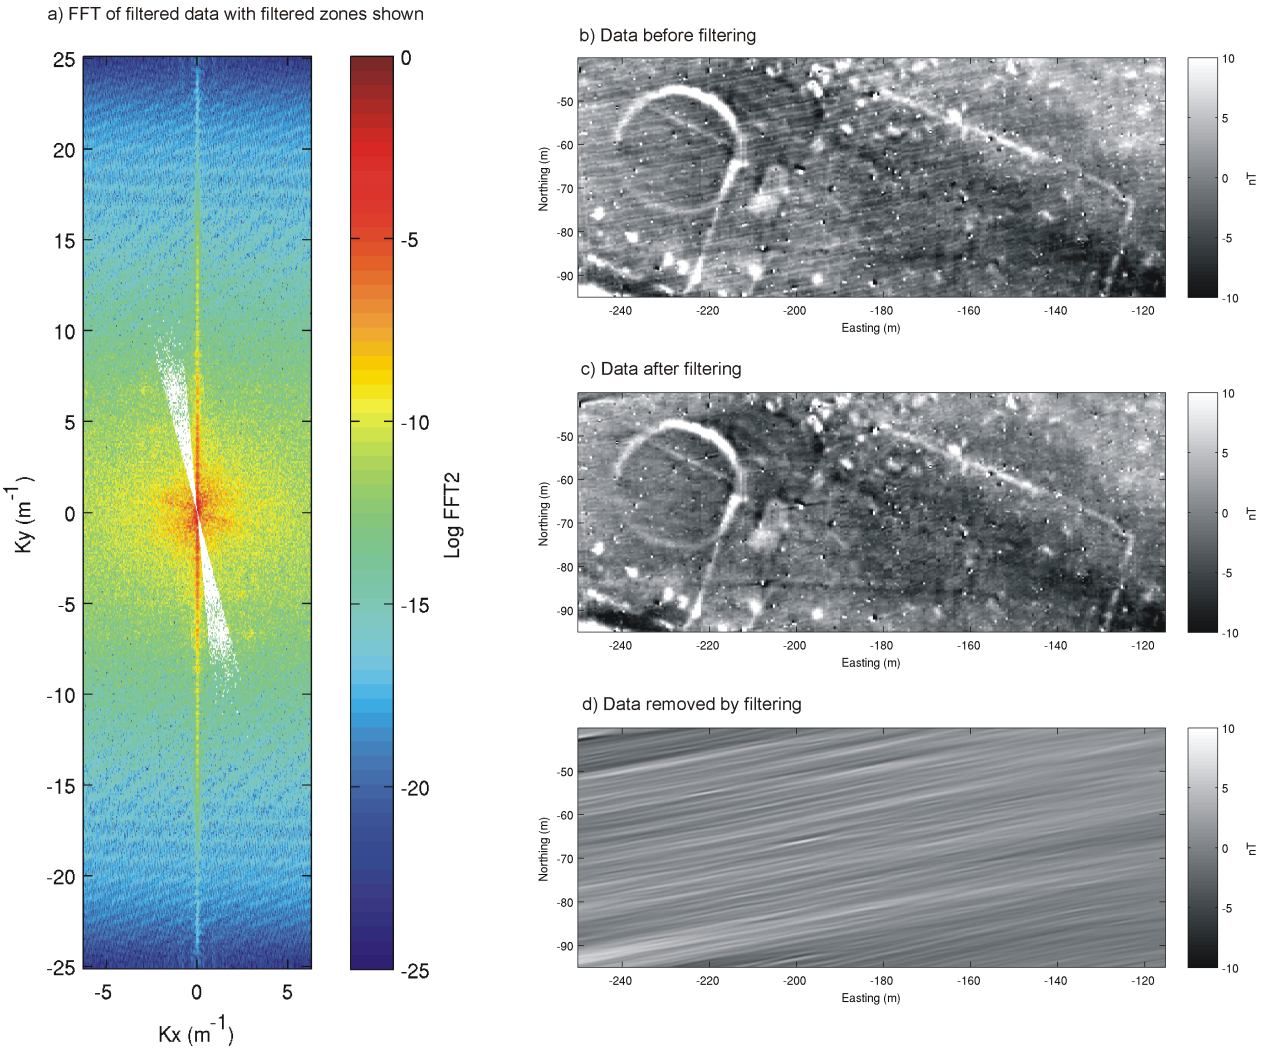

Figure 6.12: Case study from Brading Roman Villa, Isle of Wight, UK. (a) Filtered fourier domain with filtered zone shown in white. (b) Original dataset. (c) Filtered dataset. (d) Data removed by filtering.

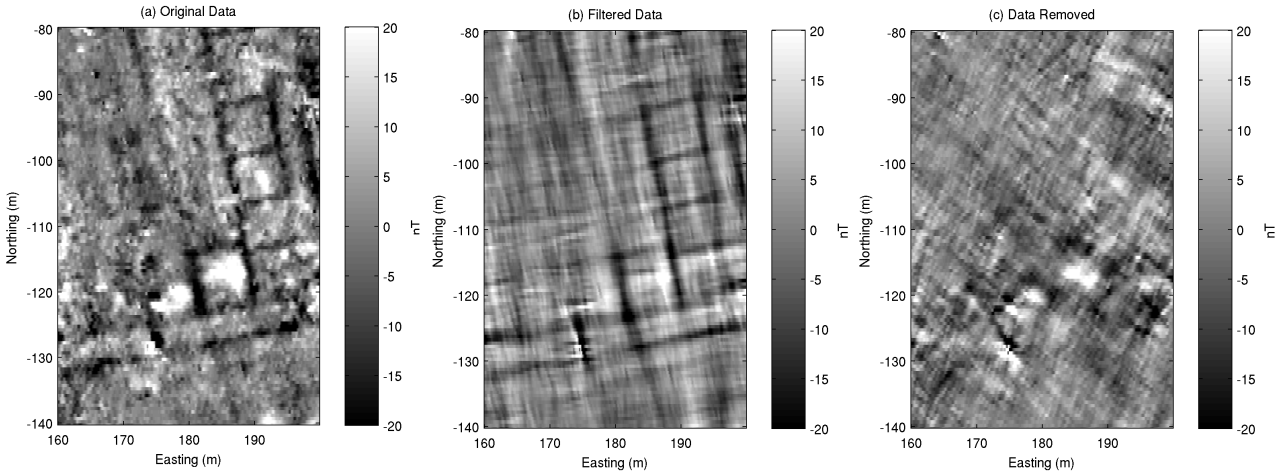

Figure 6.13: Case study from Silchester Roman Town, Forum Corner. (a) Original Dataset. (b) Filtered dataset with a large filter zone set from 0-70° and 90-160°. (c) Data removed.

dataset has improved the visual appearance of the data, and the majority of the agricultural pattern has been removed. There is little evidence of any of the archaeological anomalies being removed, therefore the amplitude of these should not have been significantly reduced due to the filtering.

6.3.7 Alternative applications

In previous sections, directional filtering has been demonstrated for removing unwanted linear features from datasets, however it could also be used to highlight linear features. Figure 6.13a is a subset of a larger dataset, and shows the south-east corner of the Roman Forum at Silchester, Hampshire, UK. The foundations of the Forum can be seen as negative linear anomalies approximately on bearings of 80° and 170°. Should these foundations features be the focus of the investigation, directional filtering can be used to separate these from the remaining data.

Figure 6.13b shows the same area once all power has been removed from wavenumbers with the 0-70° and 90-160° ranges. The data that has been removed is shown in Figure 6.13c. It would now be possible to invert the filtered dataset to produce a model that only highlights the foundation features. This type of filtering should be done with care however. A large proportion

of the dataset has been removed by this process, and artefacts within the 'pass' area of the filter are likely to be introduced due to removing power from all other directions. Interpretation of the final result should be made with this in mind.

6.4 Identifying the presence of remanent magnetisation

Section 5.6.3 highlighted the possible errors to the model that can be introduced in the presence of remanent magnetisation that is not aligned in the same direction as the inducing magnetisation vector. For the reasons discussed in Section 5.6.3 there are many added complexities to accounting for remanence during modelling, which mean it is often not feasible to do so. However, it is important to be aware if there is a strong remanent component to the magnetisation so that the final model is not incorrectly interpreted.

Roest & Pilkington (1993) devised a method for identifying the presence of remanent magnetisation. Their method involved a comparison of the analytic signal amplitude (Equation 3.4) and the horizontal gradient magnitude of the pseudogravity dataset (using Equation 3.3, but replacing T with the pseudogravity of T). For a single dataset, both of these functions were calculated for a range of inclination and declination values with the similarities identified by a cross-correlation calculation. It is stated that the analytic signal amplitude, which is largely independent of the direction of magnetisation, will correlate best with the horizontal gradient magnitude of the pseudogravity when the correct magnetisation vector has been used. The ideal magnetisation vector can therefore be determined by finding the maximum of the cross-correlation function, and should this differ from the direction of the inducing field then there is likely to be some remanent magnetisation present within the dataset.

A problem with implementing this technique in the presence of dipping features was identified by Dannemiller & Li (2006), who demonstrated that the analytic signal amplitude and horizontal gradient of the pseudogravity dataset have different rates of decay, with the fall off of the

analytic signal amplitude being $1/r^4$ and the horizontal gradient magnitude of the pseudogravity being $1/r^3$ for a point source. Therefore Dannemiller & Li (2006) presented an alternative technique, this time using the vertical gradient and the analytic signal amplitude of the reduced to the pole total-field dataset. The cross-correlation coefficient is defined as:

$$C(I_M, D_M) = \frac{\sum(\frac{\partial T}{\partial z}_j - \overline{\frac{\partial T}{\partial z}})(|A| - \overline{|A|})}{\sqrt{\sum(\frac{\partial T}{\partial z}_j - \overline{\frac{\partial T}{\partial z}})^2 \sum(|A|_j - \overline{|A|})^2}} \tag{6.7}$$

where j is the index of each quantity within the dataset, I_M and D_M are all the inclination and declinations to be tested, $\overline{\frac{\partial T}{\partial z}}$ is the mean value of the vertical gradient, and $\overline{|A|}$ is the mean value of the analytic signal amplitude.

The technique has been applied to identify the magnetisation direction for a synthetic scenario. The compound body shown in Figure 3.1 has been forward modelled with an inclination of 64° and a declination of 40°. The magnetic output from this provides the input data to be used. The dataset was reduced to the pole using all combinations of inclinations from 20 to 90° and declinations from -180 to 180°, with the cross-correlation coefficient calculated for each combination. These were then gridded and contoured, and shown in Figure 6.14a. The actual magnetisation vector is shown by the white star, and it can be seen that the maximum cross-correlation occurs in the area around the white star. The technique has therefore been successful in determining the likely magnetisation, should the scenario not be synthetic.

Unfortunately, as investigated in Chapter 3, calculating the gradients of the total field anomaly dataset causes an amplification of the high wavenumber features, which are often noise. For this reason in Figure 6.14c, where Gaussian noise with a standard deviation of 0.2 nT has been added to the data, and Figure 6.14e, where Gaussian noise with a standard deviation of 0.5 nT has been added, the vertical gradient, and analytic signal amplitude are dominated by noise (see Figure 3.4), and hence very low cross-correlation coefficients are derived, and it is not possible to identify a magnetisation direction.

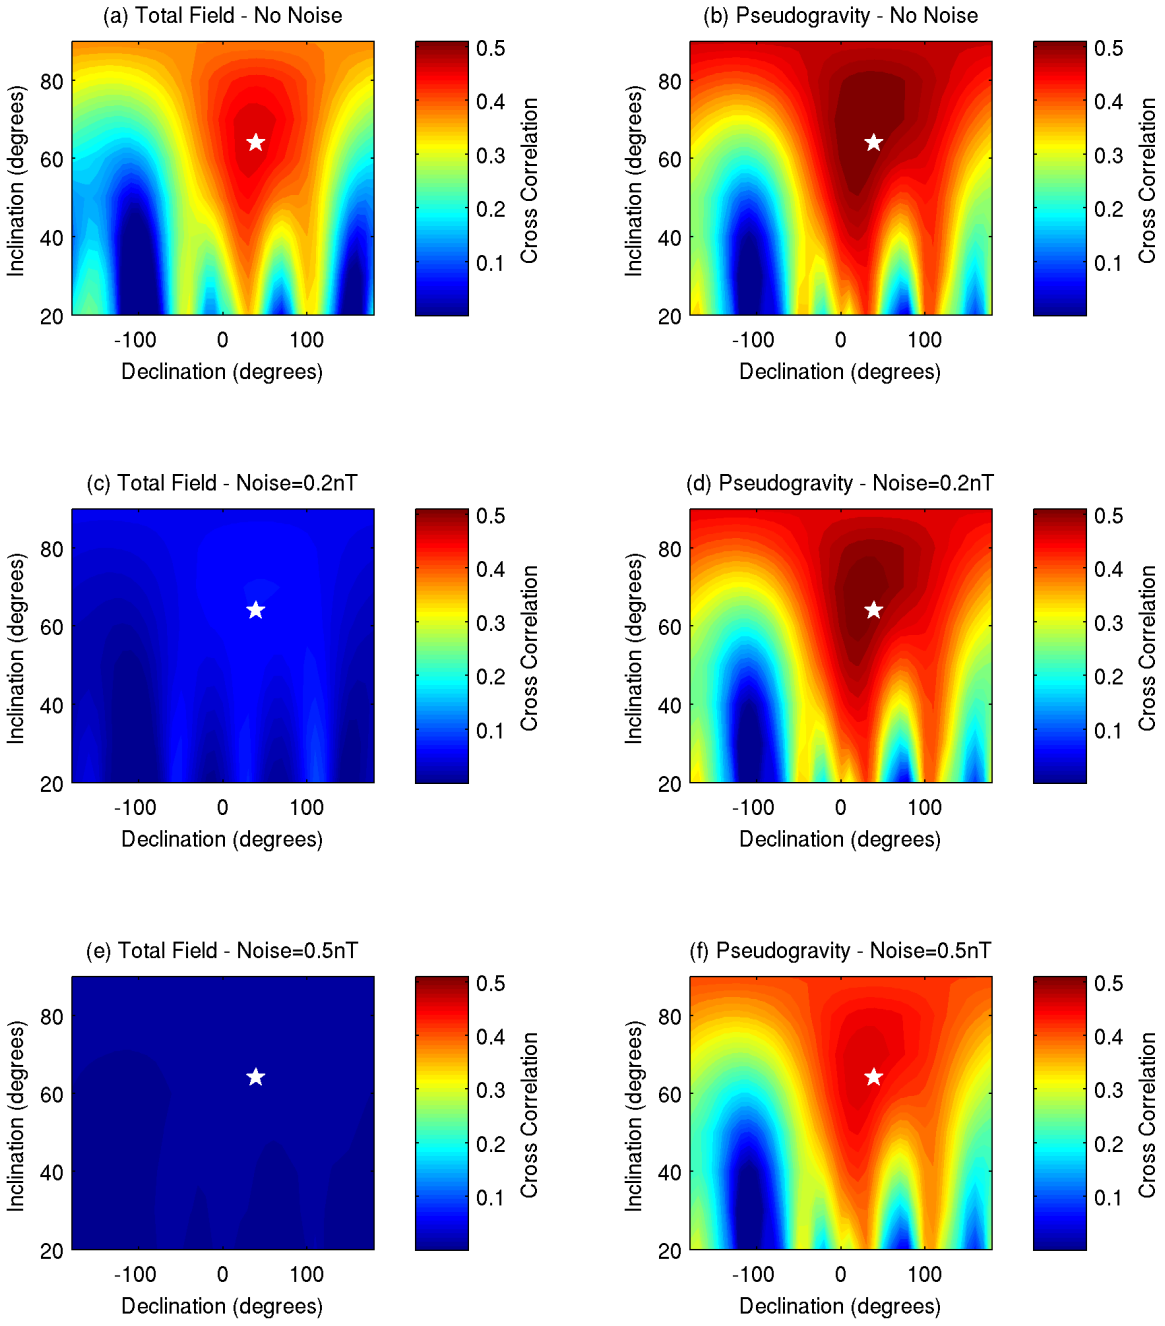

Figure 6.14: Calculation of the magnetisation vector used the method described by Dannemiller & Li (2006). (a) shows the cross-correlation calculated using the vertical gradient of the total field and the analytic signal amplitude of the total field for a synthetic scenario with an inclination of 64° and declination of 40°, with no noise added. (b) Cross-correlation generated using the vertical gradient of the pseudogravity and the analytic signal amplitude of the pseudogravity with no noise. (c) As a) but with Gaussian noise of 0.2 nT standard deviation added. (d) As b) but with Gaussian noise of 0.2 nT standard deviation added. (e) As a) but with Gaussian noise of 0.5 nT standard deviation added. (f) As b) but with Gaussian noise of 0.5 nT standard deviation added.

Once again, the pseudogravity transformation can be used to suppress the noise when calculating gradients of the total field. Instead of using the vertical gradient and analytic signal amplitude of the reduced to the pole total-field dataset, the vertical gradient and analytic signal amplitude of the pseudogravity dataset are used instead. This combination of datasets does not have the problem associated with the method described by Roest & Pilkington (1993), as both fields will decay with a rate of $1/r^3$.

The corresponding contour plots for the zero-noise, 0.2 nT and 0.5 nT Gaussian noise scenarios are shown in Figures 6.14b, 6.14d and 6.14f respectively. The maximum cross correlation coefficient is identified in the area around the actual magnetisation vector in all scenarios, and is not so heavily affected by the noise content of the data, with the maximum coefficient never falling below 0.4. This new method can therefore be used with archeo-type datasets, to identify if the magnetisation vector differs from the inducting field, and if this is found to be the case, the dataset will be known to have some quantity of remanent magnetisation present.

6.5 Conclusions

Filtering of magnetic datasets is a broad subject, with many procedures available for separating useful information from background signal. In this chapter, filtering has been considered as the steps required to make a dataset applicable for further modelling by inverse techniques.

Bandpass filtering is one of the most commonly applied type of filter in geophysics, and hence, many different techniques have been developed to fulfil this need. Here, three techniques, the simple high-pass, Butterworth and cosine roll-off filters have been used to separate regional trends originating from distant sources, from the signal generated by the archaeology of interest. It has been shown that, while statistically the high-pass filter has produced the best result for the synthetic scenario tested, any one of these methods will produce acceptable results that can

be used for subsequent modelling, and the Butterworth or cosine roll-off filters are preferred as there is less likelihood of introducing Gibb's phenomena into the filtered dataset.

As well as removing low wavenumber features, the same techniques can, and often are, applied to remove high-wavenumber noise from the data. This application has not been analysed here, as the nature of the mesh-based inverse technique means that the highest wavenumber features cannot be modelled unless very small cell sizes are placed in the inverse mesh. Where large cell sizes are allocated, shallowly buried magnetic bodies of small dimensions cannot be modelled by the inversion, and in effect a low-pass filter will naturally be applied.

Linear anomalies in near-surface geophysical data are a common problem, and often appear in archeo-magnetic surveys over agricultural land. In order to clarify the image of the deeper archaeological signal it is desired to remove these features from the data, and consequently several processing and interpretation software packages contain routines for directional filtering in the Fourier domain. The most common two of these are the directional pass/reject and cosine taper filters.

While both of these filters do greatly improve the appearance of the data, they have their own individual drawbacks. The simple filtering directional pass/reject approach can cause a large proportion of the data to be removed, while the directional cosine tapering approach fundamentally changes the amplitude of the data at all wavenumbers, not just in the filtering direction. Modern interpretation techniques are capable of making quantitative interpretations regarding depth and susceptibility based on the amplitude of the signal. While the cosine-tapering technique has produced the best image, features of the same susceptibility would be interpreted differently depending upon their orientations should this filtering technique be applied beforehand.

Here an alternative directional filtering routine is proposed. The PATH filter is based on the directional pass/reject filter, however, the problem of a large proportion of the spectrum being removed by the filter is avoided by only filtering parts of the spectrum that are over an addi-

tional energy-density criteria. This approach ensures that amplitudes of anomalies in directions outside the filter parameters remain the same, and the potential to perform quantitative interpretation techniques has not been compromised by the filtering.

Analysis of the RMS error of the three different filtering outputs shows the power/azimuth threshold filter to outperform the other two, with the resulting image the closest to the synthetic input. It has been successfully applied to real data from Brading, Isle of Wight, and shows a substantial improvement to the data.

If significant remanent magnetisation is present in the data, the results from modelling with a single inducing magnetisation vector will be invalidated. It is unlikely that enough information regarding the influence of remanent magnetisation on the data will be available to allow a filter to be designed to remove it. However, it has been shown here that the method developed by Dannemiller & Li (2006) to identify the presence of remanence, can be altered to be applied to archaeo-magnetic datasets. Once again, due to the noise-content of archaeo-magnetic data, using the derivatives of the pseudogravity is preferred to the derivatives of the total field measurements. The vertical gradient and analytic signal amplitude of the pseudogravity dataset can be used to identify the optimum magnetisation vector, and if it is found to differ significantly from the inducing field direction the dataset should not be considered for further inverse modelling.

Chapter 7

Application of quantitative interpretation techniques to archaeo-magnetic datasets: results from 3 case studies

7.1 Introduction

To test the effectiveness of the techniques investigated in the previous chapters, data are analysed from three case studies. The first, is Test Site 1, a purpose built geophysics test site of the Near Surface Geophysics Group (NSGG), situated in Leicester, UK. This site consists of a variety of geophysical targets to simulate geological, engineering and archaeological scenarios. The second is from Silchester Roman Town, Hampshire, UK. The town has been thoroughly investigated by geophysical prospecting and excavations for over 100 years, and therefore there is a wealth of data concerning the archaeological features at the site. The third case study is from land surrounding Brading Roman Villa, Isle of Wight, UK. This survey was designed to yield more information regarding the surrounding setting of the Roman villa, and therefore is the only dataset currently available. This is the scenario where 3D quantitative interpretation is most useful in providing depth and material estimates for subsequent investigation.

ID Number	Depth to base (m)	Width (m)	Height (m)	Object Type (Each 4m long)
1	1.0	0.45	0.4	Concrete, cast in-situ, non-reinforced
2	1.0	0.45	0.4	Brick, corbelled. Engineering brick.
3	1.0	0.45	0.4	Stone blocks set in mortar.
4	1.0	0.45	0.4	Aggregate, types 1 & 2
5	1.0	0.45	0.4	Sand, building sand.
6	1.0	0.45	0.4	Peat and leaf mould
7	0.6	-	-	Paving slabs 0.6×0.6 m and 0.6×0.9 m.

Table 7.1: Description of the buried walls and trenches in Area 4 of the University of Leicester Test Site 1.

7.2 Case Study 1 - Area 4, Test Site 1, Leicester, UK

The data used in this case study were obtained from a total-field survey conducted with a Geometrics G-858 magnetometer in May 1998, at the NSGG Test Site 1 located at Southmeads Road, Oadby, Leicester. The test site has been designed for training and equipment testing purposes, and provides multiple types of target, divided into ten areas. The test site is 40 m by 70 m, and the ground consists of 0.3 m of topsoil with grass coverage. Below this is 16-18 m of Boulder Clay containing clasts up to 200 mm in a clay matrix.

The G-858 gradiometer consists of two total-field caesium magnetometers, one usually placed above the other to calculate a vertical component of the total gradient of the field. For this study, only the total-field output from the bottom sensor is to be considered. The data collected by the bottom sensor are shown in Figure 7.1a. There is a large range to the amplitude of responses to the various different types of target, but of particular interest to this study is Area 4 which is outlined in Figure 7.1a, and has been designed to simulate buried foundation features for civil engineering or archaeological situations. A close up of the data collected over Area 4 is shown in Figure 7.1b, with the plans indicating the locations of buried bodies shown in Figure 7.1c. A full description of the buried targets is shown in Table 7.1.

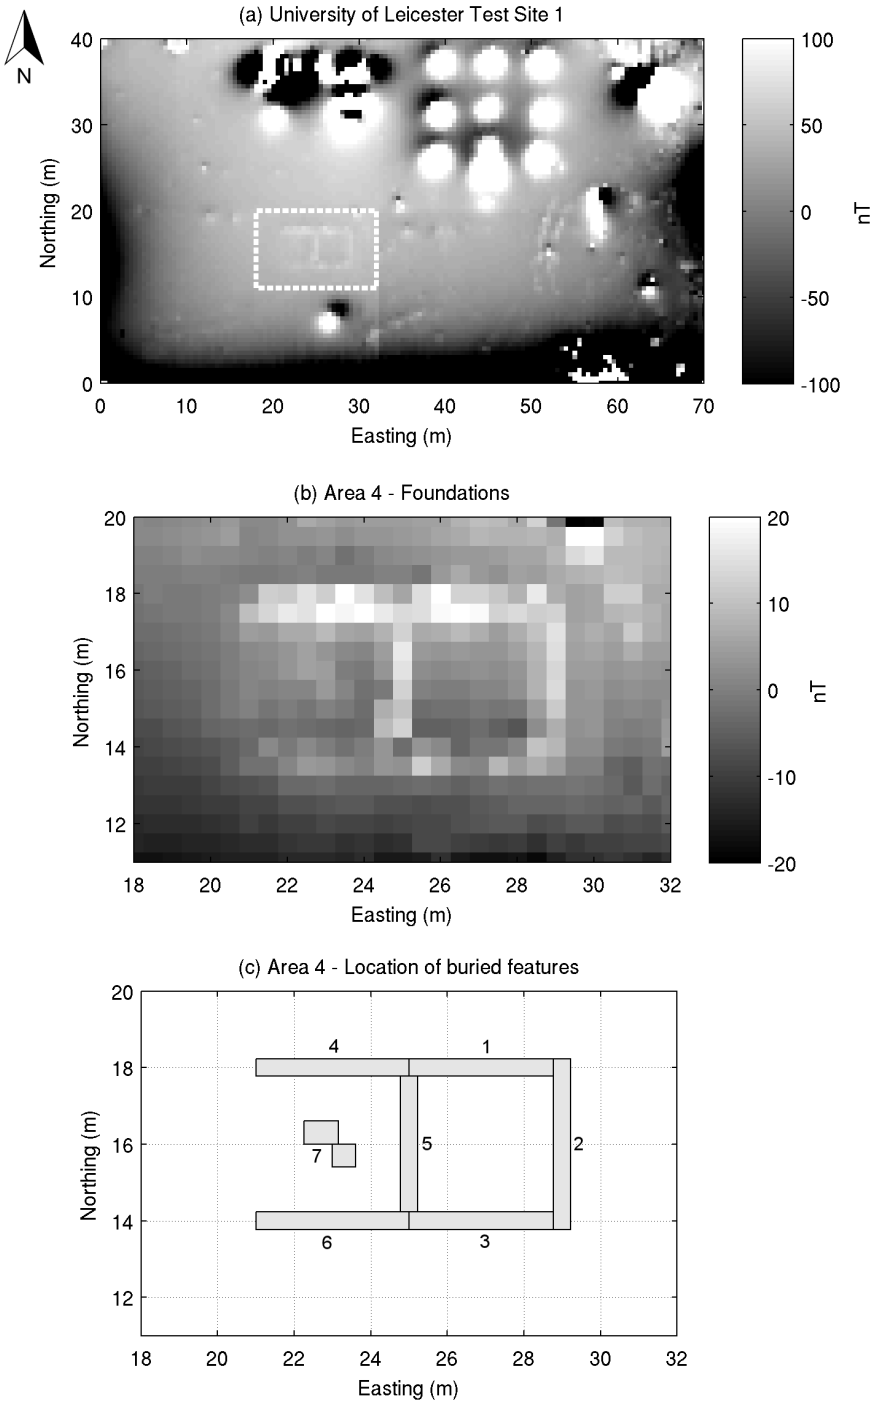

Figure 7.1: Caesium-vapour (G858) survey over University of Leicester, Test Site 2. (a) Test Site 2, has many buried targets resulting in a large range of anomalies. Area 4 is shown by the dashed area. (b) Area 4 - total field anomaly dataset. Note the change of amplitude scale. (c) Area 4 - locations of buried foundation features.

Component	Field Value	Secular Value
Declination	-4.02 degrees	10.2 arcmin/year
Inclination	67.34 degrees	0.5 arcmin/year
Horizontal Intensity	18706 nT	3.5 nT/year
North Component	18660 nT	7.4 nT/year
East Component	-1311 nT	55.1 nT/year
Vertical Intensity	44805 nT	25.7 nT/year
Total Intensity	48553 nT	25.1 nT/year

Table 7.2: IGRF output, for the NSGG Test Site 1, Leicester, May 1998.

The survey was conducted with the resolution of 0.5×0.5 m, and the bottom sensor was positioned at a height of 0.25 m above the ground surface. The subset of data clearly shows the outline of the buried features which are represented as positive anomalies ranging in amplitude between 5 and 15 nT. The largest anomaly appears as a dipole-like response of around 30 nT, not seemingly associated with the buried targets, situated towards the north-east of the area, and the lowest values are recorded towards the south-west of the area where values of approximately -20 nT are recorded.

7.2.1 Euler deconvolution

The appropriate structural index to apply to the data from Area 4 of the NSGG Test Site 1, is likely to be either $N = 1$, for a thin sheet approximation, or $N = 2$ for a line of dipoles, as the buried foundations are essentially linear features, yet with a width of 0.45 m. It is probable that the actual structural index for features of these dimensions will lie somewhere between $N = 1$ and 2.

Results for the Euler deconvolution have been obtained from data that has been upward continued by 0.2 m, using a data window of 5 samples, representing 2 m^2. The solutions for a structural index of $N = 1$ are shown in Figure 7.2a, and for $N = 2$ in Figure 7.2b.

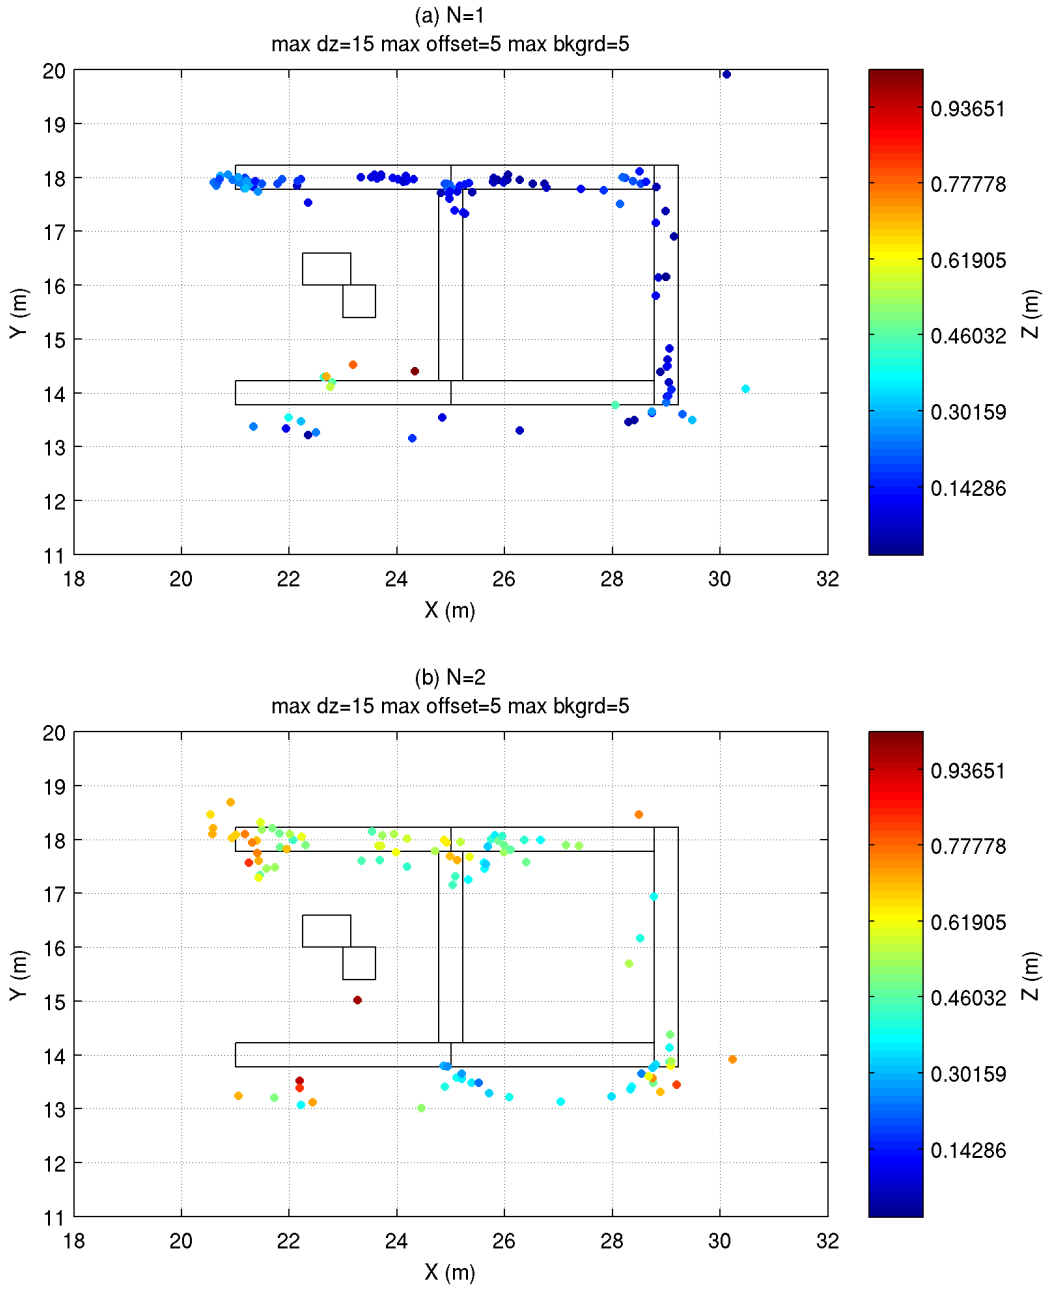

Figure 7.2: Euler deconvolution results from Area 4, NSGG Test Site 1, Leicester. (a) Applying a structural index of $N = 1$. (b) Applying a structural index of $N = 2$. The two structural indexes have yielded significantly different results.

The majority of the buried foundation features are positioned between 0.6 and 1.0 mbgl. The results from Euler deconvolution, using a structural index applicable to a thin-sheet approximation have yielded solutions that are too shallow. The average depth of the solutions shown in Figure 7.2a is just 0.12 mbgl. Applying a structural index relating to a line of dipoles has given improved solutions, with an average depth of 0.51 mbgl, however using this approximation it would be expected that the results would gather towards the mid-point of the buried body, therefore the results are significantly shallower than the expected 0.8 mbgl. A further discussion on this will be made later in comparison to the inverse model in Section 7.2.4.

7.2.2 Preparation of data

An obvious feature of the data shown in Figure 7.1b is a trend from high towards the north-east, to low in the south-west, likely to be related to iron railings along the southern side of the Test Site, and therefore related to features outside the data area. For the reasons discussed in Section 5.6.2 it is required that this trend is removed prior to modelling.

To achieve this, the data has been high-pass filtered using a Butterworth filter with the central wavenumber set to $k_0 = 0.5\ rad/m$. The degree of the filter was set to $n = 8$ to give a smooth transition between the 'accept' and 'reject' parts of the filter. The filtered dataset is shown in Figure 7.3a, with the data that has been removed by the filtering procedure shown in Figure 7.3b. The filtering has successfully removed the regional trend, which is shown varying from $\sim$-15 to 10 nT across the area. The parameters applied have successfully removed the trend without any evidence of the signal of interest in the removed dataset.

7.2.3 Remanent magnetisation

The direction of the inducing magnetic field at the Test Site in May 1998, is shown by the IGRF to have an inclination of 67.34° and a declination of -4.02° (Table 7.2). If the buried

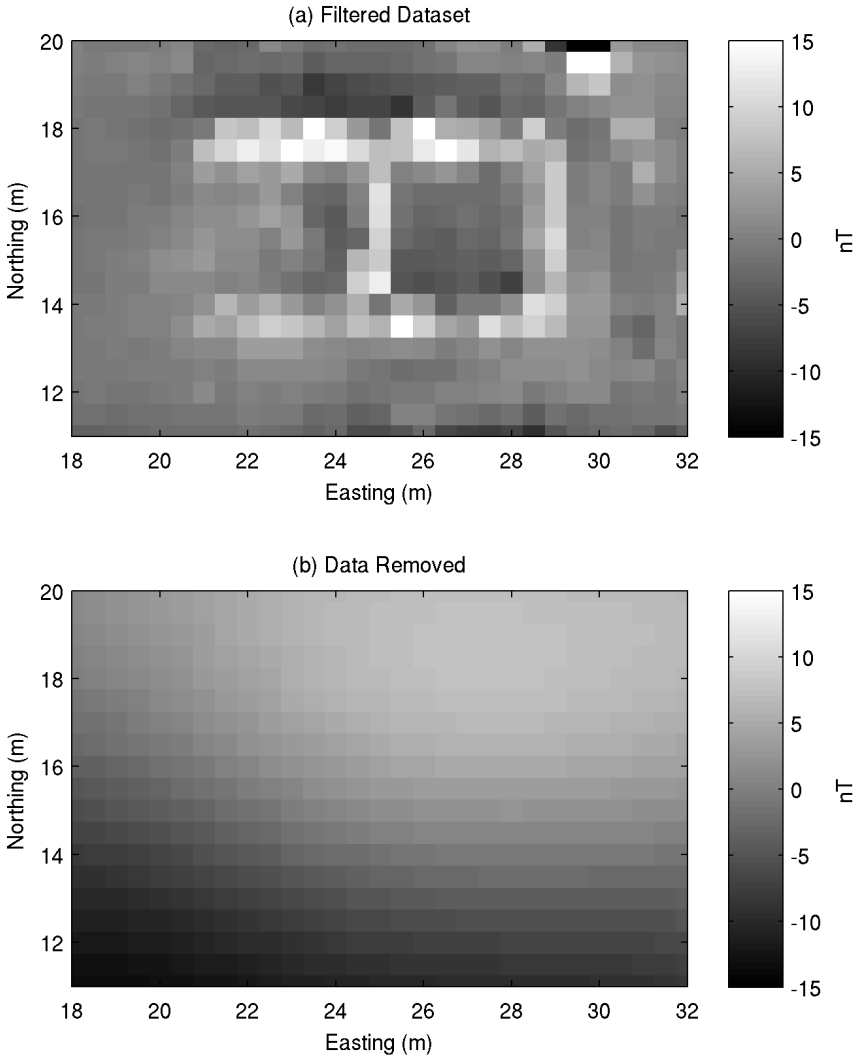

Figure 7.3: High-pass filtering of Area 4, NSGG Test Site 1, using a Butterworth Filter and $k_0 = 0.5$ and $n = 8$. (a) Data for modelling, with the regional trend across the area removed. (b) Data removed by the filter showing the regional trend from the north-east to south-west.

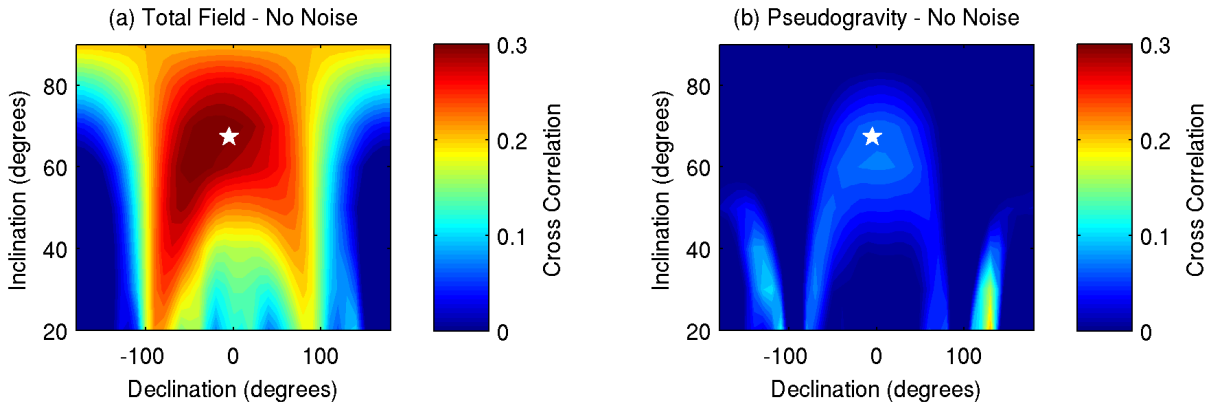

Figure 7.4: Test for the likely presence of remanent magnetisation in the dataset from Area 4, NSGG Test Site 1. (a) Using the Dannemiller & Li (2006) method of cross-correlation between the vertical gradient and the analytic signal amplitude. (b) analysing the cross-correlation between the vertical gradient of the pseudogravity and the analytic signal amplitude of the pseudogravity. The white star represents the inclination and declination derived from the IGRF.

materials in the subsurface were to have no remanent magnetisation properties, and therefore the only magnetisation is caused by the inducing field, then the magnetisation vector will be in line with the inducing vector. To test this the cross-correlation between the vertical gradient and the analytic signal of the total field, as described by Dannemiller & Li (2006), has been used to determine an optimum magnetisation vector for the dataset.

Analysis of the cross-correlation in Figure 7.4a shows a broad maximum with optimum correlations existing with inclinations between 60-70° and declinations between -50-0°. The parameters of the inducing field are shown by the star in Figure 7.4, and lies within this zone of maximum correlation. Due to the broad zone, it is not possible to rule out the presence of some remanent magnetisation in the dataset, however, as the inducing direction is within the maximum zone there is also no reason to believe that modelling using the inducing vector as the direction for all magnetisation is an invalid approach in this instance.

Cross-correlation of the vertical gradient of the pseudogravity and the analytic signal amplitude of the pseudogravity is shown in Figure 7.4b, and has resulted in lower cross-correlation coefficients than when applied to the derivatives of the total field. This result is likely to be due to the low resolution of this survey (0.5×0.5 m). This means high-wavenumber data is not cap-

tured, and therefore the effect of calculating derivatives is less pronounced than other surveys investigated here.

7.2.4 Inverse model

The inversion has been performed using a mesh 9×4×4 m in dimension, limited to the horizontal extents of the data area, and consisting of 0.5×0.5×0.1 m cells. Visual inspection of the data suggests that all the features are likely to be caused by positive susceptibility contrasts compared to the surrounding soil, therefore positivity has been imposed by limiting the modelled susceptibilities to $0 < \mathbf{m} < 1\ SI$.

The inversion has been performed multiple times with varying values of β, and the trade-off between data misfit and the model norm is shown in the L-curve in Figure 7.5. A value of $\beta = 1 \times 10^7$ was chosen as representing the maximum curvature of the L-curve. The inverse procedure was performed for both a 'small' model ($\alpha_s = 1$, $\alpha_x, \alpha_y, \alpha_z = 0$), and for a 'smooth' model ($\alpha_s = 0.0001$, $\alpha_x, \alpha_y \alpha_z = 1$). Both runs produced acceptable misfits and similar models. The results for the 'small' model are presented here.

The observed total-field data, and the calculated response to the final model are shown in Figures 7.5b and 7.5c respectively. The data misfit between these datasets is shown in Figure 7.5d. There is generally a good correlation between the observed and calculated datasets, with the largest differences associated with the anomalous high towards the north-east of the survey area, and between the low amplitude anomalies positioned along the southern border of the observed dataset.

The final model is shown in Figure 7.6, for slices with centres at 0.05, 0.25, 0.45, 0.65, 0.85 and 1.05 mbgl. The model shows the presence of high apparent susceptibilities at shallow depths, with notable anomalies present from the upper slice. In the areas above each target, susceptibilities appear to peak around 0.003 SI, at depths from the surface to 0.5 mbgl. Below

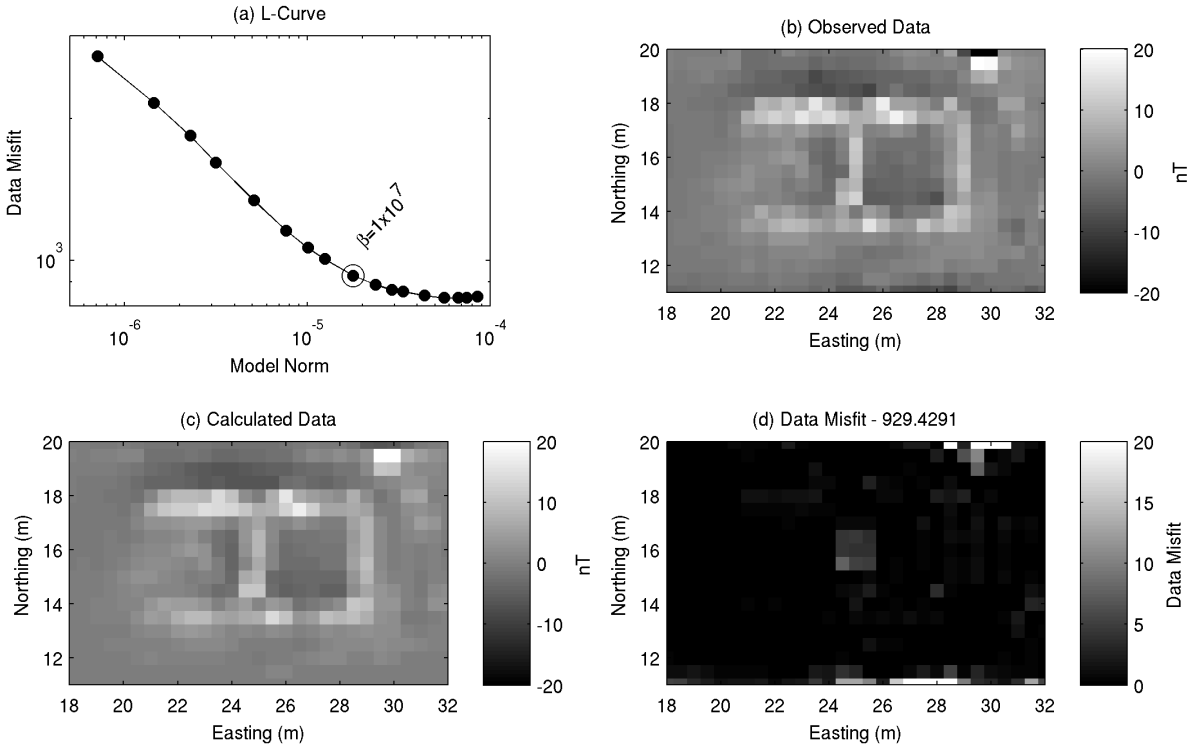

Figure 7.5: Analysis of data misfit for the inversion of Area 4 of the NSGG Test Site 1, Leicester. (a) L-curve analysis, with the chosen value of β shown. (b) Observed total-field data at a height of 0.25 m above ground level. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

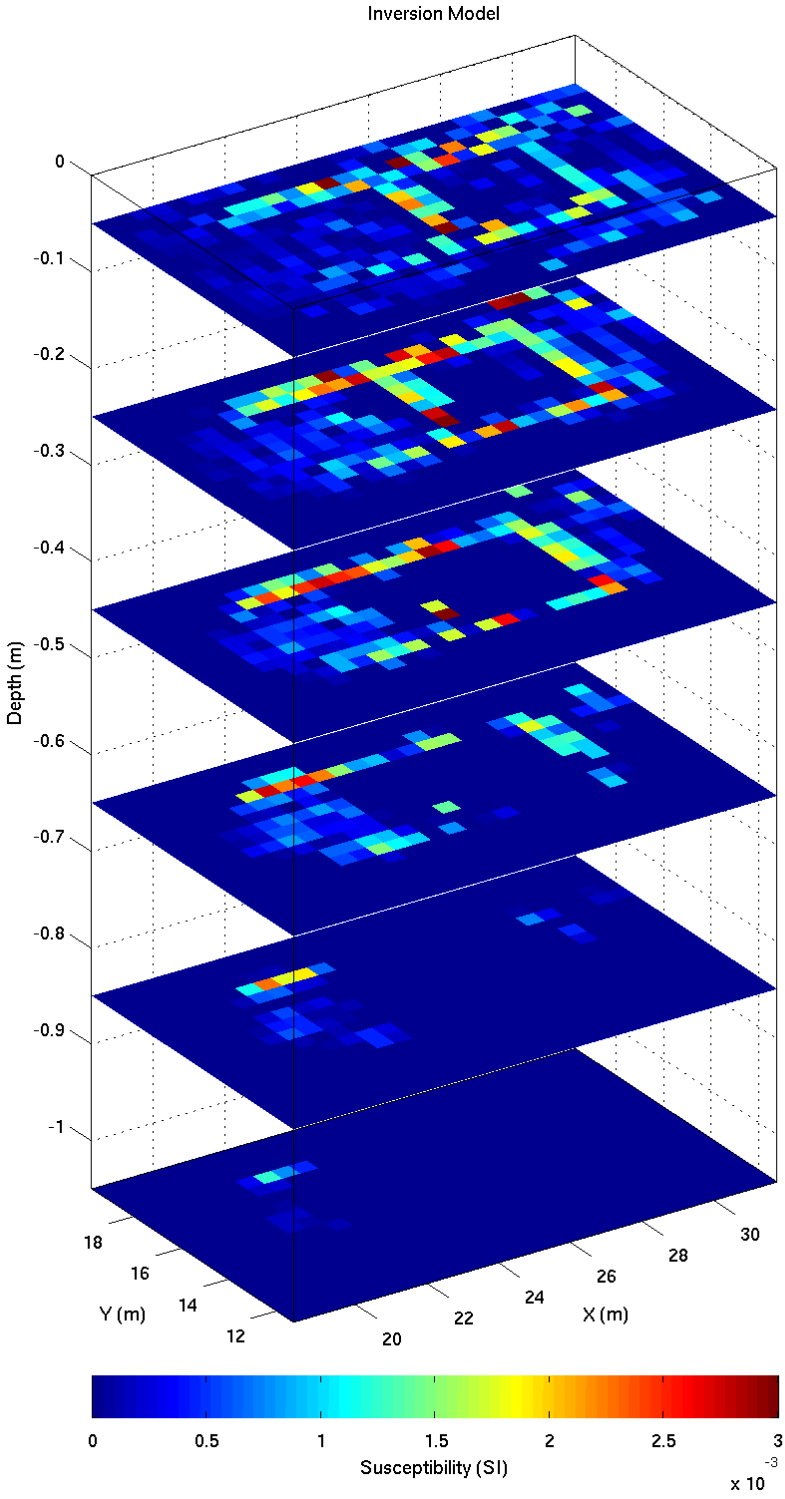

Figure 7.6: Final inversion model for Area 4 of the NSGG Test Site 1, Leicester, with positivity imposed. Slices shown have their centres at 0.05, 0.25, 0.45, 0.65, 0.85 and 1.05 mbgl.

here, a notable decrease in susceptibility is observed in all areas except for targets 2, 4, 6 and 7, which correspond to the bricks, aggregate, peat and leaf mould and paving slabs. The base of the recovered anomaly obtained for the peat and leaf mould is around 0.8 mbgl, the bricks and paving slabs can be identified to $\sim$0.9 mbgl and the aggregate to $\sim$1.1 mbgl, below which the model consists of values of $\sim$0 SI.

This result appears to be disappointing, as the majority of the foundation materials are buried between 0.6 and 1 mbgl, and therefore should appear as deeper anomalies than recovered in the model, and without the near surface high susceptibility. However, a possible explanation is that the magnetisation detected by the magnetometer, and modelled by the inversion, relate to the infill of the ditches dug to place the targets, rather than the targets themselves.

It is common for recently dug trenches to appear as a negative total-field anomaly, due to the natural remanent magnetisation being removed by the admixture of the material being excavated. Of the types of anomalies listed in Table 7.1, the bricks are most likely to consist of higher magnetisation than the surrounding material, and are also the most likely to contain a remanent magnetisation associated with the process of manufacture. The concrete, sand and stones are most likely to consist of lower susceptibilities than the surrounding matrix, and therefore relate to negative features, depending on the nature of the stones. It is therefore surprising that all the anomalies relating to the targets seen in Figure 7.1b are all positive in nature.

It is interesting to note that the bricks, which are the most likely target-type to produce a positive response, are one of the targets that appears the thickest in extent along with the aggregate, and shows high-susceptibility in the area between 0.6-1 mbgl where high susceptibility is expected. One possible explanation for the differences between the model produced in Figure 7.6 and the recorded positions of the targets in Area 4, is that the ditches dug to place the targets at the required depths, were infilled with high-susceptibility material. The magnetic response recorded over these targets is therefore positive, with little variation in amplitude, unlike what would be expected from the range of different materials. Beneath these positive ditches from

0-0.6 mbgl, anomalous susceptibilities are only being modelled for the materials which provide the positive contrasts to the surrounding soil, and therefore only the bodies relating to positive features can be observed down to depths of 1 mbgl. The fact that the ditch itself is magnetised would also help to explain the shallower than expected results from the Euler deconvolution solutions.

This theory is supported by the observation of an anomaly between 52mE, 10mN and 52mE, 20mN. This is a position away from any purposely buried targets, and was the location of tyre tracks of vehicles used in the construction of the Test Site. These tracks, as well as other areas of subsidence that later appeared across the site were infilled by the local groundsmen, using topsoil of unknown derivation (Hill, I.A, pers comm.). It is interesting to note that the anomalies relating to the tyre tracks are of a similar amplitude to the anomalies located in Area 4. This supports the theory that the origin of both the infilled tyre tracks and the anomalies identified in Area 4 are from the infill material used, rather than the target material buried.

7.3 Case Study 2 - Silchester Roman Town, Hampshire, UK

The remains of the Roman town of Calleva Atrebatum are located approximately midway between Basingstoke and Reading, near to the present village of Silchester, Hampshire. The site is known to have been an Iron Age and Roman settlement (Fulford, 2002). Visible evidence of previous occupation of this land is a defensive wall of Roman age surrounding the town and up to 4 m high in places, and Iron Age earthworks outside of this. The area located within the walls is currently pasture land with the exception of a track, mid-twelfth century church and house. Excavations have shown the site experienced dense occupation from about 25 BC until around the middle of the 5th Century (Fulford, 2002).

A survey within the town walls was conducted in June 2009 by this author assisting English Heritage. The data were collected with a 0.125×0.5 m resolution, measuring the total field

Figure 7.7: Location of data collection within the boundaries of Silchester Roman Town. The dataset is shown here underlying the interpretation of aerial photographs by the RHCME (Royal Commission on the Historical Monuments of England, 1994). The collected total field data is shown again in Figure 7.8.

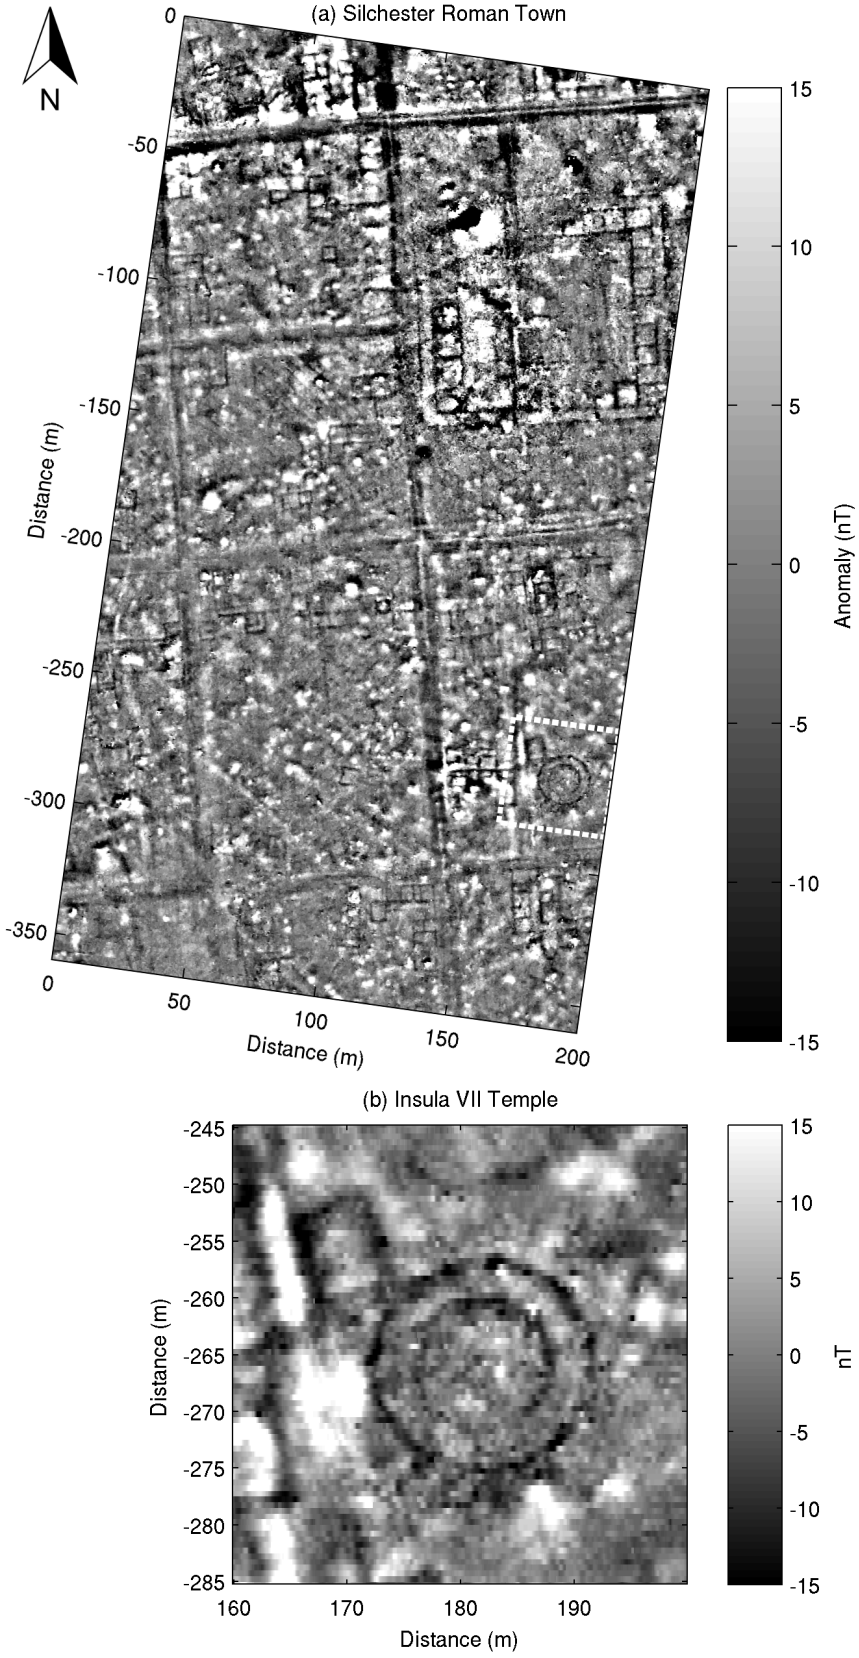

Figure 7.8: Total field magnetic suvey collected at Silchester Roman Town, Hampshire, UK. (a) Full extent of data collected. (b) Temple in Insula VII which will provide the focus of the case study.

using caesium vapour magnetometers. The extent of the data collection is shown in Figure 7.7, underlying the interpretation of aerial photography (Royal Commission on the Historical Monuments of England, 1994). The data is shown again in Figure 7.8a, where an abundance of responses indicative of Roman activity at the site can be seen. The most notable features are the linear grid pattern of roads, the Forum Basilica towards the north-east of the survey area, and the circular looking feature towards the south-east of the survey area, which has been used here as the focus of the case study. The subset is approximately 40×40 m (Figure 7.8b), and is believed to be a Romano-Celtic temple (Thomson, 1924).

The circular feature is positioned within Insula VII of the Roman town. The earliest known excavation was in 1870, by Reverend J.G. Joyce (Thomson, 1924). It was found that the foundations are not circular as they appear in Figure 7.8b, but is rather a large polygonal building consisting of 16 straight sides. Joyce however left no published account of his find, which led to the re-excavation in 1892 by the Society of Antiquaries.

The results of the Society of Antiquaries excavations are described by Fox & Hope (1894). The total diameter of the temple was found to be around 65 feet (19.81 m), consisting of two concentric rings. The foundations of outer ring are described as consisting of a slaty stone, which shows the 16-sided polygonal shape. The foundations are 3 feet (0.91 m) thick, on top of which, evidence for a wall of flint rubble masonry is found. This wall is 2 feet 5 inches (0.74 m) wide, and also shows the polygonal shape, with each section having a length of 12 feet 8 inches (3.86 m). The angles are strengthened inside and out by blocks of ironstone, which has also been used as bonds to the flint rubble masonry that make up the walls.

The foundations of the inner ring abandons the polygonal shape, and is circular, 4 feet (1.22 m) wide, also made of slaty stone. Above this stood a wall, 2 feet 6 inches (0.76 m) thick, which has the 16-sided polygonal shape on the outer side, yet is circular on the inner side. The length of each outer face is 8 feet (2.44 m). At no point during the excavation was either wall found to be greater than 1 foot 7 inches (0.48 m) above the footings.

Toward the south-west quadrant there are traces of a bed of pink cement, that have been interpreted as possibly being a bed along which gutters had been laid. It is mentioned by Thomson (1924) that the pink cement may indicate that bricks of older buildings had been ground up to make the mortar or cement.

Two photographs, believed to be showing the exposed remains of the Romano-Celtic temple, are shown in Figure 7.9. The upper photograph is believed to have been taken in 1886 and suggests that the top of the foundations were left exposed following the original excavations by Reverend J.G. Joyce. The lower photograph shows the subsequent excavation by the Society of Antiquaries in 1893. There is no visible evidence of the feature above ground at the site at the present day.

7.3.1 Edge detection

The first step in interpreting the data is to identify the lateral edge positions of the anomalies. Figure 7.10a shows the original total field data. The data has undergone a pseudogravity transformation, and then the horizontal gradient magnitude has been derived and is shown in Figure 7.10b. This transforms the data so that the peaks are positioned vertically above the edges of anomalous subsurface bodies. The theta map has also been calculated from the pseudogravity dataset and is shown in Figure 7.10c. The advantage of the theta map is again visible here, as the variable amplitude of peaks in the horizontal gradient magnitude are normalised in the theta map, however in this area there are many anomalies positioned relatively close together, so therefore both the horizontal gradient magnitude and the theta map have been used in conjunction to identify the edge positions. Using the data in Figures 7.10b and 7.10c, the location of the peaks have been interpreted manually to produce an outline of suspected buried bodies which is shown in Figure 7.10d. By identifying connected edges that relate to coherent anomalies in the data a total of 22 anomalous zones have been identified.

(a) Looking NE, photographed 19th July 1886

(b) Re-excavation by the Society of Antiquaries in 1893

Figure 7.9: Photographs of the temple in Insula VII, Silchester Roman Town. (a) Looking north-east towards St Mary's Church, taken in 1886, in between the 1872 and 1892 excavations. (b) Photograph taken during the re-excavation by the Society of Antiquaries in 1893. Photographs from Reading Borough Council (Reading Museum Service).

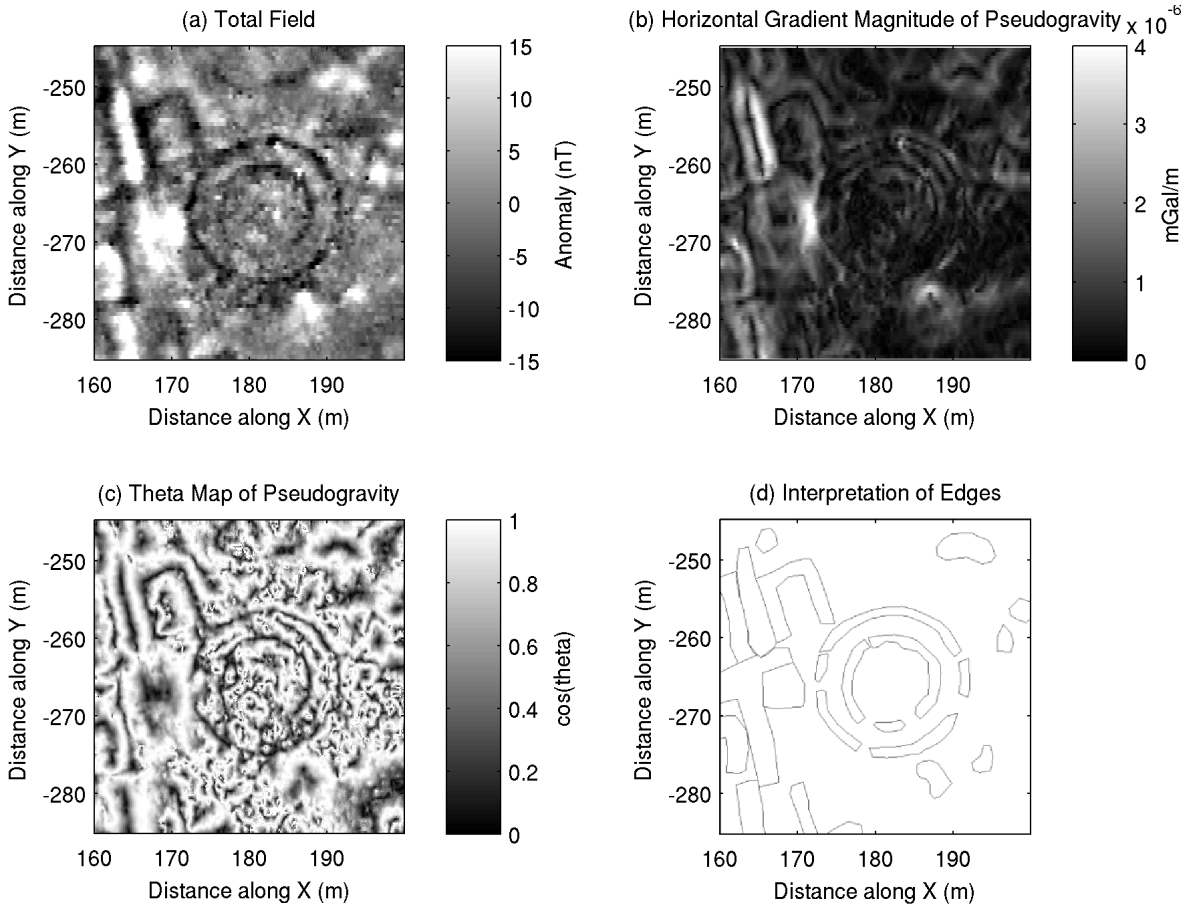

Figure 7.10: Edge detection of data from Insula VII, Silchester Roman Town. (a) The original total field data. (b) Horizontal gradient magnitude of the pseudogravity of a). (c) Theta map of the pseudogravity of a). (d) Interpretation of the edge locations based on b) and c).

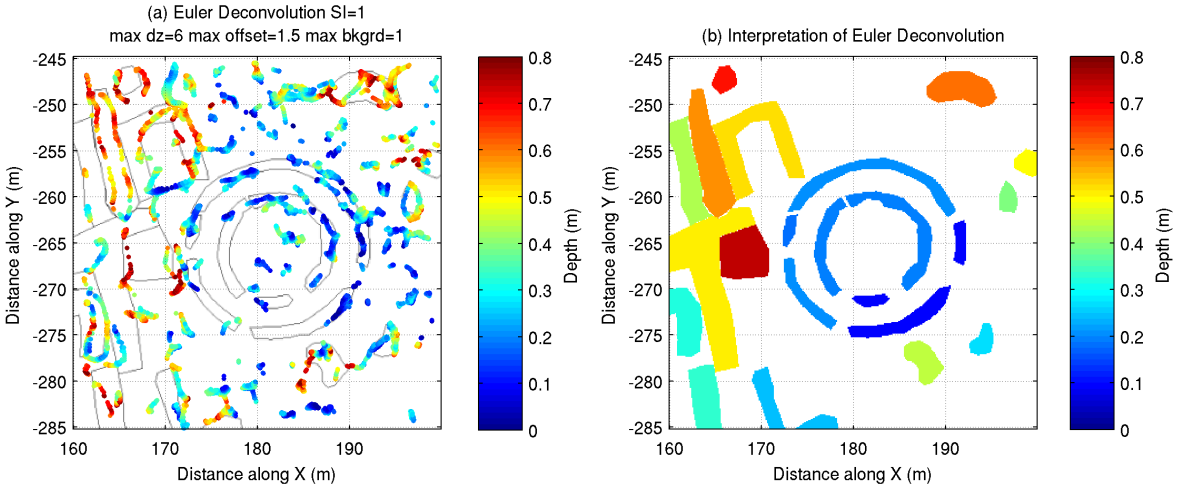

Figure 7.11: Euler deconvolution results from the Insula VII temple. (a) Solutions obtained using a structural index of 1. (b) Average depth from all solutions in each zone.

7.3.2 Euler deconvolution

It was shown in Chapter 4 that careful implementation of Euler deconvolution can provide quickly attainable estimates of depth to causative anomalies in the dataset. In a real data scenario choosing an optimum structural index for use with Euler deconvolution is difficult and often the buried bodies will not be applicable to idealistic shapes. Here, there appear to be anomalies of a variety of different shapes, therefore no single structural index is likely to be appropriate for the entire area, however Euler deconvolution has been performed with a structural index of $N = 1$, and the filtered solutions are shown in Figure 7.11. The outline of the edge detection is shown along with the Euler deconvolution solutions and it can be seen that many of the solutions are located along the lines of the determined edges. To simplify the interpretation, any solutions positioned within the lateral limits of the 22 anomalous zones have been averaged to obtain a mean depth for each anomaly zone. It is noted that this approach may not always be valid, in particular where a sloping body is located, or where the causative body has significant undulations.

Component	Field Value	Secular Value
Declination	-2.205 degrees	8.5 arcmin/year
Inclination	66.335 degrees	-0.2 arcmin/year
Horizontal Intensity	19476 nT	15.3 nT/year
North Component	19462 nT	17.2 nT/year
East Component	-749 nT	47.6 nT/year
Vertical Intensity	44441 nT	27.2 nT/year
Total Intensity	48522 nT	31.0 nT/year

Table 7.3: IGRF output, for Silchester Roman Town, June 2009

The results of the interpretation of the Euler deconvolution for the section of data surrounding the Temple in Insula VII is shown in Figure 7.11b, where blue colours represent shallow depths, and yellow and red colours represent deeper depths below the ground level. The two circular features identified as the walls of the Temple are shown to be the shallowest features in the dataset around 0.1-0.2 mbgl. Anomalies positioned to the north-west of the dataset are shown to be the deepest features, with the deepest body positioned at 0.74 mbgl.

7.3.3 Remanent magnetisation

Three-dimensional modelling of the dataset using MAG3D will only be applicable if there is no significant remanent magnetisation within the area. The IGRF for Silchester Roman Town at the time of the survey in June 2009 is shown in Table 7.3, where an inducing field with an inclination of 66.335° and a declination of -2.205° has been derived.

To test the applicability of this magnetisation vector for this dataset, multiple inclination and declination vectors have been used to reduce the data to the pole, and then the vertical gradient and analytic signal amplitude have been compared and the cross-correlation coefficient calculated as described by Dannemiller & Li (2006). This is shown in Figure 7.12a, along with the equivalent correlation based on the derivatives calculated from the pseudogravity dataset (Figure 7.12b). The white star in both cases shows the inclination and declination derived from the IGRF in Table 7.3.

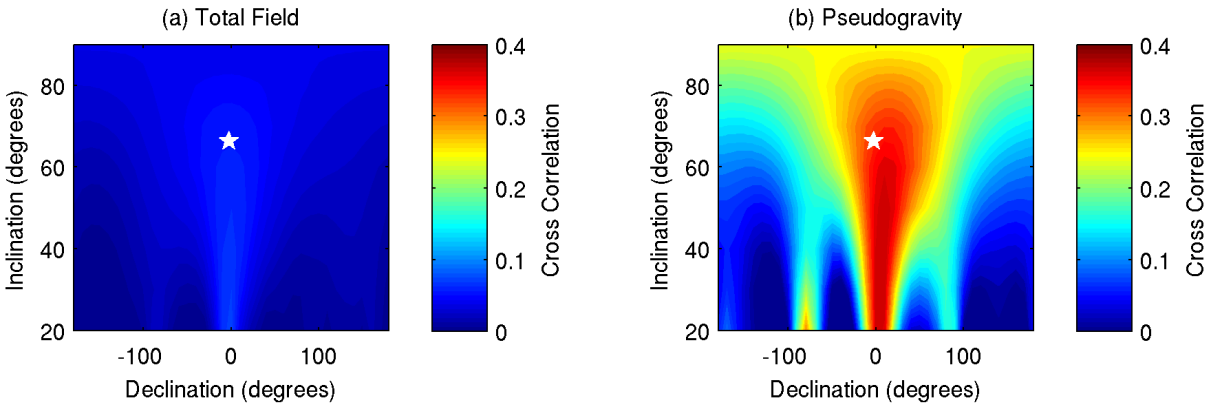

Figure 7.12: Test for the likely presence of remanent magnetisation in the subset of data from Insula VII, Silchester Roman Town. (a) Using the Dannemiller & Li (2006) method of cross-correlation between the vertical gradient and the analytic signal amplitude. (b) analysing the cross-correlation between the vertical gradient of the pseudogravity and the analytic signal amplitude of the pseudogravity. The white star represents the inclination and declination derived from the IGRF.

The high noise content present in the calculation of the derivatives from the total field data has resulted in low cross-correlation coefficients in Figure 7.12a. Higher coefficients have been derived from the derivatives of the pseudogravity, and a zone of maximum cross-correlation can be seen between inclinations of 20 to 60°, and declination of 0 to 10°. As the inducing field derived from the IGRF is outside of this zone it is possible that some remanent magnetisation could be present, and some artefacts similar to those identified in Section 5.6.3 could be present in the modelling of this dataset.

7.3.4 Initial inversions

Visual analysis of the data in Figure 7.8b suggests that the anomalies are caused by both positive and negative contrasts with the surrounding soil. In this scenario, it was shown in Chapter 5 that the inversion procedure is required to consist of three stages. During the first run a positivity constraint is placed on the recovered susceptibilities to ensure that the positive anomalies are modelled by a single positive susceptibility contrast (Figure 7.13a). For the second run a negativity constraint is imposed, so that any negative features are modelled by a negative susceptibility contrast rather than a combination of positive and negative (Figure 7.13b). For each

zone identified in the edge detection a decision is made as to whether they represent positive or negative susceptibility contrasts, and a reference model is formed by placing the susceptibilities from the relevant model into the reference model for the cells beneath each zone (Figure 7.13c). All other parts of the reference model are set to a zero susceptibility.

The positivity and negativity constrained inversions have been conducted to find both the smallest model ($\alpha_s = 1, \alpha_x, \alpha_y, \alpha_z = 0$), and the smoothest model ($\alpha_s = 0.0001, \alpha_x, \alpha_y, \alpha_z = 1$), using a mesh of $0.5 \times 0.5 \times 0.1$ m cells. There is little difference between the recovered structures of the two models. However, it is noted that in the locations where the inversions attempt to model anomalies of the opposite polarity to the susceptibility constraint, sharp discontinuities exist when the smoothness requirement was not imposed. For this reason the smoothness model was preferred. Both the positivity and negativity model were chosen from analysis of the respective L-curves, with the β representing the solution in the 'knee' of the curve chosen.

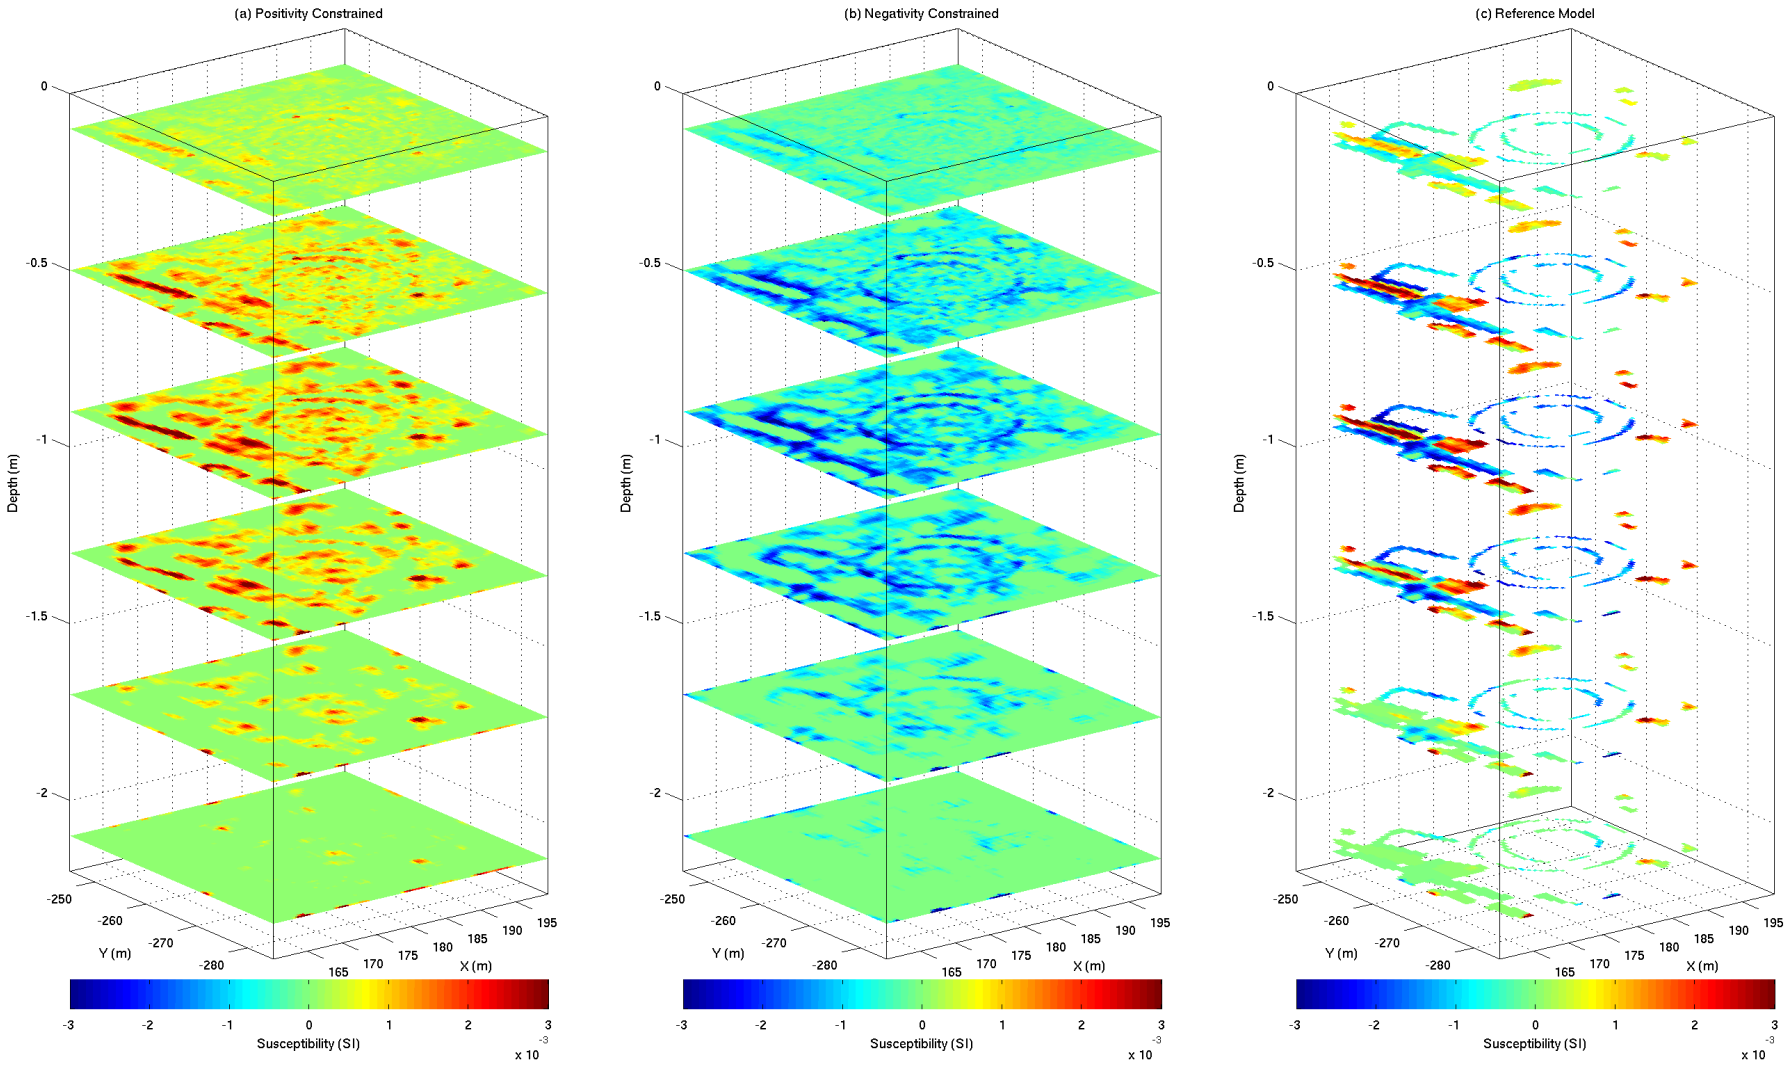

Figure 7.13: Reference model generated by selecting from the positivity and negativity-constrained inversions dependant zones from edge detection. (a) Positivity constrained model. (b) Negativity constrained model. (c) Reference model created by choosing from either the positivity or negativity model for each anomaly zone identified by the edge detection. Slices shown have their centres at 0.1, 0.5, 0.9, 1.3, 1.7 and 2.1 mbgl.

7.3.5 Inverse model

In order to obtain the final model for the dataset, a third run of the inverse procedure is required, this time relaxing the constraints on the recovered susceptibilities to $-0.5 \leq \mathbf{m} \leq 0.5$. The reference model shown in Figure 7.13c is used in the objective function in place of a zero-halfspace model. The smoothness constraint applied to the positivity and negativity inversions are no longer required, and the smallest model is sought by applying $\alpha_s = 1, \alpha_x, \alpha_y, \alpha_z = 0$.

A final L-curve was derived by varying the regularisation parameter, and an optimum value of $\beta = 5 \times 10^7$ was identified from Figure 7.14a. The calculated data, obtained by forward modelling the final model is shown in Figure 7.14c, and can be compared to the observed data in Figure 7.14b. The data misfit between these two datasets is shown in Figure 7.14d. A good correlation can be observed between the observed and calculated dataset, with only isolated peaks present in the data misfit, representing point sources that cannot be accurately modelled with the resolution of the mesh chosen.

The final model of the Insula VII temple is presented in Figure 7.15 for slices with centres at 0.1, 0.5, 0.9, 1.3, 1.7 and 2.1 mbgl. The temple is the most striking feature in the dataset represented by the double circular pattern towards the centre of the area. It is visible in the slices immediately below the ground surface, however a notable increase in the susceptibility contrast is present at 0.5 mbgl, and is visible in depth slices to 2 mbgl. The susceptibility contrast between the surrounding soil and the temple is around -0.002 SI.

7.3.6 Comparison to ground penetrating radar

The site at Silchester has been surveyed several times, with different methods. It is particularly useful to compare the results of the 3D magnetic modelling to datasets from a ground

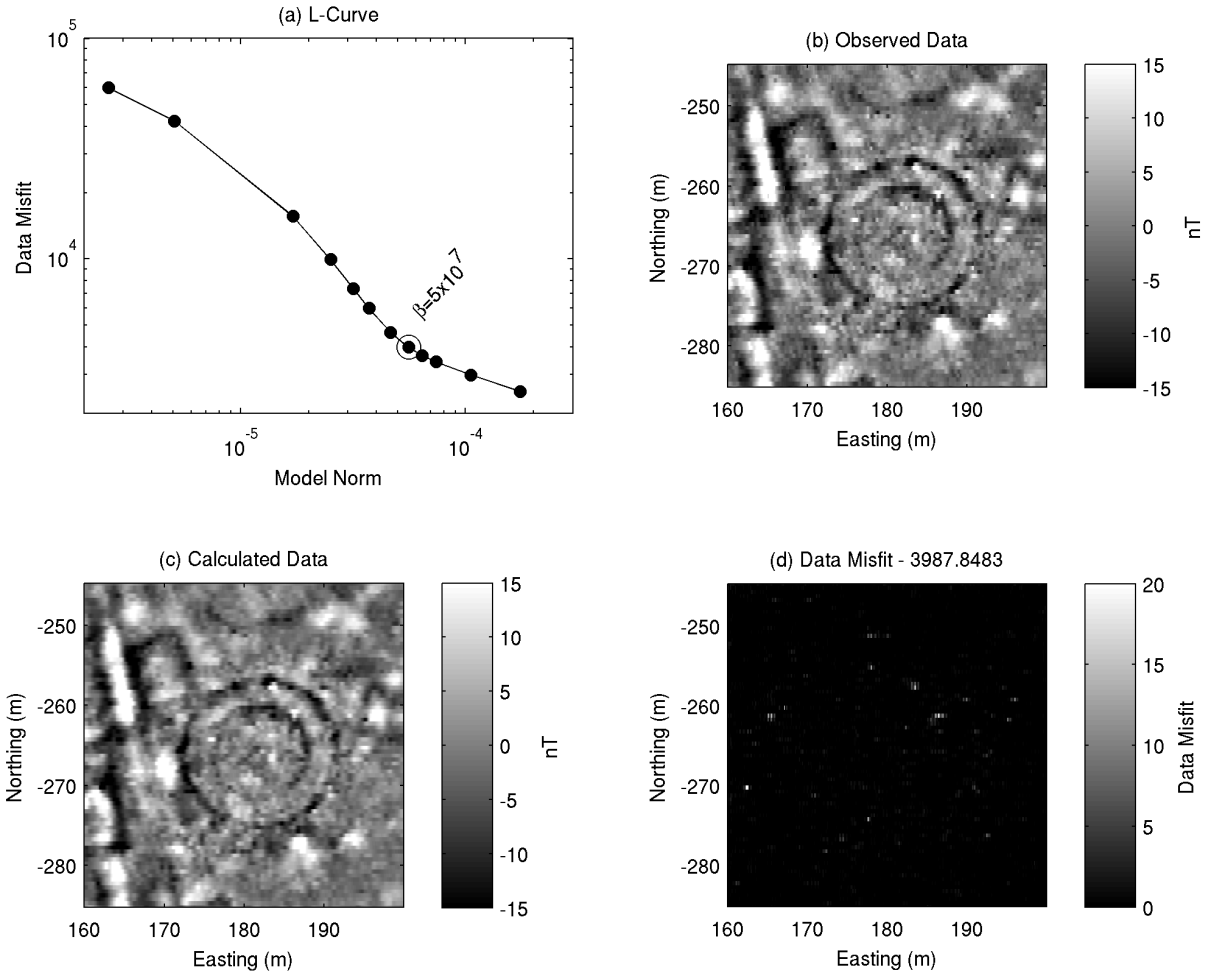

Figure 7.14: Analysis of data misfit for the inversion of a subset of data from Insula VII, Silchester Roman Town. (a) L-curve analysis, with the chosen value of β shown. (b) Observed total-field data at a height of 0.25 m above ground level. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

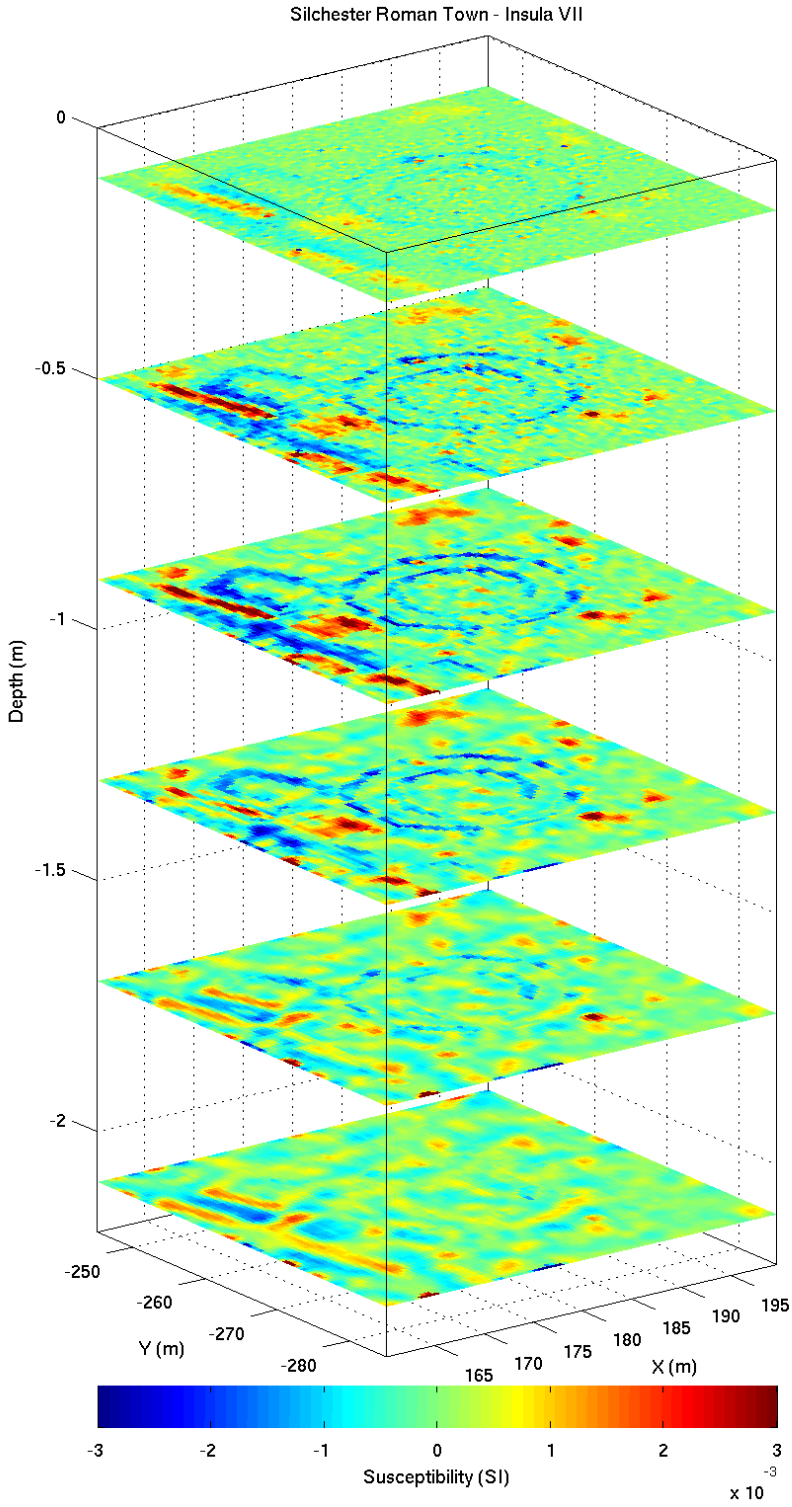

Figure 7.15: Final inversion model for the Insula VII temple at Silchester Roman Town, using the reference model shown in Figure 7.13 and relaxing the constraints on susceptibilities. Slices shown have their centres at 0.1, 0.5, 0.9, 1.3, 1.7 and 2.1 mbgl.

penetrating radar (GPR) survey, where the presentation of the final model with depth slices is commonplace for the interpretation of the data.

Here, two GPR datasets are shown, both collected by English Heritage. The June 2009 dataset (Linford *et al.*, 2010), shown in Figure 7.16b was collected at the same time as the magnetic survey interpreted above. Data collection was with a 3D-Radar GeoScope stepped-frequency continuous wave GPR. The system works by transmitting a continuous-wave signal over a series of discrete frequency steps ranging between 100 MHz and 2 GHz at each recording station. Using GPR, very high-resolution data can be collected with high frequencies, yet the depth of penetration of these frequencies are generally quite low. The advantage of the stepped-frequency system is that it allows both the high resolution data from the near-surface to be collected, as well as deriving data from the required penetration depths using lower frequencies. A horizontal resolution of 0.075×0.075 m was achieved by towing the GPR behind a vehicle.

A more traditional GPR survey was conducted in March 2000 using a Pulse Ekko 1000 impulse GPR system with a 450 MHz antenna (Linford, 2001). The sample interval derived with this system was 0.05×0.5 m and selected slices are shown in Figure 7.16c.

For comparison, the results of the 3D magnetic modelling are shown again in Figure 7.16a, this time plotted on a grayscale colour scheme. The higher resolution of the GeoScope dataset is a particularly notable feature, yet the position and extents of the temple are shown in all three datasets. The temple foundations are also visible in the shallowest slices from the GeoScope dataset (0.1 and 0.3 mbgl), suggesting that the foundations are either buried very shallowly, or that there has been significant disturbances to the ground immediately above the foundations due to the excavation work. The weaker response in the upper slices may also be indicative of a response to the remains of the walls, built on top of the foundations, and thinner than the foundations they sit on, as described by Fox & Hope (1894).

Below this, all three methods show a strong response to the temple foundations to a depth of around 1.7 mbgl, where the amplitude of the magnetic and GeoScope GPR anomaly suddenly

decreases, and the signal to noise ratio of the PulseEkko GPR causes any response to be obscured by noise.

A discrepancy between the magnetic model and the GPR data can be seen to the north-west of the temple in the suspected store of the temple (Thomson, 1924). This is a rectangular building that is clearly identifiable in the GeoScope data between 0.7 and 1.1 mbgl. This anomaly is also represented in the magnetic model, however by greater extents, as it is prominent in all depth slices down to 1.3 mbgl. One potential cause for this discrepancy is the method of separating the inversion into a three stage process, with initial positivity and negativity-constrained inversions. It is demonstrated in Chapter 5 that this process is necessary to prevent the individual magnetic contrast of a subsurface body to be represented by alternating positive and negative susceptibilities (see Figure 5.29). However, towards the north-west of the survey area of Insula VII, there are both high amplitude positive and negative anomalies situated immediately adjacent to each other. It is therefore natural in this situation, that the anomaly recorded in the total field data will be caused by interfering positive and negative anomalies, and the procedure may not be valid in this area.

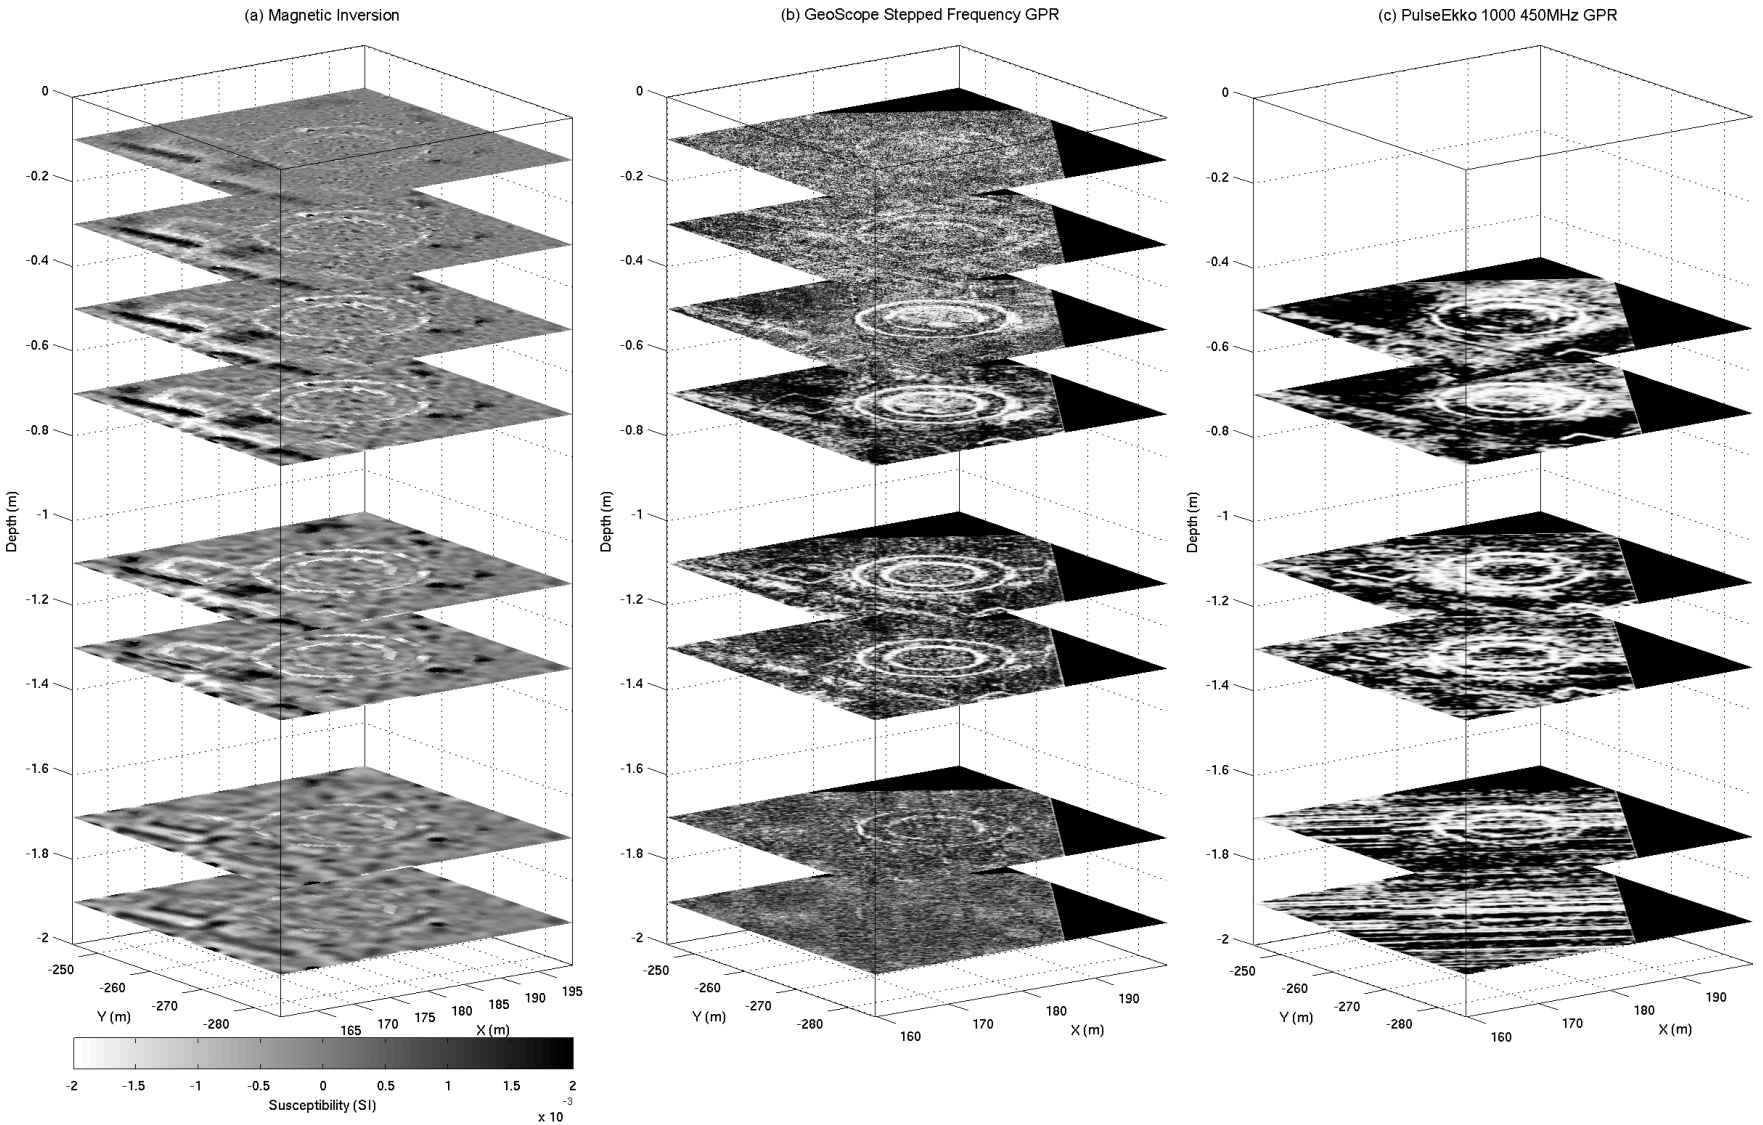

Figure 7.16: Comparison between magnetic inversion model and GPR depth slices. (a) Magnetic inversion with slices shown at 0.1, 0.3, 0.5, 0.7, 1.1, 1.3, 1.7 and 1.9 mbgl (b) GeoScope 100-1200 MHz stepped frequency (from Linford *et al.* (2010)). (c) PulseEkko 1000, 450 MHz antenna (from Linford (2001))

7.4 Case Study 3 - Brading, Isle of Wight, UK

Brading Roman Villa is situated between the towns of Brading and Sandown on the Isle of Wight, UK. The villa itself was revealed by excavations from the late 1870s until the mid-1880s, which discovered an extensive Roman villa (Price & Price, 1890). Later excavations exposed mosaics, which are now preserved in-situ, but undercover in a visitors centre built around the remains of the villa.

Total field magnetic data were collected over the surrounding grassland to the north, west and south of the villa location in February 2009 with this author assisting the English Heritage Geophysics Team. The aim of the survey was to gain information with regards to the standing of the villa in the surrounding landscape (Payne, 2009). Prior surveys to the east of the villa have found strong evidence for previous Iron Age occupation of the site (Trott, 1999).

The data collected are shown in Figure 7.17. The results show significant activity in the area surrounding the Roman villa, with numerous linear ditch systems, pit-type anomalies and three circular enclosures. These have been interpreted by Payne (2009) as most likely being indicative of the ploughed out remains of a small Bronze Age barrow cemetery, however it is speculated that due to their close relationship with the Roman villa, they may be a complex of Romano-British shrines related to the villa.

The subsection chosen for further investigation by inverse modelling is a 50×50 m area including the largest circular anomaly identified in Figure 7.17, linear anomalies appearing to connect to the circular feature, three pit-like anomalies to the south-west, and a rectangular anomaly towards the east. The area is shown by the dashed lines on Figure 7.17.

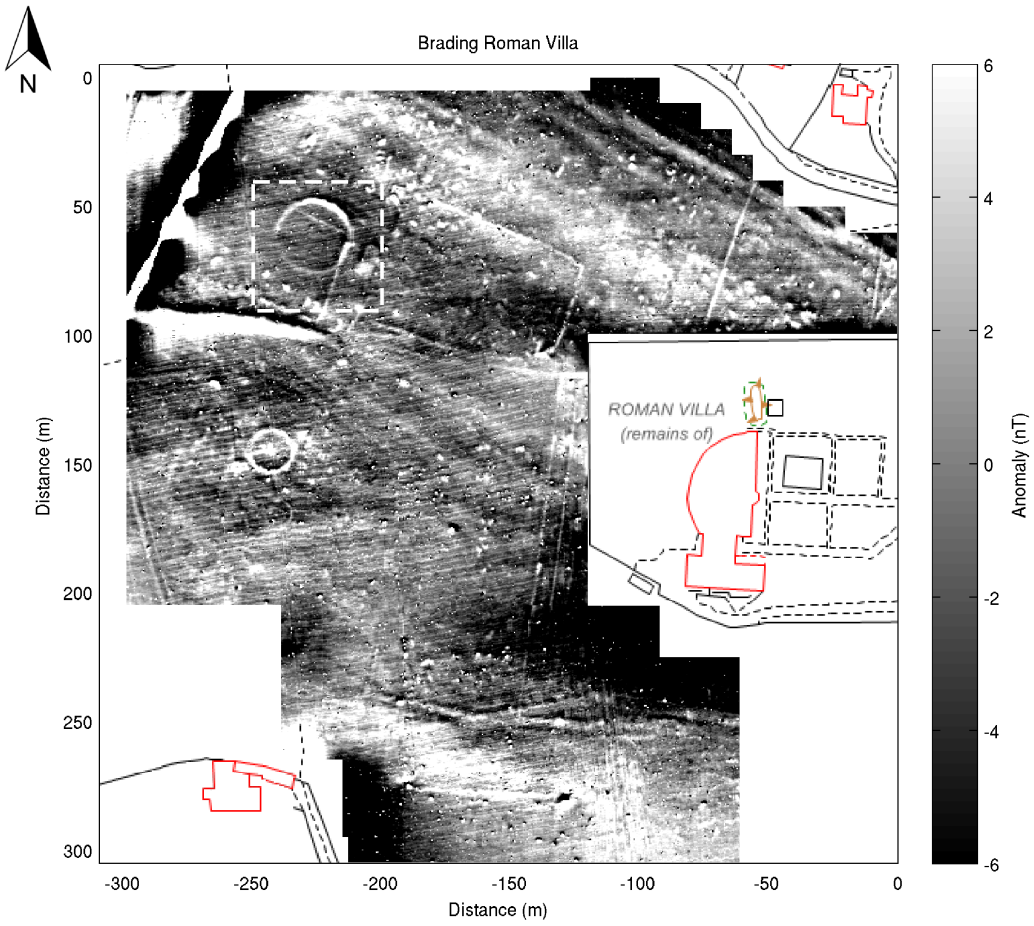

Figure 7.17: Total field magnetic survey surrounding Brading Roman Villa, Isle of Wight, UK. The dashed area shows the subset identified for further investigation.

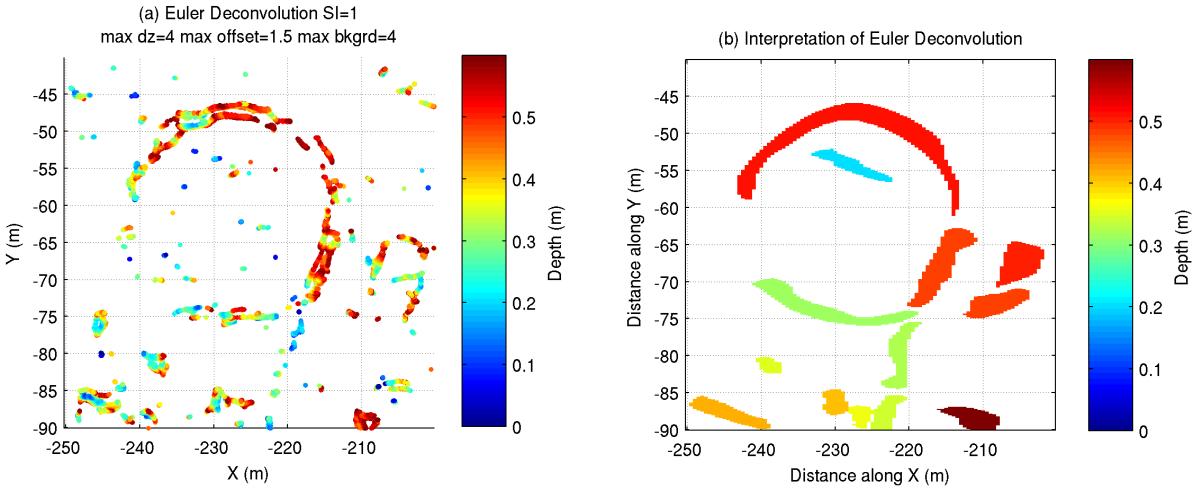

Figure 7.18: Euler deconvolution results from Brading, Isle of Wight. (a) Solutions obtained using a structural index of 1. (b) Average depth from all solutions in each zone.

7.4.1 Euler deconvolution

The main features of interest in Figure 7.17, consist of elongated anomalies with one small finite dimension, e.g. perpendicular to the suspected ditch, and one large dimension along the direction of the ditch. For Euler deconvolution purposes, this suggests a structural index of $N = 1$ or 2, for a thin sheet, or line of dipoles approximation, respectively. As the width of the response to the features appears to be up to $\sim$2 m wide, these are likely to more closely represent a thin sheet, and therefore a structural index of $N = 1$ has been chosen.

The filtered Euler deconvolution solutions are shown in Figure 7.18a. The edges of suspected anomalies have been determined from the horizontal gradient magnitude (Figure 3.17) and theta map (Figure 3.18) of the pseudogravity dataset, and all solutions positioned within the zones are used to produce the interpretation shown in Figure 7.18b.

The results show the depth to the northern part of the circular anomaly at around 0.5 mbgl, with the southern section positioned at 0.3 mgbl. The linear anomaly positioned within the circular feature is shown to be shallower at 0.25 mgbl.

Component	Field Value	Secular Value
Declination	-2.096 degrees	8.7 arcmin/year
Inclination	65.762 degrees	-0.4 arcmin/year
Horizontal Intensity	19828 nT	15.5 nT/year
North Component	19815 nT	17.4 nT/year
East Component	-725 nT	49.4 nT/year
Vertical Intensity	44042 nT	21.1 nT/year
Total Intensity	48299 nT	25.6 nT/year

Table 7.4: IGRF output, for Brading Roman Villa, February 2009.

7.4.2 Remanent magnetisation

The output from the IGRF for Brading Roman Villa in February 2009, when the survey was conducted, is shown in Table 7.4. The inclination of the inducing field is shown as $65.76°$, and the declination $-2.096°$. To identify the effect of any remanent magnetisation, the cross-correlation between the vertical gradient and the analytic signal amplitude was calculated for a range of inclinations and declinations, as described by Dannemiller & Li (2006). The results for this are shown in Figure 7.19a. The same procedure has been performed on the pseudogravity data, with the results of the cross correlation shown in Figure 7.19b.

The two plots show significantly different results. The cross-correlation using derivatives of the total field show an optimum inclination slightly less than that of the inducing field, however the maximum cross-correlation coefficient is generally low, peaking at 0.15. It is interesting to observe that the cross-correlation between the derivatives of the pseudogravity contradicts this, with an optimum inclination around that of the inducing field ($65.76°$), but with a declination towards the east of the inducing field. Therefore, it is not possible to rule in, or rule out the presence of remanence at this site, and interpretation of the inverse model should be done with caution.

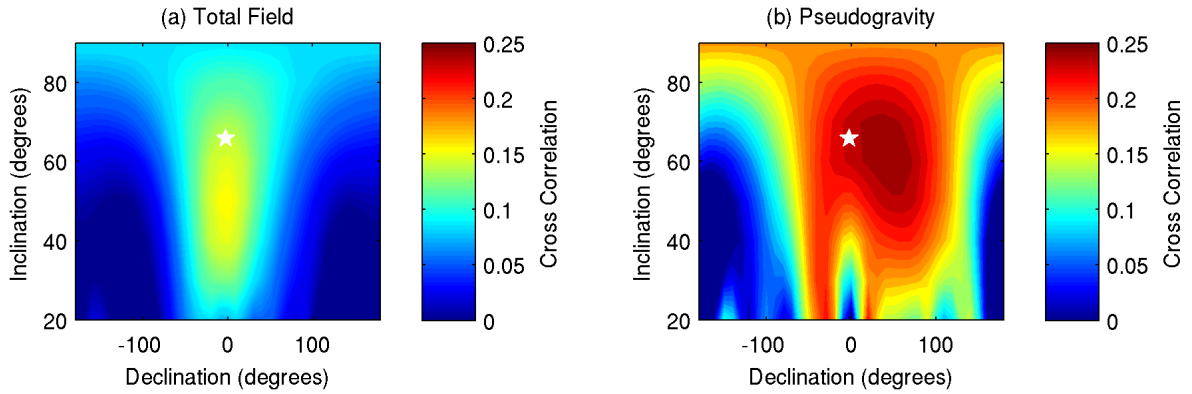

Figure 7.19: Test for the likely presence of remanent magnetisation in the subset of data from Brading, Isle of Wight. (a) Using the Dannemiller & Li (2006) method of cross-correlation between the vertical gradient and the analytic signal amplitude. (b) Analysing the cross-correlation between the vertical gradient of the pseudogravity and the analytic signal amplitude of the pseudogravity. The white star represents the inclination and declination derived from the IGRF.

7.4.3 Preparation of data

Prior to the inversion it is necessary to remove features that are either not desired in the final modelling, or cannot be modelled using the extents of the mesh provided. In this case, the data has undergone a two-stage filtering technique. First the PATH filter was applied between the angles of 75 and 85°, with a normalised energy density amplitude of >-12. This filter is the same as that applied in Figure 6.12, and has been designed to remove the agricultural pattern of tracks that results in linear anomalies across the site. Secondly, in order to remove any regional trends that may be present in the dataset, a Butterworth high-pass filter has been applied using the parameters $k_0 = 0.5$ and $n = 8$.

7.4.4 Inverse model

In order to produce a 3D model of the apparent subsurface susceptibilities, the data have been inverted with a 50×50×6 m mesh, consisting of cells of 0.5×0.5×0.2 m. As there are no distinctive negative anomalies in Figure 7.20b, a positivity constraint has been applied to the recovered susceptibilities. The smallest model has been derived by setting $\alpha_s = 1$, and $\alpha_x, \alpha_y, \alpha_z = 0$.

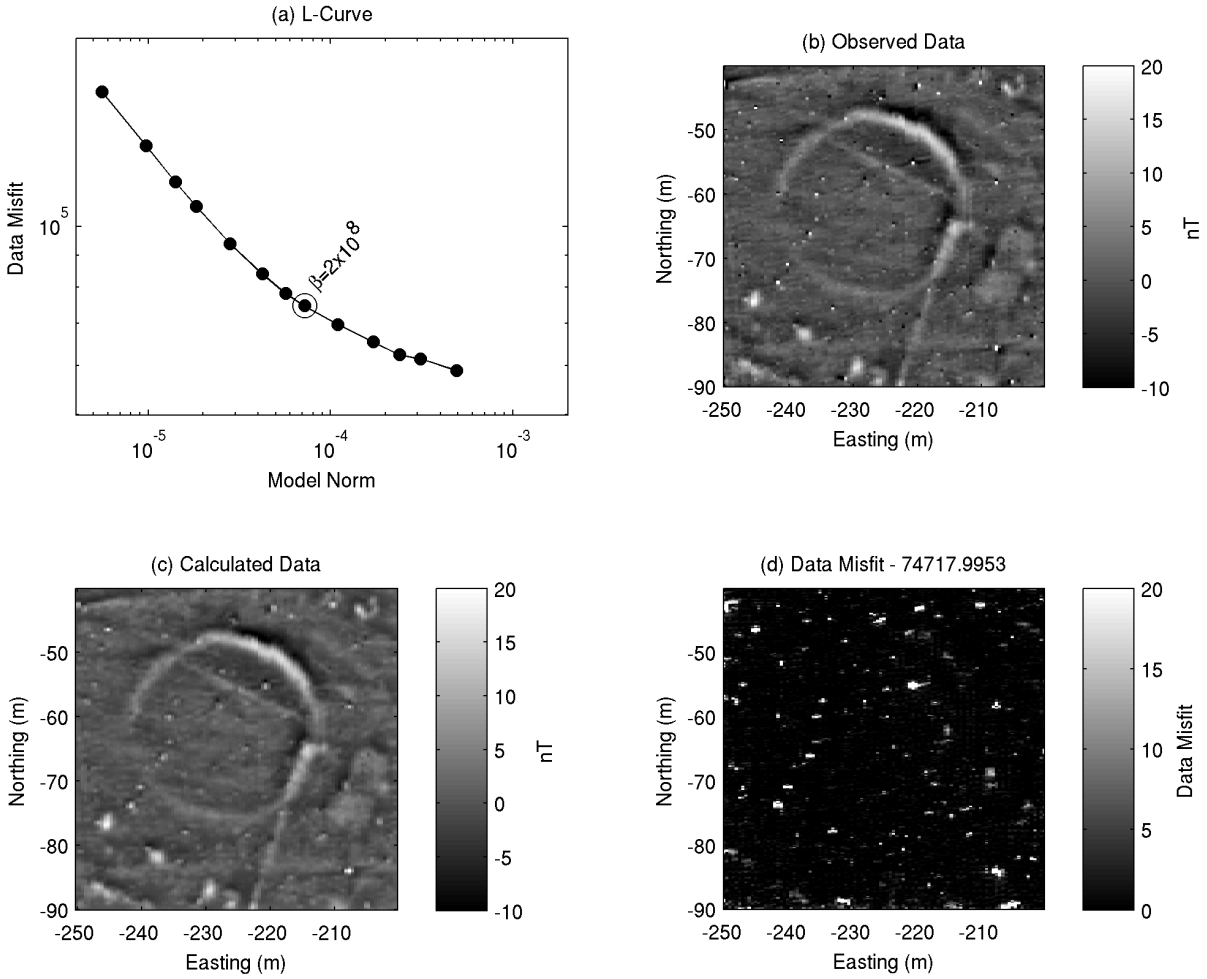

Figure 7.20: Analysis of data misfit for the inversion of a subset of data from surrounding Brading Roman Villa, Isle of Wight. (a) L-curve analysis, with the chosen value of β shown. (b) Observed total-field data at a height of 0.25 m above ground level. (c) Calculated data, generated by forward modelling the final inversion model. (d) Data Misfit calculated as shown in Equation 5.2.

The inversion has been performed for multiple values of β and the L-curve is shown in Figure 7.20a. Using analysis of the shape of the L-curve, an optimum value for β of 2×10^8 has been chosen. The calculated data and data misfit are shown in Figures 7.20c and 7.20d respectively. A good match can be seen between the observed and calculated datasets, with the data misfit showing high values only in the area correlating to dipole-like responses in the observed data that cannot be modelled accurately enough using the 0.5×0.5×0.2 m cell size.

The model produced from the inversion using $\beta = 3 \times 10^8$ is shown in Figure 7.21 for slices at 0.1, 0.5, 0.9, 1.3, 1.7 and 2.1 mbgl. The model shows a response to both the northern side of the circular body and the connecting linear features in the slice at 0.3 mbgl, which become significantly larger in the slice at 0.5 mbgl. Below this at 0.7 mbgl the southern side of the circular body becomes prominent, and there is a notable decrease in the amplitude of the linear bodies. Below 0.9 mbgl, the linear features can no longer be identified as high susceptibilities, and the model is dominated by the presence of the circular body which has a much deeper extent down to a depth of 2.1 mbgl.

The three positive anomalies positioned to the south-west of the circular body and initially interpreted as pit like features are shown to produce high susceptibilities in the slices between 0.7 and 1.7 mbgl.

The results correlate with those from the Euler deconvolution, in that the top of both the circular and linear features appear strongly in the slice centred at 0.3 mbgl. The Euler deconvolution results give no indication to the vertical extents of the body however, which have been shown in the inverse modelling to extend to depths of 1.8 m.

7.5 Conclusions

The methodologies developed over the previous chapters has been applied to data from three case study datasets here, and the following conclusions can be made about the methodology:

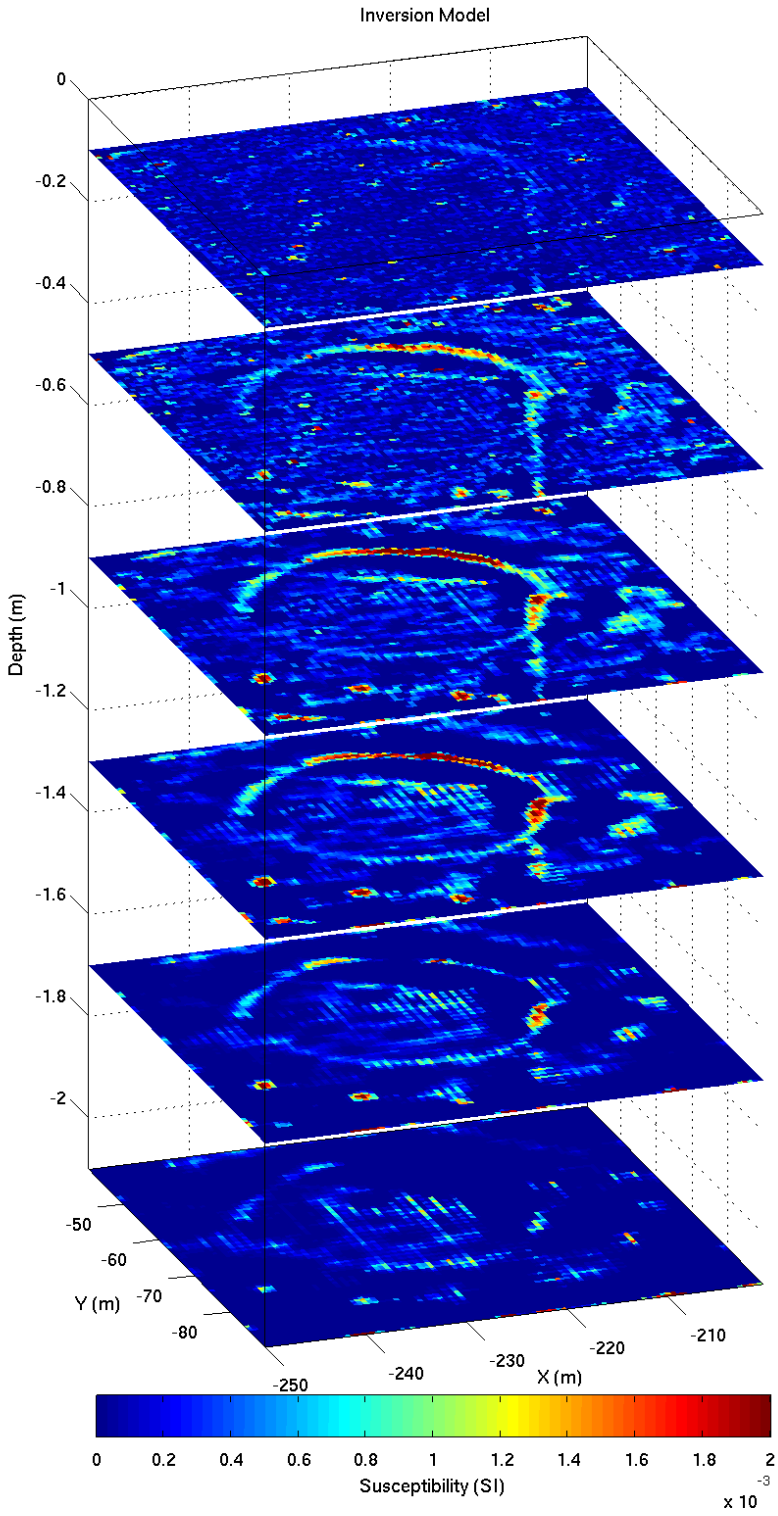

Figure 7.21: Final inversion model for a subset of data from surrounding Brading Roman Villa, Isle of Wight, with positivity imposed. Slices shown have their centres at 0.1, 0.5, 0.9, 1.3, 1.7 and 2.1 mbgl.

- Choosing an optimal structural index for use in Euler deconvolution has been problematic, even where the dimensions of the buried features are already known. Real-case scenarios are particularly unlikely to be well-represented by idealistic shapes at the depths typically encountered in archaeological site investigation.

- A high pass filter has been shown to be a simple and effective method of removing regional trends from datasets that cannot be replicated in a normal mesh extent.

- More complex filtering routines have been demonstrated to be effective at removing specific information from the dataset that is not desired in the final model. Here, the PATH filter been applied to linear anomalies present across the entire data area from recent agricultural activity.

- 3D inverse modelling of the archaeo-magnetic datasets has produced models of the magnetic susceptibility distribution in the subsurface, which can either be compared to similar 3D models attained from other geophysical techniques, or used on their own to target areas of interest for further investigation.

On a site specific basis, the following conclusions can be made regarding the interpretation of the modelling:

- The modelling results of Area 4 of the NSGG Test Site 1, reveals unexpected results considering the prior knowledge of the buried features that are present at the site. Anomalous high susceptibility material has been modelled within the upper 0.5 m of the subsurface, which appear to mask many of the target materials buried below. It is noted that the amplitude of the total field anomaly dataset above the buried features is always positive and around 10 nT in amplitude, which is surprising considering the range of foundation materials present. It is speculated that the magnetic data is primarily sensitive to the infill of the ditches and not the target materials. This is supported by the Euler deconvolution results, which also suggest shallower than expected anomalies. The inversion model shows

some evidence of high susceptibilities in the vicinity of the brick and aggregate material, yet is dominated by the anomalies in the upper near surface.

- Silchester Roman Town provides a site where the results of the 3D modelling can be readily compared to existing datasets obtained by other geophysical techniques, particularly GPR. The magnetic model of the temple situated in Insula VII has been modelled completely independently to the GPR, yet shows a good correlation with the GPR depth slices. Both datasets show the temple is buried very shallowly below the ground surface, with stronger responses being derived from the foundations of the temple walls down to a depth of $\sim$2 mbgl.

- The modelling of the dataset from Brading, Isle of Wight, shows the advantages to this 3D quantitative interpretation compared to existing magnetic interpretation methods in areas where no other data-type is at hand. The total field data in the subsection of data chosen for modelling at Brading shows overlapping responses to both a large circular feature as well as connected linear anomalies. Traditional interpretation of this dataset would not allow the user to speculate on which set of anomalies is more recent using just the data at hand. However, here the modelling shows that while the top of both bodies are situated at approximately the same depth, the vertical extent of the circular feature is significantly greater than the vertical extents of the linear pattern.

The varied case histories show added value in each case by applying the advanced interpretation and inversion.

Chapter 8

Conclusions

The objective of this thesis was to produce a 3D model of the subsurface magnetic properties, which is consistent with the characteristics of typical archaeological structures and comparable to models of other geophysical properties and excavation information. As a result, a new workflow is proposed for the processing and interpretation of archaeo-magnetic datasets, and is shown in Figure 8.1. Depending on the level of interpretation required, the workflow can be terminated at a number of points, illustrated by the different colours in Figure 8.1.

The simplest interpretation of magnetic data is generally a visual inspection of the 2D image of the data, once the data has been corrected for the artefacts introduced during the data collection process. This is shown by the *dark blue section* at the top of Figure 8.1, and was discussed in Chapter 2. The majority of this processing is standard and routinely applied within many magnetic processing software packages.

One initial improvement on the purely visual inspection of the dataset involves transforming the data into a form that allows an easier interpretation to be made, and one that should be independent of bias by the interpretor. Derivative-based methods have gained popularity in aeromagnetic processing and interpretation due to their ability to simplify large spatial datasets. Chapter 3 demonstrates that in the presence of small amounts of noise, these methods can be severely affected to the point where anomalies still identifiable in the total field data are uninter-

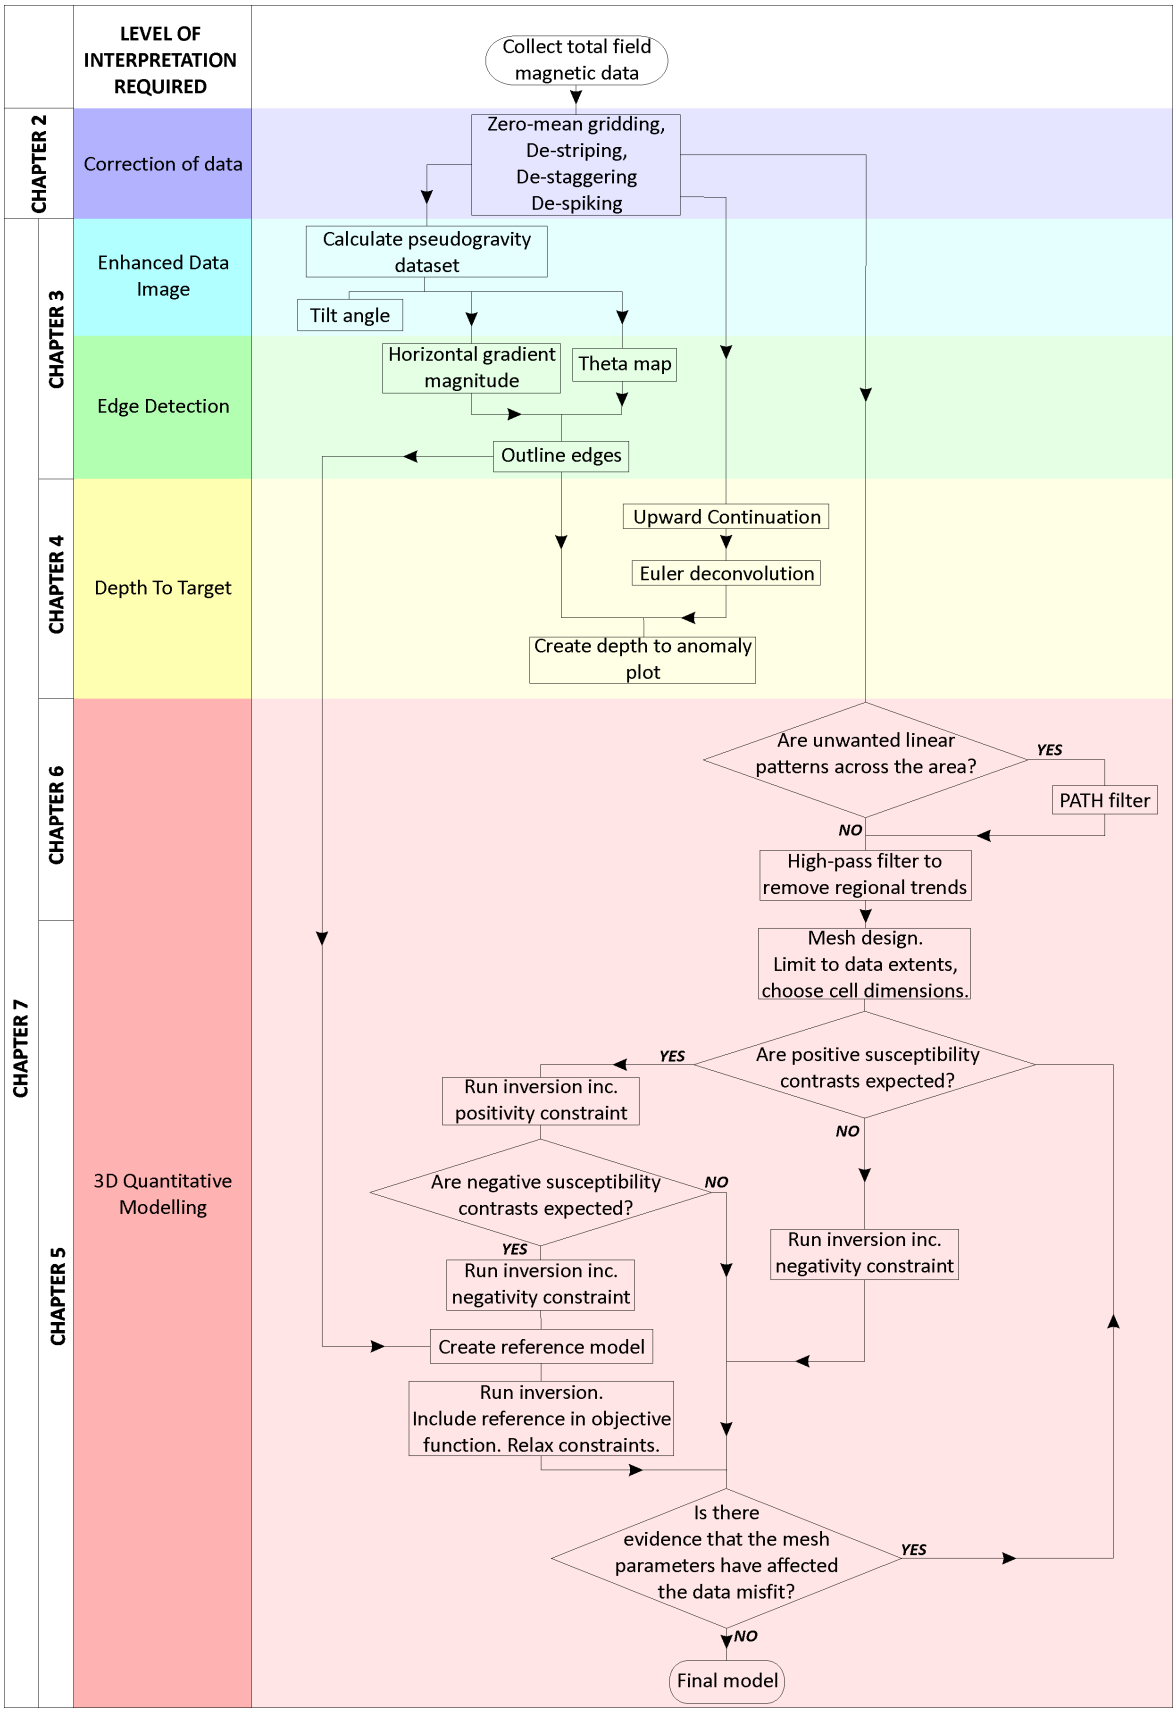

Figure 8.1: Workflow for processing and interpreting archaeo-magnetic datasets. Depending on the level of interpretation required, the workflow can be followed to any of the stages shown along the left hand side.

pretable in the derivative datasets. To overcome this problem, it is shown that the pseudogravity transformation increasingly suppresses power in higher wavenumbers, performing the opposite effect to calculating the derivatives, where the power is increased in the higher wavenumbers. Calculating derivatives of the pseudogravity data, instead of the total field data removes this bias to high wavenumber features, and results in a dataset that has a comparable power spectra to the original data, yet has all the advantages of the various derivative-based methods at enhancing the interpretation (see Cheyney *et al.* (2011)). This satisfies Objective 1, as laid out in the introduction.

Calculating the tilt angle of the pseudogravity data has been shown to be an effective way of normalising the amplitude of anomalies to improve the identification of weaker responses. This is shown by the *light blue section* of Figure 8.1. When the horizontal gradient magnitude and theta map are calculated based on the derivatives of the pseudogravity dataset, peak amplitudes are positioned above the lateral edge position of subsurface bodies, and can therefore be used to aid edge detection. This is shown by the *green section* of Figure 8.1.

Derivative-based methods can also be used to obtain estimates of the depth to simple sources, where the rate of attenuation of the magnetic field with distance is known. In Chapter 4 the tilt-depth method, a simple interpretation of the tilt angle dataset was examined for its usefulness in providing depth estimates to contact-like features, however due to the very-near surface nature of archaeo-magnetic surveys, it was found to be unlikely that an archaeological feature can be simplified to such a source shape, and therefore the method is limited in use. Other derivative-based techniques such as the SPI method show potential when applied to synthetic noise-free datasets, yet appear limited when applied to typical archaeo-magnetic situations.

Simple depth estimations can be achieved using Euler deconvolution. Although affected by the high wavenumber content of the data, a commonly applied method of suppressing the noise is to upward continue the signal prior to solving Euler's homogeneity equations. This is shown in the *yellow section* of Figure 8.1. Euler deconvolution requires the user to input a structural

index which represents the rate of attenuation of the field with distance, for the buried feature. The structural index for a range of idealistic features are known, however real archaeological anomalies are unlikely to be perfectly represented by an idealistic shape, and unlike idealistic bodies, the rate of attenuation will vary depending upon the distance between source and instrument. Chapter 4 shows that careful interpretation of Euler deconvolution results is required, and that the user must be aware of the effects on the final results of choosing an incorrect structural index, and upward continuing to a point where adjacent anomalies have an interference effect.

The methods shown in the *blue, green and yellow sections* of Figure 8.1 can be achieved relatively quickly and may provide all the interpretation required from the dataset, particularly datasets collected over large areas where a rapidly generated anomaly plan is desired. However, ideally what is desired from the survey is information about the distribution of magnetic susceptibility that produces the anomalies, so that the causative bodies, rather than their magnetic responses, can be interpreted. To achieve this the total field data is required to be inverted, and the procedure is shown by the *red section* of Figure 8.1.

In Chapters 5 and 7, 3D inversion has been conducted on synthetic and real datasets using MAG3D developed by the University of British Columbia, Geophysical Inversion Facility (UBC-GIF). This code has been primarily applied to aero-magnetic surveys of a regional-scale, and some adaptations to the procedures are required to obtain robust results from archaeomagnetic data.

Chapter 6 shows the important steps that are required to prepare the dataset prior to inversion. Typically for inversions run using MAG3D, a mesh is designed featuring a central core consisting of small cells in the volume of interest, with larger cells surrounding this outside the limits of the observation points. This type of mesh design is suggested as it allows any regional trends within the dataset, deriving from magnetic bodies outside of the mesh volume and consisting of features of low wavenumbers, to be accounted for in these 'regional' cells. This allows the core of the mesh to model the anomalies of interest. This approach has not been very successful at

recovering archaeological-type anomalies, as the regional cells, poorly constrained by nearby data have been shown to account for significant trends across the survey area, even for datasets where no regional trend is present. This has a dramatic effect on the ability to accurately recover the archaeological features. It is therefore essential that any regional trends that exist in the dataset are removed prior to inversion, and that only the core central cells of the mesh are used in the inversion to ensure that all features in the dataset are accounted for in the mesh immediately below the observation points. It is shown in Chapter 6 that a range of bandpass filters can be applied to the data allowing a high-pass which removes the low-wavenumber trends from the data.

It is also preferable to remove any anomalies caused by features that are not of interest to the final interpretation. Archaeo-magnetic surveys are often conducted in areas where much more recent activity has been present on the site, and effects of this are shown in the datasets as well as anomalies relating to features of archaeological interest. It was shown in Chapter 5 that any mis-modelling of these near surface anomalies will have a significant knock-on effect in the modelling of the features of real interest. A range of directional filters were analysed in Chapter 6, and it was found that one of the most commonly applied directional filters, the cosine-taper filter, was unsuitable if subsequent quantitative interpretation be desired as it results in a large proportion of the measured data to be altered from their original amplitude. The directional pass/reject filter, or a newly proposed PATH filter were found to be much more appropriate for removing linear features in the chosen direction without significantly affecting other parts of the Fourier spectrum of the signal.

Once the dataset is prepared the inverse procedure will vary depending on the nature of the suspected anomalies. Archaeological bodies buried within susceptible soils are often represented by both positive and negative anomalies with respect to the background level of the magnetic field. It is therefore required that the inversion can recover a model that includes both positive and negative apparent susceptibilities if required. It is demonstrated in Chapter 5 that the inversion will preferentially model dipolar anomalies with both positive and negative susceptibility

contrasts, even when only a single source polarity is present (see Figure 5.29). Using opposite polarity susceptibilities to control the wavelength of the forward modelled signal introduces uncertainty over the depth of the causative body. If positive and negative susceptibility contrasts are thought to be present, it is necessary to run the inversion twice, once with a positivity constraint applied, and again with a negativity constraint applied. Using the output from the edge detection techniques, a reference model can then be created by selecting from the positivity and negativity-constrained models, as deemed appropriate based on the polarity of the observations, keeping in mind that negative anomalies may be present adjacent to positive ones, and vice-versa, due to the dipolar nature of magnetic fields. The inversion can then be re-run for a third time, this time with the susceptibility constraints relaxed to allow both positive and negative values, and with the reference model built into the objective function to ensure the model retains a structure as close as possible to the preliminary inversions.

Modelling each inversion run with a range of values for the trade-off parameter (β) between the model-norm and the data misfit is shown to allow the selection of an optimum value in the 'knee' of the L-curve, which ensures a maximum fit to the data without excess structure being introduced into the model by attempts to fit noise induced anomalies. Using this technique does not require a prior knowledge of the noise level of the dataset.

Comparing the results to similar depth models produced by other geophysical techniques allows the magnetic susceptibility to be compared to other physical properties of the subsurface. In Chapter 7, the magnetic model is compared to ground penetrating radar where the magnetic susceptibility shows broad agreement with the structure recovered from the dielectric properties, satisfying Objective 2, as laid out in the introduction.

Objective 3 from the introduction, proposed all the above should be achievable with minimal requirement for user input, to reduce the likelihood of the result representing a bias based on the operators preference. This has been achieved to a certain extent, however it has been impossible to remove all decision making from the operator. The derivative-based edge detection

methods require no operator decisions to transform the dataset, however the interpretation of the edge locations from resulting output is user dependant. Euler deconvolution requires several user decisions, most notably the choice of an appropriate structural index to apply to the calculation. While the inversion of positive or negative-susceptibility contrasts requires little user interaction apart from the choice of appropriate mesh parameters, an inversion including both polarities of susceptibility contrast, requires the user to choose which parts of the model should be represented by each polarity.

Despite this, the added level of interpretation of archaeo-magnetic data demonstrated here, allows the physical properties of the causative bodies of the subsurface to be analysed. The data can be used to develop interpretations based on the 3-dimensional position and material composition of any archaeological bodies. Subsequent decisions as to whether to build on top of the buried features, perform more detailed surveys or commence excavation, can be made with added confidence.

8.1 Proposals for further work

Further work on this project may focus on three areas; improving the technique laid out in this thesis, expanding the technique to work with other components of the magnetic field, and developing the technique to incorporate other geophysical properties to give a more complete model of the subsurface.

Improving the technique

Two, highly significant changes in procedure are required for inverting archaeo-magnetic datasets as opposed to aero-magnetic data. These are the requirement to limit the mesh to the lateral extents of the survey area, and to constrain positivity and negativity in individual runs of the inversion, prior to making a reference model. As discussed above, it is required to prevent the

inversion from modelling the negative anomalies generated by the dipolar response to a positive body, with negative susceptibilities. The problem is highlighted once again in Figure 8.2. The positive body on the left of the model in Figure 8.2c, has generated a large positive anomaly (coloured blue in Figure 8.2a), and two small negative anomalies either side (coloured red in Figure 8.2a).

The pseudogravity data (Figure 8.2b) does not exhibit this dipolar nature, and a positive anomaly is generated above the positive susceptibility contrast, and a negative anomaly above the negative susceptibility contrast. Should the pseudogravity dataset be inverted as a gravity dataset, possibly utilising UBC-GIF's GRAV3D, then the depths of the recovered anomalies are likely to be positioned towards their actual depth, and the likelihood is that each body will be modelled as a single polarity body. The density to magnetisation ratio applied to the pseudogravity transformation could then be used to transform the density model back into a magnetic susceptibility model, for use as a reference model in a final run using the total field dataset. This will be required to fit the high-wavenumber information of the total-field data that will be suppressed in the modelling of the pseudogravity.

Expanding the uses of the technique

To expand this technique so that it may become universally adopted by archaeo-geophysicists, it is required to adapt it to work with other types of magnetic datasets. The data presented here have been measured with magnetometers with relatively low electrical noise levels compared to previous generations of equipment (Linford *et al.*, 2007), and have been collected with a dense sample spacing (0.125×0.5 m). Further investigation should be focussed on the effects of processing and modelling differing quality datasets, particularly legacy datasets collected with older equipment and coarser survey grids, as well as the potential to collect higher sampled data with higher signal to noise ratios (e.g. with SQUID magnetometers).

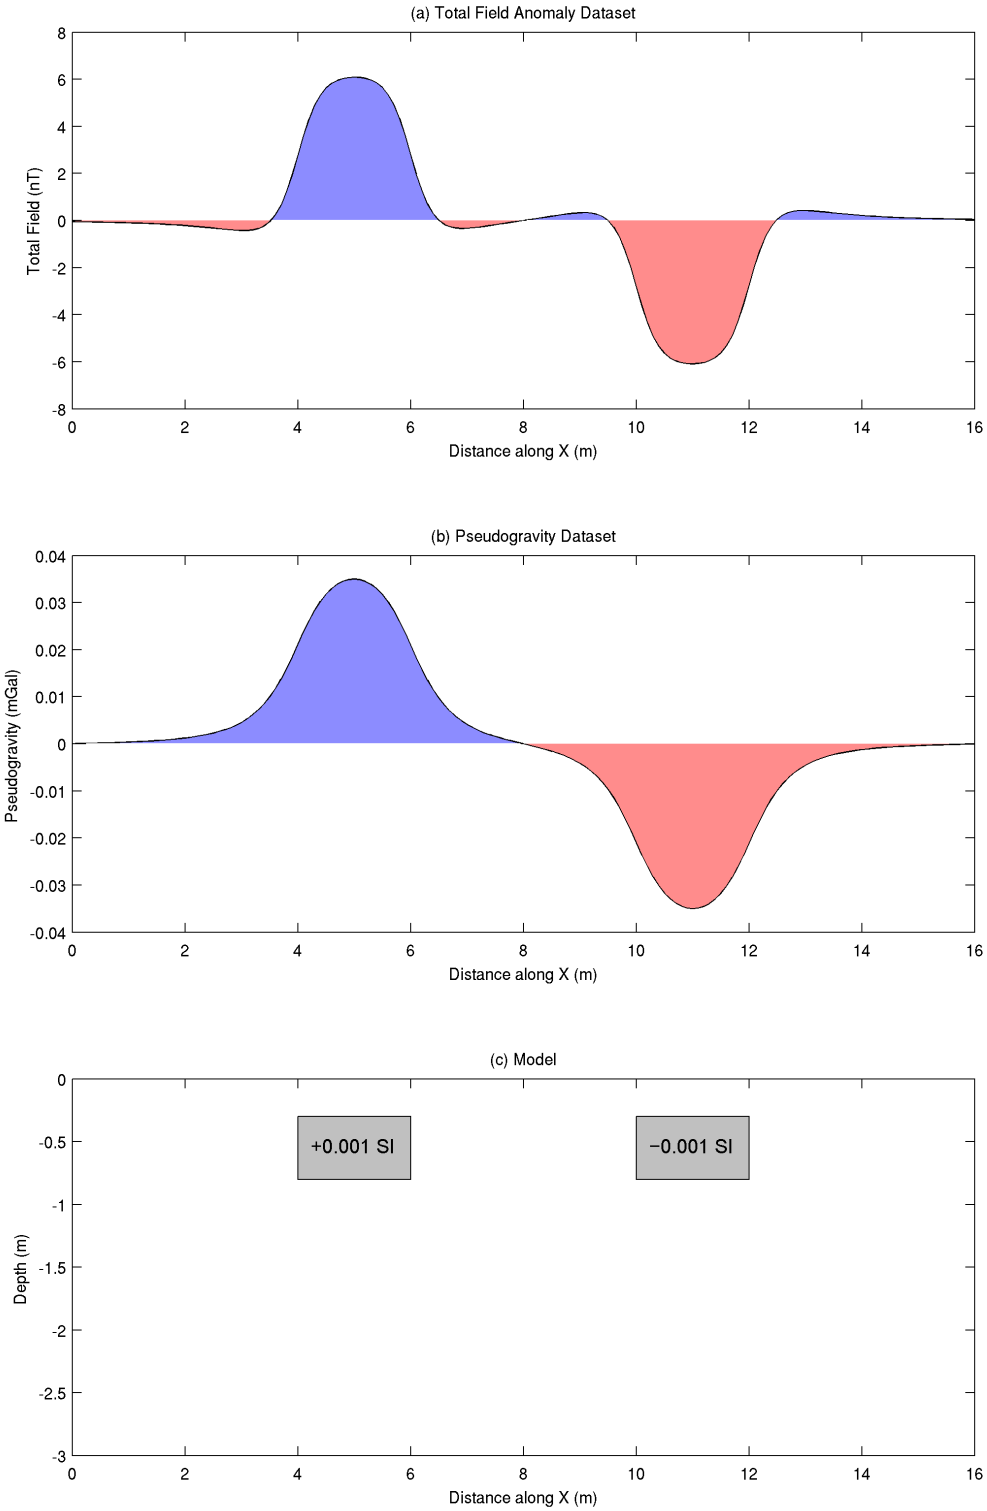

Figure 8.2: Advantages to working with pseudogravity over total field data. (a) Total field anomaly response, and (b) Pseudogravity dataset for (c) a model with both positive and negative susceptibility contracts. The polarity of the response has been coloured. The total field response has resulted in negative responses either side of the positive body, and vice versa for the negative body.

Of particular focus, should be datasets measuring the vertical component of the vertical gradient. This is the component of the magnetic field typically measured by fluxgate magnetometers. The advantages and disadvantages to fluxgate magnetometers over those that measure the total field was discussed in Chapter 2. Fluxgate magnetometers are widespread in their use in archaeological prospection (Tabbagh *et al.*, 1997), due to their relatively low price and simpler processing. Derivative-based methods require the total field to be measured, however it is likely that some useful information can be derived from fluxgate datasets, such as the horizontal gradient magnitude.

While the extra power in the higher wavenumbers has been demonstrated to be a problem with several techniques here, vertical gradient datasets have been inverted successfully using MAG3D by Piro *et al.* (2007). The authors noted that this should not have been the case. They explained the success of their result as being due to filtering steps that modified the dataset to represent a spectrum similar to one of a total field dataset. An improvement on this approach would be to modify the forward calculation code within the inversion process, to model the vertical component of the vertical gradient.

Developing the technique

To develop the technique beyond its use as a first-pass tool for initial site investigation, a future direction for research should be focused on combining the magnetic properties of the subsurface with other physical properties, particularly the results from GPR surveys to aid the characterisation of the subsurface. The dielectric properties are useful for the detection of changes in structure, and can be used to make reliable models of the subsurface, as seen in Figure 7.16. However, GPR interpretations do not reveal much information regarding the composition of the materials producing the reflections. Using GPR reflections to build the 3D boundaries of anomalous bodies in the subsurface would allow additional *a priori* constraints to be placed on the magnetic inversion, and therefore reduce the inverse problem. The final model would

reflect the confidence of deriving the structure of the subsurface using GPR, as well as allowing interpretation of the materials based on the magnetic properties.

Appendix A

Paper based on Chapter 3

This is a final draft copy of the paper published as:
Advantages to Using the Pseudogravity Transformation to Aid Edge Detection of Total Field Archaeomagnetic Datasets. *Archaeological Prospection* **18**, 81-93 (2011) doi: 10.1002/arp.408

Advantages to using the pseudogravity transformation to aid edge detection of total field archaeo-magnetic datasets

Cheyney, S.[1], Hill, I.[1], Linford, N.[2]

[1]Department of Geology, University of Leicester,
University Road, Leicester, LE1 7RH, UK (sc112@le.ac.uk),
[2]Geophysics Team, English Heritage,
Fort Cumberland, Eastney, Portsmouth, PO4 9LD, UK

July 2010

Abstract

Magnetic surveys are becoming increasingly common on archaeological sites due to the amount of data that can be collected rapidly in a non-invasive manner. Due to their relatively low cost compared to excavation they commonly provide the only dataset that covers an entire archaeological site, which can then be used to target other surveys and excavations to areas of interest. Typically interpretation is done visually using optimised images of the raw data, which when dealing with large datasets can be time-consuming and subjective.

Various derivative based methods have been developed recently to aid the interpretation of magnetic data. A particular use for these techniques is to locate the edges of subsurface magnetic bodies, and their use is gaining popularity in aeromagnetic regional and mineral exploration surveys. Despite this they are rarely used in archaeological survey interpretation. This is likely to be due to the particular challenges posed due to the low amplitude and high-wavenumber content of archaeo-magnetic surveys, where features of interest are often only slightly above the noise level.

Here, it is demonstrated that when derivative based methods are applied directly to total field data, the high-wavenumber components of the data are amplified, making datasets difficult to interpret and often proving less useful than the total field dataset alone. The pseudogravity transformation is a readily available tool for suppressing this bias to the high-wavenumber features and providing derivative based results with a power spectrum comparable to the original total field response, but with all the qualities to enhance interpretation that are obtained from using the derivative methods.

Introduction

Archaeological excavation is expensive and slow, and also damages the material excavated in an unrecoverable way. While it is essential for understanding buried historical structures, it must be applied very selectively. For these reasons, geophysical surveys have become increasingly used to guide the positioning of excavations, and magnetic surveying has now become one of the most commonly used geophysical techniques over archaeological sites, due to the ability to survey large areas rapidly whilst being able to detect subtle variations in the properties of the subsurface. For the majority of any site, these results will be the only evidence of the full spatial extent of the buried structures.

The magnetic properties of the near-surface are generally determined by the presence of small quantities (usually $<1\%$ of soil mass by weight) of iron minerals within soils and sediments (Linford, 2006), and the proportions between different forms of these minerals (Linington, 1972). Other magnetic sources of interest include ferrous metal objects which have become buried, semi-industrial and domestic processes where high temperatures have altered the iron minerals and surrounding soil, concentration of organic matter where timber structures have decayed and the disturbance of the natural magneto-stratigraphy such as that caused by the construction of a ditch or bank (see Linford, 2006).

Recent advances in magnetic surveying, including the use of several caesium (Cs) vapour magnetometer sensors mounted onto a platform, and the use of GPS for positioning, have seen a dramatic increase in the amount of data that can be collected quickly, and have thus led to higher resolution surveys becoming common, see Hill et al. (2004), Linford et al. (2007) and Dabas et al. (2010) for examples. Post-survey processing and modelling of the data is however, still usually limited to mitigating un-wanted artifacts by correcting data for offsets created by directional sensitivity of the instrument, changes in temperature and recalibration, as well as removal of data not thought to be related to archaeological sources, e.g. spikes in the dataset caused by surface metallic objects or instrument instability. Typically data are then presented and interpreted by way of a trace plot, contour plot, dot density plot or greyscale plot (English Heritage, 2008).

Several techniques are becoming increasingly popular for quickly deriving the surface locations of the edges of subsurface bodies. These can rapidly help identify the locations and shapes of anomaly sources of interest from large spatial surveys. Recent examples of use of the horizontal gradient magnitude can be seen by Sheriff et al. (2010), locating foundations over a historic town in the USA, whilst theta map and tilt angle methods have been applied by Büyüksaraç et al. (2008) to synthetic and real data from a Bronze Age burial ground in Turkey. While implementation is becoming increasingly common in mineral exploration surveys, they remain rarely implemented in archaeo-magnetic surveys. Although the data are collected and processed in fundamentally the same way in mineral and archaeological surveys, the spatial size, contrast in magnetisation, and signal-to-noise ratio of archaeological surveys can provide problems when implementing these techniques. The use of the analytic signal, used in calculating the Theta Map, has become more typical with Tabbagh (1997) and Godio & Piro (2005) providing good examples, however the

problem with increasing noise in derivative results, particularly second-order derivatives is demonstrated by Milea et al. (2010).

Synthetic data

For this study the synthetic response to a compound anomalous body was generated. The anomaly shown in Figure 1a consists of two connected parts, in order to test the sensitivity to features of differing depths and magnetic response. The top of the shallower part of the body is positioned 0.3 m below ground level, while the deeper part is 0.8 m below ground level. The base of both parts is situated at 1.2 m below ground level. The anomaly has a susceptibility of 0.001 SI, and is situated within a magnetic field of 37.5 Am^{-1}, causing a magnetisation intensity of 0.0375 Am^{-1}. The total field response has been calculated at a height of 0.2 m above ground level, at a sampling density of 0.5 x 0.125 m using the method described by Rao and Babu (1993). To replicate a dataset reduced to the pole, the response has been generated with a magnetisation inclination of 90°. This represents a well-sampled, noise free, dataset which is typical of sampling resolutions used in archaeomagnetic survey. The dataset has then been interpolated to an even 0.125 x 0.125 m grid, and can be seen in Figure 1b.

To test the effects of noise on this study, a zero-mean symmetric Gaussian distribution was generated with a standard deviations of 0.1, 0.2, 0.3, 0.5 and 1.0 nT. All scenarios show similar effects to varying degrees, and the 0.2 nT scenario is presented here to demonstrate a moderate noise case. A histogram of the distribution of noise is presented in Figure 1c. The synthetic data with noise added is shown in Figure 1d.

Methods for identifying edges of causative bodies from Total Field magnetic data

The methods evaluated here are all based on one, or multiple gradients of the original total field dataset. The application of the horizontal gradient magnitude, theta map and the tilt angle are investigated, with archaeological purposes in mind.

Horizontal Gradient Magnitude

The most commonly used tool to aid edge detection in potential field datasets is the horizontal gradient magnitude. Due to asymmetry in the magnetic response caused by the inclination of the magnetising field, to apply the horizontal gradient magnitude to a total field dataset, the data must be reduced to the pole (RTP). This performs a phase shift to make the data appear as if they were collected at the north magnetic pole. This

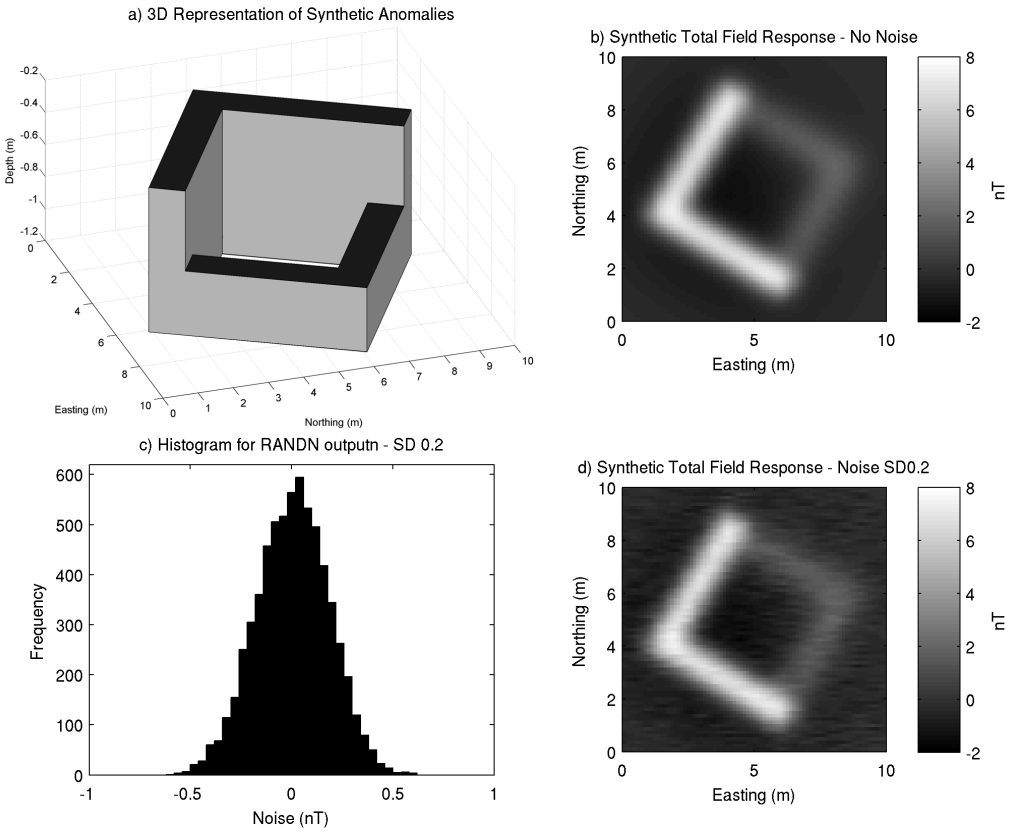

Figure 1: a) 3D model of synthetic anomaly with magnetisation of 0.0375 Am^{-1}. b) Total field data for no noise scenario. c) 0.2 nT standard deviation, zero-mean Gaussian noise. d) Total field data with noise added

has the effect of shifting anomalies laterally to be centred over their sources, and simplifies interpretation.

A procedure for outlining the edges of subsurface magnetic features using the horizontal gradient magnitude is described by Dole & Jordan (1978) and better defined by Cordell & Grauch (1985). Peaks in the horizontal gradient magnitude will overlie abrupt lateral changes in subsurface magnetisation, and will be directly above the edge of a magnetised body if the edge is a vertical contact and the data has been reduced to the pole. Where the body does not have a vertical contact, the horizontal gradient magnitude can be offset by an amount usually dependent on the depth below ground (Grauch & Cordell, 1987). In archaeological context, these offsets are therefore generally quite small. An automated approach to identifying the peaks in the horizontal gradient magnitude is demonstrated by Blakely & Simpson (1986), however this approach is often limited by the noise content found in archaeo-magnetic data.

The horizontal derivatives can be calculated in either the space or wavenumber domain. For ease of calculation, here they have been calculated in the space domain. The gradient in the x-direction is calculated using:

$$\frac{\partial T}{\partial x} = \frac{T_{(x+1,y)} - T_{(x-1,y)}}{2\Delta x} \tag{1}$$

and the y-direction using:

$$\frac{\partial T}{\partial y} = \frac{T_{(x,y+1)} - T_{(x,y-1)}}{2\Delta y} \tag{2}$$

The horizontal gradient magnitude is then calculated as:

$$\frac{\partial T}{\partial h} = \sqrt{\left(\frac{\partial T}{\partial x}\right)^2 + \left(\frac{\partial T}{\partial y}\right)^2} \tag{3}$$

The results for the synthetic case can be seen in Figure 2a. This method has successfully highlighted the vertical contacts of the anomalous bodies, which are situated beneath the maxima values of the horizontal gradient magnitude.

It can be observed in Figure 2a that the response to the lower part of the synthetic body has produced a weaker anomaly in the horizontal gradient magnitude (∼3-4 nT/m) compared with the shallower body (∼10 nT/m). This is as expected, due to a weaker response in the total field measurements and the longer-wavelength response that is typical of more deeply buried anomalies. This will make large amplitude, short wavelength anomalies easier to identify using the horizontal gradient magnitude method, which will typically show bias towards identification of shallower bodies.

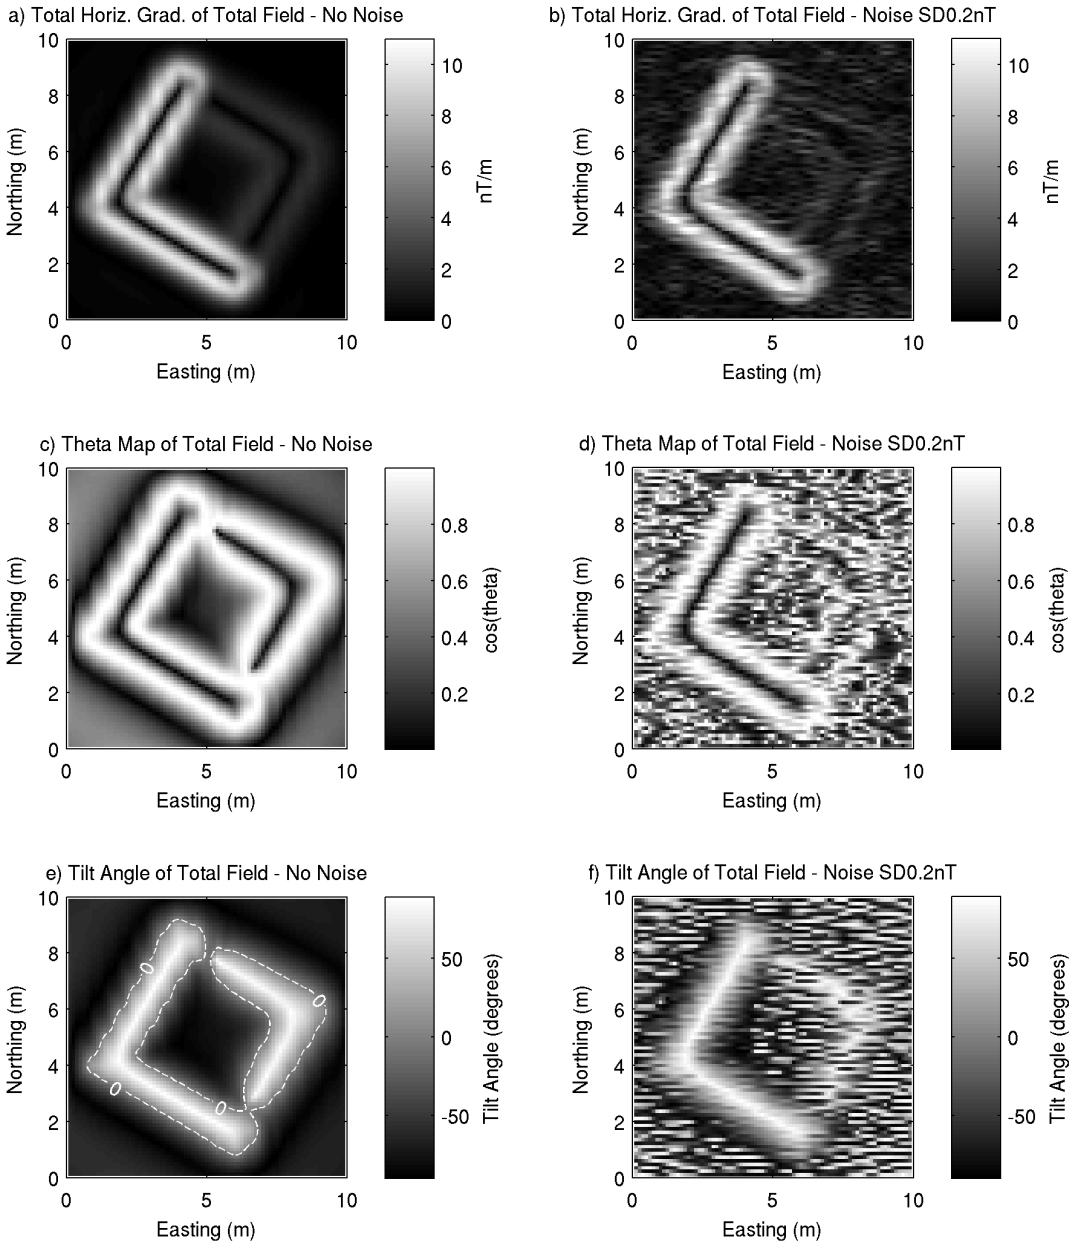

Figure 2: Edge detection methods on the total field anomaly data using synthetic data generated with no noise, and noise of 0.2 nT standard deviation. a & b) Horizontal gradient magnitude, c & d) Theta Map, e & f) Tilt Angle. Plots on the left represent the no noise scenario, plots on the right show the 0.2 nT standard deviation noise scenario.

Figure 2b shows the horizontal gradient magnitude of the synthetic model generated with zero-mean Gaussian noise with a standard deviation of 0.2 nT. Although the amplitude of the shallower anomaly means that the edges are still easily identified as maxima in the horizontal gradient magnitude, the weaker response to the deeper anomaly is hard to identify from the surrounding background noise.

Theta Map

The use of theta map for detecting edges in magnetic datasets is described by Wijns et al. (2005). Theta map combines the horizontal gradient magnitude and the analytic signal. It is defined as:

$$cos\,(\theta) = \frac{\sqrt{\left(\frac{\partial T}{\partial x}\right)^2 + \left(\frac{\partial T}{\partial y}\right)^2}}{|\mathbf{A}|} \tag{4}$$

where $|\mathbf{A}|$ is the analytic signal amplitude. This is calculated as a combination of the horizontal and vertical gradients of the magnetic field. In this respect it is often referred to as the total gradient (See Equation 5).

$$|\mathbf{A}| = \sqrt{\left(\frac{\partial T}{\partial x}\right)^2 + \left(\frac{\partial T}{\partial y}\right)^2 + \left(\frac{\partial T}{\partial z}\right)^2} \tag{5}$$

The analytic signal can also be calculated using higher numbered derivatives, where it is often referred to as the analytic signal of order (n).

The analytic signal method can be used as a edge detection method on its own. It was presented by Nabighian (1972) for 2D profile data, and Roest et al. (1992) for 3D gridded datasets. The amplitude of the analytic signal amplitude has maxima near the edges of magnetic sources, however is not considered to be as accurate as the horizontal gradient magnitude of RTP data (Li, 2006). Despite this, the analytic signal amplitude transformation can be useful in areas where reduction to the pole becomes unstable, such as at low latitudes, or where significant remanent magnetisation is thought to be present.

Theta map takes advantage of two properties of the magnetic derivatives, that the edges of a vertical contact will be situated not only below a peak in the horizontal gradient magnitude, but also at a zero crossing point in the first vertical derivative. Therefore, over an edge the amplitude of the analytic signal will be largely dependent on the amplitude of the total horizontal derivative, and the value of $\cos(\theta)$ will approach 1. Either side of this contact, the vertical derivative will be the significant component of the analytic signal, and the value of $\cos(\theta)$ will approach 0. Thus a contact is defined by a maximum value of $\cos(\theta)$ bracketed by two minimum values. Magnetically quiet areas will have a high $\cos(\theta)$ value.

The advantage of the theta map method over the horizontal gradient magnitude is that it equalises the amplitude of the anomalies across the survey grid, so maxima that appear as different amplitudes in the horizontal gradient magnitude have approximately the same amplitude in the theta map. In this respect, it is often compared to an automatic gain control (AGC) filter. The advantage of this is that subtle features in the horizontal gradient magnitude are amplified, however a disadvantage is that noise can also be amplified.

The theta map of the synthetic dataset is shown in Figure 2c and Figure 2d, for the zero-noise, and Gaussian noise examples respectively. Figure 2c shows the advantage of plotting the theta map, as the edges of both the shallow and deep anomalies are both detected by areas where the Theta Map is approaching 1. Over the centres of the anomalies the value of the theta map is close to zero. The theta map shows a termination of the deeper feature prior to the connection with the shallower feature demonstrating an interference effect where the closeness of the causative bodies cause the total field anomalies to overlap.

Figure 2d however, shows the disadvantage of using the theta map. Although the introduction of noise into the data causes some relatively small isolated anomalies in the total field data, when the theta map is calculated the noise component can be boosted to values approaching 1, the same amplitude as the features of interest. This causes Figure 2d to be difficult to interpret, and although the high-low-high pattern across the shallower body is still evident despite the noise, the lower body is now impossible to identify.

Tilt Angle

Methods using the tilt angle, and derivatives of the tilt angle have been developed to quickly identify both the horizontal location and the depth of buried anomalous bodies (see Verduzco et al., 2004 and Salem et al., 2007). The tilt angle involves the use of the horizontal and vertical gradients of total field magnetic data to produce a normalized derivative based on the ratio of the gradients. It is defined by Miller & Singh (1994) as:

$$\theta = \tan^{-1}\left[\frac{\left(\frac{\partial T}{\partial z}\right)}{\left(\frac{\partial T}{\partial h}\right)}\right] \tag{6}$$

Where θ is the tilt angle, and T is the magnitude of the total magnetic intensity of the anomaly. The first derivative of T in the horizontal direction is calculated as in Equation 3

Further work by Salem et al. (2007) showed that assuming the source structures have vertical contacts and the magnetic field is either vertical or reduced to the pole, the Tilt Angle can also be written as:

$$\theta = \tan^{-1}\left[\frac{h}{z_c}\right] \tag{7}$$

Where h is the horizontal location and z_c is the depth to the contact. Salem et al. (2007) suggest that the above equation shows that the tilt angle can be used to derive the edges of the contact, as where $h=0$, $\theta=0$, and therefore the edges of vertical contacts will be present directly below the 0° contour of the tilt angle.

An alternative way to interpret the results from the tilt angle is to calculate the horizontal gradient magnitude of the tilt angle, given by Verduzco (2004) as:

$$\frac{\partial(\theta)}{\partial h} = \sqrt{\left(\frac{\partial(\theta)}{\partial x}\right)^2 + \left(\frac{\partial(\theta)}{\partial y}\right)^2} \tag{8}$$

This is similar to the local wavenumber used in methods by Thurston & Smith (1997). Subsurface contact features should be represented beneath maxima in the horizontal gradient magnitude of the tilt angle, in a way similar to the horizontal gradient magnitude of the total field, the advantage here being that the amplitude of the anomalies have been normalised prior to the gradient being analysed.

Results from the tilt angle (Equation 6) for the synthetic scenario without noise, and with Gaussian noise added can be seen in Figures 2e and 2f respectively. In Figure 2e the total field data has been normalised to values between -90° and 90°, with the shallow and deep anomalies successfully normalised to the same amplitudes. The zero contour shows a reasonable outline of the subsurface location of the anomalous body. Figure 2f shows the tilt angle is highly susceptible to the influences of noise, and just as seen with the theta map, the noise is amplified to the same values as the signal. The continuity of the tilt angle response to the shallower body allows that to be identified, however identification of the deeper body is difficult, and tracing of the zero contour through the dataset with noise present is made impossible due to the short wavelength features of the noise causing multiple zero crossings.

The horizontal gradient magnitude of the tilt angle will be even more susceptible to noise due to being a derivative of a derivative (Cooper & Cowen, 2006). If the tilt angle dataset is not interpretable due to noise, the horizontal derivative of the tilt angle will also be uninterpretable.

Power spectra analysis of derivative-based methods

It can be seen that although the methods discussed above are successful at achieving their aims at making the no-noise synthetic data easier to interpret than the "measured" total field dataset, the introduction of noise causes a breakdown in the use of gradient-based methods, to the point where the total field data is easier to interpret by eye than the derivatives or combinations of derivatives. This is due to the accentuation of the high-wavenumber component of the data when derivatives are calculated.

Figure 3 shows the radially averaged power spectra for the synthetic dataset with the 0.2 nT standard deviation, zero-mean Gaussian noise added. The power spectrum for the total field anomaly data is represented by the solid line. Once a derivative of the total field data is calculated, high wavenumber components are accentuated relative to the low wavenumbers, as can be observed by the dashed lines in Figure 3. It is this effect that causes the noise to reach amplitudes which mask the signal of real anomalies as observed in Figure 2, and makes gradient-based methods difficult to apply to archaeo-magnetic datasets. One way to suppress the high wavenumber component of a dataset is to apply a wavenumber filter, such as the pseudogravity transformation.

Pseudogravity Transformation

Assuming the boundaries of magnetic and gravitational sources are the same, and the intensity of magnetisation and density are proportional throughout, then magnetic potential and gravitational potential are linked by Poisson's relation. This relation is shown in Robinson (1971) as:

$$A_{(x,y,z)} = \frac{I}{G\rho}\frac{\partial U_{(x,y,z)}}{\partial \alpha} \tag{9}$$

where $A_{(x,y,z)}$ is the magnetic field potential, $U_{(x,y,z)}$ is the gravitational field potential, I is uniform magnetisation intensity in direction α, and ρ is the uniform density.

Baranov (1957) made use of Poisson's relation and introduced a linear filter which is typically referred to as the pseudogravity transformation. The output from the pseudogravity transformation is a uni-polar field, and has similar properties to a gravity anomaly that would be observed should the magnetisation be accompanied with an exactly proportional density distribution. The result can simplify the interpretation of shape and location of source bodies. The transformation utility is a commonly found feature in potential field software packages, such as MagPick (Tchernychev, 2009) and Geosoft Oasis Montaj (Geosoft, 2010). The pseudogravity datasets for the synthetic case are shown in Figure 4a for the scenario with no noise, and Figure 4b with Gaussian noise added. A ratio of density to magnetisation of 1x10^2 kg m^{-3} per Am^{-1} has been used, however changing this will only act as a linear scaling factor on the output values, and therefore does not affect the wavenumber spectrum of the pseudogravity dataset.

Calculating the pseudogravity transformation of the total-field dataset has the effect of suppressing the high-wavenumber component of the data, as can be observed by comparing the dotted and solid lines on Figure 3. At progressively higher wavenumbers the power becomes progressively suppressed. This causes the pseudogravity transformation seen in Figure 4 to appear as a much smoother image than the total field data shown in Figure 1.

When the derivatives of the pseudogravity dataset are calculated, shown by the dot-dash lines in Figure 3, the resulting spectrum is much closer to the original spectrum of

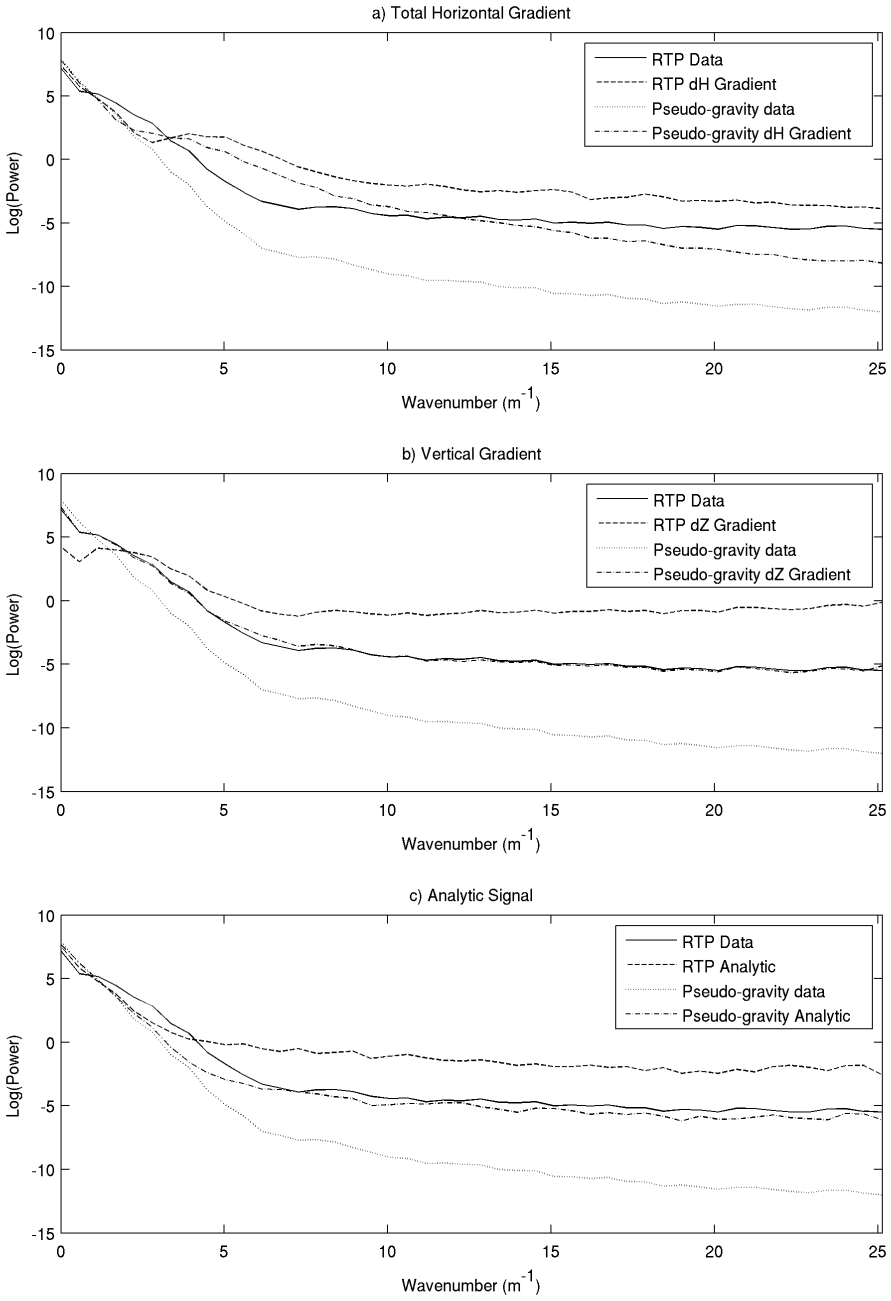

Figure 3: Radially averaged power spectra for total field magnetic data (solid line), generated with 0.2 nT Gaussian noise and reduced to the pole, and pseudogravity transformation (dotted line) are shown in all plots. For comparison, the power spectra for a) horizontal gradient magnitude, b) vertical gradient and c) analytic signal, are shown for both RTP and pseudogravity datasets.

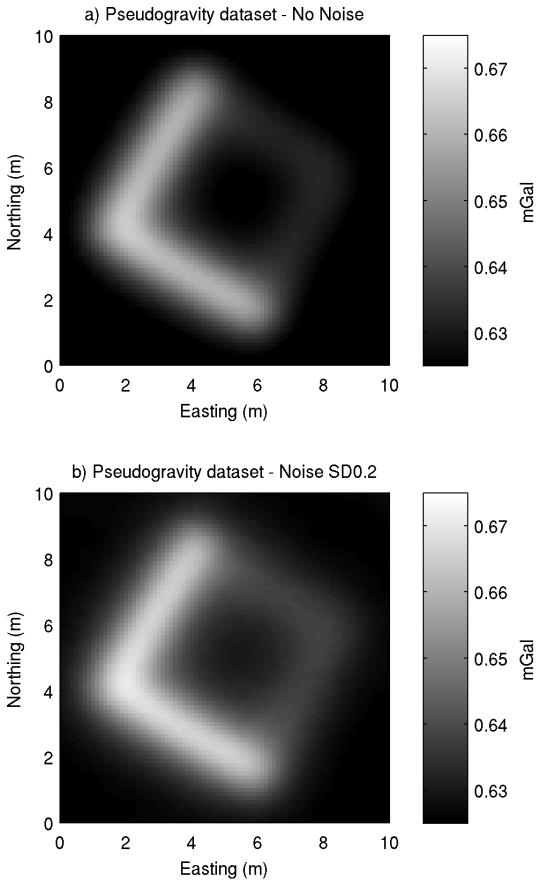

Figure 4: a) pseudogravity transformation of the synthetic data with no noise added. b) pseudogravity transformation from the total field data with 0.2nT standard deviation, zero-mean Gaussian noise added prior to transformation.

the total-field data (solid line), therefore the bias towards high-wavenumber features is removed, and the derivatives are comparable to the original total field data in terms of noise. As the pseudogravity is the vertical integral of reduction to the pole, the power spectrum for the vertical derivative of the pseudogravity dataset is equivalent to that of the RTP dataset, as observed in Figure 3b.

Edge detection using pseudogravity transformation

Derivatives calculated from the pseudogravity dataset have slightly altered properties from those calculated from the total field data (see Phillips et al., 2007), however they largely have the same desirable properties, and therefore can be used in gradient-based methods to aid edge detection without the problematic amplification of the noise identified earlier.

Figure 5a shows the horizontal gradient magnitude generated using the pseudogravity dataset. The resulting image shows similar features as seen in Figure 2a, with the shallower part of the body causing an anomaly of $\sim$0.035 mGal/m, and the deeper part of the body generating a lower amplitude response of $\sim$0.015 mGal/m. The higher wavenumber component present in the total field data causes Figure 2a to appear 'slightly sharper' than the result in Figure 5a, however largely the same interpretation can be made from the two datasets.

The big advantage to using the pseudogravity transformation can be seen in Figure 5b, where noise has been added to the original data. Whereas, in Figure 2b, the noise meant the identification of the deeper anomaly was almost impossible, the derivative generated from the pseudogravity dataset is not affected to the same degree, and appears very similar to the zero noise scenario (Figure 5a). The response of the weaker, deeper body is still clearly visible above the noise threshold.

The results, without and with noise respectively, can be seen in Figures 5c and Figures 5d for the theta map, and Figures 5e and Figures 5f for the tilt angle.

When the scenarios without noise are compared in Figure 2 and Figure 5, the results generated from the total field dataset are slightly preferable as the results derived from the pseudogravity dataset have a slightly broader signature making the crest of anomalies slightly harder to identify. However, once noise is added, the results generated from the pseudogravity dataset are much more robust than those generated from the total field data, and show all the main features of the noise-free scenarios.

Case Study

Brading Roman Villa is situated between the towns of Brading and Sandown on the Isle of Wight, UK. The site consists of a Roman villa, now incorporated into a visitor centre, and surrounding pasture land. The data used in this study is a subset of a much larger caesium

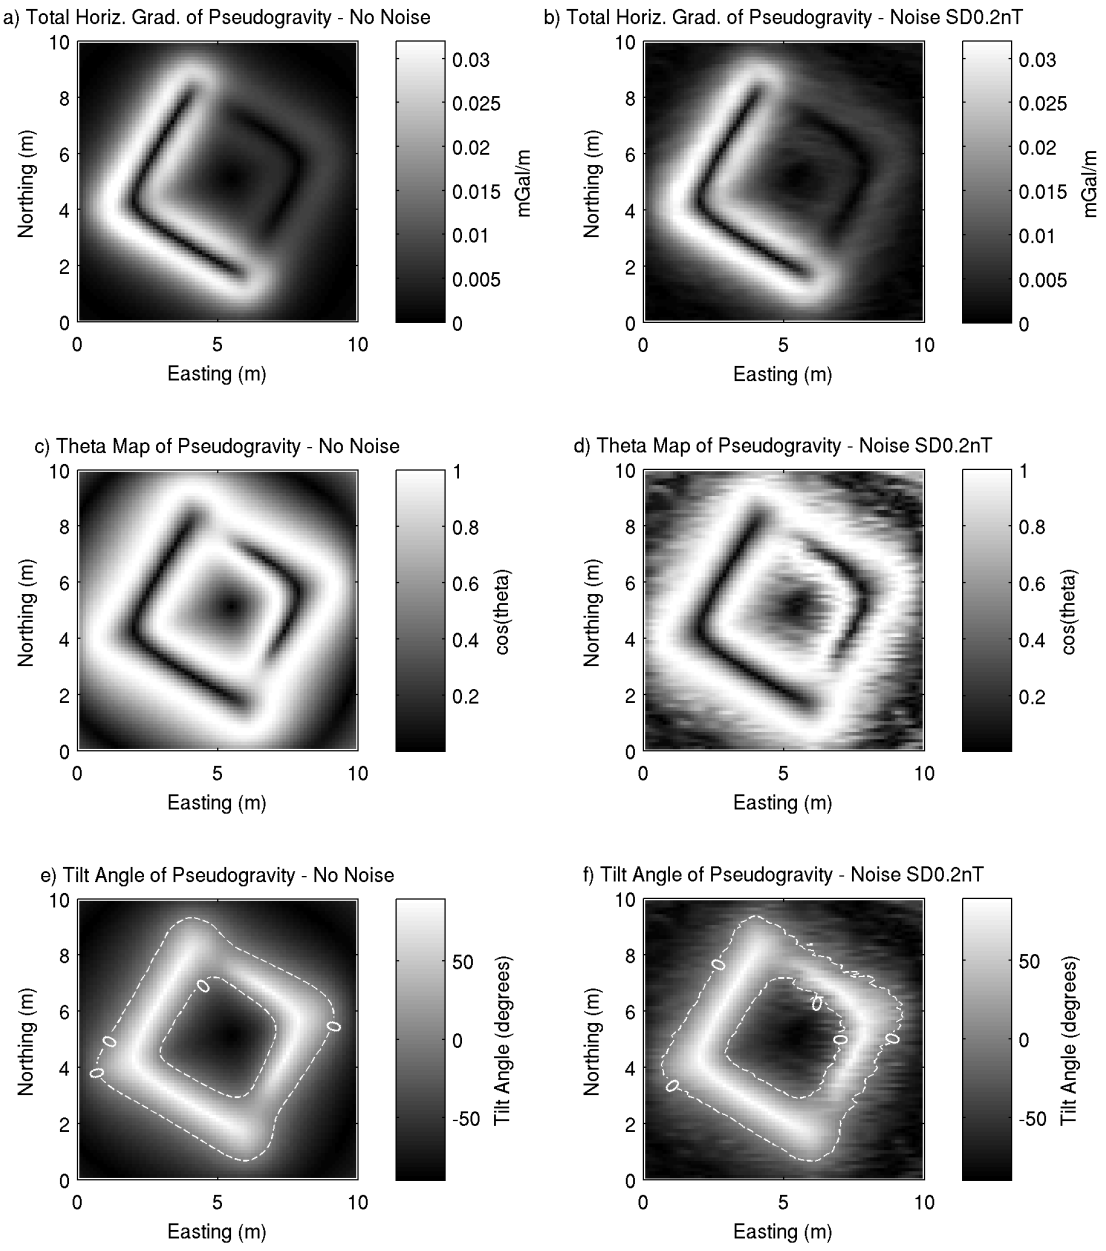

Figure 5: Edge detection methods on the pseudogravity transform data using synthetic data generated with no noise, and noise of 0.2 nT standard deviation. a & b) Horizontal gradient magnitude, c & d) Theta Map, e & f) Tilt Angle. Plots on the left represent the no noise scenario, plots on the right show the 0.2 nT standard deviation noise scenario.

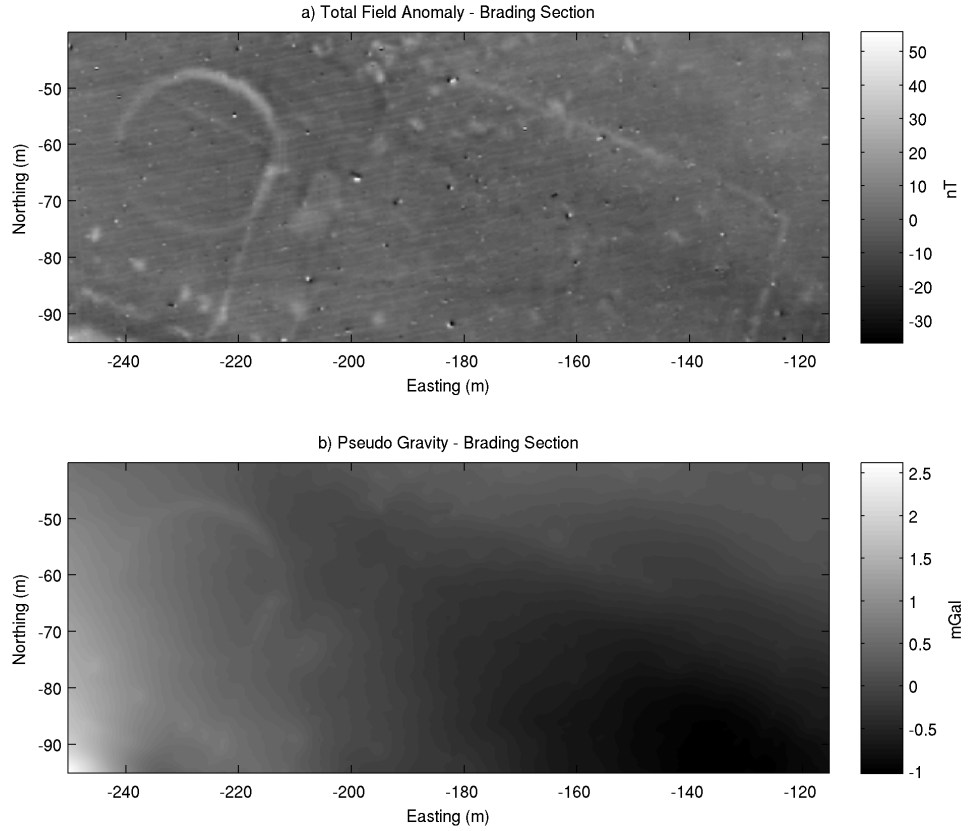

Figure 6: a) Subset of total field data from area surrounding Brading Roman Villa, Isle of Wight, UK. b) pseudogravity transformation of (a).

magnetometery survey collected in 2009 by the English Heritage Geophysics Team over the surrounding pasture land (Payne, 2009). The total field data were collected with a spatial resolution of 0.5 x 0.125 m, and have once again been interpolated to a 0.125 by 0.125 m grid.

The subset, shown in Figure 6a, is a 135 x 55 m area situated ∼100 m NW of the villa. Figure 6a shows a series of positive anomalies. The most prominent features in the data are a large circular anomaly ∼25 m in diameter towards the western side of the dataset, and a series of positive linear anomalies across the area. A ploughing pattern is prominent across the dataset, as well as a series of small positive peaks scattered across the area.

The pseudogravity transformation of the data can be seen in Figure 6b. Interpretation of Figure 6b is much harder, since the suppression of high-wavenumber components leaves only a prominent general regional trend across the area from SW to NE. The northern part of the large circular anomaly is also visible, but much harder to interpret than in the total field data.

Figure 7 shows the horizontal gradient magnitude of Figure 6. The horizontal gradient magnitude calculated from the total field data (Figure 7a) shows approximately half of

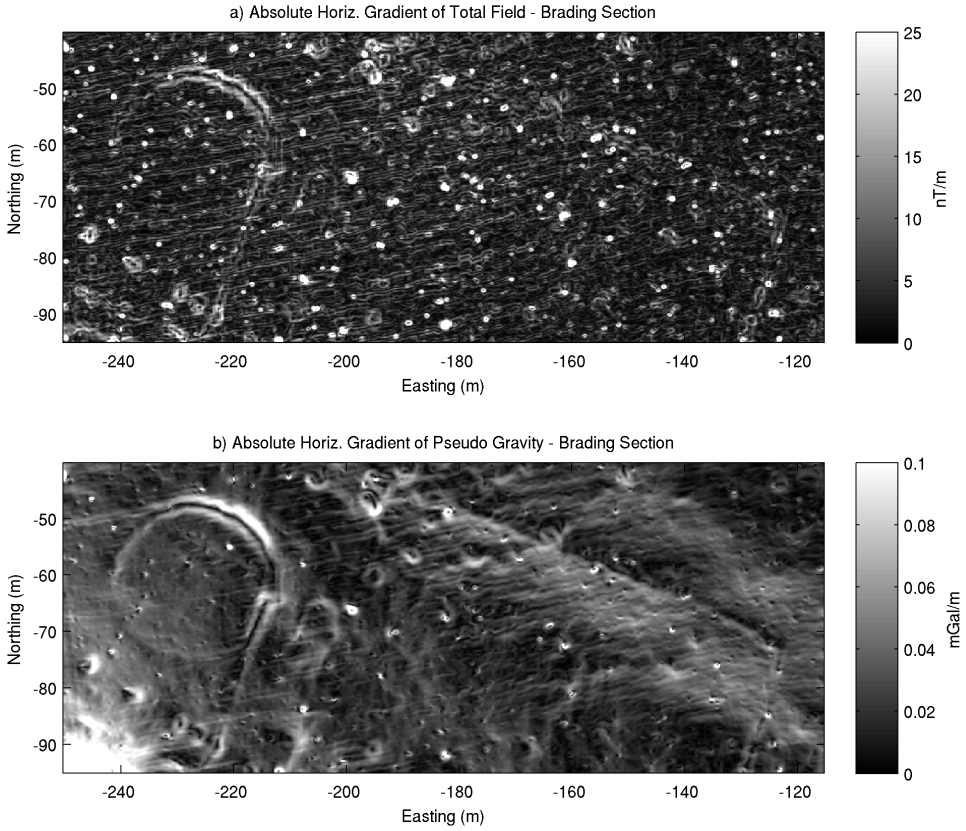

Figure 7: a) Horizontal gradient magnitude of total field data from Brading. b) Horizontal gradient magnitude of pseudogravity transformation

the circular anomaly strongly, and a possible linear continuation of this to the south. More linear anomalies can be identified towards the eastern part of the study area, and the edges of several small isolated anomalies can easily be identified. Generally though, features are harder to identify in the horizontal gradient magnitude data compared to the total field data, and many anomalies appear at least partly masked by the ∼E-W ploughing pattern.

Below this, in Figure 7b, the horizontal gradient magnitude calculated from the pseudogravity dataset can be seen. The most notable feature seen in this image is the suppression of the ploughing pattern observed in Figure 7a. This means that many more of the observable features in the raw data can be seen in the horizontal gradient magnitude data, including at least one edge of the entire circular anomaly is now identifiable, two linear tangents to the circle, a much longer and continuous anomaly running from the northern part to the eastern part of the survey area, and a more obvious segregation of the isolated anomalies from the plough pattern.

The results of the theta map are shown in Figure 8a for the total field data, and Figure 8b for the pseudogravity dataset. Once again, the improved interpretability of the results

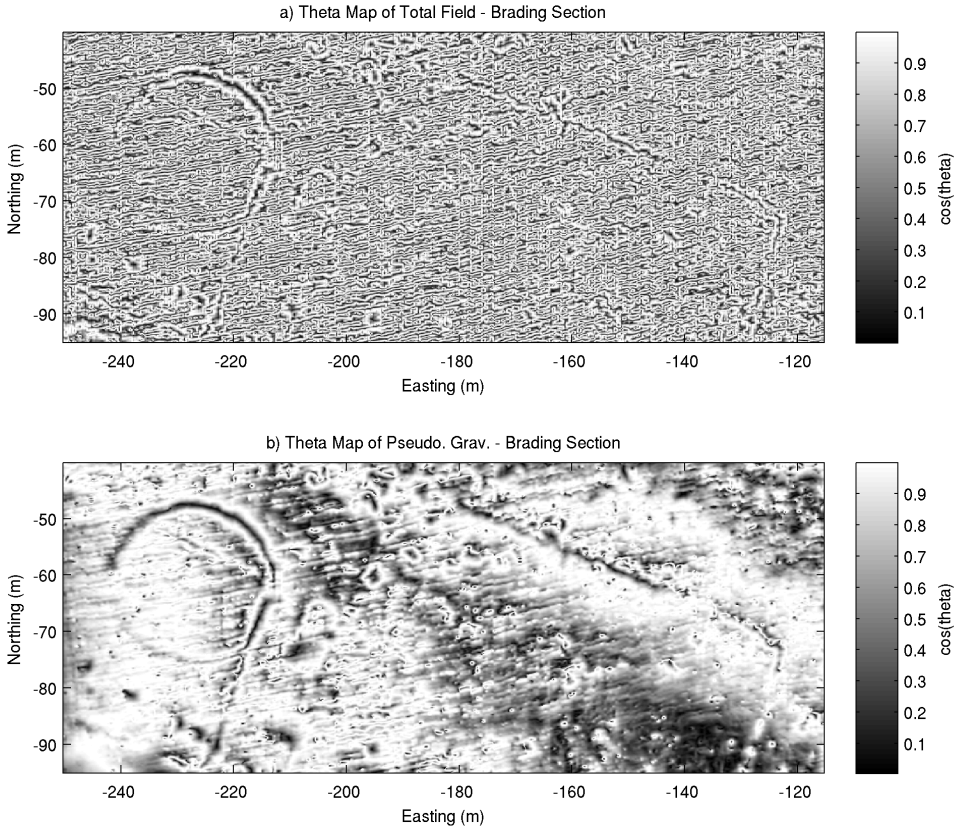

Figure 8: a) Theta Map of total field data from Brading. b) Theta Map of pseudogravity transformation

from the pseudogravity dataset over the total field is clearly observable, with Figure 8b showing far more of the anomalies observable in Figure 6a than Figure 8a, which is largely uninterpretable. The advantage of the theta map over the horizontal gradient magnitude, is that weaker anomalies are enhanced to the same amplitude as the largest anomalies, and this can be seen in the circular anomaly, where now both edges can be observed for at least part of the southern section.

Figure 9 shows the results of the tilt angle for the Brading data. The tilt angle calculated from the total field data (Figure 9a) shows the weak ploughing pattern accentuated at the expense of the other anomalies in the data, and therefore the data are hard to interpret. Figure 9b shows a result similar to the total field response but again with the advantage that the weaker linear anomalies have been boosted to the same amplitude as the circular anomaly. The ploughing pattern has caused significantly less disturbance to the features of interest than in Figure 9a.

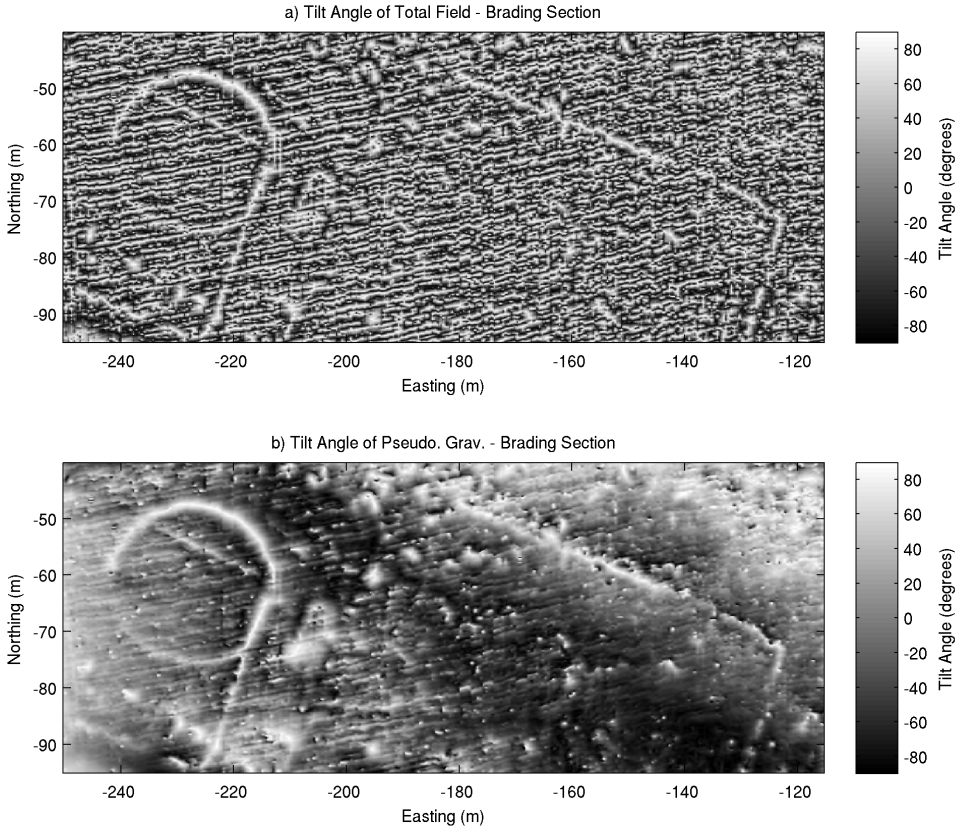

Figure 9: a) Tilt derivative of Total Field data from Brading. b) Tilt derivative of pseudogravity transformation

Discussion

It has been shown here that derivative-based methods have been successful at highlighting the edges of anomalous bodies and at normalising the amplitude of anomalies to improve the identification of weaker responses. With the synthetic anomaly designed to test these methods, the optimum results have been achieved from the derivatives taken from the total field measurements of the zero-noise scenario, as these results have shown sharp peaks in the vicinity of the edge locations. However, the introduction of small amounts of Gaussian noise has severely affected the users ability to interpret the responses, to the point where anomalies still identifiable in the total field data are uninterpretable in the derivative datasets.

The pseudogravity transformation is already a commonly available, and well used way to aid the interpretation of magnetic data. The particular feature of the pseudogravity transformation which makes it useful here is that the vertical gradient of the pseudogravity transformation shares the same distribution across its power spectrum as the measured total field response. This is because the pseudogravity transformation increasingly suppresses power in higher wavenumbers, which is the opposite effect of calculating derivatives, which boosts the power in the higher wavenumbers. When applied to the synthetic noise-added data, it is this effect of boosting the higher wavenumbers that causes interpretation to be difficult, an effect that is mirrored when applied to real archaeological data.

Conducting the derivative edge detection methods on the pseudogravity dataset removes this bias to high wavenumber features, and results in a dataset that has a comparable power spectra to the original data, yet has all the advantages of the various derivative-based methods at simplifying the interpretation. The pseudogravity transformation is being used here as a spectral shaping filter and the relative balance between low and high wavenumber features is restored once the derivatives are taken.

It is possible that a more bespoke filter can be designed to remove the effect of wavenumbers that do not relate to features of interest. This approach is likely to still further improve the clarity of the results of the derivative methods, however requires a prior level of interpretation, in that a wavenumber band of interest would already have been identified. The advantage of the using the pseudogravity transformation is that no prior interpretation is required, and no wavenumber features are compromised.

Acknowledgments

This work is funded by NERC doctoral training grant NE/G523798/1, and is CASE sponsored from Geomatrix Earth Science Ltd. Many thanks go to the English Heritage Geophysics Team, who allowed me to assist with their fieldwork, and subsequently to use the data in this study.

The advice and discussion by Stewart Fishwick, Chris Leech and two anonymous reviewers, resulted in large improvements to drafts of this paper, is also much appreciated.

References

Baranov V. 1957. A new method for interpretation of aeromagnetic maps; pseudo-gravimetric anomalies. *Geophysics* **22**(2): 359-382.

Blakely RJ, Simpson RW. 1986. Approximating Edges of Source Bodies from Magnetic Or Gravity-Anomalies. *Geophysics* **51**(7): 1494-1498.

Büyüksaraç A, Arisoy MÖ, Bektaş Ö, Koçak Ö, Çay T. 2008. Determination of Grave Locations in Dedemezari Necropolis (Western Turkey) using Magnetic Field Derivatives. *Archaeological Prospection* **15**(4): 267-283. 10.1002/arp.338

Cooper GRJ, Cowan DR. 2006. Enhancing potential field data using filters based on the local phase. *Computers & Geosciences* **32**(10): 1585-1591.

Cordell L, Grauch VJS. 1985. Mapping basement magnetization zones from aeromagnetic data in the San Juan Basin, New Mexico *in* Hinze, William J. (ed.), *The Utility of Regional Gravity and Magnetic Anomaly Maps: Society of Exploration Geophysicists*, Tulsa, OK. 181-197.

Dabas M. 2010. First Results of AMP surveys for large areas (100ha) with decimeter resolutions in archaeology. *Geophysical Research Abstracts* **12** EGU2010-15091.

Dole WE, Jordan NF. 1978. Slope Mapping. *AAPG Bulletin-American Association of Petroleum Geologists* **62**(12): 2427-2440.

English Heritage. 2008. *Geophysical Survey in Archaeological Field Evaluation.* English Heritage Publishing 514301-59.

Geosoft. 2010. Montaj Magmap Filtering: 2D frequency domain processing of potential field data. Tutorial and user guide. Geosoft Inc: Toronto, Canada.

Godio A, Piro S. 2005. Integrated data processing for archeological magnetic surveys. *Leading Edge (Tulsa, OK)* **24**(11): 1138-1144.

Grauch VJS, Cordell L. 1987. Limitations of Determining Density Or Magnetic Boundaries from the Horizontal Gradient of Gravity Or Pseudogravity Data. *Geophysics* **52**(1): 118-121.

Hill I, Grossey T, Leech C. 2004. High-resolution multisensor geophysical surveys for near-surface applications can be rapid and cost-effective. *Leading Edge (Tulsa, OK)* **23**(7): 684-688.

Li X. 2006. Understanding 3D analytic signal amplitude. *Geophysics*, **71**, L13-L16.

Linford N. 2006. The application of geophysical methods to archaeological prospection. *Reports On Progress In Physics* **69**(7): 2205-2257.

Linford N, Linford P, Martin L, Payne A. 2007. Recent results from the English Heritage caesium magnetometer system in comparison with recent fluxgate gradiometers. *Archaeological Prospection* **14**(3): 151-166.

Linington RE. 1972. A Summary of Simple Theory Applicable to Magnetic Prospecting. *Prospezioni Archeologiche* **7**: 9-59.

Milea CM, Hansen RO, Tsokas GN, Papazachos CB, Tsourlos PI. 2010. Complex Attributes of the Magnetic Signal for Multiple Sources: Application to Signals from Buried Ditches. *Archaeological Prospection* **17**(2): 89-101. 10.1002/arp.376

Miller HG, Singh V. 1994. Potential field tilt; a new concept for location of potential field sources. *Journal of Applied Geophysics* **32**(2-3): 213-217.

Nabighian MN. 1972. The analytic signal of two-dimensional magnetic bodies with polygonal cross-section; its properties and use for automated anomaly interpretation. *Geophysics* **37**(3): 507-517.

Payne A. 2009. *Brading Roman Villa, Isle of Wight. Report on Geophysical Surveys, March 1994, April 1995 and February 2009.* English Heritage Research Department Report Series 104-2009

Phillips JD, Hansen RO, and Blakely RJ. 2007. The use of curvature in potential-field interpretation. *Exploration Geophysics*, **38**, 111-119.

Rao DB, Babu NR. 1993. A Fortran-77 computer program for three-dimensional inversion of magnetic anomalies resulting from multiple prismatic bodies. *Computers & Geosciences* **19**(6): 781-801.

Robinson ES. 1971. Use of Poissons Relation for Extraction of Pseudototal Magnetic Field Intensity from Gravity Observations. *Geophysics* **36**(3): 605-608.

Roest WR, Verhoef J, Pilkington M. 1992. Magnetic interpretation using the 3-D analytic signal. *Geophysics* **57**(1): 116-125.

Salem A, Williams S, Fairhead JD, Ravat D, Smith R. 2007. Tilt-depth method: A simple depth estimation method using first-order magnetic derivatives. *Leading Edge (Tulsa, OK)* **26**(12): 1502-1505.

Sheriff SD, MacDonald D, Dick D. 2010. Decorrugation, Edge Detection, and Modelling of Total Field Magnetic Observations from a Historic Town Site, Yellowstone National Park, USA. *Archaeological Prospection* **17**(1): 49-60. 10.1002/arp.371

Tabbagh A, Desvignes G, Dabas M. 1997. Processing of Z Gradiometer Magnetic Data using Linear Transforms and Analytic Signal. *Archaeological Prospection* **4**: 1-13.

Tchernychev M. 2009. Magpick - magnetic map & profile processing. User guide. Geometrics Inc. San Jose, USA.

Thurston JB, Smith RS. 1997. Automatic conversion of magnetic data to depth, dip, and susceptibility contrast using the SPI (TM) method. *Geophysics* **62**(3): 807-813.

Verduzco B, Fairhead JD, Green CM, MacKenzie C. 2004. New insights into magnetic derivatives for structural mapping. *Leading Edge (Tulsa, OK)* **23**(2): 116-119.

Wijns C, Perez C, Kowalczyk P. 2005. Theta map; edge detection in magnetic data. *Geophysics* **70**(4): L39-L43.

Appendix B

Paper based on Chapter 6

New approach to directional filtering of near-surface magnetic data
- *Near Surface Geophysics* - in press

New approach to directional filtering of near-surface magnetic data

Cheyney, S.[1], Hill, I.[1], Linford, N.[2], Fishwick, S.[1]

[1]Department of Geology, University of Leicester,
University Road, Leicester, LE1 7RH, UK (sc112@le.ac.uk),
[2]Geophysics Team, English Heritage,
Fort Cumberland, Eastney, Portsmouth, PO4 9LD, UK

October 2011

Abstract

Topographical anomalies or compressed areas of soil caused by ploughing, or more significant relic features such as ridge and furrow produce a response in near-surface magnetic surveys that are usually identified by their repetitive, linear pattern. While they are accurate recordings of the subsurface magnetic properties and micro-topographical features of the site, it is often the anomalies due to deeper features which are the primary focus of the survey. These target anomalies can be masked by the near-surface pattern, and it is therefore often preferable to remove these from the final presentation of the data.

Two routines used for removing these features are common in commercial processing software. These are the directional pass/reject and cosine-taper filters. While these filtering techniques can dramatically improve the clarity of the data image, it is shown here that they make significant changes to the data that remains. As interpretation of near-surface magnetic data moves beyond image analysis towards more quantitative methods, it is important to ensure that the final processed dataset represents as close as possible the response to the subsurface features of interest.

Here an alternative filtering routine dependant on both the azimuth and power-content of the anomalies is proposed that overcomes the problems encountered by the traditional techniques. It is shown that patterns of agricultural linear anomalies can be removed from the data without significantly changing the properties of the desired responses, and therefore quantitative interpretation can be subsequently carried out without the data being significantly compromised by the choice of previous processing techniques.

Introduction

Relic agricultural features such as ridge and furrow, or more recent ploughing furrows and seed beds, often have a magnetic signature which is recorded in magnetic surveys over archaeological sites. These types of anomalies appear as regular linear patterns across entire sites, and while they are accurate recordings of the subsurface magnetic properties and micro-topographic features of the terrain, often they are not the primary interest of the survey. Anomalies due to deeper lying features of interest can be masked by these near-surface patterns, and therefore whilst not being noise, it is often convenient for these elements of the signal to be suppressed in the dataset.

Gaffney & Gater (2003) suggest aligning the survey direction parallel to topographic expressions of the agricultural features, using pattern identification from aerial photographs if upstanding topographic features cannot be identified on site. If this approach is used, it should be possible to remove the effect of the agricultural features using de-striping algorithms typically used to correct directional errors and sensor offsets from zigzag traverses. However, when it is not possible, or convenient to incorporate the direction of unwanted features into the survey design, it becomes necessary to identify and remove them using directional filters during post-processing.

The Fourier transform

Directional filtering is typically a process conducted in the Fourier domain of the signal (see Fuller (1967); Bezvoda *et al.* (1990); Ferraccioli *et al.* (1998); Pilkington & Roest (1998) for examples). Decorrugation of magnetic data in the direction of data collection has also been shown using wavelet transforms (Fedi & Florio, 2003), however Fourier domain methods have the advantage of being ubiquitous in potential field processing and interpretation software packages. The Fourier transformation can be used to convert magnetic field data, from a function in the space domain, into the wavenumber domain (often referred to as the Fourier domain), or vice versa. The Fourier transform is defined as (see Blakely (1995)):

The Fourier transformation can be used to convert magnetic field data, from a function in the space domain, into the wavenumber domain (often referred to as the Fourier domain), or vice versa. Directional filtering is typically a process conducted in the Fourier domain of the signal. The Fourier transform is defined as (Blakely, 1995):

$$F(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-i(k_x + k_y)} dx dy \tag{1}$$

and the inverse Fourier transform is:

$$f(x,y) = \frac{1}{4\pi^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F(k_x,k_y)e^{i(k_x+k_y)}dk_x dk_y \tag{2}$$

where *f(x,y)* are the magnetic data in the space domain, *F(k_x, k_y)* is the function of the magnetic signal in the Fourier domain, and k_x and k_y are the directional wavenumbers inversely related to wavelength (λ), by $k_x = \frac{2\pi}{\lambda_x}$ and $k_y = \frac{2\pi}{\lambda_y}$. The Fourier transform is a complex function and can be separated into two parts. The amplitude spectrum is defined as:

$$|F(k)| = \sqrt{ReF(k)^2 + ImF(k)^2} \tag{3}$$

and the phase spectrum:

$$\Theta(k) = atan\left(\frac{ImF(k)}{ReF(k)}\right) \tag{4}$$

The amplitude function can be squared to give the energy-density spectrum $|F(k)|^2$, which has been used to represent the Fourier domain in the figures presented below.

Spectra of repeating near-surface anomalies

The Fourier domain of a continually repeating sine wave, will have a peak at a particular wavenumber which is defined by the wavelength of the signal. Figure 1a shows a sine wave with a wavelength of 1 m. The Fourier domain representation of this signal is shown in Figure 1b. As expected the energy of the function peaks at 2π divided by the wavelength, therefore here is 6.28 m^{-1}.

Assuming that the input signal is composed of the combination of the magnetic field generated by the underlying archaeology of interest, and the sine wave (simulating a plough pattern) superimposed, the perfect filter to remove the effect of the sine wave in the Fourier domain would be the inverse of Figure 1b. However, this information is unlikely to be available, yet it could be possible to design a filter to suppress the energy around a peak wavenumber. In reality though, the responses from ground disturbed by agricultural processes are unlikely to produce a perfect sine wave anomaly signal, and of course, the real archaeological signal is likely to have significant energy at similar wavelengths to the plough pattern also.

Figure 1c shows the magnetic field anomaly generated by a repeating 0.5 m wide by 0.1 m deep prism. The signal appears similar to the sine wave anomaly (Figure 1a), however the energy spectrum is more complicated and consists of several peaks and troughs throughout the Fourier domain (Figure 1d). Figure 1e shows prisms similar to Figure 1c, however this

time a 1 m gap is between each pair of prisms. This is designed to replicate an scenario similar to track marks caused by a tractor. The Fourier domain of the signal is shown in Figure 1f. As in the previous example, the energy spectrum is made up of a number of peaks and troughs at a range of wavenumbers.

It can be seen in these 1D examples that filtering of realistic signals from agricultural sources is much more complicated than for the sine wave example, as the energy of the anomaly is spread out across the entire spectrum.

Synthetic dataset

In order to test the effectiveness of various directional filtering algorithms, a synthetic dataset has been derived by calculating the response to a compound anomalous body, representing a typical archaeological structure. The anomaly shown in Figure 2a consists of two connected parts, in order to test the sensitivity to features of differing depths and magnetic response. The top of the shallower part of the body is positioned 0.3 m below ground level, while the deeper part is 0.8 m below ground level. The base of both parts is situated at 1.2 m below ground level. The anomaly has a susceptibility of 0.001 SI, and is situated within a magnetic field of 37.5 Am^{-1}, causing a magnetisation intensity of 0.0375 Am^{-1}. The total field response has been calculated at a height of 0.2 m above ground level, at a sampling density of 0.1 x 0.1 m, and magnetic inclination of 64°, using the method described by Rao & Babu (1993). This represents a well-sampled, noise-free dataset, at a mid-latitude site, and is hereafter referred to as the archaeological signal (Figure 2a).

Agricultural signal is generated as described in Figure 1e, and has a bearing of 170°. The magnetisation intensity of the agricultural signal is also 0.0375 Am^{-1}, resulting in responses between -1.5 and 2 nT (Figure 2b). The synthetic dataset used in this study is generated by the addition of the archaeological and agricultural signal, and is shown in the space domain in Figure 2c and the Fourier domain in Figure 2d. Although it is not actually "noise" the agricultural signal is referred to here as the agricultural noise, as it is the undesired part of the dataset.

The aim of the directional filtering routines presented here, is to remove the agricultural noise in order to produce a dataset that most closely represents the archaeological signal. The effect of the agricultural noise in the Fourier domain can be seen as a linear peak in Figure 2d running from the origin to ∼30 Kx, 5 Ky and mirrored in the negative Kx domain.

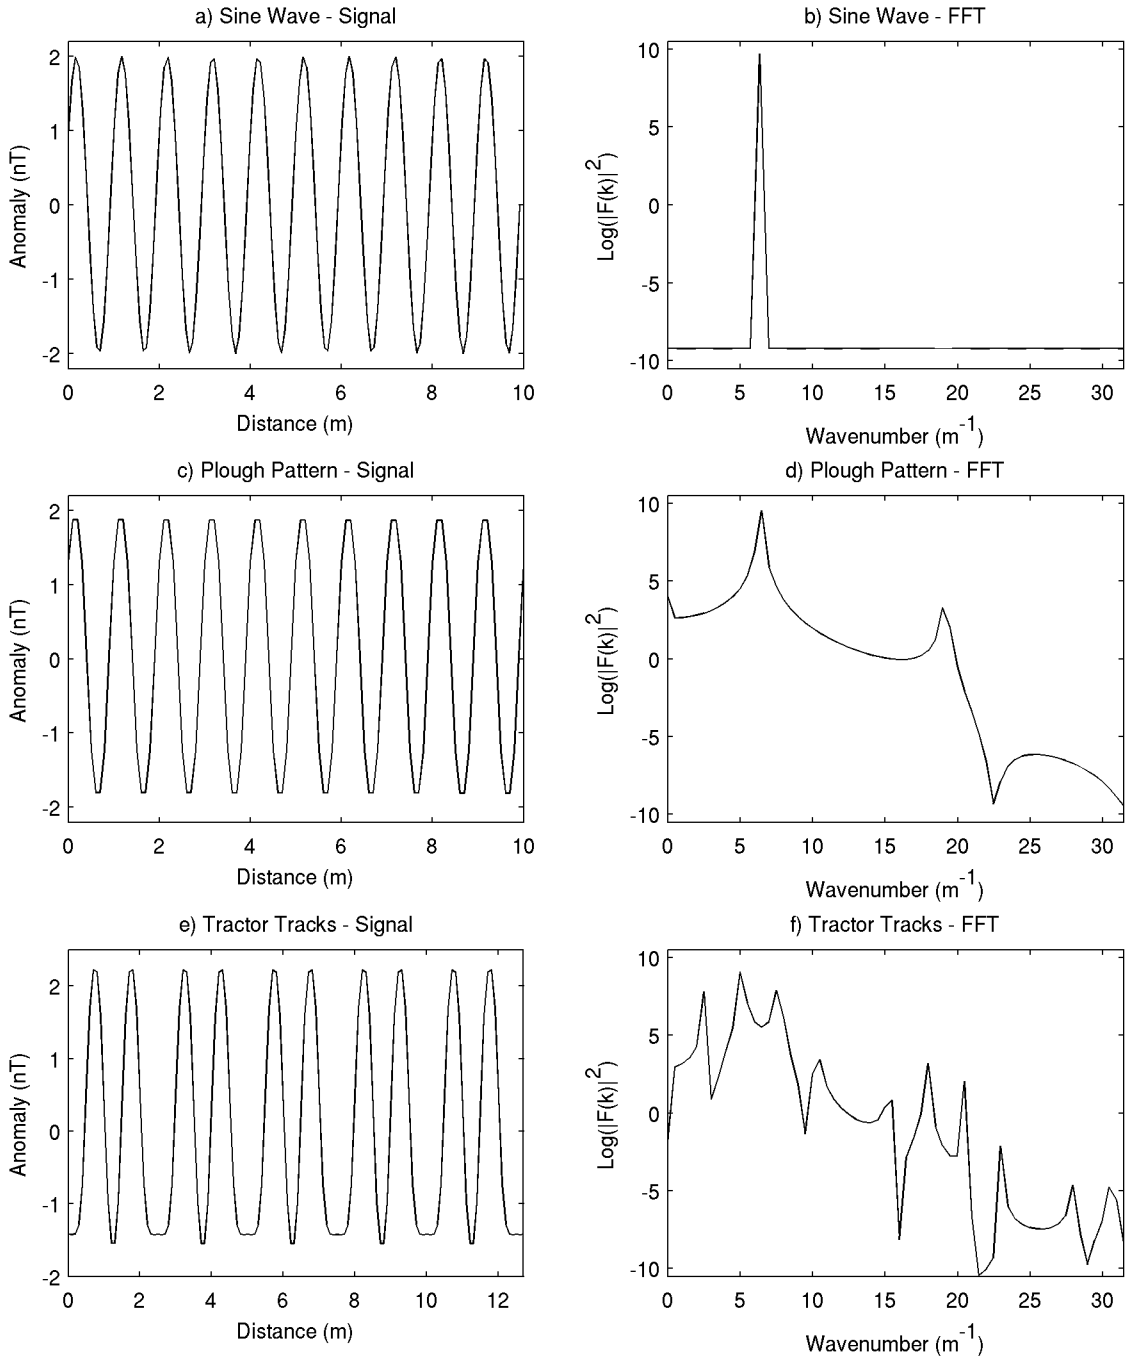

Figure 1: (a) Sine wave of wavelength=1m. (b) FFT of a). (c) Repeated "plough" anomaly of prisms 0.5m wide, 0.1m deep and infinite length positioned 0.5m apart. (d) FFT of c) (e) Repeating "tracks" anomaly pattern of pairs of ∞ x 0.5 x 0.1m prisms positioned 0.5m apart, separated by 1m. (f) FFT of e)

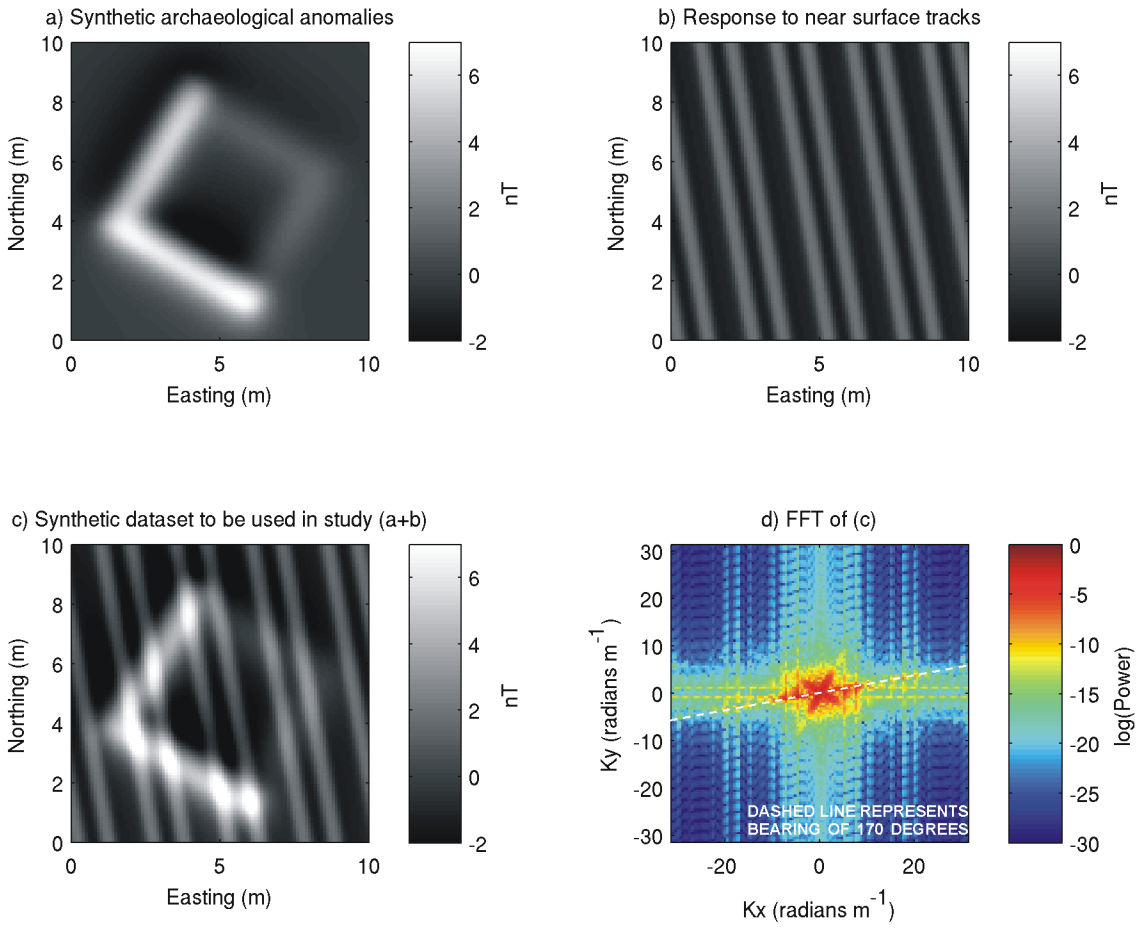

Figure 2: (a) Synthetic archaeological signal required from the data (b) Near-surface agricultural "tracks" required to be filtered out by the directional filtering. (c) Synthetic dataset to be used, combining a) and b). (d) 2D FFT of c), showing the effect of the agricultural noise as a linear peak running roughly left to right.

Directional pass/reject filter

Once the bearing of the features to be removed is identified, it is possible to isolate and remove all power from the wavenumbers representing this bearing. This method is presented by Fuller (1967) and Bezvoda *et al.* (1990), and typically is done by selecting a minimum and maximum bearing. Any wavenumbers within this zone are set to zero in the Fourier domain, all other wavenumbers remaining the same. The inverse Fast Fourier Transform (FFT) is then performed and the power of lineations along the selected bearings will be reduced.

In reality, editing the spectra to zero within the filter zone would create a sharp discontinuity within the spectra, and increases the likelihood of Gibbs phenomena, where ringing is often observed once the spectra is transformed back into the space domain. In order to reduce the effects of Gibbs phenomena, the filter function should be allowed to change from 1 to 0 in a smooth way. Scollar (1970) discusses several methods for this, one of which is to cosine taper the filter function. This is the method that has been used here.

Linear anomalies present in the space domain can be identified in the Fourier domain along an azimuth rotated 90° from the azimuth observed in the space domain. For example, a N-S trending pattern, will appear as peaks in the left to right direction in the Fourier domain, and vice versa. Figure 3a shows the filter to be multiplied by the Fourier domain function. The minimum azimuth is 160° and the maximum is 180°. Between these two angles the filter equals zero, and no power will be retained in this area. In order to reduce Gibbs phenomenon, tapering has been conducted between 155° and 160° and 180° and 185°. The filter is multiplied by the Fourier domain and the filtered spectrum is shown in Figure 3b. The inverse FFT is then performed and the filtered dataset in the space domain is presented in Figure 3c. The filter has performed a good job at removing the majority of the agricultural noise, as little of the 170° pattern remains in the filtered dataset. In order to identify the data that has been removed, the filtered dataset is subtracted from the original dataset, and the data removed by the filtering is shown in Figure 3d. It can be seen that a small amount of the amplitude of the larger archaeological anomaly has been removed from the data, and is shown as a higher amplitude in the removed data towards the SW corner.

Cosine tapering of filtered zone

An alternative to completely removing the power from wavenumbers relating to the filtering direction, is to increasingly suppress the power depending on how close the data are to the filtering direction. This is the approach of directional cosine tapering. The algorithm is stated by Geosoft (2010) as:

$$L(\theta) = |cos^n \left(\alpha - \theta + \frac{\pi}{2} \right)| \tag{5}$$

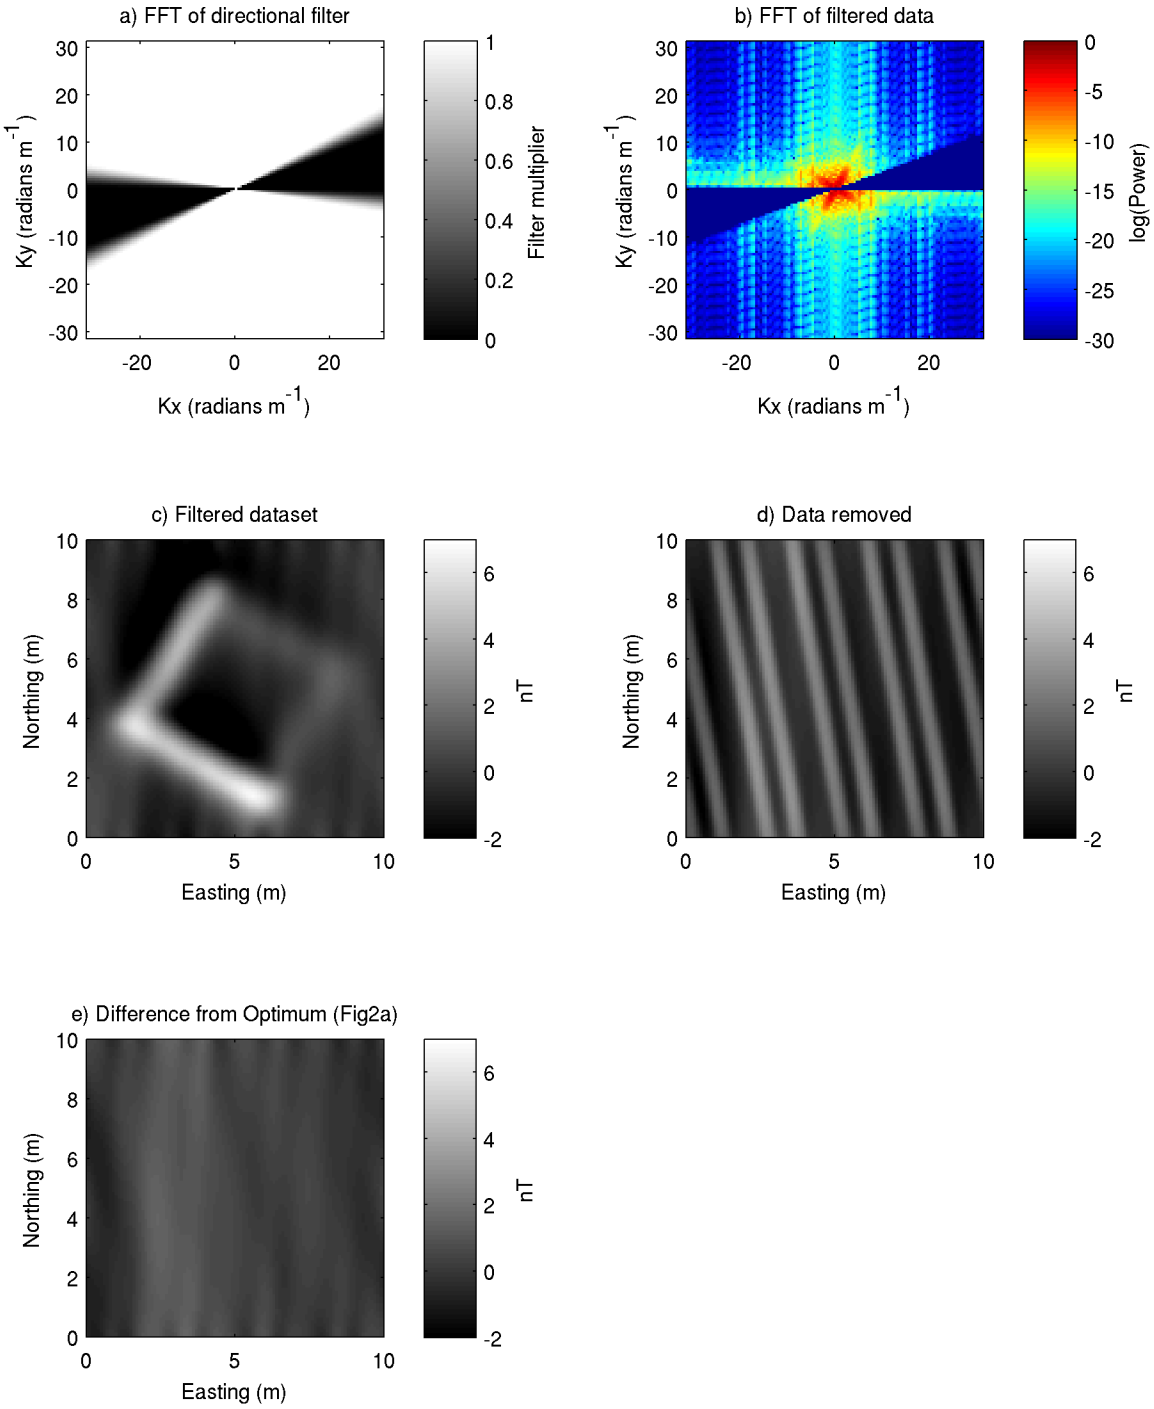

Figure 3: (a) Filter to be applied to the Fourier domain. (b) FFT of filtered data. Note the reduction in power between 160° and 170°. (c) Post-filtering dataset. (d) Data removed from the original data by filtering. (e) Difference from the optimum dataset (Fig2a minus Fig3c).

where α is the filtering direction, and the value of n is the degree of the cosine function, that can be varied to alter the shape of the function. Directional cosine filtering is commonly used and can be seen included in decorrugation techniques by Ferraccioli *et al.* (1998) and Pilkington & Roest (1998), as well as archaeological examples such as Linford *et al.* (2007).

Three scenarios have been tested here, with varying values for n. Figures 4a-c, show the filter to be applied to the Fourier domain for *n=0.5*, *n=1.0* and *n=2.0* respectively. Figures 4d-f, show the filtered Fourier domain, Figures 4g-i show the final filtered dataset, and Figures 4j-l show the data that has been removed by the filtering process. It can be seen in the Fourier domain plots that the obvious difference between the Cosine tapering technique, and the previous technique is that very little of the spectrum is actually removed, instead it is increasingly suppressed depending upon how close the data is to the filter azimuth.

Figure 4g shows that the *n=0.5* value has not been sufficient to remove all the agricultural signal, as striping can still be observed across the dataset. It is notable however, that part of the archaeological anomaly has been filtered out, as Figure 4j shows evidence of this in the removed data. This will result in a decrease in amplitude of the signal of interest in the filtered dataset. This is a problem that increases as the value of n is increased. Whereas Figures 4h and 4i show little evidence of the agricultural noise, the amplitude of the archaeological signal is increasingly reduced as more of the signal of interest is removed from the dataset. In Figure 4i, the response of the upper archaeological anomaly is $\sim$4 nT, instead of the $\sim$7 nT seen in Figure 2a. Whilst in this scenario both the archaeological bodies are still evident in the results, smaller amplitude anomalies would also have been reduced and may go undetected. This is particularly a problem for archaeological surveys, where signals are typically only a few nT in amplitude. Furthermore any subsequent quantitative interpretation techniques, which rely on the amplitude of the total field signal will be compromised.

Power/azimuth threshold (PATH) filter

An alternative method to the directional pass/reject filter and cosine taper filter, would be to add an additional amplitude constraint to the filter. A minimum and maximum azimuth is defined by the operator, as well as an amplitude threshold. Anything in the Fourier domain between the two angles and above the threshold is removed. All other data remains. The procedure is demonstrated in Figure 5. This techniques solves two of the drawbacks identified in the previous techniques. 1) The amount of the spectrum that is removed by filtering will be small, and not all the power at certain wavenumbers will be removed/reduced. 2) All data outside of the azimuths declared will remain the same, causing no loss of amplitude in features along different azimuths to the filter.

This filtering can seen in Figure 6a. A minimum azimuth of 160°, maximum of 180° and a normalised energy density threshold of -14 has been selected. The choice of the threshold

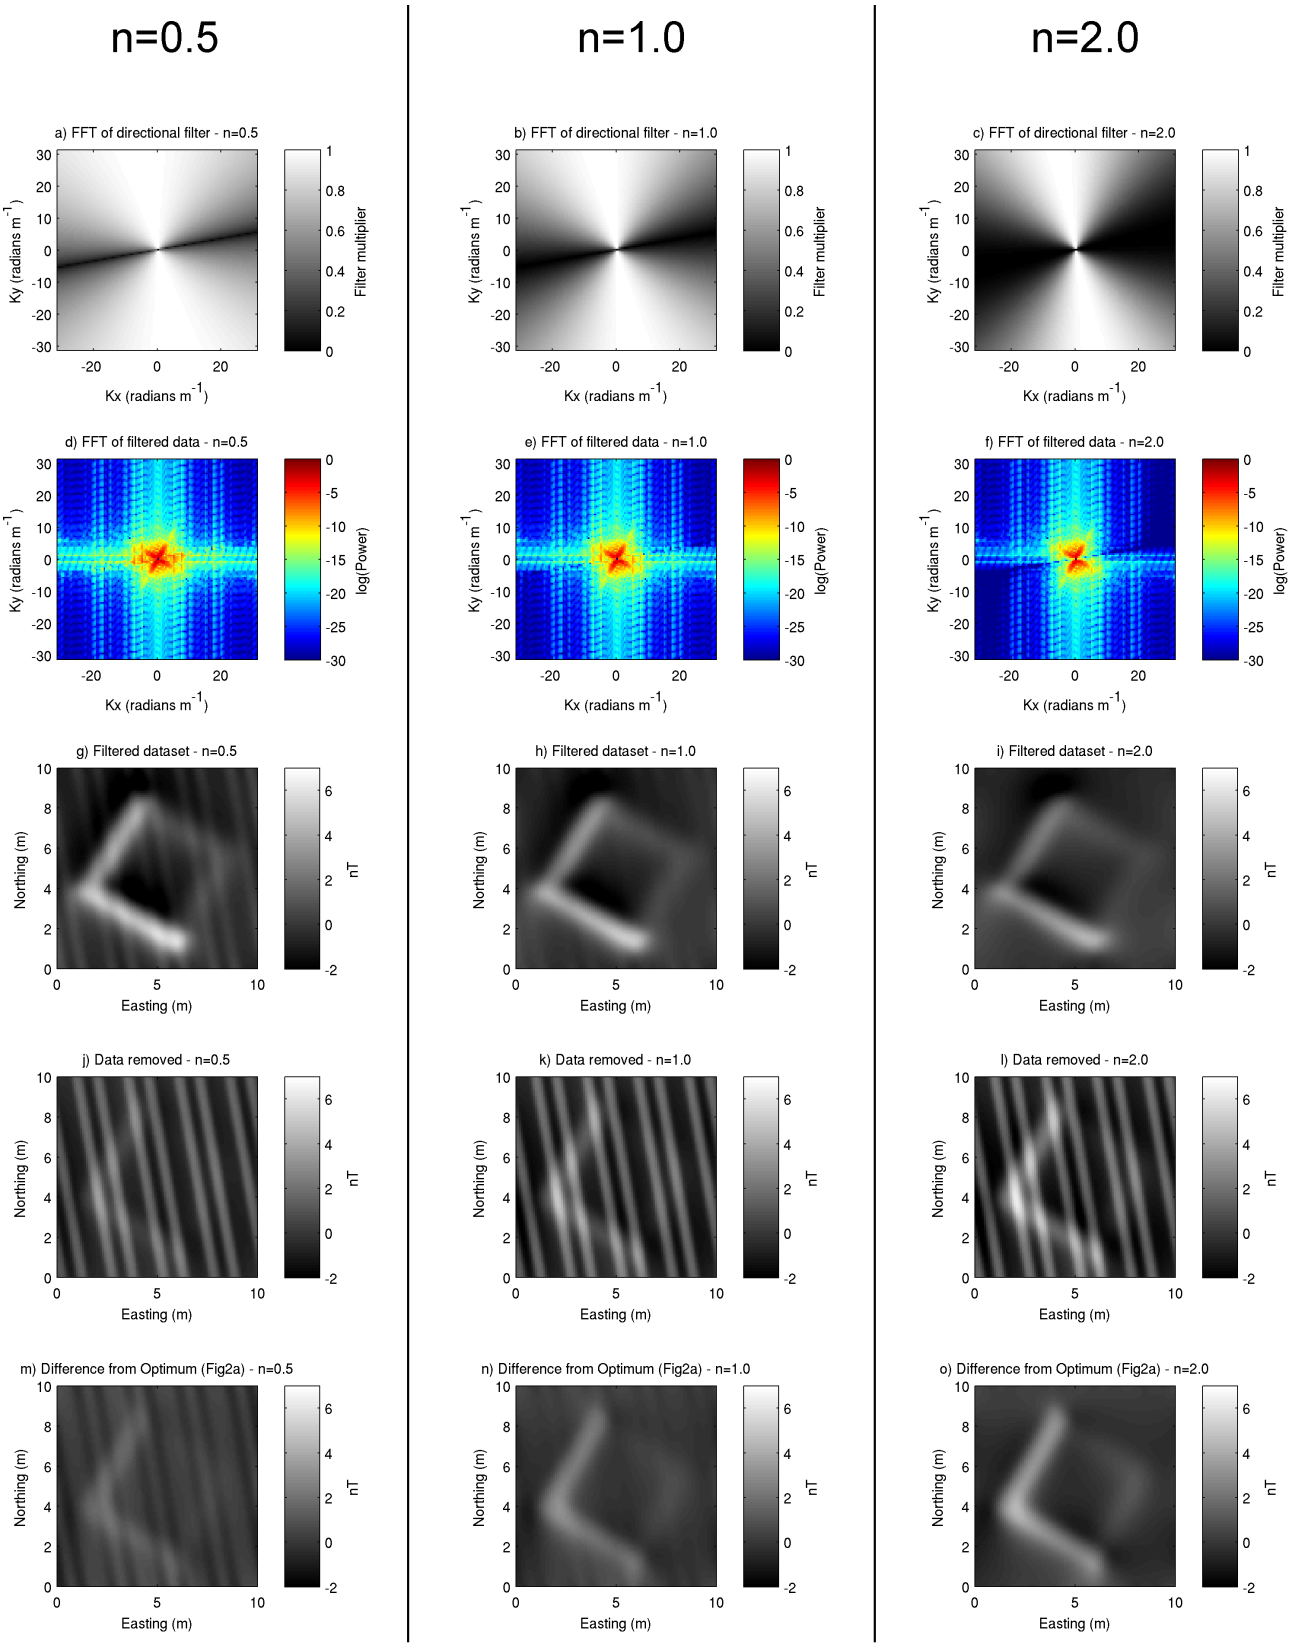

Figure 4: Illustration of the effect of the cosine taper filter, using α=*170°*. For *n=0.5*, the filter to be applied to the Fourier domain (a), filtered Fourier domain (d), filtered dataset (g), data removed by filtering (j), and difference from the optimum dataset (Fig2a minus Fig4g) are shown (m) . The respective plots for n=1 are shown in figures (b), (e), (h), (k) and (n), for n=2 in (c), (f), (i), (l) and (o).

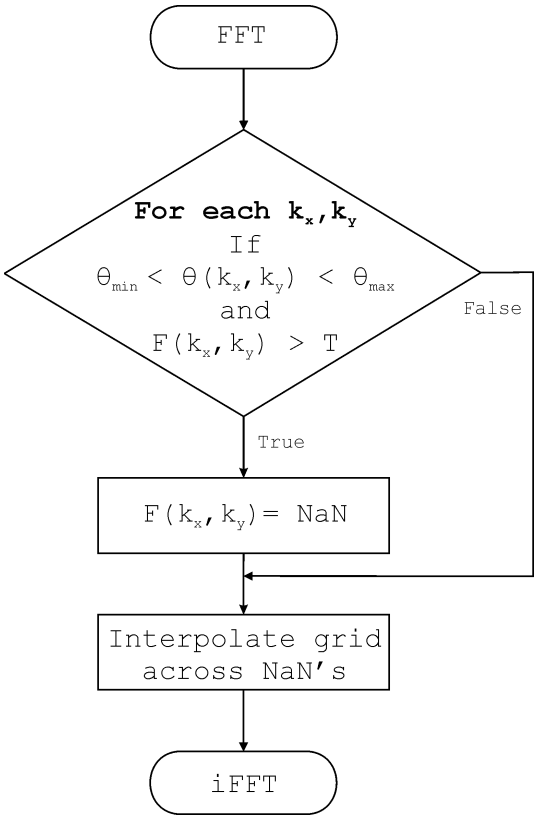

Figure 5: Flow chart outlining the procedure of the PATH filter, where F(kx,ky) is the function of the magnetic data represented in the Fourier domain, and the user inputs, min, max and T are the minimum and maximum angles to be filtered, and the minimum normalised power threshold, respectively.

value is important to the final result, and will dependent upon the spectrum of the specific problem. Selecting a threshold value that is too low will result in a larger proportion of the spectrum becoming filtered than required, and the result will tend towards that obtained by the directional pass/reject filter. Choosing a threshold value that is too high will filter too little of the spectrum, and the undesired linear patterns will remain in the filtered dataset. It is therefore ideal to select the highest possible threshold that removes the undesired features.

Any data within these constraints are removed and shown by the areas which are not white in Figure 6a. The Fourier domain to subsequently interpolated across the removed data to avoid sharp discontinuities and Gibb's phenomena. Comparatively, much less data has been filtered than Figures 3a and 4a-c. The filtered Fourier domain can be seen in Figure 6b. The filtered data (Figure 6c) shows good removal of the 170° agricultural noise. The removed data (Figure 6d) clearly shows the agricultural noise, with little evidence of the archaeological anomaly.

Comparing the filtering techniques

In order to evaluate the outputs of the filtering techniques described above, it is necessary to compute how closely the filtered data matches the original synthetic input dataset (Figure 2a). In order to do this, a simple root mean squared (RMS) approach has been used to quantify the difference between the 'ideal' and filtered output. The RMS error is defined as:

$$RMSerror = \sqrt{\frac{\sum_{i=1}^{n}(T_i - \bar{T}_i)^2}{n}} \tag{6}$$

where T is the data from Figure 2a and $\bar{T}$ is the filtered output.

It can be seen in Table 1 that the power/azimuth threshold filtering has produced a much lower RMS error than any other technique. The second best technique based on the value of the sum of squares is cosine tapering with an n-value of 0.5, however, as demonstrated in Figure 4 this has left evidence of the linear anomalies present in the image. As n is increased, the value of the sum of squares is dramatically increased, showing that although the images look the clearest, a large proportion of the desired archaeological signal has been removed leading to a dataset that is the furthest away from the desired output.

The power/azimuth threshold filtering requires the filtered parts of the spectrum to be interpolated to remove the problem of sharp discontinuities in the spectrum. To test this, three interpolation techniques, triangle-based linear, triangle-based cubic and nearest neighbour interpolation have been applied to the data, as well as a scenario where no interpolation has been applied. In this case, Table 1 shows the linear interpolation to be the most effective, while the un-interpolated scenario is the second best. In addition to the increased likelihood of Gibbs phenomenon, when the spectrum is not interpolated

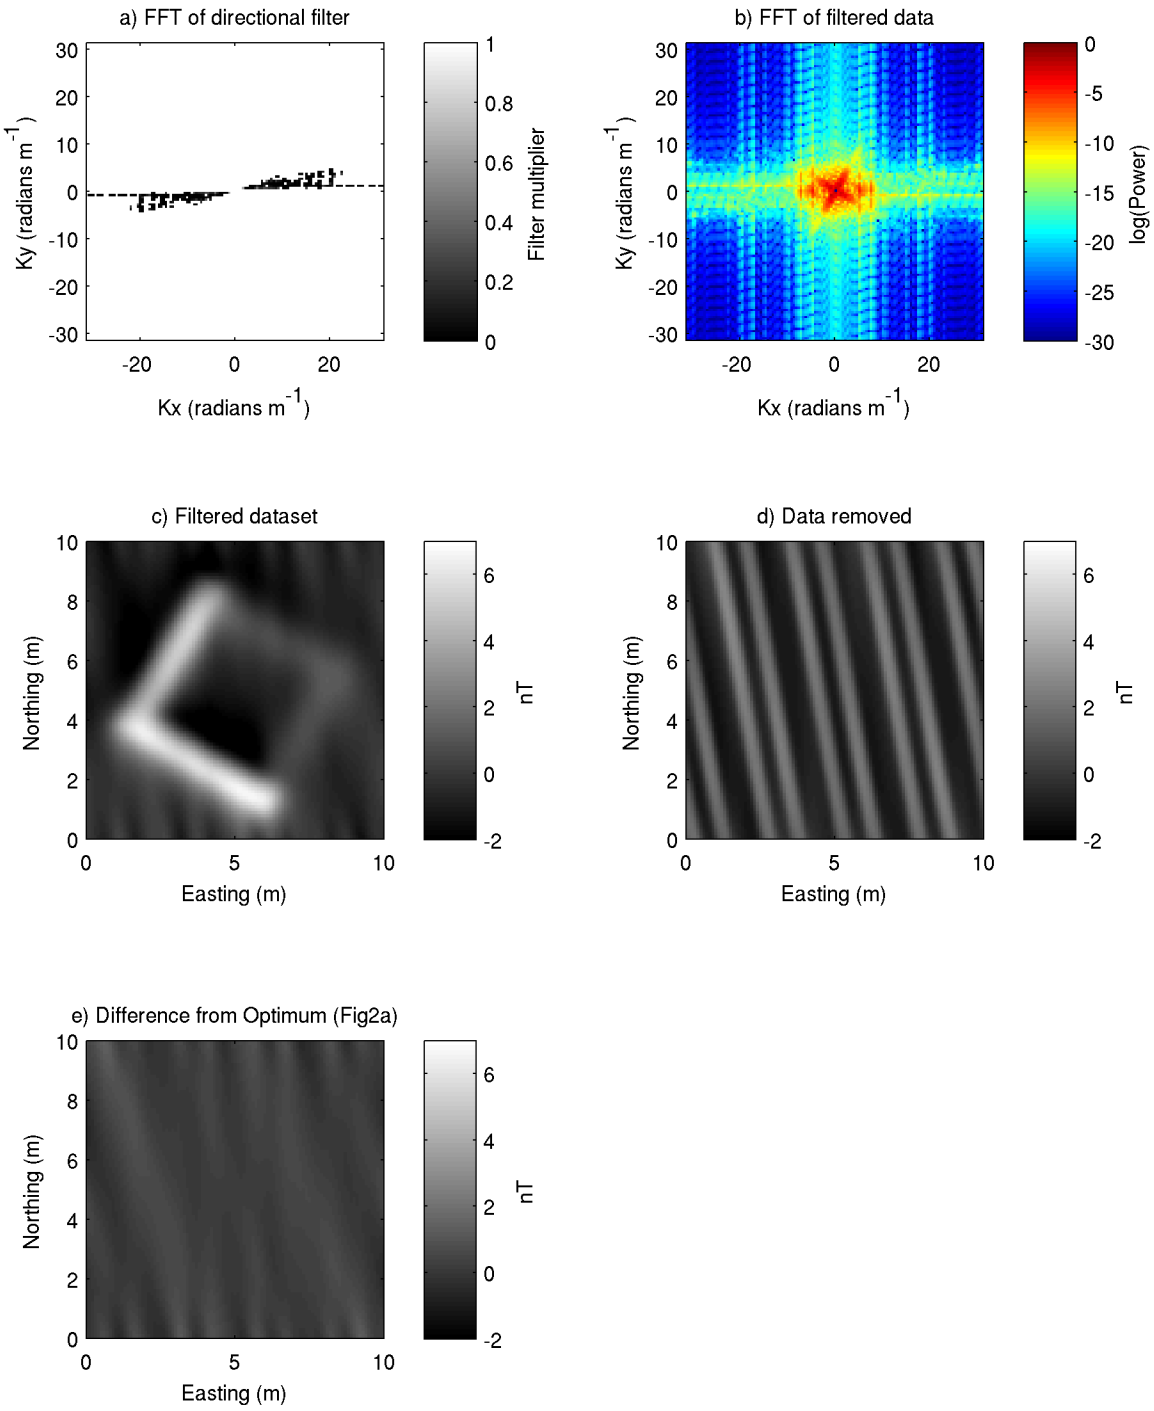

Figure 6: (a) Filter to be applied to data 160> α >180 and normalised energy density >-14. (b) Filtered data in the Fourier domain. (c) Filtered dataset. (d) Data removed by filtering and (e) Difference from the optimum dataset (Fig2a minus Fig6c).

Filtering method	Additional information	RMS Error
Directional pass/reject		0.63 nT
Cosine tapering	n=0.5	0.59 nT
Cosine tapering	n=1.0	0.79 nT
Cosine tapering	n=2.0	1.07 nT
Power/azimuth threshold	Linear interpolation	0.37 nT
Power/azimuth threshold	Cubic interpolation	0.38 nT
Power/azimuth threshold	Nearest neighbour interpolation	0.46 nT
Power/azimuth threshold	No interpolation	0.38 nT

Table 1: Root mean square of the differences between the desired data and the filtering output.

prior to the inverse Fourier transform, previously high-powered wavenumbers are altered to low-power, and a striping pattern can again be present, this time with negative polarity. This is illustrated in Figure 7.

Case study

Brading Roman Villa is situated between the towns of Brading and Sandown on the Isle of Wight, UK. The site consists of a Roman villa, now incorporated into a visitor centre, and surrounding pasture land. The data used in this study is a subset of a much larger Caesium magnetometery survey collected in 2009 by the English Heritage Geophysics Team over the surrounding pasture land (Payne, 2009). The data were collected with a spacial resolution of 0.125 x 0.5 m.

The subset, shown in Figure 8b, is a 135 x 55 m area situated ∼100 m NW of the villa, showing a series of positive anomalies. The most prominent features in the data are a large circular anomaly ∼25 m in diameter towards the western side of the dataset, and a series of positive linear anomalies across the area. An agricultural pattern is prominent across the dataset, as well as a series of small positive peaks scattered across the area.

The Fourier domain representation of the data is shown in Figure 8a, highlighting the areas where the filtering has been applied in white. The reject criteria for the filter were an angle >75 and <85°, with a normalised energy density amplitude of >-12. The post-processed inverse FFT is shown in Figure 8c, and the removed data in Figure 8d. It can be seen that the filtered dataset has improved the visual appearance of the data, and the majority of the agricultural pattern has been removed. There is a little evidence of any of the archaeological anomalies being removed, therefore the amplitude of these should not have been significantly reduced due to the filtering.

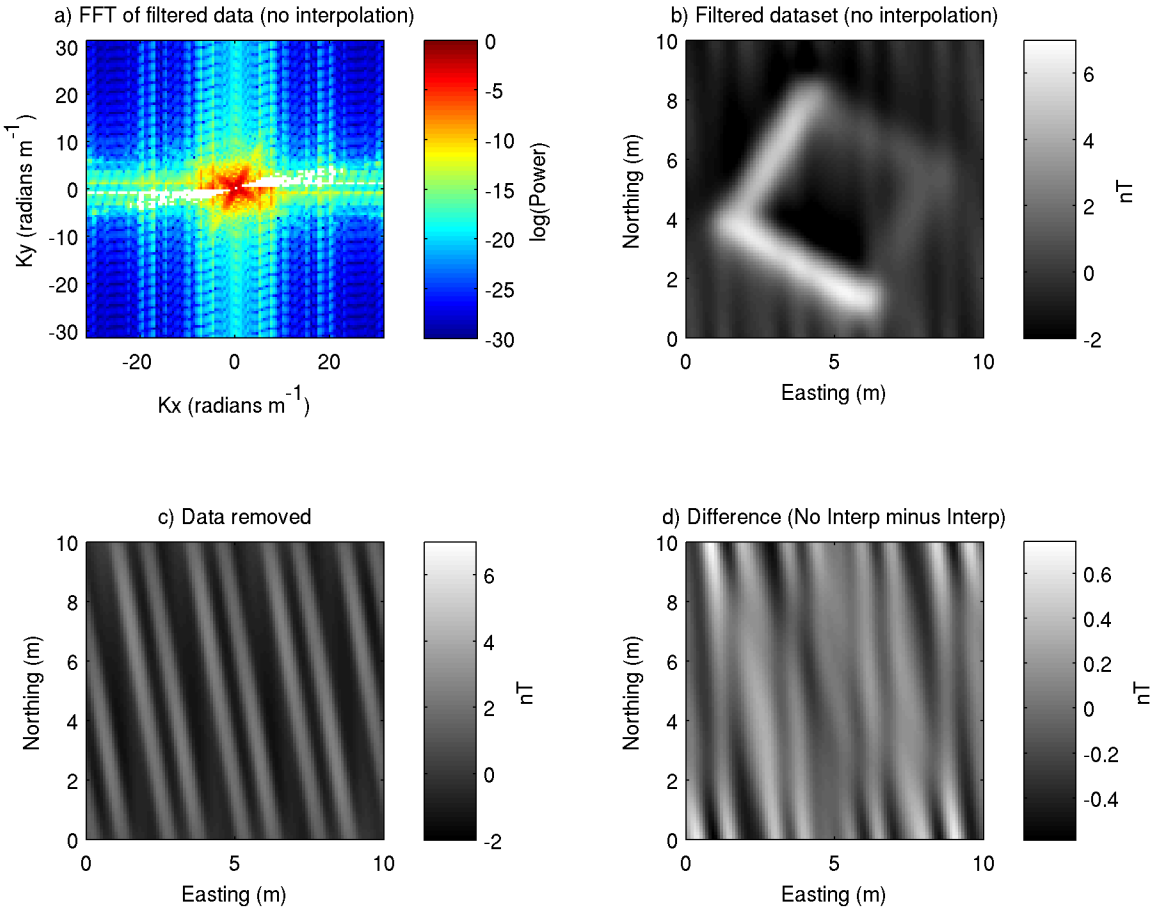

Figure 7: Illustration of the effect of interpolating across the filtered spectrum. (a) Filtered data in the Fourier domain (no interpolation). (b) Filtered dataset. (c) Data removed by filtering. (d) Non-interpolated data minus interpolated data (Figure 7b minus Figure 6c). The effect of not interpolating the Fourier domain prior to iFFT has created a set of linear anomalies closer to a N-S direction. Note the scale axis is much smaller on this plot, as these differences are much smaller amplitude.

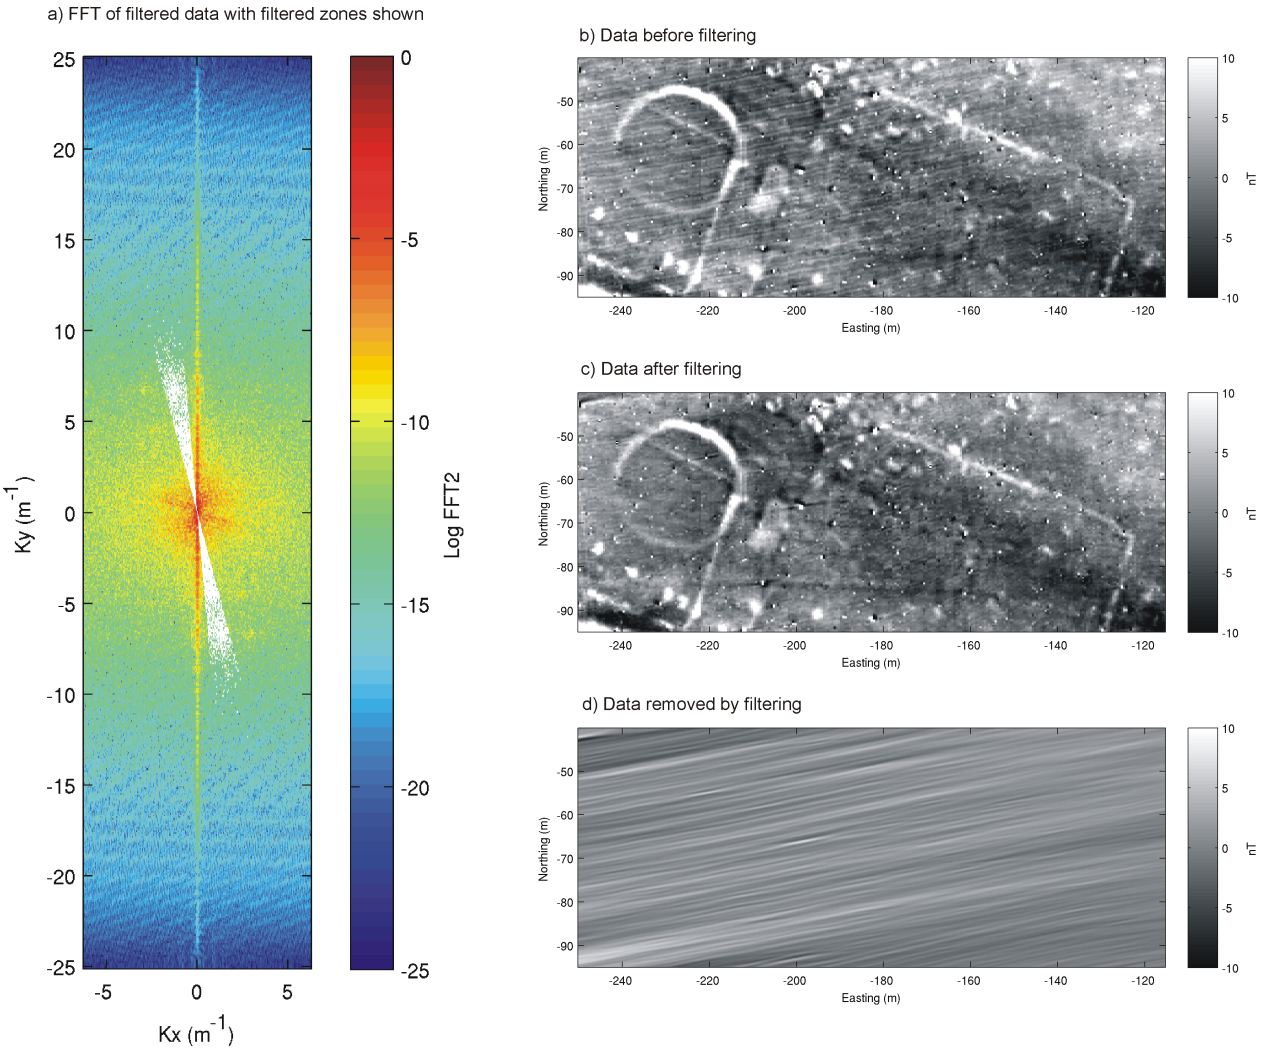

Figure 8: Case study from Brading Roman Villa, Isle of Wight, UK. (a) Filtered fourier domain with filtered zone shown in white. (b) Original dataset. (c) Filtered dataset. (d) Data removed by filtering.

Discussion and conclusions

Linear anomalies in near-surface geophysical data are a common problem, and often appear in archeo-magnetic surveys over current agricultural land. In order to clarify the image of the deeper archaeological signal it is desired to remove these features from the data, and consequently several processing and interpretation software packages contain routines for directional filtering in the Fourier domain. The most common two of these are the directional pass/reject and cosine taper filters.

It is shown here, that while both of these filters do greatly improve the appearance of the data, they have their own individual drawbacks. The simple filtering directional pass/reject approach can cause a large proportion of the data to be removed, while the directional cosine tapering approach fundamentally changes the amplitude of the data at all wavenumbers, not just in the filtering direction. Modern interpretation techniques are capable of making quantitative interpretations regarding depth and susceptibility based on the amplitude of the signal. While the cosine-tapering technique has produced the best image, features of the same susceptibility would be interpreted differently depending upon their orientations should this filtering technique be applied beforehand.

Here an alternative filtering routine is proposed. The PATH filter is based on the directional pass/reject filter, however, the problem of a large proportion of the spectrum being removed by the filter is avoided by only filtering parts of the spectrum that are over an additional energy-density criteria. This approach ensures that amplitudes of anomalies in directions outside the filter parameters remain the same, and the potential to perform quantitative interpretation techniques has not been compromised by the filtering.

Analysis of the RMS error of the three different filtering outputs shows the power/azimuth threshold filter to outperform the other two, with the resulting image the closest to the synthetic input. It has been successfully applied to real data from Brading, Isle of Wight, and shows a substantial improvement to the data.

Acknowledgments

This work is funded by NERC doctoral training grant NE/G523798/1, and is CASE sponsored from Geomatrix Earth Science Ltd. Many thanks to the English Heritage Geophysics Team, who allowed me to assist with the fieldwork, and then use the data from Brading Roman Villa, IOW, UK.

References

Bezvoda, V., Jezek, J. & Segeth, K., 1990. Fredpack-a program package for linear filtering in the frequency domain. *Computers and Geosciences*, **16**(8), 1123–1154.

Blakely, R.J., 1995. *Potential theory in gravity and magnetic applications*. Cambridge, Cambridge University Press, United Kingdom.

Fedi, M. & Florio, G., 2003. Decorrugation and removal of directional trends of magnetic fields by the wavelet transform: application to archaeological areas. *Geophysical Prospecting*, **51**(4), 261–272.

Ferraccioli, F., Gambetta, M. & Bozzo, E., 1998. Microlevelling procedures applied to regional aeromagnetic data: an example from the transantarctic mountains (antarctica). *Geophysical Prospecting*, **46**(2), 177–196.

Fuller, B.D., 1967. *Two-dimensional frequency analysis and the design of grid operators*. Mining Geophysics Vol 2: Theory, Society of Exploration Geophysicists, Tulsa, Oklahoma, 1 edition, pp. 658–709.

Gaffney, C. & Gater, J., 2003. *Revealing the buried past: Geophysics for archaeologists*. Tempus Publishing, UK, 1 edition.

Geosoft, 2010. *Montaj Mapmap Filtering: 2D Frequency Domain Processing Of Potential Field Data. Tutorial and User Guide*. Geosoft Inc, Toronto, Canada.

Linford, N., Linford, P., Martin, L. & Payne, A., 2007. Recent results from the english heritage caesium magnetometer system in comparison with recent fluxgate gradiometers. *Archaeological Prospection*, **14**, 151–151–166.

Payne, A., 2009. Brading roman villa, isle of wight. report on geophysical surveys, march 1994, april 1995 and february 2009. Technical Report 104-2009, English Heritage.

Pilkington, M. & Roest, W.R., 1998. Removing varying directional trends in aeromagnetic data. *Geophysics*, **63**(2), 446–453.

Rao, D.B. & Babu, N.R., 1993. A fortran-77 computer program for three-dimensional inversion of magnetic anomalies resulting from multiple prismatic bodies. *Computers & Geosciences*, **19**(6), 781–801.

Scollar, I., 1970. Fourier transform methods for the evaluation of magnetic maps. *Prospezioni Archeologiche*, **5**, 9–41.

Appendix C

MATLAB scripts and synthetic datasets (on CD)

MATLAB scripts

ID	Method	Filename
C.1.1	Analytic Signal Amplitude	*analytic_signal_amplitude.m*
C.1.2	Butterworth Filter	*butterworth_filter.m*
C.1.3	Canny Edge Detection	*canny_edge_detection.m*
C.1.4	Cosine Rolloff Filter	*cosine_rolloff_filter.m*
C.1.5	Identifying remanent magnetisation using pseudogravity	*dannemiller_pgrav*
C.1.6	Identifying remanent magnetisation using total field	*dannemiller_rtp*
C.1.7	Directional cosine filter	*directional_cosine_filter.m*
C.1.8	Directional pass/reject filter	*directional_pass_reject_filter.m*
C.1.9	Directional PATH filter	*directional_PATH_filter.m*
C.1.10	Downward continuation	*downward_continuation.m*
C.1.11	Forward modelling	*forward_modelling.m*
C.1.12	High pass filter	*high_pass_filter.m*
C.1.13	Horizontal gradient magnitude	*horizontal_gradient_magnitude.m*
C.1.14	iSPI	*ispi.m*
C.1.15	Pseudogravity transformation	*pseudograv.m*
C.1.16	Reduction to the pole	*reduction_to_pole.m*
C.1.17	SPI	*spi.m*
C.1.18	Theta map	*theta_map.m*
C.1.19	Tilt angle	*tilt_angle.m*
C.1.20	Various plotting functions	*plotting_functions*

The scripts presented here are written in MATLAB 7.8.0.347 (R2009a), running on openSUSE 11.4 (i586).

Synthetic datasets

ID	Name	Image
C.2.1	Compound	Figure 3.1
C.2.2	Prism	Figure 5.2b
C.2.3	Prism - depths	Figure 4.1
C.2.4	Prisms - positive/negative contrasts	Figure 5.26
C.2.5	Prisms - remanence	Figure 5.22
C.2.6	Regional trend	Figure 5.14a
C.2.7	Tracks pattern	Figure 5.19a

Synthetic datasets are provided with, and without noise, in a gridded format, and as an XYZ file. An PNG image of each dataset is also present.

References

Aspinall, A., Gaffney, C. & Schmidt, A., 2008. *Magnetometry for Archaeologists*. AltaMira Press, U.S.A., first edition.

Baranov, V., 1957. A new method for interpretation of aeromagnetic maps; pseudo-gravimetric anomalies. *Geophysics*, **22**(2), 359–382.

Bescoby, D.J., Cawley, G.C. & Chroston, P.N., 2006. Enhanced interpretation of magnetic survey data from archaeological sites using artificial neural networks. *Geophysics*, **71**(5), H45–H53.

Bevan, B., 1994. The magnetic anomaly of brick foundations. *Archaeological Prospection*, **1**, 93–104.

Bezvoda, V., Jezek, J. & Segeth, K., 1990. Fredpack-a program package for linear filtering in the frequency domain. *Computers and Geosciences*, **16**(8), 1123–1154.

Blakely, R.J. & Simpson, R.W., 1986. Approximating edges of source bodies from magnetic or gravity-anomalies. *Geophysics*, **51**(7), 1494–1498.

Blakely, R.J., 1995. *Potential theory in gravity and magnetic applications*. Cambridge, Cambridge University Press, United Kingdom.

Bueyueksarac, A., Arisoy, M.O., Bektas, O., Kocak, O. & Cay, T., 2008. Determination of grave locations in dedemezari necropolis (western turkey) using magnetic field derivatives. *Archaeological Prospection*, **15**(4), 267–283.

Canny, J., 1986. A computational approach to edge-detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **8**(6), 679–698.

Cheyney, S., Hill, I. & Linford, N., 2011. Advantages to using the pseudogravity transformation to aid edge detection of total field archaeomagnetic datasets. *Archaeological Prospection*, **18**, 81–93.

Ciminale, M. & Loddo, M., 2001. Aspects of magnetic data processing. *Archaeological Prospection*, **8**, 239–246.

Clark, A.J., Tarling, D.H. & Noel, M., 1988. Developments in archaeomagnetic dating in britain. *Journal of Archaeological Science*, **15**(6), 645–667.

Cordell, L. & Grauch, V.J.S., 1985. *Mapping basement magnetization zones from aeromagnetic data in the San Juan Basin, New Mexico.* The utility of regional gravity and magnetic maps, Society of Exploration Geophysicists, Tulsa, USA, 1 edition, pp. 181–197.

Dabas, M., 2010. First results of amp surveys for large areas (¿100ha) with decimeter resolutions in archaeology.

Dannemiller, N. & Li, Y., 2006. A new method for determination of magnetization direction. *Geophysics*, **71**(6), L69–L73.

Desvignes, G., Tabbagh, A. & Benech, C., 1999. The determination of the depth of magnetic anomaly sources. *Archaeological Prospection*, **6**, 85–85–105.

Dittmer, J.K. & Szymanski, J.E., 1995. The stochastic inversion of magnetics and resistivity data using the simulated annealing algorithm. *Geophysical Prospecting*, **43**(3), 397–416.

Dole, W.E. & Jordan, N.F., 1978. Slope mapping. *Aapg Bulletin-American Association of Petroleum Geologists*, **62**(12), 2427–2440.

Eder-Hinterleitner, A., Neubauer, W. & Melichar, P., 1995. *Reconstruction of archaeological structures using magnetic prospection*, *Interfacing the past: Computer Applications and Quantitative Methods in Archaeology*, volume 1. Analecta Praehistorica Leidensia 28, University of Leiden, caa95 edition, pp. 131–137.

English Heritage, 2008. Geophysical survey in archaeological field evaluation. Technical Report 51430, English Heritage Publishing.

Fassbinder, J., Stanjek, H. & Vali, H., 1990. Occurrence of magnetic bacteria in soil. *Nature*, **343**(6254), 161–163.

Ferraccioli, F., Gambetta, M. & Bozzo, E., 1998. Microlevelling procedures applied to regional aeromagnetic data: an example from the transantarctic mountains (antarctica). *Geophysical Prospecting*, **46**(2), 177–196.

Finlay, C.C., Maus, S., Beggan, C.D., Bondar, T.N., Chambodut, A., Chernova, T.A., Chulliat, A., Golovkov, V.P., Hamilton, B., Hamoudi, M., Holme, R., Hulot, G., Kuang, W., Langlais,

B., Lesur, V., Lowes, F.J., Luehr, H., Macmillan, S., Mandea, M., McLean, S., Manoj, C., Menvielle, M., Michaelis, I., Olsen, N., Rauberg, J., Rother, M., Sabaka, T.J., Tangborn, A., Toffner-Clausen, L., Thebault, E., Thomson, A.W.P., Wardinski, I., Wei, Z., Zvereva, T.I. & Wo, I.A.G.A., 2010. International geomagnetic reference field: the eleventh generation rid a-5747-2011 rid a-5627-2011 rid h-1822-2011 rid a-5670-2011. *Geophysical Journal International*, **183**(3), 1216–1230.

Fox, G.E. & Hope, W.H.S., 1894. Excavations on the site of the roman city at silchester, hants, in 1893. *Archaeologia (Second Series)*, **54**, 199–238.

Fulford, M., 2002. *A Guide to Silchester: The Roman Town of Calleva Atrebatum*. Calleva Trust, Stroud, U.K., second edition.

Fuller, B.D., 1967. *Two-dimensional frequency analysis and the design of grid operators*. Mining Geophysics Vol 2: Theory, Society of Exploration Geophysicists, Tulsa, Oklahoma, 1 edition, pp. 658–709.

Gaffney, C. & Gater, J., 2003. *Revealing the buried past: Geophysics for archaeologists*. Tempus Publishing, UK, 1 edition.

Geosoft, 2010a. *Montaj Grav/Mag Interpretation: Processing, Analysis and Visualization System for 3D Inversion of Potential Field Data for Oasis montaj v7.1. Tutorial and User Guide*. Geosoft Inc, Toronto, Canada.

Geosoft, 2010b. *Montaj Mapmap Filtering: 2D Frequency Domain Processing Of Potential Field Data. Tutorial and User Guide*. Geosoft Inc, Toronto, Canada.

Gerovska, D. & Arauzo-Bravo, M.J., 2003. Automatic interpretation of magnetic data based on euler deconvolution with unprescribed structural index. *Computers & Geosciences*, **29**(8), 949–960.

Gharibi, M. & Pedersen, L.B., 2000. Removal of dc power-line magnetic-field effects from airborne total magnetic-field measurements. *Geophysical Prospecting*, **48**(3), 617–627.

Godio, A. & Piro, S., 2005. Integrated data processing for archeological magnetic surveys. *Leading Edge (Tulsa, OK)*, **24**(11), 1138–1144.

Golub, G.H., Heath, M. & Wahba, G., 1979. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, **21**(2), 215–223.

Grant, F.S. & West, G.F., 1965. *Interpretation theory in applied geophysics*. McGraw-Hill, New York, U.S.A.

Grauch, V.J.S. & Cordell, L., 1987. Limitations of determining density or magnetic boundaries from the horizontal gradient of gravity or pseudogravity data. *Geophysics*, **52**(1), 118–121.

Haigh, J.G.B., 1992. *Automatic grid balancing in geophysical survey*. Computer applications and quantitative methods in archeology, British Archaeological Reports, United Kingdom, 1991 edition, pp. 191–196.

Hansen, P.C., 1992. Analysis of discrete ill-posed problems by means of the l-curve. *SIAM Review*, **34**(4), 561–580.

Hansen, P.C. & O'Leary, D.P., 1993. The use of the l-curve in the regularization of discrete ill-posed problems. *Siam Journal on Scientific Computing*, **14**(6), 1487–1503.

Hansen, R.O., Racic, L. & Grauch, V.J.S., 2005. *Magnetic methods in near-surface geophysics*. Near-surface geophysics, Society of Exploration Geophysicists, U.S.A., first edition, pp. 151–176.

Herwanger, J., Maurer, H., Green, A.G. & Leckebusch, J., 2000. 3-d inversions of magnetic gradiometer data in archaeological prospecting; possibilities and limitations. *Geophysics*, **65**(3), 849–860.

Hesse, A., Barba, L., Link, K. & Ortiz, A., 1997. A magnetic and electrical study of archaeological structures at loma alta, michoacan, mexico. *Archaeological Prospection*, **4**, 53–67.

Hill, I., Grossey, T. & Leech, C., 2004. High-resolution multisensor geophysical surveys for near-surface applications can be rapid and cost-effective. *Leading Edge (Tulsa, OK)*, **23**(7), 684–688.

Lagarias, J., Reeds, J., Wright, M. & Wright, P., 1998. Convergence properties of the nelder-mead simplex method in low dimensions. *Siam Journal on Optimization*, **9**(1), 112–147.

Lawson, C.L. & Hanson, R.J., 1974. *Solving Least Squares Problems*. Prentice-Hall Inc, New Jersey, USA, 1 edition.

Le Borgne, E., 1955. Susceptibilit magntique anormale du sol superficiel. *Annales de Geophysique*, **11**, 399–419.

Li, X., 2006. Understanding 3d analytic signal amplitude. *Geophysics*, **71**(2), L13–L16.

Li, Y. & Oldenburg, D.W., 1996. 3-d inversion of magnetic data. *Geophysics*, **61**(2), 394–408.

Linford, N., 2001. Silchester roman town, hampshire; report on ground penetrating radar survey, march 2000. Technical Report Centre for Archaeology Report Number 9/2001, English Heritage.

Linford, N., 2004. Magnetic ghosts: Mineral magnetic measurements on roman and anglo-saxon graves. *Archaeological Prospection*, **11**, 167–180.

Linford, N., 2006. The application of geophysical methods to archaeological prospection. *Reports On Progress In Physics*, **69**(7), 2205–2205–2257.

Linford, N., Linford, P., Martin, L. & Payne, A., 2007. Recent results from the english heritage caesium magnetometer system in comparison with recent fluxgate gradiometers. *Archaeological Prospection*, **14**, 151–151–166.

Linford, N., Linford, P., Martin, L. & Payne, A., 2010. Stepped frequency ground-penetrating radar survey with a mulit-element array antenna: Results from field application on archaeological sites. *Archaeological Prospection*, **17**, 187–193.

Maher, B.A. & Taylor, R.M., 1988. Formation of ultrafine-grained magnetite in soils. *Nature*, **336**, 368–371.

Milea, C.M., Hansen, R.O., Tsokas, G.N., Papazachos, C.B. & Tsourlos, P.I., 2010. Complex attributes of the magnetic signal for multiple sources: Application to signals from buried ditches. *Archaeological Prospection*, **17**(2), 89–101.

Miller, H.G. & Singh, V., 1994. Potential field tilt; a new concept for location of potential field sources. *Journal of Applied Geophysics*, **32**(2-3), 213–217.

Milsom, J. & Eriksen, A., 2011. *Field Geophysics*. Wiley-Blackwell, United Kingdom, fourth edition.

Murdie, R.E., Styles, P., Upton, P., Eardley, P. & Cassidy, N.J., 1999. Euler deconvolution methods used to determine the depth to archaeological features. *Geoarchaeology; exploration, environments, resources Geological Society Special Publications*, **165**, 35–40.

Nabighian, M.N., 1972. The analytic signal of two-dimensional magnetic bodies with polygonal cross-section; its properties and use for automated anomaly interpretation. *Geophysics*, **37**(3), 507–517.

Oldenburg, D.W., McGillivray, P.R. & Ellis, R.G., 1993. Generalized subspace methods for large-scale inverse problems. *Geophysical Journal International*, **114**(1), 12–20.

Oldenburg, D.W. & Li, Y., 2005. *Inversion for applied geophysics: A tutorial*. Near Surface Geophysics: SEG Investigations in Geophysics Series No. 13, SEG, Tulsa, USA, pp. 89–150.

Payne, A., 2009. Brading roman villa, isle of wight. report on geophysical surveys, march 1994, april 1995 and february 2009. Technical Report 104-2009, English Heritage.

Peters, L.J., 1949. The direct approach to magnetic interpretation and its practical application. *Geophysics*, **14**, 290–320.

Phillips, J.D., Hansen, R.O. & Blakely, R.J., 2007. The use of curvature in potential-field interpretation. *Exploration Geophysics*, **38**(2), 111–119.

Pilkington, M., 1997. 3-d magnetic imaging using conjugate gradients. *Geophysics*, **62**, 1132–1142.

Pilkington, M. & Roest, W.R., 1998. Removing varying directional trends in aeromagnetic data. *Geophysics*, **63**(2), 446–453.

Piro, S., Sambuelli, L., Godio, A. & Taormina, R., 2007. Beyond image analysis in processing archaeomagnetic geophysical data: case studies of chamber tombs with dromos. *Near Surface Geophysics*, **5**(6), 405–414.

Powell, D.W., 1967. Fitting observed profiles to a magnetized dyke or fault-step model. *Geophysical Prospecting*, **15**, 208–220.

Price, J.E. & Price, F.G.H., 1890. *A guide to the Roman villa recently discovered at Morton, between Sandown & Brading, Isle of Wight.* Briddon, United Kingdom, 14th edition.

Rao, D.B. & Babu, N.R., 1993. A fortran-77 computer program for three-dimensional inversion of magnetic anomalies resulting from multiple prismatic bodies. *Computers & Geosciences*, **19**(6), 781–801.

Ravat, D., 1996. Analysis of the euler method and its applicability in environmental magnetic investigations. *Journal of Environmental and Engineering Geophysics*, **1**(3), 229–238.

Reid, A.B., 1980. Aeromagnetic survey design. *Geophysics*, **45**(5), 973–976.

Reid, A.B., Allsop, J.M., Granser, H., Millett, A.J. & Somerton, I.W., 1990. Magnetic interpretation in 3 dimensions using euler deconvolution. *Geophysics*, **55**(1), 80–91.

Robinson, E.S., 1971. Use of poissons relation for extraction of pseudototal magnetic field intensity from gravity observations. *Geophysics*, **36**(3), 605–&.

Roest, W.R. & Pilkington, M., 1993. Identifying remanent magnetization effects in magnetic data. *Geophysics*, **58**(5), 653–659.

Roest, W.R., Verhoef, J. & Pilkington, M., 1992. Magnetic interpretation using the 3-d analytic signal. *Geophysics*, **57**(1), 116–125.

Royal Commission on the Historical Monuments of England, 1994. Calleva atrebatum, silchester roman town; archaeological survey, air photographic interpretation and transcript of the town and environs. Technical Report 1059939, Royal Commission on the Historical Monuments of England.

Salem, A. & Ravat, D., 2003. A combined analytic signal and euler method (an-eul) for automatic interpretation of magnetic data. *Geophysics*, **68**(6), 1952–1961.

Salem, A., Williams, S., Fairhead, J.D., Ravat, D. & Smith, R., 2007. Tilt-depth method: A simple depth estimation method using first-order magnetic derivatives. *Leading Edge (Tulsa, OK)*, **26**(12), 1502–1505.

Scollar, I., 1970. Fourier transform methods for the evaluation of magnetic maps. *Prospezioni Archeologiche*, **5**, 9–41.

Scollar, I., Tabbagh, A., Hesse, A. & Herzog, I., 1990. *Archaeological prospecting and remote sensing*. Cambridge University Press, Cambridge, U.K., first edition.

Scollar, I., Weidner, B. & Segeth, K., 1986. Display of archaeological magnetic data. *Geophysics*, **51**(3), 623–633.

Sheen, N.P., 1998. Automatic interpretation of archaeological gradiometer data, using a hybrid neural network.

Sheriff, S.D., MacDonald, D. & Dick, D., 2010. Decorrugation, edge detection, and modelling of total field magnetic observations from a historic town site, yellowstone national park, usa. *Archaeological Prospection*, **17**(1), 49–60.

Smith, R.S., Thurston, J.B., Dai, T.F. & MacLeod, I.N., 1998. ispi (super tm) ; the improved source parameter imaging method. *Geophysical Prospecting*, **46**(2), 141–151.

Spector, A. & Grant, F.S., 1970. Statistical models for interpreting aeromagnetic data. *Geophysics*, **35**(2), 293–302.

Tabbagh, A., Desvignes, G. & Dabas, M., 1997. Processing of z gradiometer magnetic data using linear transforms and analytic signal. *Archaeological Prospection*, **4**, 1–1–13.

Tabbagh, J., 2003. Total field magnetic prospection: Are vertical gradometer measurements preferable to single sensor survey? *Archaeological Prospection*, **10**, 75–81.

Tchernychev, M., 2009. *Magpick - Magnetic Map and Profile Processing. User Guide.* Geometrics Inc., San Jose, U.S.A.

Telford, W.M., Geldart, L.P. & Sheriff, R.E., 1990. *Applied geophysics*. Cambridge University Press, U.K., second edition.

Thompson, D.T., 1982. Euldph - a new technique for making computer-assisted depth estimates from magnetic data. *Geophysics*, **47**(1), 31–37.

Thomson, J.C., 1924. *A Great Free City: The Book of Silchester*. Simpkin, Marshall, Hamilton, Kent & Co. Ltd, London, volume 1 edition.

Thurston, J.B. & Smith, R.S., 1997. Automatic conversion of magnetic data to depth, dip, and susceptibility contrast using the spi (tm) method. *Geophysics*, **62**(3), 807–813.

Trott, K., 1999. A rescue excavation at the brading roman villa coach park, isle of wight. *Proceedings of the Hampshire Field Club & Archaeological Society*, **54**, 189–215.

UBC-GIF, 2005. Mag3d: A program library for forward modelling and inversion of magnetic data over 3d structures. version 4.0.

Ugalde, H. & Morris, W.A., 2010. Cluster analysis of euler deconvolution solutions: New filtering techniques and geologic strike determination. *Geophysics*, **75**(3), L61–L70.

Verduzco, B., Fairhead, J.D., Green, C.M. & MacKenzie, C., 2004. New insights into magnetic derivatives for structural mapping. *Leading Edge (Tulsa, OK)*, **23**(2), 116–119.

Weston, D.G., 2002. Soil and susceptibility: Aspects of thermally induced magnetism within the dynamic pedological system. *Archaeological Prospection*, **9**, 207–215.

Wijns, C., Perez, C. & Kowalczyk, P., 2005. Theta map; edge detection in magnetic data. *Geophysics*, **70**(4), L39–L39–L43.

Williams, N.C., 2008. *Geologically-constrained UBC-GIF gravity and magnetic inversions with examples from the Agnew-Wiluna greenstone belt, Western Australia*. Ph.D. thesis, University of British Columbia.