GEOLOGICA ULTRAIECTINA

Mededelingen van de
Faculteit Geowetenschappen
Universiteit Utrecht

No. 260

Implicit structural inversion for lithology using a gridded model

Arnout Tim van Zon

Promotoren:	Prof. dr. W. Spakman Prof. dr. J.C. Mondt*
Co-promotoren:	Dr. K. Roy Chowdhury Dr. P. Ditmar†

* Shell International Exploration and Production, Rijswijk, The Netherlands
† Delft Institute for Earth Observation and Space Systems, TU Delft, The Netherlands

The research described in this thesis was conducted under the programme of the Vening Meinesz Research School of Geodynamics, and was funded by the Geodynamics Research Institute (GOI, UU). It was carried out at:

Department of Earth Sciences
Utrecht University
Budapestlaan 4
3584 CD Utrecht
The Netherlands

http://www.geo.uu.nl/Research/ExplorGeophysics/

Members of the dissertation committee:

Prof. dr. ir. J. T. Fokkema	Civil Engineering and Geosciences Delft university of Technology, The Netherlands
Prof. dr. ing. habil. R. Klees	Aerospace Engineering, Delft university of Technology, The Netherlands
Prof. dr. J. Trampert	Department of Earth Sciences, Utrecht University
Dr J. W. Bredewout	Department of Earth Sciences (retired), Utrecht University

Printed by Wöhrmann Print Service, Zutphen

ISBN-10: 90-5744-124-1
ISBN-13: 978-90-5744-124-0

Implicit structural inversion for lithology using a gridded model

Impliciete structurele inversie voor gesteentesoort, gebruik makende van roostervormige modellen van de ondergrond

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor
aan de Universiteit Utrecht
op gezag van de rector magnificus, prof.dr. W.H. Gispen,
ingevolge het besluit van het college voor promoties
in het openbaar te verdedigen
op dinsdag 9 mei 2006 des ochtends te 10.30 uur

door

Arnout Tim van Zon

geboren op 21 juni 1976, te Amsterdam, Nederland

Promotoren:	Prof. dr. W. Spakman
	Prof. dr. J.C. Mondt
Co-promotoren:	Dr. K. Roy Chowdhury
	Dr. P. Ditmar

Contents

Chapter 1

Introduction

An important goal of geophysical exploration is the delineation of the rock structures present in the sub-surface. Lithologies are used to simplify the description of the sub-surface; a lithology is a classification of rock type, based on gross properties, like grain size, colour and mineral composition. This leads to various average physical properties of lithologies e.g. density, magnetic susceptibility, shear speed, bulk speed and electrical conductivity, which can be determined in the laboratory or in situ in bore-holes.

Measurements of these physical properties can, conversely, be used to characterize and/or differentiate lithologies in the earth. It should be noted though, that these properties may vary considerably within the same lithology found at different sites or at different depths. Furthermore, several rock types could have overlapping values for the same characteristic physical property, viz. density or susceptibility, thus making the assignment of the lithology, based purely on the former, difficult. However, if two lithologies have distinctly different values of some physical property, then geophysical techniques can be used to delineate the boundary between them. Examples are salt-structures (density contrast), magnetic ores (susceptibility contrast), aquifers and ore-bodies (resistivity contrast) etc.

The aim of this thesis is to develop methods to derive sharp images of the sub-surface, which show the geometry of the structures in the region of interest. In the region of interest, the sub-surface is taken to contain two lithologies with distinctive values of some physical property. In an exploration setting, such a class of models includes - among others - (heavy) ore-bodies and (light) salt structures. Throughout this study, the sub-surface is described by 3 dimensional (3-D) homogeneous, horizontal prisms, with a rectangular cross section in the vertical plane. These are used to construct 3-D and 2.5-D models.

The results of the research carried out are being presented in a number of chapters, all of which are intended for eventual publication. Chapter 2 has already been accepted and Chapters 3 & 4 have been submitted. This introduction is therefore primarily meant for putting the different chapters into context and not as an introduction for the entire research. Also, a certain amount of overlap is therefor unavoidable.

1.1 Gravity measurements

Masses attract each other due to gravitation. The gravitational attraction at the Earth's surface due to mass distributions inside the Earth is conveniently expressed as a volume integral over the Earth's density distribution. Internal density variations lead to gravitational anomalies. Such gravity accelerations are typically of the order of 10^{-4} m/s^2. In Earth Sciences, the commonly used unit is "Gal" (in honor of Galileo); the practical unit of 1 milliGal (mGal) equals 10^{-5} m/s^2.

Recently, Nabighian et al. (2005) published an historical review of using the gravity method in exploration. See also Telford et al. (1990). Observations of the vertical component of the gravitational acceleration - henceforth called gravity - either the absolute values or relative differences, can be carried out on, below or above the surface of the earth. Several instruments have been developed throughout history for this purpose. The oldest is the *pendulum*; its period of oscillation depends on the absolute gravity. Another (*free fall*) apparatus measures the time needed for a mass to drop through a fixed distance, which depends on absolute gravity. The *torsion balance* is a gravimeter with two masses on the ends of a bar. A horizontal gradient in the gravity field will cause a rotation of the bar.

A modern land gravimeter is the *spring gravimeter*, which is based on a mass attached to a metal or quartz spring. Either the displacement of the mass is measured (via a lever system) or the electrical force needed to maintain the position of the mass. These systems have also been modified for marine and airborne measurements. *Full Tensor Gradient* gravimeters have recently become available commercially. They measure gradients of the gravitational field to obtain the 5 independent components (out of 9) of the gravity gradient tensor.

Generally, measurements are corrected for the acceleration due to the Earth's rotation, ellipticity and height of observation. In exploration geophysics, other unwanted effects present in measurements (like Earth tides, topography) have also to be corrected for. Their influence on the measurements could be much larger than the signal of interest to the exploration geophysicist.

1.2 Seismic measurements

Seismic velocity is also an important diagnostic characteristic of lithology. In addition, it is also needed for the processing of reflection seismic data to obtain accurate depth images of the sub-surface. In practice, the term *slowness*, the reciprocal of velocity, is frequently used.

Besides direct measurements in laboratories and in bore-holes (sonic logs), information regarding wave speeds can also be obtained indirectly from surface seismic measurements. Measuring travel-times between sources and receivers situated in bore-holes is a technique for deriving dependable information regarding the slowness distribution.

Unlike gravitational attraction, which is linear with density, seismic travel-times are non-linearly related to the velocity field. Obtaining information about the latter from the travel-times thus presents extra problems.

1.3 Inversion

The goal of inversion in geophysics is to obtain information about (part of) the Earth, from measurements made outside it. It is well known that these inverse problems can lead to non-unique solutions, or solutions with large variances, notwithstanding the quality and spatial distribution of the observations; if that is the case, the problem is said to be ill-posed. Therefore, only a limited amount of knowledge regarding the sub-surface can be obtained confidently from the data. Much research has been dedicated to the topic of retrieving reliable information; several authors investigated the possibility to compute bounds on ill-posed problems, see for instance Parker (1974), Parker (1975), Sabatier (1977a), Sabatier (1977b), Matsu'ura and Hirata (1982), Tarantola and Valette (1982), Richard et al. (1984), Huestis (1996), Silva et al. (2002) and Barbosa et al. (2002).

Using a suitable discretisation (called *parameterisation*) of the modeled sub-surface is an important part of any strategy for geophysical inversion. A classic approach to parameterisation is to divide the region of interest in 3-D homogeneous prisms/cells, or a in a 2-D grid, infinite in the strike direction. The parameters to be estimated from the observations are the values of the chosen physical property (density, seismic velocities, electrical conductivity, etc) of each cell. The parameterisation with cells is an example of a parameterisation with local basis functions, for which the coefficients (the parameters) are sought. Parameterisation with non-local functions, like (spherical) harmonics is common in large-scale (global) inverse problems.

Earlier approaches to determine the model parameters often started with an initial guess. The model response was then calculated (forward calculation) and was visually compared with the observations. The geophysicist then changed the model (guided by experience) and iterated this process till a reasonable match was obtained.

This process has now-a-days largely given way to a computed update of the model, in order to fit the data adequately. The model parameters for which the data are most likely, given a noise distribution for the observations, is taken to be the solution. The different update algorithms lead to different inversion techniques. Menke (1984) and Press et al. (1992) explained the Maximum Likelihood method, which is often used for data corrupted with Gaussian noise.

For gravity data inversion, theoretically (even with noise-free data), a unique solution does not exist. The usual lack of data only complicates matters. In order to select a solution, constraints are imposed, which lead to, so called, *regularisation* of the inverse problem. The nature of these additional constraints and their physical/geological relevance, are areas of active research.

1.4 Smooth inversion & structural inversion

We use the phrase *smooth inversion* to denote schemes that use a gridded parameterisation and a rule that some spatial derivative of the parameters (the properties of the cells) is small, while fitting the observations properly. The resulting models will have a certain *smoothness* (Tikhonov and Arsenin (1977)).

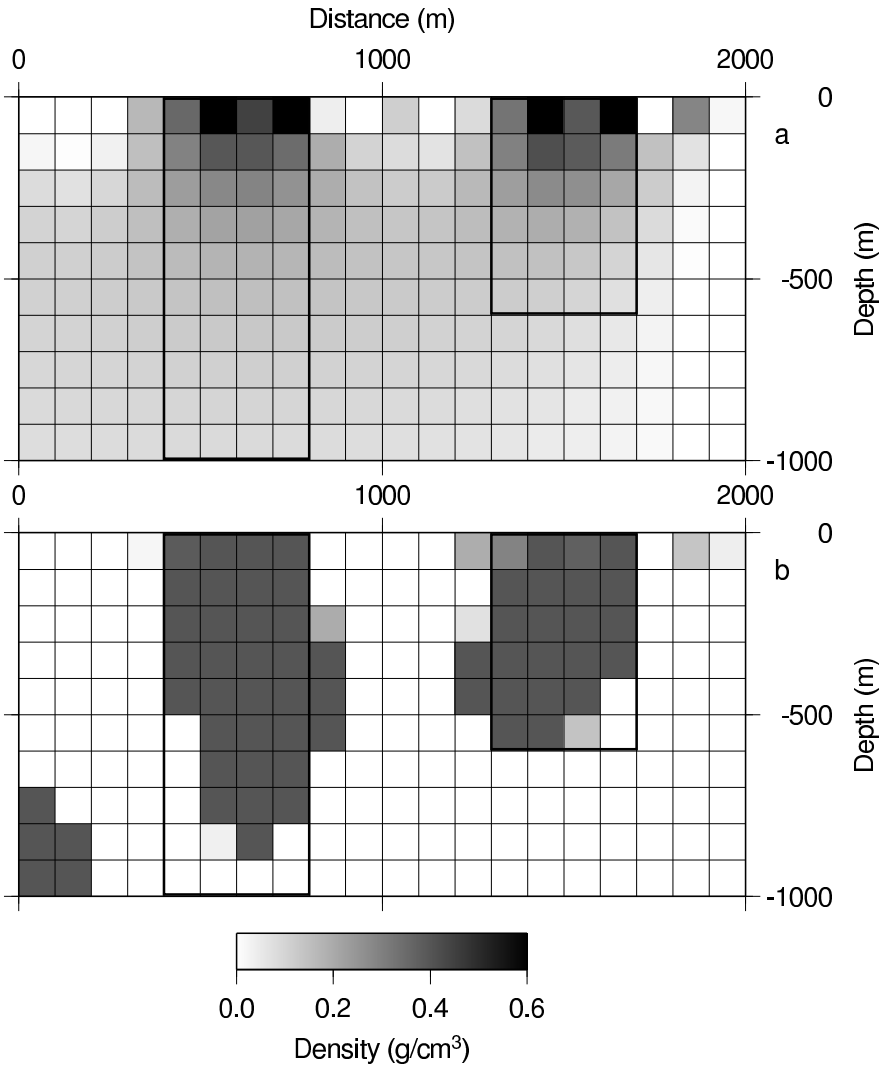

Figure 1.1: A smooth inversion result and a structural inversion result of a synthetic gravity data set, perturbed with 16% Gaussian noise. The smooth inversion result, panel (a) is obtained with the Truncated Singular Value Decomposition. Linear Programming is used for the structural inversion, shown in panel (b). The two source bodies have a density contrast of 0.4 g/cm^3 with respect to some background density (white). The outline of these two source bodies are shown with the bold black lines. Details of these methods are given in **Chapter 2**.

Smooth models were also found by Jackson (1972), who carried out the singular value decomposition (SVD) of the linear operator that relates the parameters to the observed data. The influence on the predicted data of a combination of parameters with a large Singular Value (SV) is large. It is relatively small for small SVs or none at all for zero SVs. A solution (model) can then be obtained by ignoring SVs below a chosen threshold. The latter are linear combinations of the model parameters, which are hardly influenced by the data (measurements) and can, in turn, be poorly determined from it: the case of zero-valued SVs corresponds to eigenvectors spanning the so-called null-space. Neglecting the relatively small SVs results in a robust solution with small variances for the parameters (depending on the threshold) and yields robust models.

An advantage of the smooth inversion is the ease of solving the problem; a gridded medium (consisting of many cells) is defined and the physical property value of each homogeneous cell is inverted for. With the Maximum Likelihood method, a *cost-function* is used whose minimum is taken to be the solution of the inverse problem. The cost function is usually the sum of two terms, one for the data misfit and another for regularisation (subjective).

Another advantage of smooth inversion is the possibility of using efficient algorithms, e.g., (sparse) conjugate gradient solvers or subspace methods (Oldenburg and Ellis (1993), Oldenburg et al. (1993), Brown and Saad (1990), Sasaki (2001)). Furthermore, a gridded model also allows arbitrary geometrical complexities to be represented. Figure 1.1a shows an example of a smooth inversion result of noisy synthetic gravity data.

The main problem with smooth inversion results are the smeared transitions between areas with different values of rock properties. Many economically interesting geological structures (ore-bodies, salt-flanks etc.) are not smooth spatially; on the contrary, they usually have quite distinct boundaries. Smooth inversion results are difficult to interpret in terms of sharp boundaries between the lithologies.

In contrast to smooth inversion, *explicit structural inversion* aims at finding the shape of a limited number of anomalous bodies, while fitting the observations adequately. Sometimes, the value of the physical rock property is inverted for as well. The outcome of an explicit structural inversion is a clear spatial distribution of the rock property and can easily be interpreted as a lithology image of the sub-surface.

An example of explicit structural inversion is the work of Jupp and Vozoff (1975). They inverted electrical resistivity data for the distribution of the conductivity of the sub-surface. They parameterized the sub-surface with 5 layers (infinitely long in the horizontal plane) and inverted for the thickness of each layer (and thus also its depth) and its conductivity.

Another example is the study by Ditmar (2002), in which the aim of the inversion was to find the shape of an isolated 2-D body of anomalous density (assumed known). The outline of the body is given by connected nodes. In an iterative scheme, the nodes are moved, in order to find the configuration that fits the data best. Starting with a seed, René (1999) and Camacho et al. (2000) found the shape of isolated bodies by iteratively adding cells of anomalous density to their external borders, until a satisfactory datafit was achieved.

Mundim et al. (1998) parameterized the anomalous volume with a number of vertical prisms, with a fixed and known density contrast. They inverted for the depth extent of the prisms. Like in many other cases of the explicit structural inversion, this is a non-linear problem, and they used a stochastic search method (generalized simulated annealing), to find the global minimum of the cost-function. Stochastic search methods are, however, time consuming and the inversion allows for only a few parameters. See also Gidas (1985), Tsallis and Stariolo (1996) and Nagihara and Hall (2001).

The explicit structural inversion techniques described above have certain drawbacks. Firstly, they require rather detailed a priori information about the medium. The geophysicist should at least know the number of anomalous bodies and their approximate locations. Secondly, manipulating a model where structural information is included explicitly may be difficult, especially in a 3-D situation. Lastly, computation of the Jacobian matrix (the matrix composed of the partial derivatives of a function with respect to the parameters) for models of this class may be difficult, especially if the computation is based on numerically solving differential equations, as happens, say, for the inversion of electrical conductivity data.

The drawbacks of explicit structural inversions might be overcome if the sub-surface is represented by means of a grid of suitably oriented rectangular prisms. Structural constraints are implemented in an inversion of this type in an implicit form. For example by allowing a high contrast in properties between neighboring cells, or "encouraging" parameter values close to the physical properties of specific lithologies. We call this approach *implicit structural inversion*. The aim of the implicit structural inversion is to fit the observations properly, while the resulting model should have clear structures, arising from the data and the structural constraints. This should make the interpretation of the result, in terms of delineation of lithologies, easier than with a smooth inversion. For example, compare the results of a smooth inversion (Figure 1.1a) with those of a structural inversion (Figure 1.1b).

Several authors, e.g., Last and Kubik (1983), Guillen and Menichetti (1984), Acar and Vogel (1994), Li and Oldenburg (1998), Zhdanov and Hursan (2000), Portniaguine and Zhdanov (2002), Farquharson and Oldenburg (2003) have used this approach for obtaining non-smooth models from different types of geophysical data. A further example of such an implicit structural inversion is the work of Bertete-Aguirre et al. (2002). They obtained non-smooth models with a stabilizing functional called 'total variation penalisation', which minimizes the L1-norm of the differences between cells. This regularisation function does not penalize sharp contrasts between neighboring geologic structures as much as traditional smoothing does, resulting in solutions with more contrast. A drawback of the technique is that the starting model has to be 'rather close' (difference less than 25%) to the true model, in order to find the correct minimum of the cost function. This extent of a priori knowledge may not be available.

Another example of implicit structural inversion, in this case of gravity data, was presented by Portniaguine and Zhdanov (1999). They minimized the area with non-zero spatial gradient in density (minimum gradient support). Additionally, they set a limit on the maximum and minimum density values to prevent geologically unrealistic models; they were able to obtain structured images.

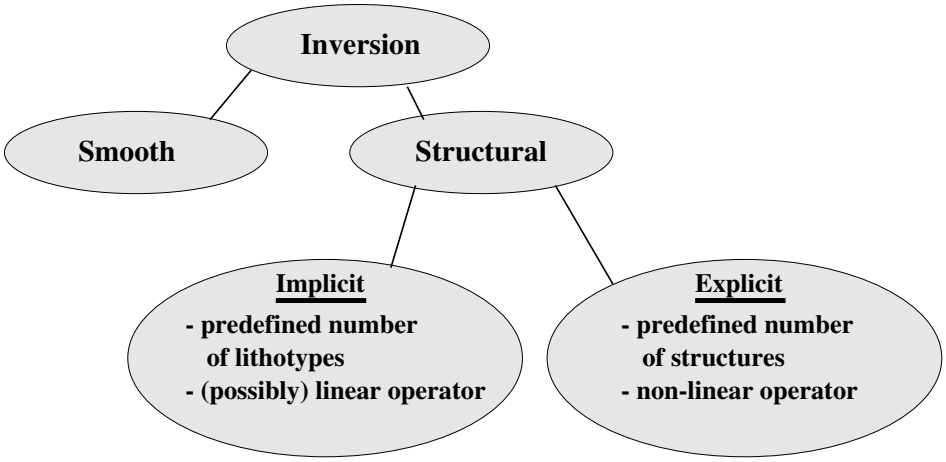

Figure 1.2: Our categorisation of inversion strategies as used in the field of geophysics. The structural inversion comprises implicit inversions and explicit inversions. The operator in an explicit structural inversion is typically non-linear, but not necessarily so.

The above methods use the L_2-norm, but structural inversion using the L_1-norm has also been investigated. Oldenburg and Ellis (1991) presented a scheme for an explicit structural inversion using Linear Programming (LP). LP is a mathematical technique used to optimize a linear objective function, while satisfying (linear) constraints on the variables. They applied LP to optimize the L_1-norm of the data misfit to solve a 1-D magnetotelluric inverse problem, utilizing an approximate inverse mapping. Within an iterative scheme, they used LP to calculate a smooth update for their layered conductivity model. Their objective function contains the absolute value of the data residuals and a weighted second term with the absolute value of the differences in conductivity between all neighboring layers. Obviously, the weight parameter must be chosen to correctly balance the data misfit with the size of the model update. Methods and applications for estimation of proper values for the weight parameter are, for instance, given by Golub et al. (1979), Hansen (1992),Kusche and Mayer-Guerr (2001), Kilmer and O'Leary (2001), O'Leary (2001), Kusche and Klees (2002).

In order to reduce some of the uncertainties, it may be desirable to perform a sequential inversion or a joint (structural) inversion; more than one independent set of observations could be inverted in sequence or together (Vozoff and Jupp (1975), Lines et al. (1988), Bosch (1999), Bosch (2001), Sasaki (1989), O'Brien et al. (2005), Berge et al. (1997)). Gravity data could - for example - be combined with first-arrival seismic traveltimes (Haber and Oldenburg (1997), Vernant et al. (2002)).

The above is summarized in Figure 1.2. Geophysical inversion methods can be divided in 2 large subgroups, namely smooth inversions (yielding smooth models of property values) and structural inversions (providing models with sharp boundaries), the latter

can be further subdivided into explicit (inversion for geometry) and implicit (inversion for property values) categories. An important advantage of the latter is the ease to delineate the structures in the model obtained.

1.5 Outline of the thesis

In this thesis, we propose 2 new approaches for implicit structural inversion of geophysical data. Both methods are formulated with their extension to joint inversion in mind. All the studies are carried out using a gridded cell parameterisation of the sub-surface, where the latter is described by homogeneous, horizontal 2.5-D or 3-D prisms with a rectangular cross section.

First, in **Chapter 2**, we investigate Linear Programming as a tool for implicit structural inversion of synthetic gravity data (corrupted with noise). Specifically, we use the L_1-norm of the data misfit to invert for density values associated with the cells of a gridded 2.5-D model of the sub-surface using constraints appropriate for a bi-modal lithological setting. Schematic models of salt diapirs and dikes are used for the study. The results are compared with those obtained using Truncated Singular Value Decomposition.

This LP technique is used in **Chapter 3** to invert a real data set. The regional-scale data set, courtesy of Shell, was acquired over a basin. Extensions to the method make it possible to work with the differences between observations, instead of the absolute value of the gravity (floating reference) and to find a residual linear trend in the data. Solving for a linear trend helps to prevent a leakage of the regional field into the result of the local inversion.

In **Chapter 4** , we explore another approach for implicit structural inversion. This technique is based on the L_2-norm of the data misfit and a priori information regarding the density difference between two lithologies assumed present in the volume of interest. This is implemented as a dynamic (iteratively updated) cost function using three terms: data misfit, smoothing and quadratic structural regularisation. Besides the gridded parameterisation, the only a priori information used is the difference in density between the background and the suspected anomalous zone.

Preliminary results from the application of the method developed in Chapter 2 to synthetic cross-well seismic first-arrival times to obtain structural images of seismic velocity are presented in **Chapter 5**. One problem is the non-linearity of seismic travel-time inversion, requiring an iterative approach, even when the velocity contrast is modest. Finally, joint inversion of gravity data and seismic data is also investigated briefly in this chapter. An assumption is made to couple the inversions: the spatial structure (variation) of the two causative properties (density and seismic pressure wave velocity (V_p)) is the same. This clearly demonstrates the improvement possible by joint inversion. Some suggestions for future research on trying to handle large velocity contrasts and 3 lithologies are provided.

Forward modelling forms an important part of geophysical inversion and the research reported here had necessitated the development of several software tools. An industry-sponsored investigation offered a chance to apply the software to perform a feasibility

study for investigating the internal structure of a blast furnace using micro gravity measurements. This is not geophysical inversion, but the results are of general interest.

The iron manufacturer Corus is interested in quantifying the wear of the heat-resistant walls of a blastfurnace with time. Furthermore, the thickness of layers of solid iron, which form on the bottom of the furnace due to incomplete drainage of liquid iron, is important information for estimating the time needed for repairs. Our aim was to investigate the possibility of resolving the resulting features using micro gravity measurements. We compared different data-acquisition geometries and carried out sensitivity analysis for the different parameters. This material (a technical report submitted to Corus) has been included as **Appendix A**.

Chapter 2

Structural Inversion of Gravity Data using Linear Programming

Structural inversion of gravity data - deriving robust images of the subsurface by delineating lithology boundaries using density contrasts - is an important goal in a range of exploration settings (e.g. ore bodies, salt flanks). Application of conventional inversion techniques in such cases, using L_2-norms and a smoothing regularisation, produces smooth results, and is thus sub-optimal. We investigate an L_1-norm based approach which yields structural images without the need for explicit regularisation. The density distribution of the subsurface is modeled with a uniform grid of cells. The density of each cell is inverted for by minimising the L_1-norm of the data misfit using Linear Programming (LP) while satisfying a priori density constraints. The estimate of the noise level in a given data set is used to qualitatively determine an appropriate parameterisation. The 2.5-D and 3-D synthetic tests reconstruct the structure of the test models adequately. The quality of inversion depends upon a good prior estimation of the minimum depth of the anomalous body. A comparison of our results with one using Truncated Singular Value Decomposition (TSVD) on a noisy synthetic data set favors the LP-based method.

There are two advantages in using Linear Programming for structural inversion of gravity data. Firstly, it offers a natural way to incorporate a priori information regarding the model parameters. Secondly, it produces sub-surface images with sharp boundaries (structure).

2.1 Introduction

An important aspect of geophysical exploration is the delineation of the lithologies present in the sub-surface. Contrasts in specific physical properties are used as discriminants for

[1]This chapter is a slightly modified version of a manuscript that has been accepted for publication in *Geophysics* (2006). Authors: Tim van Zon and Kabir Roy Chowdhury, Applied geophysics, Utrecht University

this purpose. The value of physical properties of rocks, e.g., density, magnetic susceptibility, acoustic velocities and electrical conductivity can be measured in the laboratory or in-situ in boreholes. These values can then be used to characterise and/or differentiate lithologies. These property values may, however, vary within the same lithology at different sites or at different depths. Different rock types can, moreover, have similar values (partially overlapping range of values) for, e.g., density or susceptibility, thus making the discrimination of the lithology based on these physical properties difficult. For lithologies with distinctly different physical property values, appropriate geophysical techniques can be used to delineate the boundaries among them. Some examples of such properties are density contrast (salt structures, heavy ore bodies), susceptibility contrast (magnetic ore bodies and dikes) and resistivity contrast (aquifers, ore bodies). Determining geological boundaries from geophysical measurements is a type of geophysical inversion.

A classic approach for L_2-norm based geophysical inversion, parameterises the region of interest with many homogeneous 2/3-D prisms/cells and estimates the value of the physical property (density, acoustic velocity, electrical conductivity, etc) of each cell. The solution is obtained using the Maximum Likelihood (ML) method (Menke 1984; Press et al. 1992), which finds the model for which the data are most likely, given a noise distribution for the latter.

Sometimes the estimates of the model parameters, given by the ML method, can have large variances; small perturbations in the data will cause large changes in the estimates. To reduce these variances, additional constraints are imposed on the model parameters. For example, a regularisation, often used in smooth inversions, is the rule that the parameters should be as small as possible with respect to some background (starting) model, or that some spatial derivative of the parameters should be small, i.e., the model should have a certain smoothness (Tikhonov regularisation, Tikhonov and Arsenin 1977). We use the phrase *smooth inversion* for such procedures that produce smooth images of the subsurface.

Jackson (1972) took another approach, using the Singular Value Decomposition (SVD) of the forward operator, which maps the values of the model parameters into data-points. The influence of a combination of parameters with a large Singular Value (SV) on the predicted data is large, relatively small for a small SV or none at all for a zero SV. A solution (model) can then be obtained by ignoring the relatively small SVs below a chosen threshold. The eigenvectors corresponding to the neglected SVs are linear combinations of the model parameters, which hardly influence the (predicted) data and can, in turn, be poorly determined from the observations; the case of zero-valued SVs corresponds to eigenvectors spanning the so-called null-space. Neglecting the relatively small SVs results in a robust solution with small variances for the parameters.

An advantage of the smooth inversion is the ease of solving the problem; a gridded medium (consisting of many homogeneous cells) is used and the physical property value of each cell is inverted for. With the maximum likelihood method, a quadratic *cost-function* is used. The latter is a function of the model parameters; the set of parameter values corresponding to the minimum of the cost-function is the solution of the problem. The cost function is usually the sum of two terms, a term for the data misfit (say, the L_2-norm) and a regularisation term. The minimum is efficiently found by (sparse)

conjugate gradient solvers. Furthermore, gridded models allow easy manipulation and representation of arbitrary complexities.

The main disadvantage of the results from the inversion procedures mentioned above are the smooth transitions between areas with different values of rock properties (i.e. smearing). Sharp boundaries in the real medium get smeared out in the reconstructions. Many economically interesting geological structures (ore bodies, salt-flanks, etc.) are not smooth spatially; they have relatively sharp boundaries. Hence, results of smooth inversions may be difficult to interpret in terms of distinct boundaries between lithologies.

An alternate approach - structural inversion (SI) - attempts to counter this disadvantage by 'encouraging' models with relatively sharp transitions. 'Explicit' structural inversion aims at explaining the observed data in terms of a limited number of anomalous bodies and finding their shape. Sometimes, the values of the rock properties are solved for as well. The outcome of an explicit SI is a clear spatial distribution of the rock property values which can easily be interpreted as a lithological image of the subsurface. See for instance Jupp and Vozoff (1975) or Ditmar (2002).

The 'explicit' SI techniques mentioned above have certain drawbacks too. Firstly, they require a rather detailed a priori information about the medium; at least the number of anomalous bodies and their approximate locations should be known. Secondly, manipulating a model with structural information explicitly included may be difficult, especially in a 3-D situation. Lastly, computation of the model response (the forward problem) and of the Jacobian for such models may also be difficult, especially if the computation is based on numerically solving differential equations, e.g. for the inversion of electrical conductivity data.

The drawbacks of the two approaches for inversion described above (smooth inversion and explicit structural inversion) could be overcome by combining them. In this approach, the subsurface is represented by a rectangular grid. Structural constraints can be implemented in a model of this type in an implicit form - by "encouraging" a high contrast in parameter values between neighboring cells. This should make the interpretation of the result, in terms of delineation of lithologies, easier than those obtained using smooth inversion.

Structural inversion can also be carried out 'implicitly'. Several authors, e.g., Last and Kubik (1983); Guillen and Menichetti (1984); Acar and Vogel (1994); Li and Oldenburg (1998); Zhdanov and Hursan (2000); Portniaguine and Zhdanov (2002); Farquharson and Oldenburg (2003) have used this approach for obtaining non-smooth models from different types of geophysical data. Portniaguine and Zhdanov (1999) used this approach for inversion of gravity data. They minimised the area with non-zero spatial gradient in density (minimum gradient support), while limiting the influence of large density contrasts to a constant value. Additionally, they set a limit on the maximum and minimum density values to prevent geologically unrealistic models. They were able to regularise a gravity inversion and obtain structured images. Li and Oldenburg (1998) devised an objective function in which an a priori model could be specified which the result should resemble. Furthermore, the level of smoothness of the result could be specified in all 3 dimensions. With the aid of a properly chosen weighting function, they obtained good results.

The above mentioned approaches minimised a quadratic function (L_2-norm) of the

data residuals. The use of other norms for SI has also been investigated. Oldenburg and Ellis (1991) presented a scheme for an explicit SI using Linear Programming (LP). They applied LP to a 1-D magnetotelluric inverse problem, utilising an approximate inverse mapping. Within an iterative scheme, they used LP to calculate a smooth update for their layered conductivity model. The objective function used by them contains two terms: the absolute value of the data residuals and the absolute value of the differences in conductivity between all neighboring layers. The relative importance of these terms is given by weighting coefficients; the latter must be chosen to properly balance the data misfit with the size of the model update.

Linear programming can also be used to perform an implicit SI. Esparza and Gómez-Treviño (1997) gave an example for the interpretation of electrical sounding data. They applied LP to invert for the resistivities of a large number of thin layers (a mathematical model). The aim was to obtain a non-smooth result with only a few different values for the resistivity, i.e., a few layers (a geological model).

Here, we propose a method for an implicit structural inversion of geophysical data in a two-lithology setting. We illustrate the technique using synthetic gravity data (with noise), but the approach can be used for other types of geophysical data as well (with a linear or linearisable forward operator). We invert for densities associated with the homogeneous cells of a gridded model of the sub-surface, using the Linear Programming to minimise the L_1-norm of the data residual with an assumed maximum density contrast.

2.2 Linear programming for geophysical inversion

2.2.1 Linear Programming: A brief introduction

Linear programming is a method to perform constrained optimisation, see for instance Brown and Brown (1992) or Press et al. (1992). Here, the constraints are linear (in)equalities in the variables and the function to optimise is also a linear function of the variables.

In general, a set of (in)equalities has no solution, one solution, or more than one solution. If the constraints on some variable, x, are mutually inconsistent, e.g., $x > 5$, and $x < 3$, then obviously no solution exists. However, usually there are many combinations of the variables which satisfy all the constraints. The set of all these combinations is called the *feasible set.*

A graphical example of a feasible set in two variables x and y is given in Figure 2.1, the latter also contains contours of an objective function $\phi = y - x/2$. Non-negativity ($x, y \geq 0$) restricts the set to the positive quadrant. The two solid lines represent some additional constraints given by $y \leq -0.5x + 5$ and $y \leq -3x + 15$. The feasible set is the shaded area; the model-space for x, y, for which all the constraints are satisfied. The dashed lines are contours of the objective function. The minimum of this objective function, satisfying all the constraints, is at $(5, 0)$, at a corner of the feasible set, with $\phi = -2.5$.

The location of the extremum (maximum or minimum) of the linear objective function is normally on a corner of the (convex) feasible set. In the special case that a constraint

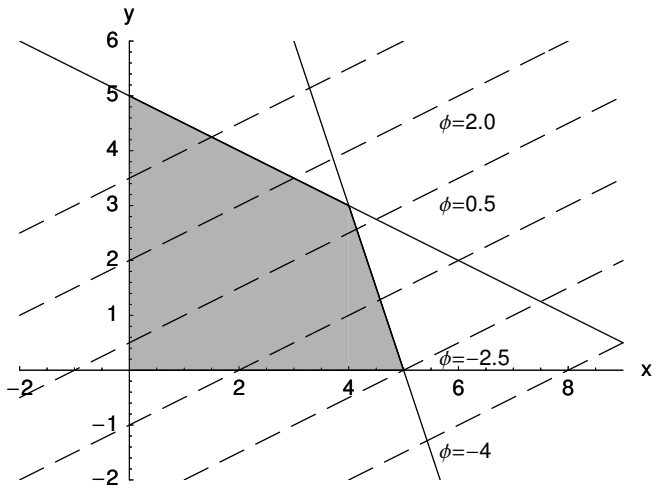

Figure 2.1: A 2-D example of a feasible set in two variables x and y, both of which are expected to be positive. The solid lines $y \leq -0.5x + 5$ and $y \leq -3x + 15$ represent some further constraints. The feasible set is the shaded area. The dashed lines are contours of the objective function $\phi = y - x/2$.

is parallel to the objective function, the extremum of the objective function can be a line segment or a (hyper) plane, which is then the model null space.

A linear programming problem can be solved with the Simplex method. The implementation used here is the one from Press et al. (1992).

2.2.2 Application of LP to gravity inversion

For the application of linear programming to exploration geophysics, an objective function needs to be defined and constraints need to be set. In the case of gravity anomalies, the model vector, $\mathbf{m}$, contains the density contrasts of the prisms with respect to some background value. The first set of constraints needed on the unknowns, m_j, the components of the model vector, $\mathbf{m}$, are the non-negativity constraints: $m_j \geq 0$. The use of additional constraints, and the choice of the function to be optimised, have been investigated by various authors, e.g., Safon et al. (1977) and Cuer and Bayer (1980). The former have used the constraints that the difference between predicted data from a trial density distribution and the actual measurements must be smaller than a certain error, ϵ, specified individually for each observation. Below, we take a closer look at their approach.

The sub-surface is parameterised with prisms, for which the density contrasts are sought. Nagy (1966) derived a formula (by performing the volume integration) for the vertical component of the gravitational attraction of a buried prism of unit density at any measuring station:

$$f = \mathcal{G} \left| \left| \left| x \ln(y+r) + y \ln(x+r) - z \arcsin\left(\frac{z^2+y^2+yr}{(y+r)\sqrt{y^2+z^2}} \right) \right|_{z=z_1}^{z=z_2} \right|_{y=y_1}^{y=y_2} \Bigg|_{x=x_1}^{x=x_2}, \tag{2.1}$$

$\mathcal{G}$ being the universal gravitation constant and $r = \sqrt{x^2+y^2+z^2}$. We follow the notation used by Nagy on the right hand side, i.e., $|f(t)|_{t=t_1}^{t=t_2}$ should be read as $(f(t_2) - f(t_1))$. The eight corners of the prism, relative to the measuring position, are given by (x_u, y_v, z_w), with $u, v, w = 1, 2$. The formula is valid for points outside or on the border of the prism. See Nagy (1966) for details.

Using equation 2.1 and writing G_{ij} as the contribution of the j-th prism at the i-th measurement station, the total gravity, d_i, can be written as

$$d_i = \sum_{j=1}^{M} G_{ij} m_j, \tag{2.2}$$

m_j being the densities of the prisms. This can be written in a matrix notation as

$$\mathbf{Gm} = \mathbf{d}. \tag{2.3}$$

The column vector $\mathbf{m}$, of length M, contains the densities of the prisms. The data are contained in the column vector $\mathbf{d}$, of length N. The $(N \times M)$ matrix $\mathbf{G}$ has elements G_{ij}.

The constraints of Safon et al. (1977) limit the residuals, $\mathbf{r}$. The residual for datum, i, is defined as $r_i = \sum_j G_{ij} m_j - d_i$. Writing all the constraints, $r_i \leq \epsilon_i$, in matrix form gives

$$\mathbf{Gm} \leq \mathbf{d} + \boldsymbol{\epsilon}, \tag{2.4}$$

where the vector $\boldsymbol{\epsilon}$ has positive components that need to be properly chosen, e.g., on the basis of the noise characteristics. The notation with the $\leq$ sign in this equation (and in the following), should be read as an inequality equation for each row.

Similarly, the "greater-than" constraint equations are given by:

$$\begin{aligned} \mathbf{Gm} &\geq \mathbf{d} - \boldsymbol{\epsilon} \qquad \text{and} \\ \mathbf{m} &\geq \mathbf{0}. \end{aligned} \tag{2.5}$$

Safon et al. (1977) used Linear Programming to optimise various moments of the densities of 2D prisms, to get lower and upper bounds for the total mass of the anomalous body, an estimate for the location of its center of mass and limits for its horizontal extent. This can be considered to be an explicit structural inversion.

The drawback of the approach of Safon et al. (1977) and of Esparza and Gómez-Treviño (1997) is that in equations (2.4) and (2.5) they do not allow deviations larger than ϵ between the forward modeling results and the measured data. The maximum deviation

(positive and negative) is fixed. Some of our tests showed that in the case of Gaussian noise, or worse, noise distributions with many outliers, this can cause incompatible constraints, i.e., there may be no feasible set and thus no solution.

Furthermore, the "inner workings" of the algorithm will cause some of the data-imposed "greater-than" type constraints be satisfied as equalities. The data will then not be fitted properly, but will be fitted as datum plus (or minus) the allowed deviation.

The objective functions used by Oldenburg and Ellis (1991) and Esparza and Gómez-Treviño (1997) were L_1-norm objective functions, which are not linear. These absolute value functions can be handled within the LP formalism by introducing some extra variables and constraints, as will be shown in the next section.

2.2.3 Implementing L_1-norm objective functions for LP

Like Oldenburg and Ellis (1991), we will minimise the absolute value of the data residuals; minimising the L_1-norm of the difference between the data and the forward modeling response. Furthermore, we set a maximum for the density difference between the background and the anomalous body (lithological contrast). This is implicit SI; the desired structural characteristics result from the fact that the solution of an LP-problem is always at a corner of the feasible set.

The L_1-norm objective function to be minimised is

$$\phi = \sum_{i=1}^{N} |\sum_{j=1}^{M} G_{ij} m_j - d_i|. \tag{2.6}$$

It may be noted that assumed measurement errors have been used here to modify the matrix $\mathbf{G}$ and the vector $\mathbf{d}$. This is implemented by dividing each of the N measurements by its assumed measurement error and by scaling the elements of the matrix $\mathbf{G}$ accordingly so as to produce dimensionless model responses. The resulting objective function is thus also dimensionless. The data vector, $\mathbf{d}$, now contains N dimensionless numbers. The vector, $\mathbf{m}$, is the model vector with M unknown components. The unknowns are the density contrast of each homogeneous prism relative to the background density, assumed to be known a priori. The constraints are:

$$\mathbf{0} \leq \mathbf{m} \leq \boldsymbol{\rho}_{max}, \tag{2.7}$$

with $\boldsymbol{\rho}_{max}$ the column vector with the positive upper limits for the density contrasts, assumed to be known a priori. In this study, all the components of $\boldsymbol{\rho}_{max}$ are the same, i.e. $0 \leq m_j \leq \rho_{max}$ (scalar).

The LP formalism can, in most cases, be used to optimise an L_1-norm objective function. This is done by introducing some new variables and by adding additional constraints. For this, we introduce two column vectors, $\mathbf{y}$ and $\mathbf{z}$, each of length N, as additional dummy variables; the LP problem is then to minimise the objective function

$\phi = \sum_{i=1}^{N}(y_i + z_i)$ with the constraints:

$$\begin{aligned} \mathbf{y} &\geq \mathbf{Gm} - \mathbf{d}, \\ \mathbf{z} &\geq \mathbf{d} - \mathbf{Gm}, \\ \mathbf{m}, \mathbf{y}, \mathbf{z} &\geq 0, \\ \mathbf{m} &\leq \boldsymbol{\rho_{max}}, \end{aligned} \tag{2.8}$$

This LP problem can be solved with the Simplex method, see Press et al. (1992).

In the following we will first investigate the problem of choosing a proper parameterisation (cell size) in a 2.5-D case. (With 2.5-D we mean that the length of the cells in the strike dimension is much larger than their width and height.) Then, we test the unregularised, L_1-norm based, inversion of gravity data for robustness against noise. Next, we compare the results of the inversions of the under-determined system of equations (2.3) using LP with those obtained using Truncated SVD (TSVD), which is well known to be robust against noise in the data. We will then test the robustness of the algorithm further by perturbing the data set with several noise simulations and performing the inversion of all perturbed data sets. The averages and standard deviations of the retrieved model parameters gives some insight into the well-resolved features of the model. Finally, we will apply our approach to a synthetic 3-D situation.

2.3 Results

2.3.1 Choice of parameterisation

In real applications the parameterisation cannot describe all the details of the true (unknown) earth. To make the method robust against noise, it is best to work with as few unknowns as possible. Care should also be taken to avoid fitting noise present in the real data during inversion; the aim should, instead, be to fit the data to the expected level of noise. We propose to start to parameterise with large cells and decrease the cell size until a satisfactory data misfit is achieved.

For the parameterisation test we used an earth model containing a dipping dike-shaped intrusion with a known *minimum* depth-to-the-top. There are no anomalies in this top layer, so that this layer does not need to be parameterised and its gravity effect can be corrected for in a pre-inversion step.

The width of the 2.5-D model is 2 km and the depth is 1 km. The length of the prismatic cells along strike (into the plane of the figure) is 1 km. The true earth is made up of 40 by 20 prisms, each of 30x30 m section. (Note, that the parameterisation of the earth model does not fill the area for which the density distribution is sought.) The data set consists of 200 measurements, 2 m above the surface, with 10 m separation between each location, the leftmost station is located 10 m to the left of the earth model. The data have been scaled with an assumed measurement error, $\theta = 0.04$ mGal, to obtain dimensionless numbers. The scaled data were then corrupted with unit-variance, zero mean Gaussian noise. The maximum permissible density contrast (ρ_{max}) was set at the correct value of 0.4 g/cm^3. The earth model and the data are plotted in Figure 2.2a.

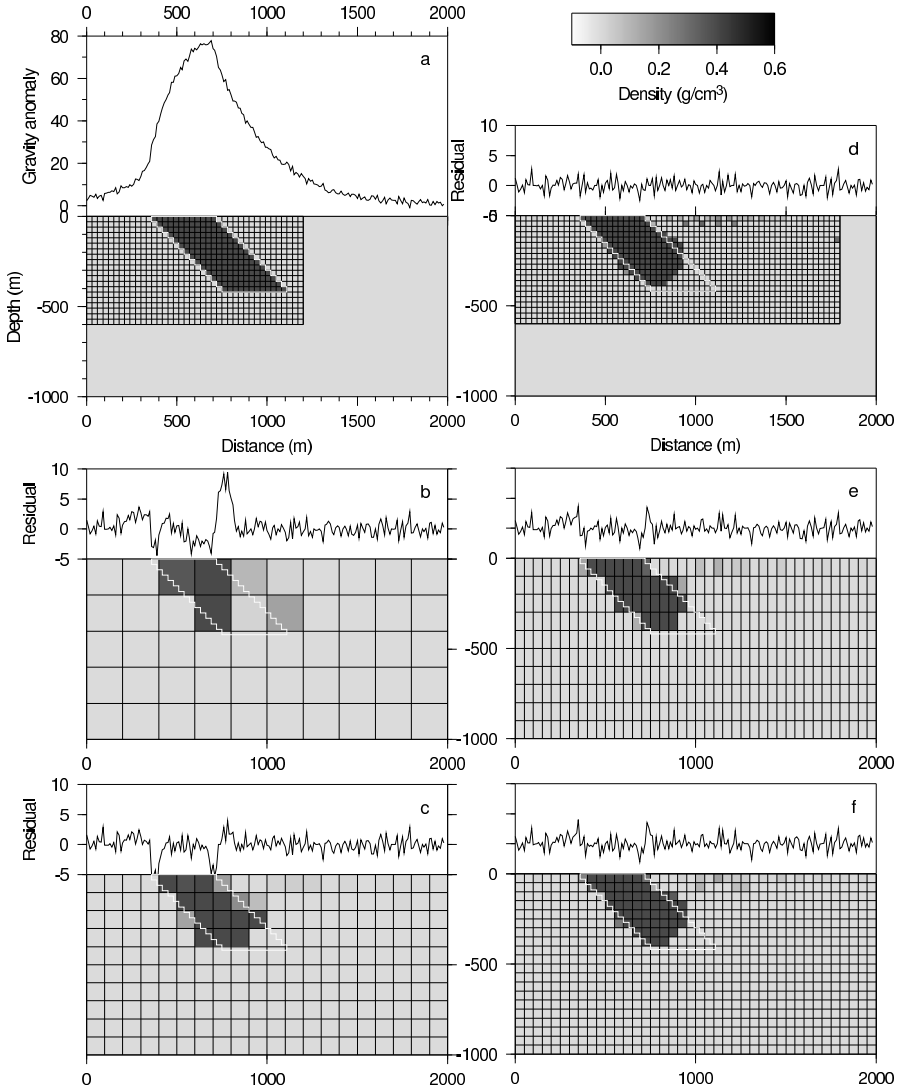

Figure 2.2: Effect of parameterisation on inversion. Panel (a) shows the (true) earth model used and the dimensionless noisy data. The scale bar denotes the density contrast in g/cm^3. The other 5 panels show typical inversion results, and the respective residuals, for different cell sizes used. The number of cells in the grid are: b (10x5), c (20x10), d (60x20), e (40x10) and f (40x20). Only the case (d) shows no apparent signal in the residuals.

This noisy data set was inverted several times with different parameterisations of the sub-surface using prismatic cells. Some of the results are shown in Figures 2.2b-f, along with the residuals. There is a clear signal remaining in the residuals if the parameterisation is too coarse. This effect is very pronounced in the case with only 50 (10 by 5) prisms each of 200x200 m cross section (Figure 2.2b), but also visible in the 400 (40 by 10) cell model with 50x100 m cross section, (Figure 2.2e) and even in the 800 (40 by 20) cell model with 50x50 m cross section (Figure 2.2f). This means that the data has features that cannot be modeled with these parameterisations. We also tested three parameterisations with 800, 1000 and 1200 prisms, each of 30x30 m section respectively. These were special, because the parameterisation was perfect. The inversion of the noisy data for the first two of these (not shown) showed edge-effects on the right, whereas no such effect is visible for the case with 1200 prisms (Figure 2.2d). The question now arises, as to which of the parameterisations reported in this paragraph is preferable. Our approach is not to be guided solely by the goodness of the data misfit, but also to take the noise level into account.

For this, we start with the noiseless synthetic data, given by: $\mathbf{d^{theor}} = \mathbf{G m^{true}}$, where $\mathbf{m^{true}}$ is the true causative body. If zero mean, unit variance Gaussian noise, $\mathbf{n}$, is added to this to obtain the noisy data, $\mathbf{d}$, the L_2-norm of the data residuals, corresponding to a model, $\mathbf{m}$, is

$$\sum_{i=1}^{N}(d_i - G_{ij}m_j)^2 = \sum_{i=1}^{N}(d_i^{theor} + n_i - G_{ij}m_j)^2. \tag{2.9}$$

Hence, even if the true earth model, $\mathbf{m^{true}}$, could be exactly recovered, the expected value of the L_2-norm of the data misfit for N data points would be

$$< \sum_{i=1}^{N}(d_i^{theor} + n_i - G_{ij}m_j^{true})^2 > \tag{2.10}$$

$$=< \sum_{i=1}^{N} n_i^2 > \quad = \quad N < n_i^2 >= N(= 200 \text{ in this case }).$$

This is because the expected value of n_i^2 is the variance, which is 1. In other words, this is the best one should aim for. If the data misfit for an earth model obtained is less than N, it is obvious that noise characteristics have partly been fitted. Similarly, the expected value of the L_1-norm of the data misfit is

$$< \sum_{i=1}^{N} |d_i - G_{ij}m_j| >= N < |n_i| > . \tag{2.11}$$

With the definition of the expected value $< x >= \int xP(x)dx$, with $P(x)$ the probability density function, we obtain

$$N < |n_i| >= N \int |n_i| \frac{1}{\sqrt{2\pi}} e^{-0.5(n_i)^2} dn_i = \frac{2N}{\sqrt{2\pi}} \approx 160. \tag{2.12}$$

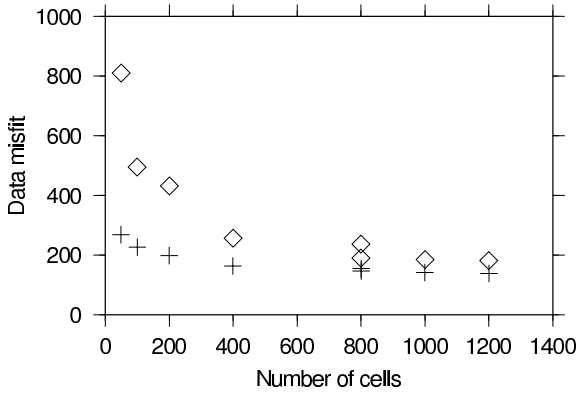

Figure 2.3: Comparing L_1-norm data misfit (plusses) and L_2-norm data misfit (diamonds) resulting from inverting a noisy data set with 200 measurements, as a function of the number of cells in the models. There were 2 models with 40 by 20 cells in this test; the first one with the cell-size 50 by 50 m (inversion result 2.2f), representing a 2000 by 1000 m subsurface (upper diamond). The second one had 30 by 30 m cells (not shown in Figure 2.2) representing a 1200 by 600 m subsurface (lower diamond). Note, that this parameterisation fits the true model perfectly.

The L_2 and the L_1-norms of the data misfits of several tests, including all those shown in Figure 2.2, are summarised in Figure 2.3, as a function of the number of prisms used to parameterise the sub-surface. As expected, the data misfit decreases as the number of cells in the models increases.

The model with 400 cells (40 by 10), each 50x100 m (Figure 2.2e), gives an L_1-data misfit of 164, very close to the expected misfit, while the L_2-data misfit is still rather high at 256. The parameterisation with 800 cells (40 by 20), each 50x50 m (Figure 2.2f), produces an L_1-norm data misfit of 155 indicating a slight over-fit (expected optimum value: 160); the corresponding L_2-norm data misfit is 237, indicating an under-fit (expected optimum value: 200).

Combining the knowledge of the expected data misfit, the (strength of) remaining signal in the residuals, and the philosophy of using as few parameters as possible, we prefer the 400 prism parameterisation of the subsurface, using prisms with a cross section of 50x100 m for this data set. The corresponding inversion result is shown in Figure 2.2e.

2.3.2 Comparison with inversion using TSVD

For testing the robustness of inversion of noisy gravity data using the above approach, we use the earth model shown in Figure 2.4d, which schematically represents two diapirs. For the inversion, we use the same parameterisation as the earth model, i.e. the parameterisation fits the model perfectly. There are 200 cells in the model and only 100 measurements; the system of equations is thus under-determined.

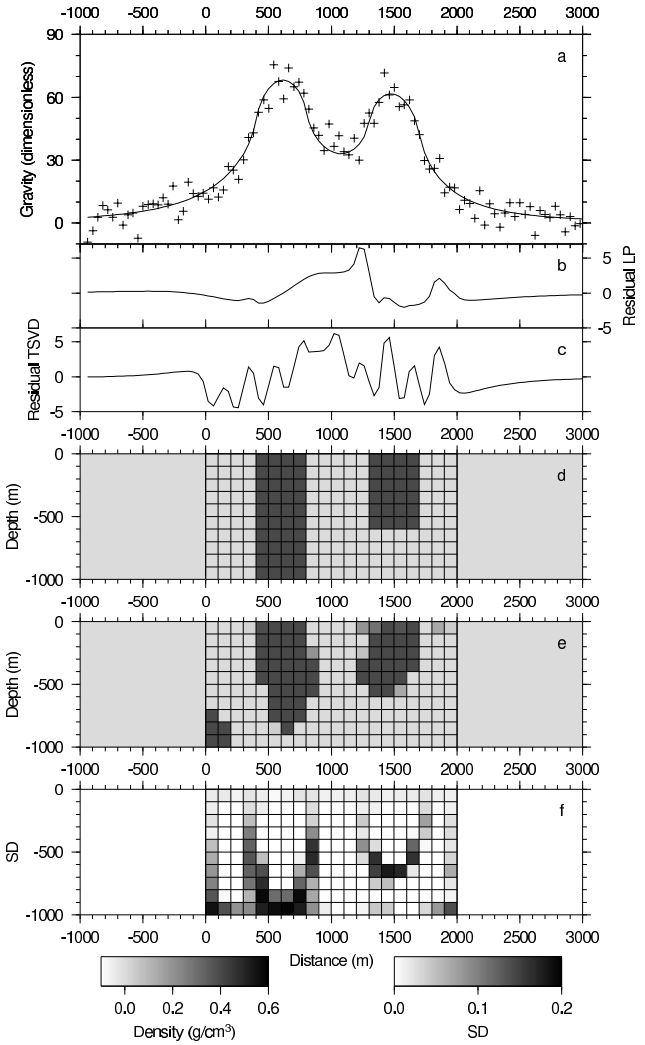

Figure 2.4: LP-based inversion results for a double diapir model (panel d). Panel (a) shows 100 noise free data (solid line) and their perturbed version (crosses), with 16% noise added, which is actually used for the inversions. The 2.5-D medium is parameterised with 20 by 10 cells of 100x100 m (panel (d)) that are 2 km long in the y-direction. Panel (e) shows the result of an LP-based inversion, with the dimensionless residuals (compared with the noise free data) shown in the panel (b). Panel (c) shows the dimensionless residuals (again noise free) from an inversion (Figure 1.1a) using TSVD. Panel (f) shows the standard deviation (SD) of the results of 100 LP-based inversions of the data each time perturbed with about 3% noise.

A remarkable result of our LP-based inversion scheme is that, if the correct ρ_{max} is used and the data are noise free, then the true earth model is *perfectly* retrieved (see discussion at the end). In the case of a well-posed system of equations with as many unknowns as measurements, this could also have been achieved with the Singular Value Decomposition (SVD) by including *all* singular values. (See Jackson (1972) or Press et al. (1992) for further discussion of SVD.)

It is well known, however, that constructing a model using singular values that are much smaller than the largest one will lead to unstable inversion results for real (noisy) data; the variances of the estimated parameters being proportional to the variance of the data and inversely proportional to the square of the smallest singular value. Hence, if very small SVs are included, a little noise can dramatically change the outcome of the inversion. To make the inversion more robust, only the larger singular values should be used to calculate the inverse. An appropriate cut-off needs, therefore, to be set - usually after some experimentation.

Below, we will compare the robustness of the LP-method with one using TSVD for an under-determined inversion problem, but first some definitions and explanations. We generate synthetic Gaussian noise with zero mean and unit variance and add this to the (previously scaled) measurement values. The signal-to-noise ratio $\mathcal{S}/\mathcal{N}$ is defined as

$$\mathcal{S}/\mathcal{N} = \sqrt{\frac{\sum_i^N d_i^2}{\sum_i^N n_i^2}}, \tag{2.13}$$

where d_i is the scaled i-th theoretical measurement and n_i is the i-th Gaussian noise sample. For generating a noisy data set with a higher level of noise, the noise is drawn from the same distribution, but its amplitude is multiplied.

The model misfit (as a percentage) is defined as:

$$\text{model misfit} = \frac{100\%}{M\rho_{an}} \sum_j^M |m_j - m_j^{true}|, \tag{2.14}$$

where $\mathbf{m^{true}}$ is the reference (true earth) model, $\mathbf{m}$ is the inversion result, M is the number of cells, and ρ_{an} is the expected density difference between the lithologies. This represents the average error (as a percentage) associated with the inverted densities for the cells. With this definition, if the density of most cells are either equal to zero or ρ_{an}, then the model misfit is approximately equal to the percentage of cells with the wrong density.

The cut-off value for the truncation of the singular values was chosen as a percentage of the largest one, and all eigenvectors with eigenvalues smaller than the cut-off value were excluded from the inversion. For the test, this cut-off value was chosen by inverting several synthetic data sets, corrupted with zero mean, unit variance Gaussian noise, each with a different cut-off for the singular values. We selected a cut-off of 5%, which yielded the best model misfit.

The implementation of Linear Programming used assumes that all variables are positive. It is therefore not possible to use negative values for the measurements. For this

reason we chose to change the negative gravity anomalies, that may have resulted after adding simulated noise, to small positive values.

The earth model as shown in Figure 2.4d was used in this test. In Figure 2.4a, the noise-free data are shown, together with the corrupted version (noise added). The $\mathcal{S}/\mathcal{N}$ ratio for this realisation is 6.3 (or approx. 16% noise). Figure 2.4e shows the inversion result with the linear programming method. Panels b and c show the residuals corresponding to the inversion results with LP and TSVD. As expected, the LP approach results in a structural model, the visual fit to the true earth model (Figure 2.4d) is to be put in the context of the rather large noise.

The inversion results of both methods are quite different. The result using LP shows 2 bodies with roughly the correct depth extent and a small, deep anomaly at the lower left border of the grid. The result using TSVD (Figure 1.1a) could also be interpreted to have two anomalous bodies, but the depth extent is very poorly resolved. TSVD concentrates the density anomaly at shallower levels. The data residual for the LP-based inversion is much smoother than that for the TSVD result, although the maximum difference with the noise free data is almost the same, i.e. the curves in Figures 2.4b and 2.4c have almost the same amplitude.

The above test was done with one realisation of a noise sequence. A different realisation would almost surely have given a different result. For a statistical comparison of the two approaches, the test needs to be repeated with many realisations of the simulated noise. For this, we first generate 100 noise sequences from a normal distribution with zero mean and unit variance (each sequence containing 100 samples to fit the length of our data). Five noise data sets (with different levels of noise) are then obtained by multiplying each sample of these 100 sequences respectively by the factors 1, 3, 5, 7 or 10. These five sets of noise were used to corrupt the theoretically computed data. This resulted in five groups of data sets corrupted by different levels of noise, each containing 100 data sets to be inverted. These groups are each categorised (labeled) with their respective average $\mathcal{S}/\mathcal{N}$ ratio ($\overline{\mathcal{S}/\mathcal{N}}$), and the corresponding standard deviation ($= \sqrt{\frac{1}{P}\sum_p \{(\mathcal{S}/\mathcal{N})_p - \overline{\mathcal{S}/\mathcal{N}}\}^2}$, the sum being over P, the number of data sets used). Note, that the $\mathcal{S}/\mathcal{N}$ ratios are computed *after* the statistical noise generation, *directly* from the actual realisations.

Each "noisy" data set was inverted separately, using truncated singular values and LP. Figure 2.5 summarises the results from these 1000 inversions (2 methods, each with 100 data sets corrupted with 5 different noise levels). The model misfits corresponding to the inversion results are plotted against the noise level of its data set.

This test shows that inversion using truncated singular values is relatively robust against noise; the standard deviation in model misfit is small. However, even the models resulting from data with little or no noise have a poor model misfit. A noise free data set was also inverted, with all the non-zero singular values included. This gave a better result, namely with a model misfit of 16.2%. The true model cannot be retrieved, because the system is under-determined (there will be SVs equal to zero) - there will be a null space. When a model is constructed, truncation is used to minimise the amount of null space in the solution. This has a smoothing effect on the resulting image.

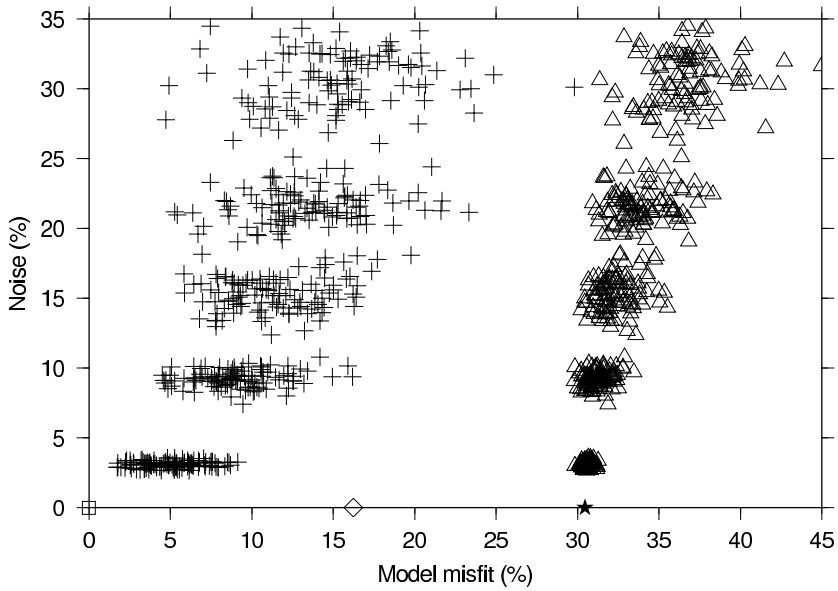

Figure 2.5: Statistical comparison of inversion results based on L_1-norm (using LP), and L_2-norm (using SVD). The vertical axis shows the noise level, 100%/($\mathcal{S}/\mathcal{N}$-ratio), in the perturbed data to be inverted, the horizontal axis the model misfit (Equation 2.14). The 500 LP inversions are marked with a plus-sign (left clusters), the 500 using truncated singular values (with the 5% cut-off) are marked with a triangle. The open square (0,0) is the perfect result with LP for noise free data, the diamond (16.2,0) is the SV inversion result with all non-zero SVs included in the inversion, and the filled star (30.4,0) is the truncated SV result, with the cut-off at 5%.

The LP performs better in this noise test. The model misfits are much better for all noise levels. The standard deviations are larger though. The noise free inversion yielded, as already mentioned before, the true earth model.

Instead of showing the results of each different noise realisation, we will summarise them with the average density and the associated standard deviation, SD, per cell (SD $= \sqrt{\frac{1}{P}\sum_p (m_p - \bar{m})^2}$, the summation being, once again, over the number of inversions). We calculated the average density per cell for the models with the noise level around 3% (not shown). The standard deviation of these resulting models is shown in Figure 2.4f. The LP inversions show variations around the edges of the bodies and some at the border of the grid. The good spatial resolution and the poor depth resolution, inherently present for the gravitational problem, are reflected in this figure.

2.3.3 Robustness against noise

Geophysical data is always contaminated with noise. One normally tries to remove (minimise) correlated noise by some type of preprocessing, but spatially or temporally incoherent noise gets inverted along with the true-earth response. An inversion scheme, therefore, has to be robust with respect to noise.

We used the dike model as in the parameterisation test (Figure 2.2a). The corresponding 'contaminated' raw data (at 200 surface locations), also shown in the same figure, is used for the inversion. In order to study which parts of the inverted model are robust, we would like to investigate if and where the inversion result is affected if the 'raw' data are perturbed. Also, we would like to investigate the effect of parameterisation (cell size) on the inversion of the 'perturbed' data.

In ten separate experiments we added zero mean, unit variance Gaussian noise to the 'raw' data. These data sets were inverted for the two best parameterisations of the subsurface obtained from the earlier test (Figure 2.2e and f). Finally, some statistics were derived from the results: the mean density of each cell and the standard deviation of the densities per cell. The results are summarised in Figure 2.6.

Panels (a) and (b) show the average models recovered for the two parameterisations and their theoretical responses. Panels (c) and (d) show the standard deviations associated with the cells of the average model. In general, the result is robust against small perturbations of the data; only a small number of cells have a non-zero standard deviations, which means that most cells ended up with the same density after each inversion.

To be more precise, the interior of the body seems to be well determined from the data except for the deepest part on the right, which is not imaged at all. All large standard deviations are found near the edges of the body and/or near the upper part of the subsurface. The 800 cell model has more ambiguous cells near the upper part of the model. This is a further confirmation that the choice of the 400 cell parameterisation is preferable, and validates our approach to determine the optimum cell size for inversion as outlined earlier. Note that the two inversions from the parameterisation test (Figure 2.2e and 2.2f) resemble these averaged models very well.

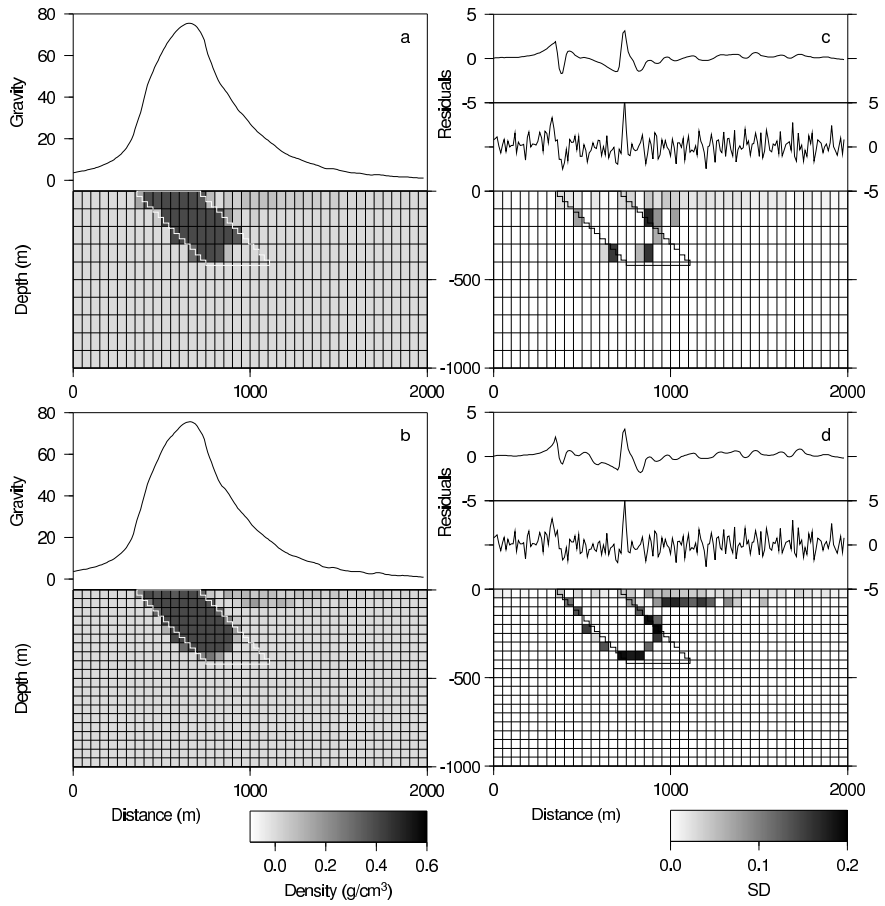

Figure 2.6: Robustness of LP-based scheme against noise using different parameterisations. The true model (also shown in Figure 2.2a) is outlined in all the panels. Figure (a) shows the average density of each cell, obtained from 10 inversions using perturbed data sets and the parameterisation of Figure 2.2e; Figure (b) shows the same with the parameterisation used in Figure 2.2f. The upper part of both these two Figures show the response of the average models. The gravity data, the responses and the residuals are dimensionless. The grayscale for the average density is shown on the lower left. Figures (c) and (d) show the standard deviations (SD) of the cell densities retrieved in the 10 inversions averaged in the panels (a) and (b). The corresponding grayscale is plotted in the lower right. The upper parts of the Figures (c) and (d) show the dimensionless residuals of the average model-responses with respect to the 'unknown' noise-free data (top) and the 'known' noisy data (bottom).

2.3.4 3-D example

Encouraged by the results reported above, we tested our approach for a more realistic 3-D situation. We used a dike-like 3-D body embedded for this purpose in a 2000 cell (20x10x10) subsurface, consisting of homogeneous cubic cells, with sides equal to 50 m. The anomalous density with respect to the background is 0.4 g/cm^3. The results are summarised in Figure 2.7. The shape of the actual anomaly is outlined in all the slices in white. The intrusion ranges from y=150, to y=350, corresponding to the slices (f) through (i) in Figure 2.7. Note that there is no obvious strike direction. This is also evident from the dimensionless "noisy" data shown in the panel (a).

A synthetic 2-D data set is generated at 800 locations on a regular horizontal 40 by 20 grid. After scaling and addition of noise, the dimensionless "noisy" data set is shown by means of contours in Figure 2.7a. The maximum (scaled) data value is 42 and the value of the $\mathcal{S}/\mathcal{N}$-ratio turns out to be 14.1. Again, note the absence of any obvious symmetry of the density distribution with respect to the parameterisation of the subsurface.

The problem of inverting for 2000 elements - given 800 data points - is clearly an under-determined one. The LP-based inversion is carried out with the maximum density contrast set at the correct value of 0.4 g/cm^3. The contoured data corresponding to the inversion result is shown in Figure 2.7b. The data misfit is 524, whereas theoretically we would expect 684 (see earlier discussion regarding the expected value of data misfit for noisy data). The data is thus slightly over-fitted. This inversion took about 14 minutes on a Pentium 4 processor.

The inversion has placed the anomalous mass at the correct lateral position in x and y, even though there was no favorable symmetry in the geometry of the 'acquisition' of the synthetic data. Very little mass is placed outside the slices containing the intrusion. The deeper part of the intrusion is - as before - not recovered properly. Some small density perturbations can be seen in the cells near the surface, above the lower part of the intrusion.

2.3.5 Limitation

Tests showed that the top of the source body does need to be known a priori. To illustrate this point clearly, two inversion results for a buried spherical body are shown in Figure 2.8. The extent and the parameterisation of the subsurface inverted for is shown with the cells. In one test, the anomalous body is actually buried in the middle of the parameterised area at an 'unknown' depth. The inversion result, as shown in Figure 2.8b, is completely wrong. Figure 2.8c shows the result when the top of the anomalous body coincides with the top of the parameterisation. The more the top of the source body is near the top of the parameterised area, the better the inversion result - this is to be expected, given the nature of the gravity field and the way the algorithm works.

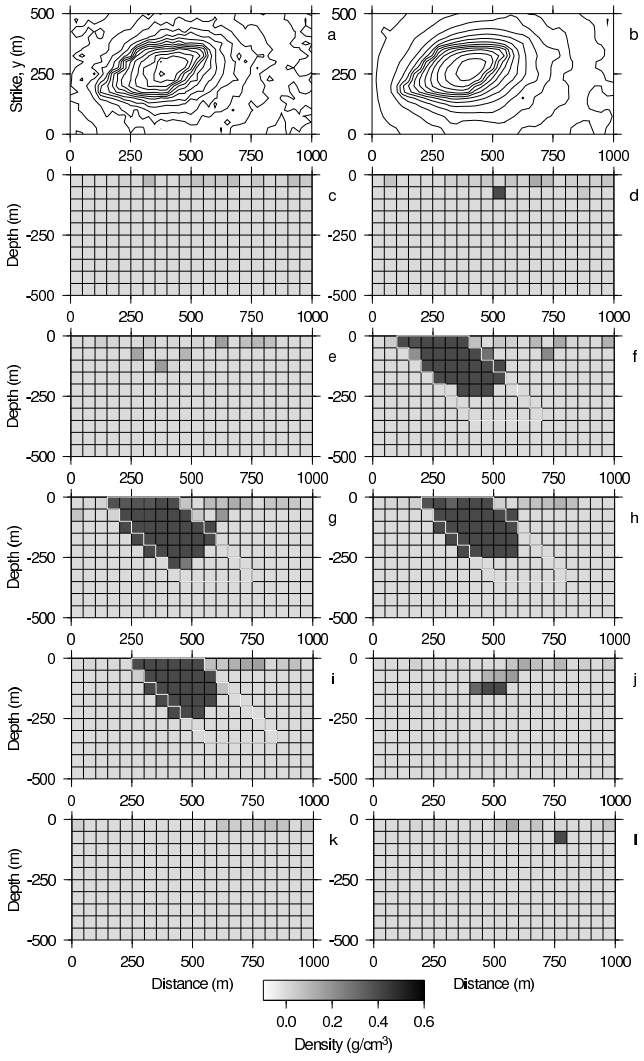

Figure 2.7: An example of a 3-D inversion. The sub-surface grid used for inversion consists of 2000 cubical cells, 50x50x50 m each. The realistic (dimensionless) data, computed at 800 stations in the x-y plane and corrupted with Gaussian noise ($\mathcal{S}/\mathcal{N}$-ratio=14.1), is shown in panel (a). The contours are drawn every 3 dimensionless units. The contoured (dimensionless) data corresponding to the inversion result is shown in panel (b). Panels (c) through (l) show vertical slices of the gridded model, obtained after LP-inversion, at $y = 25, y = 75, \ldots, y = 475$ m respectively. The asymmetrical dike-like source body is outlined in white in all slices.

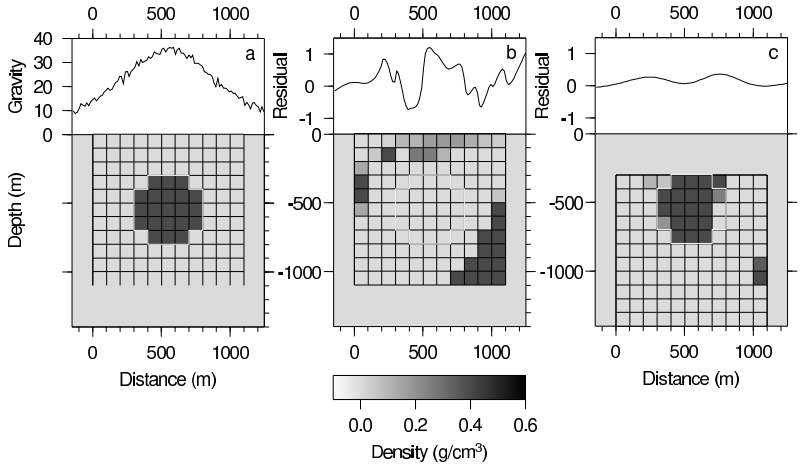

Figure 2.8: Importance of top-depth in LP-based inversion. Panel (a) shows the anomalous body and the corresponding dimensionless data (scaled and perturbed with noise). Panel (b) shows the inversion result when the top of the body is unknown and it is located in the middle of the parameterised area. The outline of anomalous body is plotted in white. Panel (c) shows the inversion result with the parameterisation chosen so that the top of the anomalous body coincides with the top of the parameterised area. The residuals in (b) and (c) are calculated with respect to the dimensionless noise free data.

2.4 Discussion

Perfect inversion results could be achieved with noise free data, even for an (originally) under-determined set of equations. This is due to three reasons: 1) the constraints on the density differences limit the model-space and add enough information to make this an over-determined problem, 2) the parameterisation used could perfectly represent the true earth, and, 3) the minimum of a linear cost function is always situated on a corner (or, in a degenerate case, a side, or a face) of the convex set of the allowed models in the model-space. The way the algorithm works prevents the search ending up in a local minimum somewhere inside the convex set. Such a minimum could exist due to noise in the data, in combination with an ill-posed inverse problem. Because the interior of the convex set of allowed models is not evaluated, possible local minima of the cost function over the domain of the model-space cannot be obtained as solutions. An implementation of the conjugate gradient algorithm could, on the other hand, get stuck in such a local minimum while minimising a quadratic cost function. The need to know the top of the anomalous body does, at first, appear to be a serious shortcoming of our approach. However, our focus is on delineating the structure laterally (sides of dikes, flanks of salt-structures, etc.); information regarding the tops of such structures may often be available from other geophysical techniques or from geological considerations. We conclude that the better the top of the source body is known, the better the inversion results are.

An advantage of our L_1-norm based inversion using LP is that a weighting function and many other tunable parameters are not needed. A weighting function is often used in other approaches to balance the influence of the individual parameters on a regularisation term. Secondly, if the maximum density contrast is set at (or close to) the true density contrast, our formalism does not need a smoothing term.

The only tunable parameter needed in the method presented here is the maximum density contrast, but even that need not be a hard constraint. A possible alternative approach would require an extra parameter to be added to the cost function. This extra parameter to be minimised would be smaller than the absolute boundary set by the maximum density contrast while larger than all other densities. In effect, this upper limit parameter (ρ_{ul}) enforces a flexible upper bound which is to be chosen as small as possible. In this case, the objective function would become $\phi = \beta\rho_{ul} + \sum_{i=1}^{N}(y_i + z_i)$, where β is a weight factor needed to balance the need to make ρ_{ul} small, while still fitting the data adequately. Of course the constraints $\rho_{ul} \geq 0$ and $\rho_{ul} \leq \rho_{max}$ need to be added. However in this case a trade-off value β between data misfit and importance of a small flexible upper limit must also be chosen; we have not investigated this further.

Other additions to the objective function could be a term to minimise or maximise the mass of the system, in a trade-off with the data misfit. The noise test, however, showed that this type of extra information is not necessary to obtain robust results.

2.5 Conclusions

The tests indicate that our L_1-norm based inversion-scheme using LP yields robust 'structural' 2.5-D models from a limited amount of 1-D data, even in the presence of substantial noise. One 2-D data set contaminated with noise also was inverted successfully for a 3-D model. This can be of interest in exploration settings like salt flanks, ore bodies, dikes etc. with sharp density contrasts.

The noise test indicates that the LP based approach is less robust against noise than the one using truncated SVs, but the performance of the former, in terms of model misfit, is much better.

Advantages of our LP-based scheme are the absence of a weighting function or other tunable parameters. Only the maximum density (contrast) is needed a priori. The results are, however, sensitive to a proper estimation of the minimum depth to the causative anomalous density.

A suitable parameterisation can be found by trial and error using the datafit. The criteria being the size of the residuals given the noise distribution, and the absence of large, spatially correlated, patterns in the former. Furthermore, the robustness of the result can be assessed by perturbing the data set several times and comparing the inversion results. A suitable parameterisation should yield robust inversion results.

2.6 Acknowledgments

The manuscript has benefitted from the comments/suggestions made by the reviewers. This research was made possible by the award of a Ph.D. position to van Zon by the Geodynamic Research Institute, Utrecht University, The Netherlands.

Chapter 3

Structural Inversion of Gravity Data using Linear Programming: a field data example

Regional scale gravity data has been inverted to infer the structure (topography) of the top of the basement. The 2.5-D volume of interest is parameterised with homogeneous horizontal prisms and a two-lithology medium is assumed. Linear Programming is used to minimise the L_1-norm of the data misfit, relative to a floating reference.

Given a known density contrast between the lithologies, an inversion using Linear Programming has the intrinsic advantage that a structural model is retrieved instead of a smooth one. The model recovered is almost bi-modal and its general features seem to be robust with respect to the parameterisation scenarios investigated. The floating reference and a linear trend in the data were also retrieved simultaneously. The results are comparable to those obtained earlier based upon detailed 2-D forward modeling (subjective) using many narrow, near-vertical prisms.

3.1 Introduction

An important goal of geophysical exploration is the delineation of the lithologies present in the sub-surface. Contrasts in specific physical properties are used as discriminants for this purpose. The value of physical properties of rocks, e.g., density, magnetic susceptibility, wave speed, and conductivity can be measured in the laboratory or in-situ in boreholes. These values can then be used to characterise and/or differentiate lithologies.

The more distinct the lithologies are with respect to some physical property, the easier it is for geophysical techniques based upon the latter to delineate the boundary between

[1]This chapter is a slightly modified version of a manuscript that has been submitted for publication in *Geophysics*. Authors: Tim van Zon and Kabir Roy Chowdhury, Applied geophysics, Utrecht University

the former. Examples of such physical properties - and associated geological situations - include density contrasts (light salt structures, heavy ore bodies), susceptibility contrasts (magnetic ores) and resistivity contrasts (aquifers, ore bodies). In Chapter 2, we have developed a method to obtain "geologically reasonable" models from surface measurements of gravity, and have demonstrated its effectiveness using numerical simulations of "diapirs" and 2.5/3-D "inclined dikes".

In this paper we extend this technique to make it applicable for an implicit structural inversion of geophysical data in a field situation; as before, we deal with an anomalous mass-distribution embedded within a homogeneous background, i.e. a two-lithology setting. A regional-scale, marine gravity data set, provided by the industry, will be used for this purpose. An interpretation of this data, taking other available information also into account, was made available to us and provides a benchmark for comparison

We group under *structural inversion* procedures that result in images of the subsurface that are easily interpretable in terms of lithology. Many authors have performed such an inversion by solving for the thicknesses of layers and/or the lateral change in depth of the layer interface, e.g., Pilkington and Crossley (1986), Guspi (1993), Smith et al. (1999), Auken and Christiansen (2004) and de Groot-Hedlin and Constable (2004).

Another approach is to parameterise the region of interest with vertical prisms of constant density, whose depth extend is solved for. This has been performed by using either discrete depth steps, or continuous variations. See for instance Mirzaei (1996), Barbosa et al. (1997), Mundim et al. (1998) Rao and Prakash (1997) and Gómez-Treviño et al. (2002). These are non-linear inverse problems, which require linearisation or stochastic search methods.

Assuming a 2.5-D situation, we use horizontal rectangular prisms, of finite extent in the strike direction, to represent the sub-surface structure. In the "dip" (cut-away) section, the prisms appear as rectangular cells, and the inverse problem to be solved consists of determining the densities associated with the cells of this gridded model. This is a linear problem. We use the L_1-norm for the data misfit with respect to a floating reference, and assume the a priori knowledge of the maximum density contrast. The problem is then solved with an implementation of Linear Programming (LP) as given by Press et al. (1992).

3.2 Methodology

Below, we highlight salient features of our approach; for details, reference may be made to Chapter 2.

Linear Programming (LP) is a method to perform constrained optimisation. Here, the constraints are linear (in)equalities in the variables and the function to be optimised is also a linear function of the variables. In general, a set of (in)equalities may have no solution, one solution, or more than one solution. If the constraints (which can be of 'greater-or-equal', 'less-or-equal', or 'equal' type) are mutually inconsistent like $c \geq 5$ and $c \leq 3$, then there is obviously no solution to the problem. Usually, however, there are many combinations of the unknown values, all of which satisfy the constraints. The set

of these combinations is called the *feasible set.*

For the application of LP to exploration geophysics, a suitable objective function needs to be defined and constraints need to be set. Given the linear constraints, and an objective function, LP determines the feasible set, and tries to find a solution that optimises the objective function. Positivity constraints are implied and the use of additional constraints and the choice of the function to optimise have been investigated by some authors. See for instance Sabatier (1977a), Sabatier (1977b), Safon et al. (1977), Cuer and Bayer (1980), Oldenburg and Ellis (1991), Huestis (1996) and Esparza and Gómez-Treviño (1997).

In our approach, the target sub-surface is parameterised with 2.5-D homogeneous horizontal prisms, with a rectangular cross-section in the dip direction, for which the density contrast is sought with respect to a homogeneous background medium. We use the L_1-norm of the data-fit, the difference between the observed data and the response of the model being investigated, as the objective function to be minimised. Advantages of our L_1-norm based inversion using LP include the absence of a weighting function, and of many other tunable parameters. A weighting function is often needed in other approaches to balance the influence of the individual parameters with a regularisation term, the latter providing for some desirable properties in the solution, e.g. having a certain smoothness, honoring other a priori information etc.

Most importantly, our goal is structural inversion, and as shown in Chapter 2, LP yields robust 'structural' 2.5-D models from a limited amount of 1-D data, even in the presence of substantial noise. This is achieved because 1) the constraints on the density differences limit the model-space (feasible set) and 2) the minimum of a linear cost function is always situated on a corner (or, in a degenerate case, a side, or a face) of the convex set of allowed models in the model-space. If the feasible set is limited in one direction by a density constraint on a parameter, then that parameter will be assigned the limiting value, provided that it corresponds to the optimal value of the cost function.

The cost function to be optimised depends upon the model (to be determined). The latter can be represented by a vector $\mathbf{m}$, containing the density contrasts of the (homogeneous) cells with respect to some background value.

With the gridded parameterisation of the sub-surface, as mentioned above, any datum, d_i, can be calculated from the following relation - linear in the density contrasts of the cells:

$$d_i = \sum_{j=1}^{M} G_{ij} m_j, \tag{3.1}$$

with M being the number of prisms and G_{ij} the vertical component of the gravitational attraction at the measurement station i, due to j-th rectangular prism of unit density. The density contrast of the j-th cell is denoted by m_j. Nagy (1966) derived an explicit formula for G_{ij} and we adapted it in Chapter 2.

3.2.1 Special considerations for real data

Floating reference

Gravity data is usually acquired as differences with respect to a base station. Thereafter, corrections for time-varying effects like instrument drift and tides are performed. Further, extraneous elevation- and spatial effects are removed as much as possible. The measurements are usually relative differences, the absolute value of gravitational attraction being unknown.

Thus the method of Chapter 2, which uses absolute gravity values, is not applicable to real data and has to be modified. By using a floating reference, i.e. allowing for an unknown offset (constant bias) in the data, one could (locally) invert for the relative gravity signal instead. If the floating reference is positive, this could be implemented directly within the LP-formalism.

Linear trend

After the standard pre-processing of the gravity data for temporal effects (tides, instrument drift) and spatial effects (free-air, terrain, Bouguer), a regional long wavelength trend is usually still present in the data. This trend can, for instance, be caused by non-horizontal structures deeper than the volume of interest, or dipping layers above it. In the case of a properly chosen - and pre-processed- regional scale data set, this long wavelength variation can - for the most part - be represented by a linear trend. If not accounted for, such a component of the input data will corrupt the inversion. To prevent such artifacts in the inversion result, a possible linear trend in the data should, therefore, also be included in the formalism and inverted for. Ideally, the direction of the linear trend should not be assumed to be known, before the inversion, especially if the trend were to be a weak one. This poses an extra problem because the Linear Programming algorithm used (Press et al. (1992)) requires the variables to be positive, so a possible negative slope cannot be implemented directly.

However, as Figure 3.1 shows, both the floating reference and the contribution due to the effective trend, can be incorporated by allowing for a possible "left trend" and a "right trend". The contribution of the left trend is $\alpha_l \delta_{l,i}$, where $\delta_{l,i}$ is the absolute value of the distance of station i from the leftmost station and α_l is the (positive) slope as defined in Figure 3.1. Similarly, the contribution of right trend is $\alpha_r \delta_{r,i}$, where $\delta_{r,i}$ is the absolute value of the distance of station i from the rightmost station and α_r is the (positive) slope. Thus, if $\alpha_l < \alpha_r$, (as in Figure 3.1) the effective trend is $(\alpha_r - \alpha_l)\delta_{r,i}$. The maximum correction for this trend is thus at the leftmost station. By inverting for the positive quantities α_l and α_r, the floating reference and the contribution to each individual station due to any effective linear trend can be found using LP.

The objective function originally used in Chapter 2 is the absolute value function:

$$\phi = \sum_{i=1}^{N} |d_i - \sum_j G_{ij} m_j|. \tag{3.2}$$

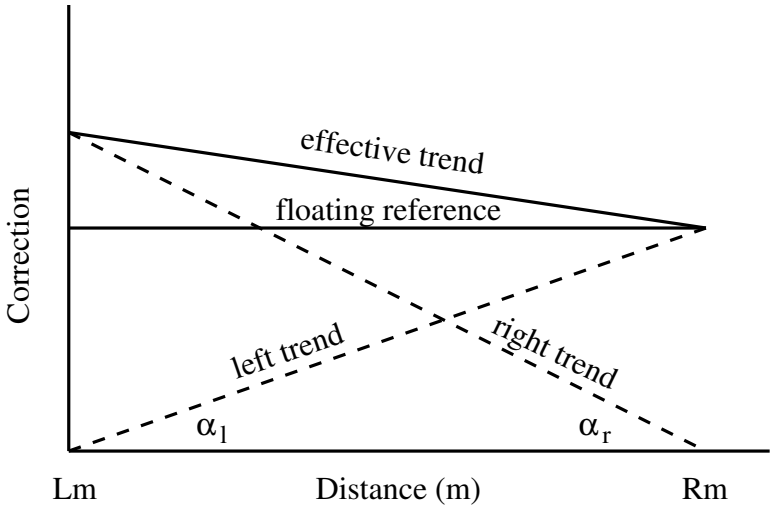

Figure 3.1: Synthesis of an arbitrary linear trend and the floating reference in the data. The left trend line and the right trend line are shown (dashed). Their respective slopes are α_l and α_r, both positive with respect to their respective sides. Adding both trend lines gives the *effective trend* line and the constant *floating reference*. *Lm* is the location of the left most station of the data set and *Rm* is the location of the right most station.

To invert the data (with respect to a floating reference) for the mass distribution and for a possible linear trend, we define

$$\mathcal{D}_i = d_i - (\alpha_l \delta_{l,i} + \alpha_r \delta_{r,i}), \tag{3.3}$$

and modify the original objective function (3.2) to

$$\phi = \sum_{i=1}^{N} |\mathcal{D}_i - \sum_j G_{ij} m_j|. \tag{3.4}$$

The $\mathcal{D}_i$'s are the values of the N measurements corrected for the effective trend (left and right combined) and the floating reference.

Note, in the Equations 3.2 to 3.4 all the quantities on the r.h.s. are dimensionless. This is achieved by dividing the N observations with an assumed measurement error. The elements of the forward operator $\mathbf{G}$ are also scaled accordingly. Thus, the objective function is dimensionless too.

3.2.2 Implementation

The objective function given in Equation 3.4, can be written in compact notation as:

$$\phi = \sum_{i=1}^{N} |\mathcal{D}_i - \sum_j G_{ij} m_j| = |\boldsymbol{\mathcal{D}} - \mathbf{Gm}|, \tag{3.5}$$

where $\mathbf{m}$ is the M-dimensional column vector of the unknowns, the homogeneous density contrast of the cells with respect to the background density, which is assumed to be known. The $(N \times M)$-matrix $\mathbf{G}$ is the forward operator.

As described in Chapter 2, we need two column vectors $\mathbf{y}$ and $\mathbf{z}$, both of length N, as additional dummy variables, to be able to solve the absolute value function (Eq. 3.5) with Linear Programming.

The LP problem is to minimise the objective function

$$\phi = \sum_{i=1}^{N} (y_i + z_i), \tag{3.6}$$

with the constraints

$$\begin{aligned} \mathbf{y} &\geq \mathbf{Gm} - \boldsymbol{\mathcal{D}}, \\ \mathbf{z} &\geq \boldsymbol{\mathcal{D}} - \mathbf{Gm}, \\ \mathbf{m}, \mathbf{y}, \mathbf{z} &\geq 0, \\ \mathbf{m} &\leq \boldsymbol{\rho_{max}}, \\ \alpha_l, \alpha_r &\geq 0. \end{aligned} \tag{3.7}$$

Most of the inequality relationships above are between column vectors; the inequalities should then be understood as relationships between corresponding elements (row) of the vectors on both sides. Note that the objective function in $\mathbf{y}$ and $\mathbf{z}$ is equivalent to Eq. 3.5, when ϕ is minimal. Note, that one of the relationships implements the constraint regarding the maximum allowable density contrast, ρ_{max}.

The formalism described in Chapter 2 has been modified to simultaneously invert for α_l and α_r and the densities of the cells m_j, by minimising the objective function given by Eq. 3.5. This LP-problem can be solved with the Simplex method as given by Press et al. (1992).

3.3 The gravity data and other information

The gravity data used in the present study has been provided by Shell and has been extracted from a regional data set. It consists of 61 gravity measurements, which have been corrected for time-varying effects. This data set is shown in Figure 3.2. It may be noted that the gravity values shown are with respect to some reference level (datum). This, while important in a large-scale context, does not directly affect the problem at hand; here we are interested in the gravity differences only.

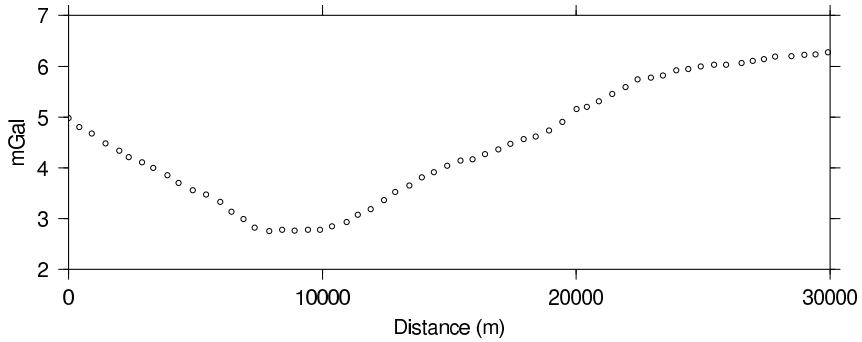

Figure 3.2: The gravity anomaly in mGal. (Courtesy of Shell, published with permission.)

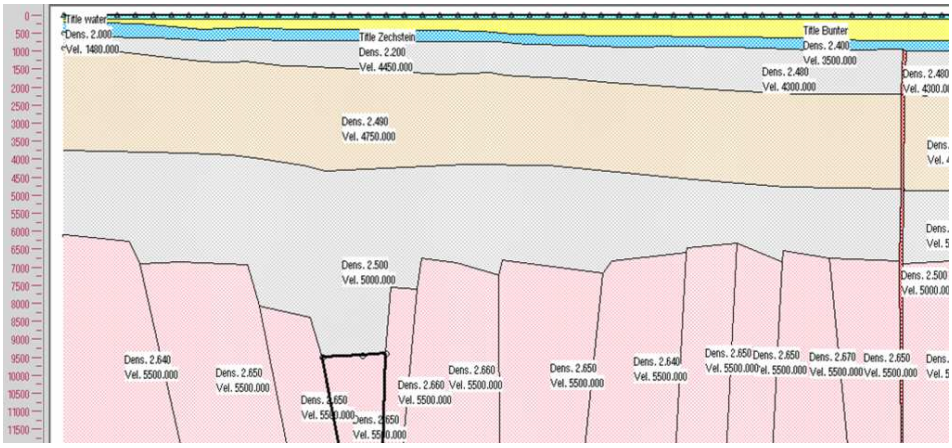

Figure 3.3: Modeling result. (Courtesy of Shell, published with permission.)

An important geological question in this region concerns the structure of the basement which has implications for the development of the overlying basins. The anomaly clearly indicates an asymmetrical mass deficiency in the sub-surface, the aim is to relate this to the basement topography. This information is expected to be useful for a better interpretation of the seismic reflection data recorded in this area.

Gravity data is - as is well known - not easy to interpret in terms of sub-horizontal structures; furthermore, in the context of inversion, the lateral resolution of the eventual presence of multiple basins is also a challenging problem. The data was earlier interpreted by detailed forward modeling using thin near vertical prismatic bodies; the results (Figure 3.3) indicated the presence of depressions in the basement. For this rather detailed analysis, inputs from additional sources (magnetics, seismics) were also used.

The earlier forward-modeling of the gravity and magnetic data indicated that the near-surface - above the basement - consisted of several horizontal to gently-dipping layers (also indicated in seismic section), and also provided an estimate of the minimum depth to the basement. We would like to carry out a structural inversion of this gravity data (anomaly) in terms of the topography of the basement, with a known minimum depth. We assume sub-horizontal layering above it and a two-lithology setting.

Before inversion, some choices needed to be made and the digitised data set needed to be appropriately manipulated to make it suitable for LP. These are summarised below:

- Information regarding measurement errors were not available to us. We have therefore chosen to scale the data with an assumed (constant) measurement error, σ (=0.1 mGal), to make the former dimensionless. Elements of the forward operator were also scaled similarly.

- The sub-surface is supposed to consist of a homogeneous background medium (half-space, strictly: laterally invariant), in which the 2.5-D anomalous volume with a rectangular cross-section is embedded. Because a 2-D description seems to be geologically appropriate here, a strike length of 90,000 m - (much) larger than the depth and width of the model - is assumed and the 3-D software is used.

  The density contrast is assumed to be 0.15 g/cm^3. This value is used as the maximum density contrast allowed within the LP formalism. The part of the background medium above the region of interest is not homogeneous, but consists of a few sub-horizontal layers. These dipping layers break the assumption of lateral invariance of density. To account for this, the inversion also solves for a linear trend in the data. The top of the basement, and thus the top of the region of interest, is assumed to be at a depth of 6,000 m.

- The inversion assumes that all anomalous mass causing the anomaly is within the anomalous volume (Figure 3.4g); the basement-layer outside is set to zero-density and the data is corrected accordingly.

These pre-processing steps involved in order to make the data suitable for the LP-based inversion are indicated in the Figure 3.4. This figure also shows the conceptual model to be used for the inversion. The input data is first modified by removing the effect

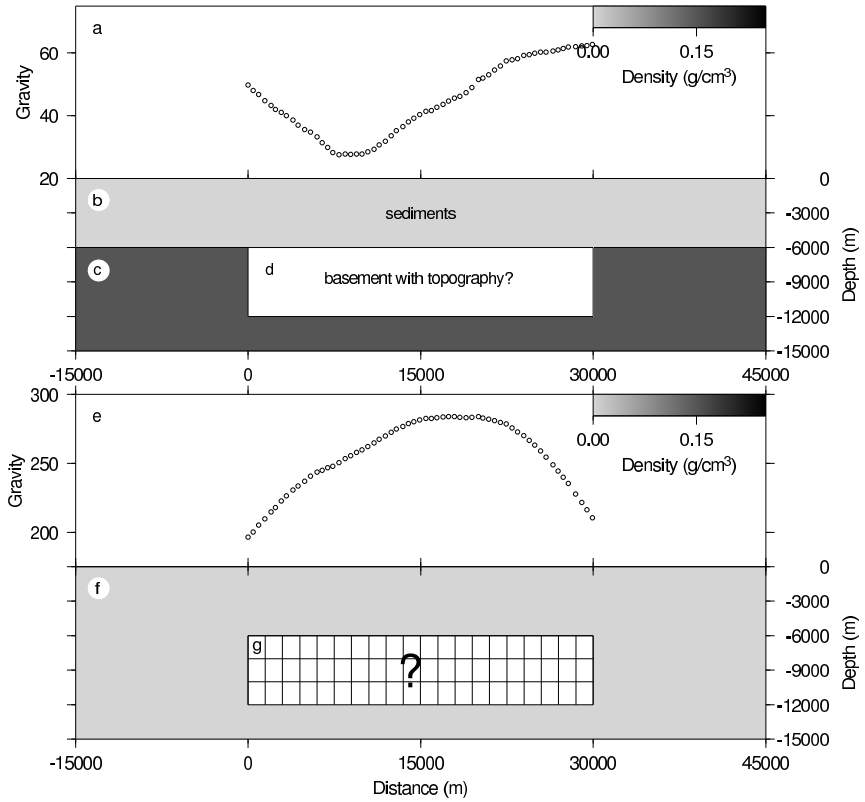

Figure 3.4: Processing the observed data, prior to LP-based inversion. Panel a shows the scaled (dimensionless) gravity data. Panels b, c and d show the conceptual model to be used for inversion. Panel b represents the overburden, assumed to be laterally infinite and sub-horizontal with a uniform background density. Panel c represents the basement (and deeper structures), assumed to be laterally infinite and with a known excess density relative to the overburden. Panel d is that part of c which will be parameterised, and inverted for, to explain the observed data. The values shown in panel a are modified by removing the effect of the region c (except for the region d) due to its excess density. Panel e shows the input data for the inversion, wherein the effects of the heavier basement (region c, except for the region d) have been removed. This anomaly needs to be explained by a density distribution in the area g, which is now embedded in the background medium (region f) of constant density.

of the basement except for target area (Figure 3.4c). This is done using the assumed density contrast of 0.15 g/cm^3 between the overburden and the basement. To ensure a positive floating reference, an arbitrary constant value is then added to each data point. The resulting positive anomaly (Figure 3.4e) is then to be explained in terms of density distribution in the (gridded) target area (Figure 3.4g) plus a (regional) linear trend that may be present in the data.

3.4 Results

3.4.1 Choice of parameterisation

As in every inverse problem, the number and the size of the cells to describe the target need first to be chosen. In Chapter 2, we described a method to select the optimal parameterisation. The approach is to use a parameterisation with as few cells as possible, while fitting the data adequately. This is based on the assumption that the noise in the data is random.

The data residual, r_i, for station i is the difference between the corrected, dimensionless data and the forward modeled data from the inversion result, corrected for the trend and the offset. Hence, using Equation 3.3,

$$r_i = \sum_j G_{ij} m_j - \mathcal{D}_i = \sum_j G_{ij} m_j + \alpha_l \delta_{l,i} + \alpha_r \delta_{r,i} - d_i. \tag{3.8}$$

A typical set of residuals is shown in Figure 3.5b. These residuals are not purely random noise, but seem to contain some signal. Due to the depth (6,000 m) to the top, even a severely over-parameterised model cannot fit the data. Thus, the parameterisation selection procedure described in Chapter 2, needs to be modified. As a sensible alternative, we suggest inverting for roughly as many cells, say 60, as data points (61).

Inversions with several parameterisations describing the volume of interest have been performed, from as few as 45 cells (15x3), up to as many as 180 cells (30x6). The results of these inversions are shown in Figure 3.6. They give some feeling for the robustness of the features in the inversion results. The results of these experiments - the L_1-data misfit, the floating reference, and the maximum of the effective trend - are also summarised in the columns 3,4 and 5 of Table 3.1.

Note, that figure 3.6 contains, among others, two inversion results with the preferred number of cells (60) viz. panel c (20x3) and panel g (15x4).

3.4.2 Additional constraint

Some of the models shown in Figure 3.6, contain 'heavier' cells overlying 'lighter' ones (local density inversion), especially on the left margin of the parameterised subsurface. This is geologically unreasonable in the context of the situation under study. Extra constraints on the variation of density with depth may, therefore, be considered. The Linear Programming approach is very convenient for applying such linear constraints. Thus we

Panel	# cells x × z	L_1 misfit	fl. ref.	max. trend	+L_1 misfit	+fl. ref.	+ max. trend
a	15 × 3	30.4	84.1	3.1	30.5	84.2	3.6
b	15 × 6	29.2	84.1	4.4	29.3	81.4	5.0
c	20 × 3	28.1	87.4	4.5	28.1	87.4	4.5
d	30 × 3	27.3	85.4	6.7	27.5	85.8	5.0
e	30 × 6	26.8	84.1	5.1	26.9	83.9	4.2
f	30 × 4	26.9	80.5	6.5	27.0	80.5	6.1
g	15 × 4	29.8	81.9	6.0	29.9	81.5	5.6

Table 3.1: Comparison of inversions with different parameterisations/constraints. Column 1 refers to the panels in the figures 3.6 and 3.7. Column 2 indicates the number of cells in the models shown in those figures in the lateral and vertical directions. The remaining pair of three columns each are all in dimensionless units. The columns marked with 'L_1 misfit' give the L_1-norm data misfit. Both the trend columns give the maximum correction, i.e. the correction for the leftmost station, because here the effective trend turns out to be a right trend. The rightmost three columns (marked with '+') were obtained using the extra constraint, that density may not decrease with depth (see section 3.4.2).

can easily implement an additional constraint that the density of each cell be greater or equal to its neighboring cell above.

The results of the inversions with these extra constraints on the density are shown in Figure 3.7. The corresponding data misfit, trend and offset are shown in the rightmost three columns of Table 3.1. The data misfit with the extra constraints is somewhat worse. This is to be expected, because now the inversion is more constrained - leading to a sub-optimal data-fit. This holds for all inversions, except for the model shown in Figure 3.6c and Figure 3.7c. These results are identical, because - fortuitously - the result in Figure 3.6c, without the explicit extra constraints, had no density inversion.

3.5 Discussion and Conclusions

The inversion results shown in Figures 3.6 and 3.7 all display two main features; the large depression of the basement at about 9,000 m and a smaller one at 17,000 m. These seem to be robust features, although the maximum depth varies somewhat with the parameterisation used. The larger depression of the basement was also noticed during the earlier forward-modeling by industry. The second - smaller - depression, however, was not found consistently.

A feature of all inversions is the near bi-modality of the resulting model. There are only a few cells that do not have either the basement density or the background (sediment) density. The inversion shows clear structures, and thus confirms the results from the earlier tests (Chapter 2).

All the inversion results in Figure 3.7 suffer from edge effects on both lateral borders of the gridded target area. Gravity effects of anomalous density distributions outside the

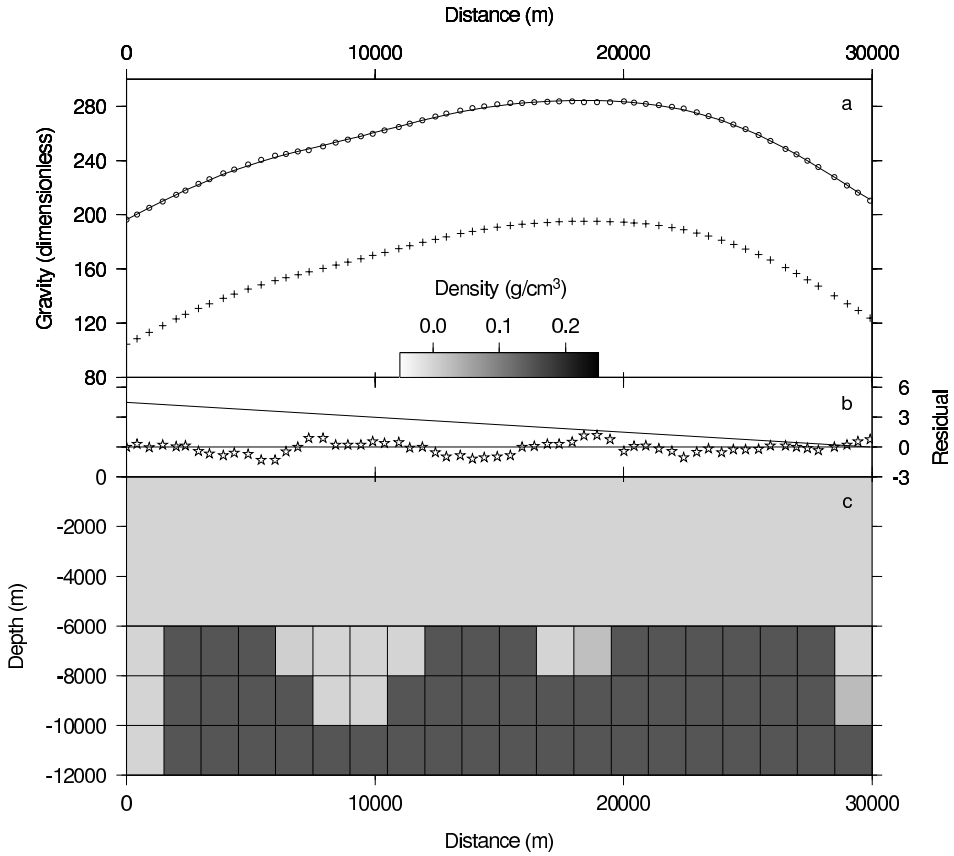

Figure 3.5: Result of an LP-based inversion. The parameterisation used (c) is a grid of(20x3) cells of 1500 by 2000 m starting at a depth of 6000 m. covering the area shown earlier in Figure 3.4g. The input data (Figure 3.4e) is shown in panel a with circles, as in Figure 3.4. The computed gravity effect due the (inverted) model is shown with crosses. The solid line is this forward modeled data plus the corrections for the effective trend (shown in b) and the offset. Panel b also shows the data residuals, i.e. the difference between the solid line and the circles in the panel a.

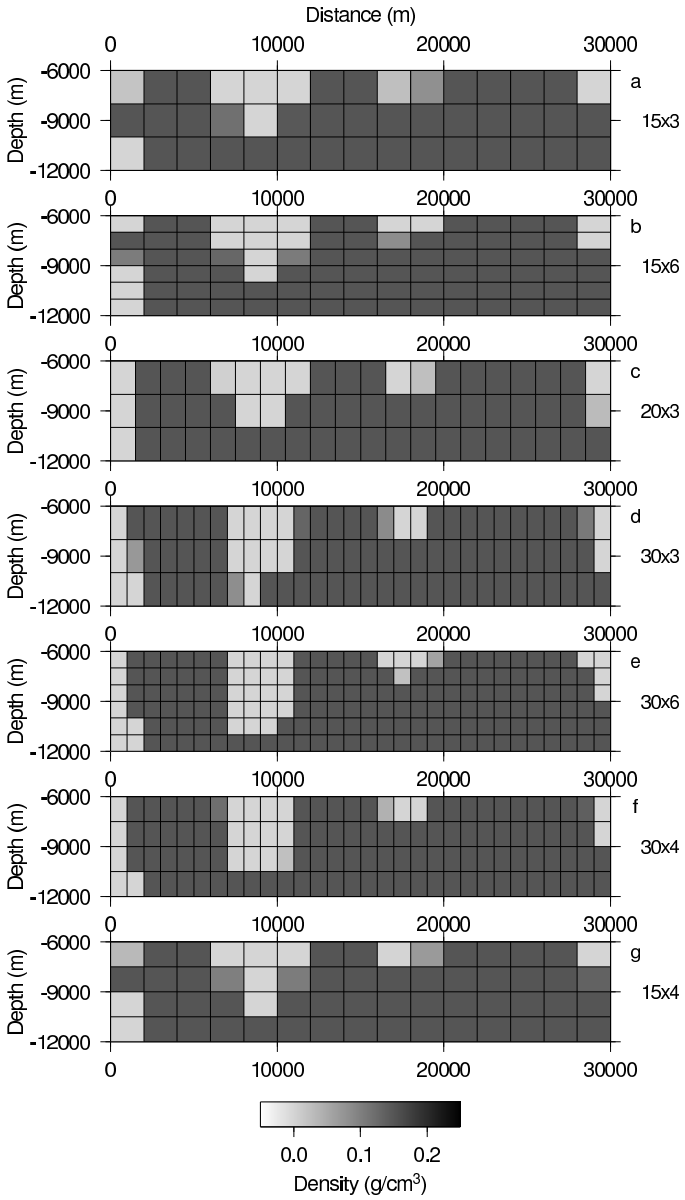

Figure 3.6: Inversion results with different parameterisations of the target area. The number of cells (lateral x vertical) are, respectively: a: (15x3), b: (15x6), c: (20x3), d: (30x3), e: (30x6), f: (30x4), g: (15x4).

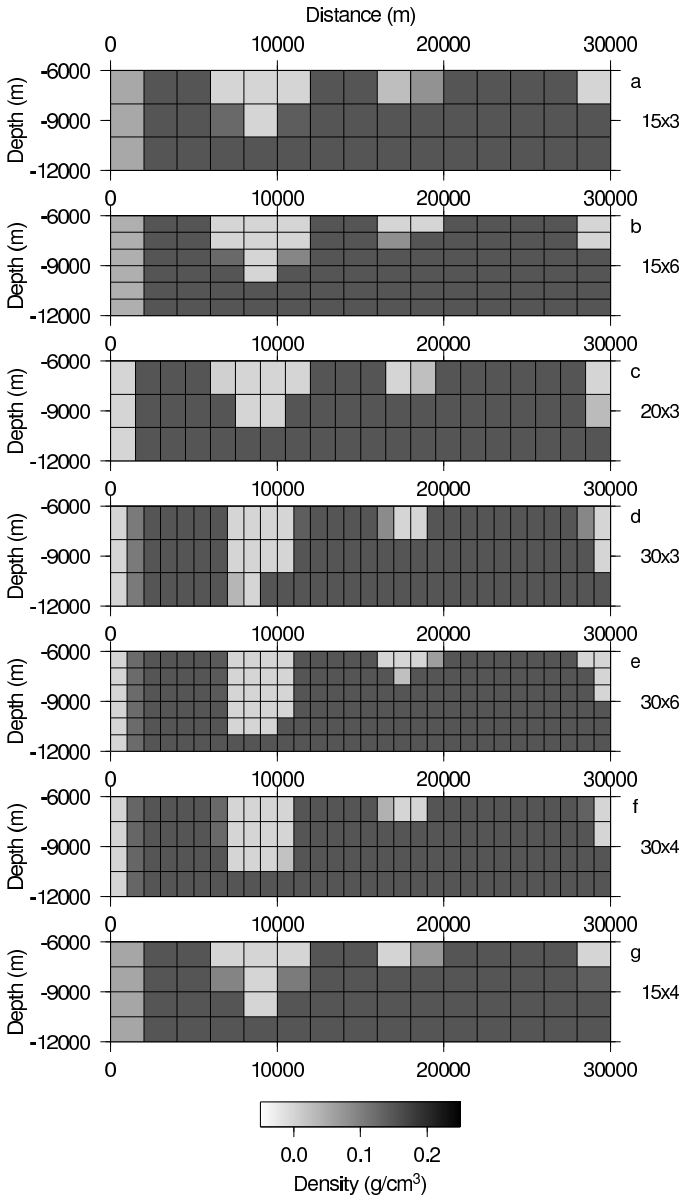

Figure 3.7: The inversion experiments shown in the Figure 3.6 repeated with the additional constraint that density cannot decrease with depth. The parameterisations used are, as in Figure 3.6: a: (15x3), b: (15x6), c: (20x3), d: (30x3), e: (30x6), f: (30x4), g: (15x4).

parameterised area are, no doubt, present in the data. This could be due to the limited horizontal extent of the line of measurement chosen, or incorrect values assumed for the densities or the strike length (3-D-effects). During inversion, these effects get mapped into the region of interest. The sides of the model are most susceptible for such errors, because the cells near a side of the grid influence only a limited number of measurements significantly and are thus not well constrained by the latter.

The linear trends retrieved from the data, as given in Table 3.1, all have their maximum at the left of the line of measurements ($\alpha_r > \alpha_l$). This is in agreement with the overall dip of the layers in the subsurface above the region of interest. Note that this region was modeled as homogeneous, hence, gently dipping layers that are present, give rise to a linear trend in the data. Older (deeper) layers have a higher density and they are closer to the surface at the left of the measurement line. Thus a larger correction is necessary at the left side, in accordance with the trends found.

In our formalism, we also solve for a floating reference. This term, which may be of potential importance in a large-scale study, combines - in the present case - contributions from several sources, e.g. the unknown datum that was used earlier for processing the original data and an offset that has been added to the data before the LP-based inversion to honor the positivity requirements of the latter. It also includes - implicitly - the effect of the background (Panel f in Figure 3.4). In our treatment, this "Bouguer-like" slab has been neglected; its density being equated to zero. When these individual components of the floating reference, as obtained from the LP-based inversion, are important, extra steps should be taken for their separation. Extensive numerical tests with synthetic data (with noise added) have shown that there is no trade-off in the inversion results between the effective trend and the floating reference.

An important piece of a priori information is the depth of the grid. This point has been discussed in detail in Chapter 2. The minimum depth of the basement was known for this study, but in experiments, wherein the grid was positioned at a shallower depth, the results were markedly different (not shown). The two depressions were still present, but their maximum depth extent was much less. This is expected because, given the same density contrast, shallower structures have a larger gravity anomaly. Overall, the method seems readily applicable to real data and it produced structured inversion results of the transition between sediments and basement. The main features were reproducible with several parameterisations and the earlier modeling, giving some confidence in the results.

3.6 Acknowledgments

We would like to thank Jaap Mondt at Shell International Exploration and Production for providing the data and Shell for allowing its publication.

Chapter 4

Lithotype discrimination within a gridded model in the context of gravity data inversion

In this chapter we present a method which allows delineation of geologic structures in a bi-modal lithology setting. We propose to use gravity data in combination with a priori information about the density contrast between the two lithologies. Without the need of an a priori density model, the iterative method produces structural images of the subsurface. The depth to the top of the inhomogeneity is also retrieved.

4.1 Introduction

The gravity field of the Earth can be used in the applied earth sciences, to extract information about the subsurface at various length-scales. Some examples are density constraints for the mantle, delineation of tectonic structures, locating ore-deposits, detection of cavities, etc. However, the same gravitational field can be produced by an infinite number of different density distributions. The inverse problem is thus ill-posed and a priori information needs to be included to get stable results.

Many authors have studied the problem of geophysical data inversion in the presence of non-uniqueness in the solutions (Jackson (1972),Sabatier (1977a), Menke (1984)). A classic approach is to parameterize the region of interest with many cubic prisms or cells and estimate the value of the physical properties (density, acoustic velocity, conductivity, etc) of each cell. We use the phrase *smooth inversion* for schemes using a gridded parameterisation and a rule that some spatial derivative of the parameters (the properties of the

[1]This chapter is a slightly modified version of a manuscript that has been submitted for publication in *Geophysical Journal International* (2006). Authors: Tim van Zon, Pavel Ditmar, Kabir Roy Chowdhury and Jaap Mondt

cells) is small, while fitting the observations properly. The resulting models will have a certain smoothness.

However, interesting geological structures are frequently not smooth spatially; they have distinct boundaries. This could be modeled by prescribing a structure. For instance, Jupp and Vozoff (1975) assumed a set of layers, whose thickness and physical property values needed to be determined. Another example is the inversion performed by Mirzaei et al. (1995) for the shape of the basin assuming a known density difference between the basin-fill and the surrounding basement.

Another example of explicit structural inversion is the study by Ditmar (2002), with the aim to find the shape of an isolated 2D body of anomalous density (assumed known). The outline of the body is given by connected nodes. In an iterative scheme, the nodes are moved, in order to find the configuration that fits the data.

The 'explicit' structural inversion techniques mentioned above have certain drawbacks too. Firstly, they require rather detailed a priori information about the medium; at least the number of anomalous bodies and their approximate locations should be known. Secondly, manipulating a model with structural information explicitly included may be difficult, especially in a 3-D situation. Lastly, computation of the response of such models (the forward problem) may also be difficult, especially if the computation is based on numerical solution of differential equations, e.g. for the inversion of electrical conductivity data. It might be necessary to employ stochastic search methods, which are time consuming; then the inversion allows for only a few unknown parameters.

The drawbacks of the two approaches described above (smooth inversion and explicit structural inversion) could be overcome by combining them. In this approach, the subsurface is represented by a Cartesian grid. Structural constraints can be implemented in a model of this type in an implicit form - by "encouraging" or "allowing" a high contrast in property values between neighboring cells. This should make the interpretation of the result, in terms of delineation of lithologies, easier than in the conventional smooth inversions.

Several authors, e.g., Last and Kubik (1983); Guillen and Menichetti (1984); Acar and Vogel (1994); Li and Oldenburg (1998); Zhdanov and Hursan (2000); Portniaguine and Zhdanov (2002); Farquharson and Oldenburg (2003) have used this approach for obtaining non-smooth models from different types of geophysical data. Portniaguine and Zhdanov (1999) inverted gravity data by minimizing the area with non-zero spatial gradient in density (minimum gradient support), while limiting the influence of large density contrasts to a constant value. Additionally, they set a limit on the maximum and minimum density values to prevent geologically unrealistic models. They were able to regularise a gravity inversion and obtain structured images.

Li and Oldenburg (1998) devised an objective function in which an a priori model could be specified, which the result should resemble. Furthermore, the level of smoothness of the solutions could be specified separately in all 3 dimensions. With the aid of a properly chosen weighting function, they obtained good results.

Recently, van Zon and Roy Chowdhury (2006) (Chapter 2) used Linear Programming for structural inversion of gravity data, based on a gridded model. Given a density contrast, they obtained results with structure by minimizing the L_1-norm of the data residuals.

The method is robust against noise, but the minimum depth to the top needed to be known a priori.

Here, we propose a method to invert gravity data for the density of cells on a grid, while using lithological constraints. The lithological constraints are data of the density contrast of the lithologies assumed present in the subsurface. During an iterative scheme, we encourage that the value of the physical property of the cells of intermediate models are (increasingly) in agreement with the lithological data, via the lithological constraints. These constraints are implemented as quadratic functions, that are iteratively updated. During each update, the previous inversion result is used to deduce the lithology (and thus the density) of each cell. A function encouraging a density value close to the density of its deduced lithology is implemented for the next iteration.

Our method has 5 tunable constants, but the interpreter need not set those; only a proper range is required. The strategy is given for the 3-D case, but the examples shown are 2.5-D; the cells have a finite extend in the strike direction.

4.2 Method

The subsurface is parameterised with a grid of 3-D, identically shaped, prisms, for which the density contrasts are sought. Each observation, d_i, can be calculated from the following relation, linear in the density contrasts of the prisms:

$$d_i = \sum_{j=0}^{M} m_j G_{ij}, \tag{4.1}$$

where M is the number of prisms, m_j the density contrast of the j-th prism (and the j-th component of the vector $\mathbf{m}$) and G_{ij} is the vertical component of the gravitational attraction at the measurement station i due to the j-th rectangular prism (of unit density). Nagy (1966) derived a formula for the calculation of the components G_{ij}.

Equation 4.1 can be written in a matrix notation:

$$\mathbf{Gm} = \mathbf{d} \tag{4.2}$$

with $\mathbf{m}$ the M-dimensional column vector with the M model parameters, m_j; $\mathbf{d}$ is the N-dimensional column vector composed of the N measurements, d_i, and $\mathbf{G}$ is the $N \times M$ forward operator matrix (or Jacobian matrix), containing the elements G_{ij}.

The inversion uses an iteratively updated cost function of the parameters, $\mathbf{m}$, consisting of 3 terms, each meant to promote desired features in the solutions:

$$\psi_k(\mathbf{m}) = \text{data misfit}(\mathbf{m}) + \beta_k^2 \times \text{structural regularisation}(\mathbf{m}) + \gamma^2 \times \text{smoothing}(\mathbf{m}), \tag{4.3}$$

with β_k and γ being weight factors. β_k (the weight for the structural regularisation term) increases with each iteration, while γ remains constant. This will ensure that the solution will be easily interpretable in terms of lithology. The result of the inversion are the values of the parameters for which the cost function is minimal. We use an implementation of the

conjugate gradient method, as given by Press et al. (1992), to minimize the cost function. The three terms of the right hand side will be treated in more detail below.

The first term of the right hand side of the cost function is the data misfit term. This is the L_2-norm of the data residuals ($= \chi^2$) given by

$$\text{data misfit}(\mathbf{d}, \mathbf{m}) = (\mathbf{d} - \mathbf{G}\mathbf{m})^T (1/\sigma^2)(\mathbf{d} - \mathbf{G}\mathbf{m}), \tag{4.4}$$

where T denotes the transpose and σ^2 is the variance of the data noise.

The second term in Equation 4.3, the structural regularisation, is chosen as

$$R_{struc}(\mathbf{m}, \mathbf{w}, \mathbf{m^{prev}}) = \sum_{j=1}^{M} w_j^2 \phi_j(m_j, m_j^{prev}), \tag{4.5}$$

where ϕ_j is given in Equation 4.7 below. The M-dimensional vector $\mathbf{m^{prev}}$ is either the initial model (at the first iteration), or the result obtained at the previous iteration (at later iterations).

As Li and Oldenburg (2000), we use a sensitivity-based weighting function. The components of the weight vector, $\mathbf{w}$, reflect the differences in influence of the cells on the data and are defined to be:

$$w_j = \frac{1}{N} \sum_{i=1}^{N} |G_{ij}|^P, \tag{4.6}$$

with $|G_{ij}|$ the absolute value of the element in row i and column j of the forward operator matrix $\mathbf{G}$. P is a scalar that needs to be chosen properly.

The function ϕ, used in Equation 4.5, is defined as

$$\phi_j(m_j, m_j^{prev}) = \begin{cases} t(m_j^{prev}) m_j^2 & \text{if} \quad |m_j^{prev}| \leq |\mu| \\ t(m_j^{prev})(m_j - \rho_{an})^2 & \text{if} \quad |m_j^{prev}| > |\mu| \end{cases}, \tag{4.7}$$

where μ is a constant between 0 and ρ_{an}, the latter being the density contrast between the background density and the anomalous density. $t(m_j^{prev})$ is a function that determines the steepness of the parabola used for the structural regularisation for cell j (see below Equation 4.8).

At each iteration the initial value m_j^{prev} is fixed, so $\phi_j(m_j, m_j^{prev})$ is a quadratic function of the unknown parameters, m_j. A contour plot of the structural regularisation, $\phi(m_j^{prev}, m_j)$ is shown in Figure 4.1. Furthermore, two typical parabolas are plotted as well. The minimum of each parabola is either at the background density (0) or the anomalous density (ρ_{an}). In the case of a salt-dome embedded in sandstone the density difference, $\rho_{an} \approx -0.4$ g/cm^3.

We have implemented the function $t(m_j^{prev})$, as

$$t(m_j^{prev}) = \frac{(m_j^{prev} - \mu)^2}{\mu^2}, \qquad (\mu \neq 0). \tag{4.8}$$

The rationale for this choice is that 'in between' the lithologies, there is no preference for either one, so $t(m_j^{prev} = \mu) = 0$. So the structural regularisation term is zero and the choice of m_j is driven by the data.

In summary, the structural regularisation is used to favor models containing only the background density or the anomalous density. When the initial density of a cell is in the middle between the lithologies it is not clear to which lithology a cell is likely to belong. Then the function t reduces its influence on the structural regularisation term. At each particular iteration, the structural regularisation term reduces to a zero-order Tikhonov regularisation. The relative weight of the structural regularisation term, β_k, increases with the iteration number, to steer towards a clear structure. This corresponds to an increased trust in the solution with iteration. (Later models should be trusted more than the initial homogeneous model.)

The third term of the right hand side of the cost function (Equation 4.3) is a smoothing function. We choose a weighted Second Order Tikhonov smoothing. The 3-D implementation is:

$$S = (\mathbf{wDm})^T(\mathbf{wDm}), \tag{4.9}$$

with $\mathbf{D}$ being a numerical difference operator: $\mathbf{wDm}$ is the M dimensional vector with components:

$$\begin{aligned}\{wDm\}_j \quad = w_{p,q,r} \quad (6m_{p,q,r} - m_{p-1,q,r} - m_{p+1,q,r} - m_{p,q-1,r} \\ -m_{p,q+1,r} - m_{p,q,r-1} - m_{p,q,r+1}).\end{aligned} \tag{4.10}$$

For clarity we did not use the column vector notation in this formula, but denoted each cell in the 3-D grid with indices p, q, r that represent the cell indices in the x, y, z-direction respectively. In this notation, cell j (column vector) is cell (p, q, r) and the weight $w_{p,q,r}$ is w_j. Outside the grid, the density is assumed to be that of the background, so a 'parameter' like $m_{-1,1,1}$ is zero by definition.

In summary, the smoothing term of a cell is zero when all the neighboring cells have the same density contrast. The purpose of the smoothing term is to prevent density distributions with many small anomalous bodies and to promote a limited number of compact structures. Furthermore it regularises the solution when the influence of the structural regularisation is small.

At each iteration in the algorithm, the value for β_k in the cost function is increased. This has the effect of making the structural term relatively more important, so models with increasingly more structure will be produced. We use two criteria to stop the iterations; there is a pre-determined, maximum number of iterations, but when β_k becomes larger than a threshold, the iterative process is also terminated. When β_k is large, we can assume that the mass distribution has much structure and that cells cannot change lithology anymore. Then, we apply a final discrimination:

$$\begin{aligned}&\text{if} \quad |m_j| > |\mu| \quad \text{then} \quad m_j = \rho_{an} \\ &\text{if} \quad |m_j| \leq |\mu| \quad \text{then} \quad m_j = 0.\end{aligned} \tag{4.11}$$

The final product is an image of the subsurface containing only 2 distinct lithologies.

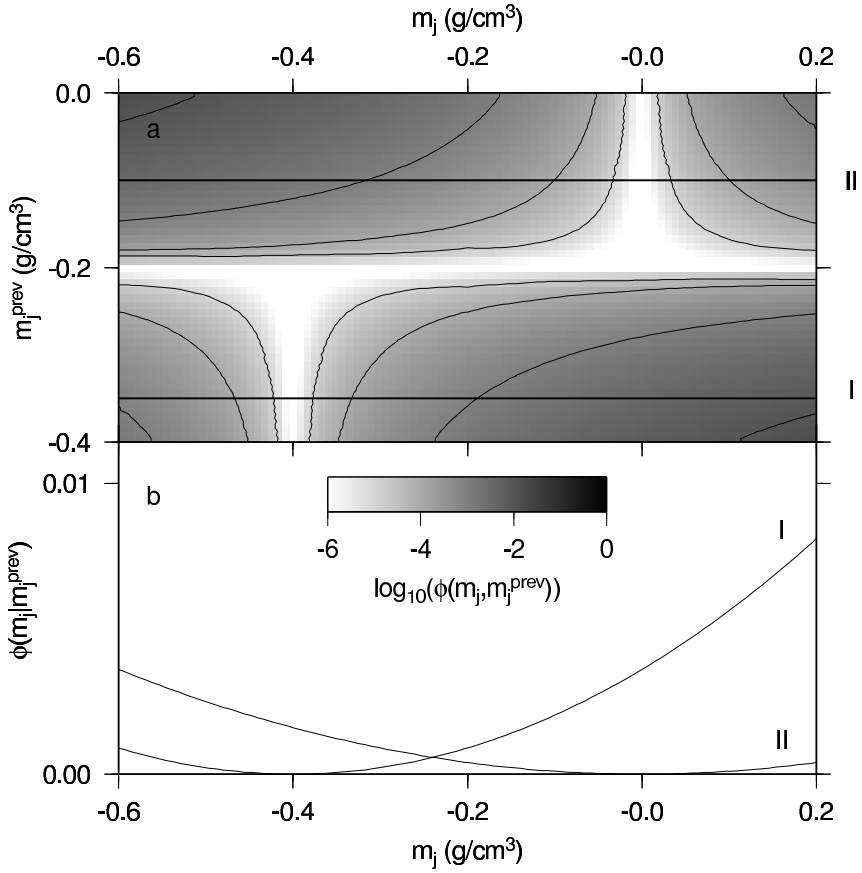

Figure 4.1: Panel (a) shows a contour plot of the structural regularisation $\phi(m_j, m_j^{prev})$. The contours corresponding to the function values 0.00001, 0.0001, 0.001 and 0.01 are plotted. The grayscale is logarithmic. In panel (b) two parabolas, I and II, are shown. These are the functions $\phi(m_j|m_j^{prev} = I)$ and $\phi(m_j|m_j^{prev} = II)$. These two chosen m_j^{prev}, (-0.35 and -0.1) are shown with the straight lines in panel (a), marked with the I and II on the right. In this figure, $\mu = -0.2$. Note that for each iteration, m_j^{prev} is fixed and that the structural regularisation is thus a 1-D parabola in m_j, with its minimum either at 0 (the background density) or at -0.4 (the anomalous density contrast).

4.2.1 Strategy

The cost function (Equation 4.3) has 5 tunable parameters, namely γ, P, μ, β_0 and $d\beta$. Here, β_0 and $d\beta$ are, respectively, the initial value and the multiplicative factor - per iteration - for β (i.e. $\beta_{k+1} = \beta_k \times d\beta$). The effect of increasing β_k with iteration is that the model is smooth in the beginning and that more structure appears in the course of the iterative process, when it becomes clearer as to which lithology a cell belongs. A large β_k implies a large influence of the model obtained at the previous iteration.

The threshold μ has 3 effects: Its obvious role is to promote the final discrimination according to Equation 4.11. Another role is played in Equation 4.8, where μ determines the slope of the parabolas. Lastly, the values of the density at the first iteration (initial model) are set equal to μ. This implies that the structural regularisation term is zero in the first iteration because t (Equation 4.8) is zero. Parameter γ controls the degree of smoothness at all the iterations. Finally, the constant P determines the weight vector $\mathbf{w}$ (Equation 4.6). A large value for P favors mass distributions deeper in the subsurface, whereas smaller P-values encourage shallower mass concentrations.

In the synthetic tests below, the true earth model is known, so the results obtained for several combinations of the tune parameters, can be compared. For this, we introduce the model misfit (as a percentage) as an objective measure of 'likeliness' between the inversion result and the true earth model:

$$\text{model misfit} = \frac{100\%}{M\rho_{an}} \sum_{j}^{M} |m_j - m_j^{true}|, \tag{4.12}$$

where $\mathbf{m^{true}}$ is the true earth model, $\mathbf{m}$ the inversion result, M the number of cells and ρ_{an} the expected density difference between the lithologies. This represents the average error associated with the computed densities in the cells. This definition, in combination with the final discretisation, makes the model misfit equal to the percentage of cells with the wrong density.

Our approach to tuning the 5 pre-determined parameters is to try several combinations of them and calculate both the data misfit (Equation 4.4) and the model misfit (Equation 4.12). To choose these 5 tunable parameters, we use the Neighborhood Algorithm (NA), as developed by Sambridge (1999). The neighborhood algorithm, in its original form, can be used to stochastically search a model parameter space. Regions of the parameter space with a good data misfit are sampled more densely, to efficiently find combinations of the model parameters that give a good data misfit. Another advantage of this method is that the resulting ensemble of models can be used to compute statistical properties of the model parameters.

We did not use the Neighborhood Algorithm to find the density contrasts for each cell (our model parameters) because we have too many parameters and the search would take too much time. Instead we use the NA to identify good combinations of the 5 tunable parameters. Their influence on the inversion result can be highly non linear and a stochastic search is necessary to properly investigate this parameter space.

The NA can be tuned to perform a broad search, including many points in poorly fitting regions, or to perform a narrow search, which densely samples the better fitting

parameter	minimum	maximum	optimal
β_0	$1.0 \cdot 10^{-6}$	1.0	$8.4 \cdot 10^{-4}$
$d\beta$	1.0	1.5	1.2
γ	$1.0 \cdot 10^{-4}$	10.0	3.1
P	0.2	1.5	1.26
μ	-0.35	-0.05	-0.21

Table 4.1: The ranges for the 5 tunable parameters in the cost function. In the fourth column the optimal values, as found with the Neighborhood Algorithm, are given.

regions. A broad search needs more time to produce good fitting models, while a narrow search might end up investigating a local minimum. We set the algorithm to perform a broad search.

To quantify the amount of noise added to the data, we define the signal to noise ratio as:

$$\mathcal{S}/\mathcal{N} = \sqrt{\frac{\sum_{i=1}^{N} d_i^2}{\sum_{i=1}^{N} n_i^2}}, \tag{4.13}$$

where d_i is the scaled i-th theoretical measurement and n_i is the i-th Gaussian noise sample. In the present case, the expected data misfit, given Gaussian noise statistics, is 200 (dimensionless) - equal to the number of measurements.

4.3 Results

4.3.1 Buried pillar

Our first 2.5-D test model contains a buried pillar, and is shown in Figure 4.2c. Its top is 200 m below the surface and its strike length is 1000 m. The pillar has a density contrast of -0.4 g/cm^3 with respect to the background medium. The subsurface is parameterised with 40×20 cells, of 50×1000×50 m each. This parameterisation is shown with the mesh in Figure 4.2. Synthetic data was calculated at 200 locations, evenly spaced between -490 m and 2495 m, at a height of 2 m above the surface. The line of measurements crosses the inhomogeneity half way. To make the data dimensionless, each measurement is scaled with an assumed measurement error, σ, of 0.04 mGal. Zero-mean, unit variance non-correlated Gaussian noise was added to the data. This dimensionless data set is shown in Figure 4.2a. The forward operator is scaled accordingly.

This 'noisy' data set ($\mathcal{S}/\mathcal{N}$-ratio = 25.7) was inverted 1250 times, using different values for the 5 tunable parameters. The search ranges (parameter space) for these constants are tabulated in Table 4.1. The Neighborhood Algorithm was used to find their optimal combination, using the data misfit to steer the search. The total time needed to perform the 1250 inversions is about 9 hours on a Pentium IV, 3 GHz CPU.

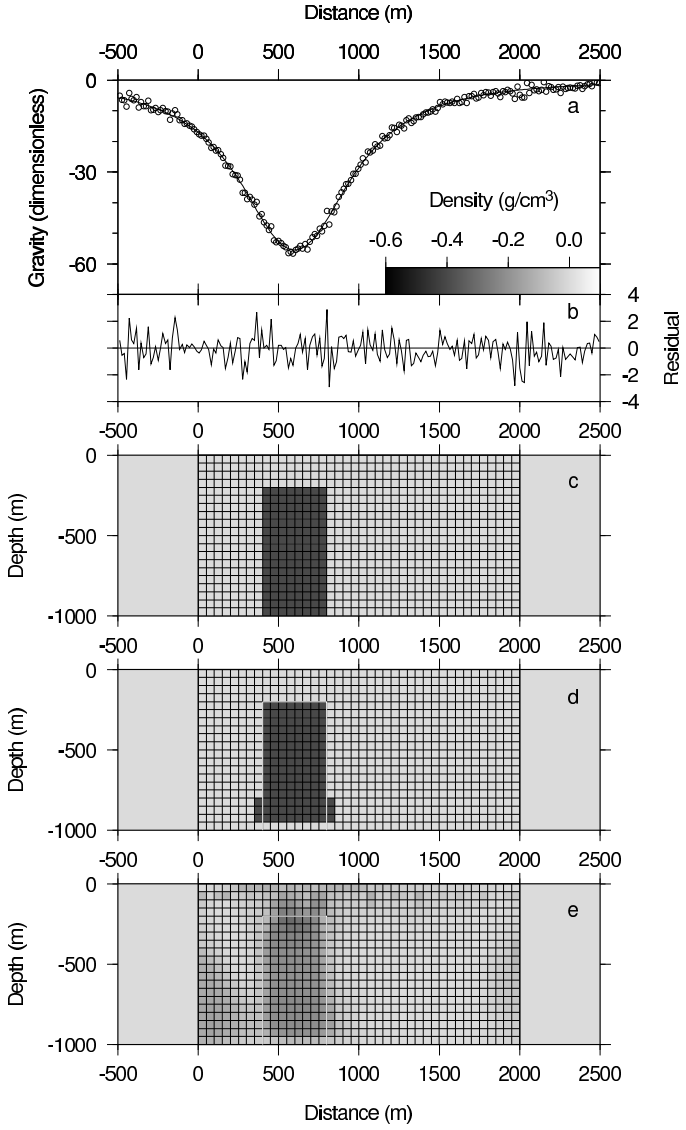

Figure 4.2: The dimensionless (scaled) synthetic data with added Gaussian noise is shown with the open circles in panel (a). Panel (c) shows the source body and the parameterisation used. The data corresponding to the inversion result (d) is shown with the solid line in panel (a). The dimensionless data residuals are presented in (b). (e) shows the average per cell of the last 500 models found during the parameter search with the neighborhood algorithm. The white line in (d) and (e) is the outline of the true anomalous mass distribution.

Figure 4.2d shows the resulting model with the best data misfit from the 1250 inversion results. The data misfit is calculated after the final discretisation and is 200.5, very close to the expected value of 200. The data corresponding to this mass distribution is plotted with the solid line in Figure 4.2a. The data residuals are shown in Figure 4.2b. The resulting model has a model misfit (Equation 4.12) of 1.75%. To satisfy the stopping criterion, 75 iterations were needed.

The effect of increasing the importance of the structural regularisation term in the cost function can be seen in Figure 4.3. The initial model is homogeneous (Figure 4.3a). Figures 4.3b through 4.3i show intermediate models; the results after 35, 40, 45, ... , 70 iterations. The cost functions used for the first 45 iterations (small structural regularisation) produce very smooth models. When the structural regularisation is made more important, clear structure emerges. Note that the model after 35 iterations (Figure 4.3b) does not look like the final result at all.

All terms of the cost function for each iteration are shown in Figure 4.4. Each term has been multiplied by its weight factor, β_k^2 or γ^2. The last point in each curve is the value after the final discretisation. This automatically makes the structural term zero, because then each parameter is either background or anomaly. Note that for the early iterations, the data misfit is below 200, so the data is being over-fitted.

We have looked at the best fitting result from the 1250 inversion in detail now, what about the others? Of great importance is the ability of a strategy to find models that properly represent the true subsurface. In these synthetic experiments the true model is known and the model misfit can be calculated (Equation 4.12). In Figure 4.5, these model misfits are plotted against the data misfit of the result. Models with a data misfit larger than 500 are not represented in this figure. (Models with a data misfit as bad as $7.1 \cdot 10^5$ were retrieved.) Our approach does well; the models with a good data misfit also have a good model misfit.

The last 500 inversion results (of the total of 1250 inversions) were averaged and are shown in Figure 4.2e. Due to the broad search with the Neighborhood Algorithm, the search has not converged yet, but the search should have zoomed in on reasonable regions of the search ranges. Some areas of the subsurface appear to be much more likely to contain mass (darker gray) than others (light gray). The true density anomaly can be recognised easily, which means that many combinations of the tunable parameters recover the same features of the subsurface. The lower left corner has some edge effects and sometimes the top layers seem to accommodate noise features.

4.3.2 Bedrock topography

The methodology is tested further using the schematic model shown in Figure 4.6b; it is 10 km wide, 5 km deep and the strike length is taken to be 20 km. It includes a basinal structure. The density distribution outside the grid is supposed to be known and its effect on the data has been removed. 200 observations are simulated, equally distributed along a line crossing the basin half way along strike. Unit variance, zero mean Gaussian noise has been added to the measurements (Figure 4.6a), after scaling them with the assumed measurement error of 0.04 mGal. (A typical noise level is 0.02 mGal for land data and

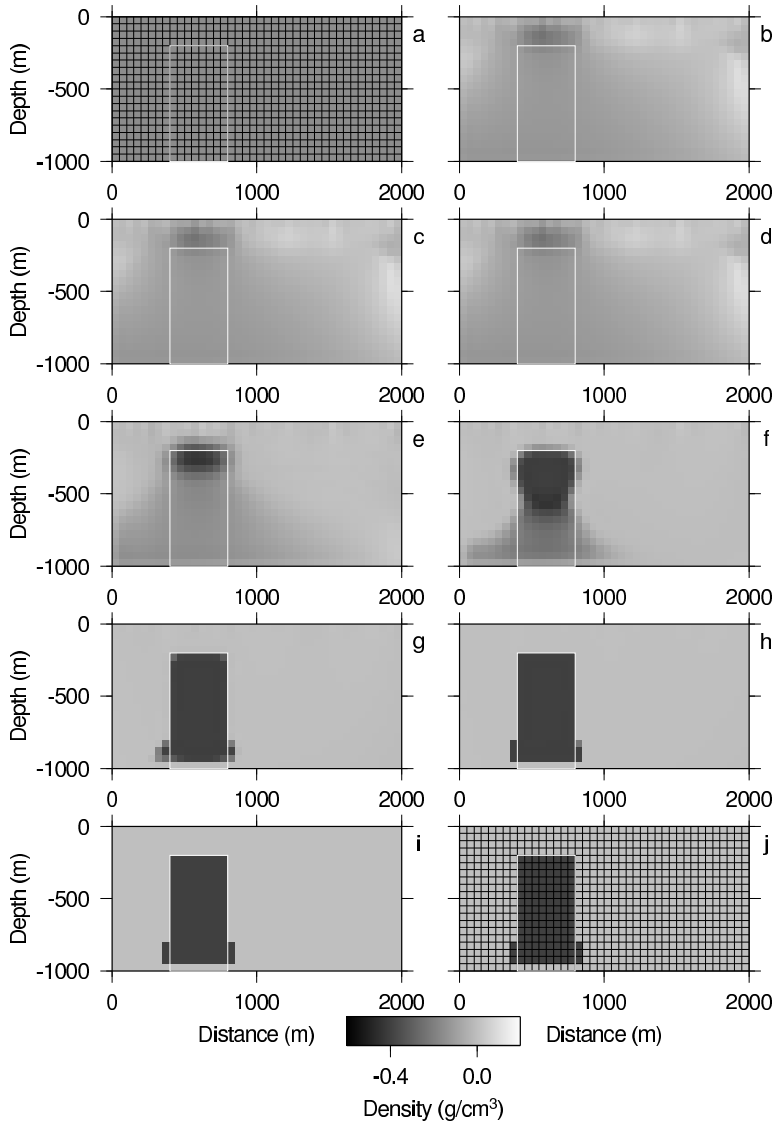

Figure 4.3: Some of the intermediate models of the inversion of the synthetic data shown in Figure 4.2a, with the optimal constants for the cost function, as found during the search with NA. Panels (b) through (i) show the intermediate models after, 35, 40, 45, ... , 70 iterations, respectively. (a) shows the homogeneous start model (with density value μ) and (j) is the result after the final discretisation. The white line is the outline of the true source body. For clarity, the parameterisation used is only shown in panels (a) and (j).

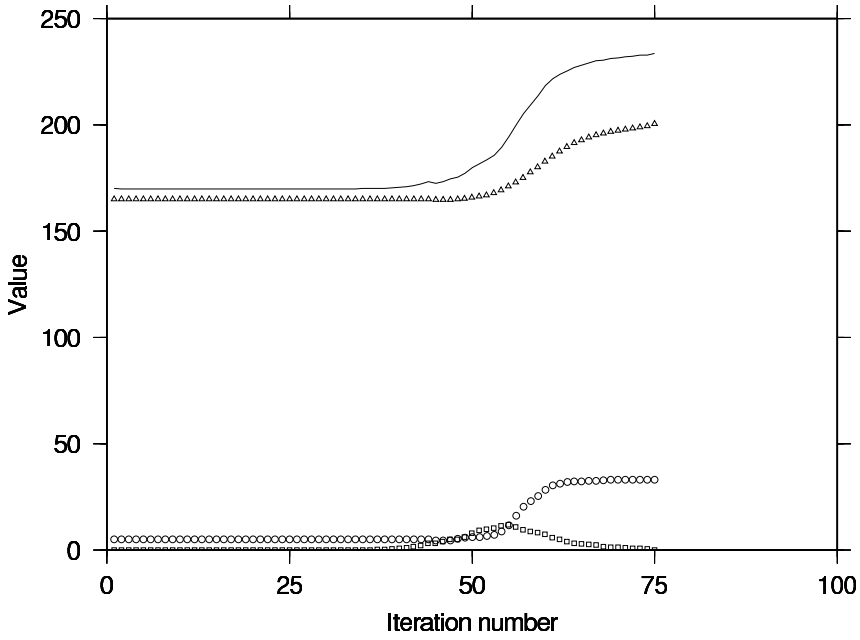

Figure 4.4: The individual terms of the cost function (Equation 4.3) at each iteration. The total value is shown with the solid line, the Tikhonov regularisation by the open circles, the structural regularisation by the open squares and the data misfit by the triangles. The first two terms have been multiplied with their weight constant. The last point in each curve is the value after the final discretisation.

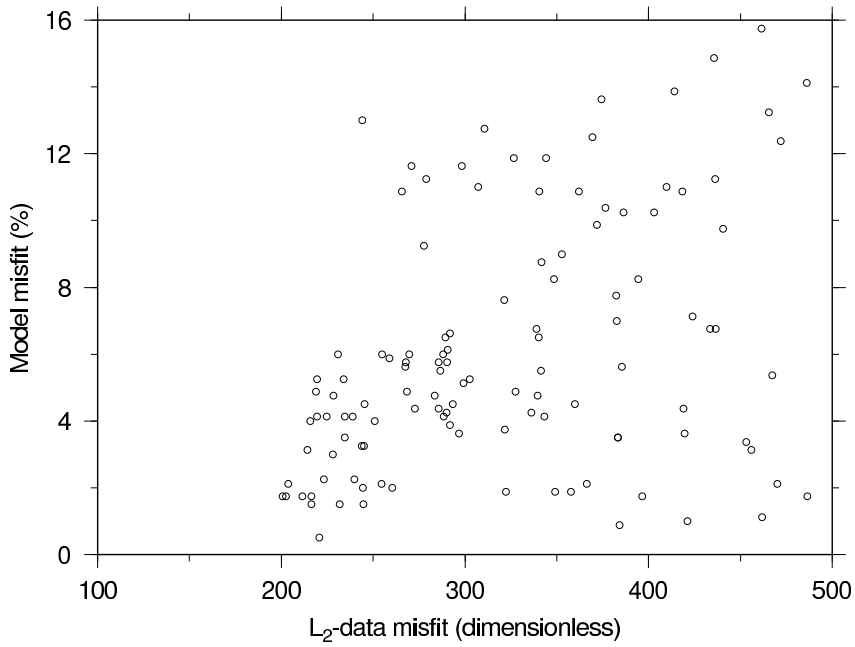

Figure 4.5: A scatter plot of some of the combinations of the model misfit and the data misfit as found during the search with the neighborhood algorithm for the buried pillar model. The expected data misfit is 200. There are many solutions with a data misfit worse than 500, these are not shown.

parameter	minimum	maximum	optimal
β_0	$1.0 \cdot 10^{-6}$	1.0	$1.71 \cdot 10^{-3}$
$d\beta$	1.0	1.5	1.29
γ	$1.0 \cdot 10^{-4}$	10.0	$1.03 \cdot 10^{-4}$
P	1.5	3.0	2.33
μ	0.05	0.35	0.25

Table 4.2: The ranges for the 5 tunable parameters used for the inversion for the bedrock topography. In the fourth column, the values leading to the solution with the minimum data misfit (as found with the NA search), are given.

0.1 mGal for marine data.) The $\mathcal{S}/\mathcal{N}$-ratio of the anomaly equals 154. The search ranges for the 5 constants are tabulated in Table 4.2.

The inversion result with the best data misfit (= 238) is shown in Figure 4.6c. The data misfit is somewhat higher than the expected value of 200. The model misfit (Equation 4.12) of the result is 1.50%. The maximum depth of the depression is recovered, as is the depth of the bedrock on the sides. The asymmetry of the basin is not well resolved.

Lets review the ensemble of solutions. The average of the last 500 inversion results (Figure 4.6d) shows rather sharp boundaries, which implies that similar models were recovered during the later stages of the parameter search. The sides of the depression show some variations, as do the cells on the borders of the grid. Again, the models with a good data misfit also have a good model misfit (not shown).

4.4 Discussion

The cost function has 3 terms, the data misfit, the structural regularisation and the smoothing term. We investigated what happens when β_k, the weight of the structural regularisation term, is very small: the resulting models are very smooth. Examples are the models shown in Figure 4.3, for the early iterations, when β_k is still small. Increasing β_k with iterations leads to structured results.

The smoothing term in the cost function seems at first illogical because it seems to counteract the structural term. To investigate the influence of this term, we performed the experiment for the buried pillar model with the search ranges as given in Table 4.1, except that now $\gamma = 1 \cdot 10^{-10}$. The four models with the best data misfit, out of the 1250 models, are shown in Figures 4.7a through 4.7d. The inversion results are not really compact, but fit the data reasonably well (a data misfit of about 300 instead of 200). These are the best fitting models, but the average over the last 500 results of the NA search looks nothing like the source body (Fig. 4.7e). This indicates that the inversion is not very robust.

The model misfit and data misfit are not as good as when the proper smoothing is applied. The addition of the smoothing term seems to favor compact bodies and the source body was compact in this case, so it is to be expected that the inversion with smoothing gives better models.

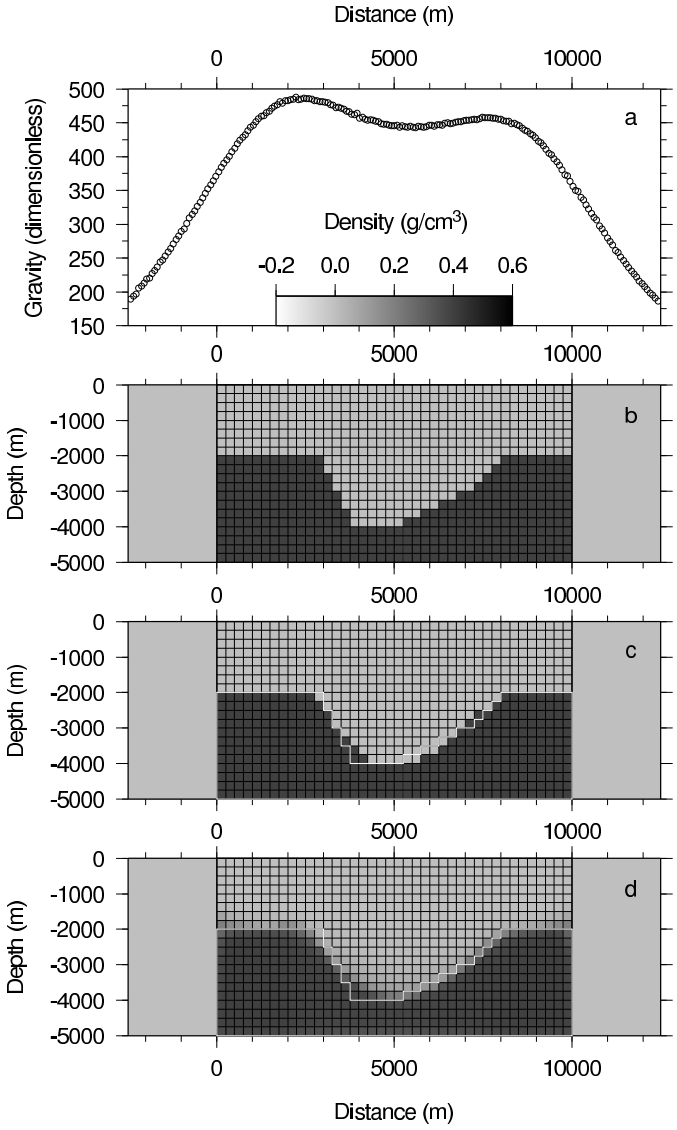

Figure 4.6: The dimensionless (scaled) synthetic data with added Gaussian noise is shown with the open circles in panel (a). Panel (b) shows the source body and the parameterisation used. The inversion result with the best data fit is shown in (c). (d) shows the average per cell of the last 500 models found during the parameter search with the neighborhood algorithm. The white line in (c) and (d) is the outline of the true anomalous mass distribution.

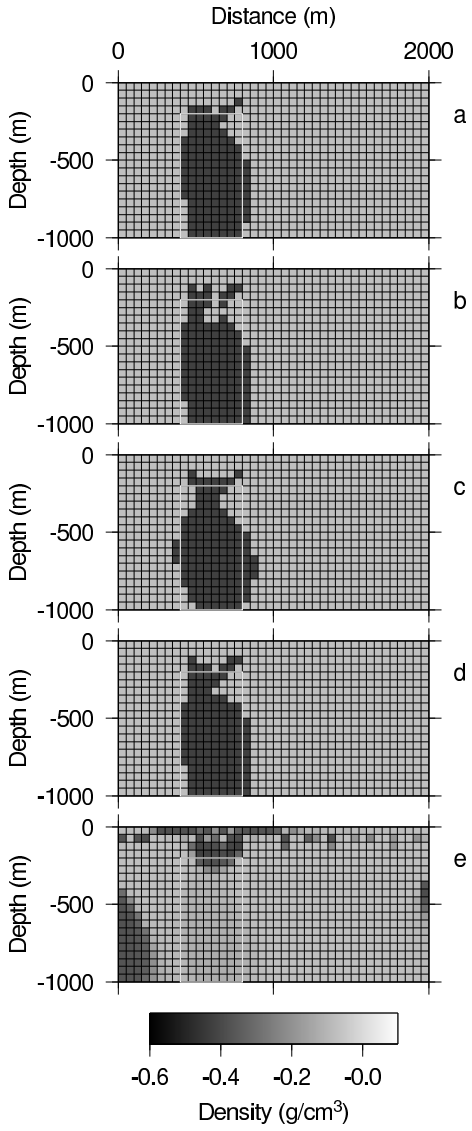

Figure 4.7: Example of the four inversion results with the best data misfit with $\gamma = 1 \cdot 10^{-10}$. The source body (outline in white) and data are the same as in Figure 4.2. The results in the panels (a,b,c,d) have a data misfit of respectively 264, 294, 300 and 306 (dimensionless units). Panel (e) is the average over the last 500 inversion results.

We made two observations, which indicate that the inversion results are robust. First, we examined the 9 models with the best data misfit, obtained with each NA search; they looked alike and the data misfits and model misfits were comparable. Secondly, we took the values for the 5 constants for the best fitting model and used them to invert the data of the same source body, but corrupted with other noise realisations. We obtained very similar inversion results; almost the same model misfit and reasonable data misfits, although worse than the result optimised with NA.

The search ranges for the constants β_0, $d\beta$, γ and μ do not influence the inversion results much, as long as they are chosen wide enough. The optimal values (Tables 4.1 and 4.2) are not at the boundaries of the parameter space. If the constants leading to the best data misfits are very close to the extremes of the range, the range should be extended, especially if the best found data misfit is much larger than expected.

However, the search range for P does influence the inversion results. We observed that when allowing too large values for P, there were many models with a good data misfit (slightly over-fitting the data), which had a very poor model misfit.

We have also experimented with an alternate function for $t(m_j^{prev})$, to be used in the structural regularisation (instead of Equation 4.8), namely

$$t_{lin}(m_j^{prev}) = |\frac{-m_j^{prev}}{\mu} + 1|, \qquad (\mu \neq 0). \tag{4.14}$$

Like the other equation for t, this function is zero 'in between' the lithologies, when $m_j^{prev} = \mu$. The parameter search with NA was about twice as fast compared with the search using the other implementation for t (Equation 4.8), but the resulting models were poor.

4.5 Conclusions

The proposed strategy can robustly retrieve density models of the subsurface with clear structure, from gravity data corrupted with Gaussian noise, given the density contrast. The depth to the top of an inhomogeneity can also be retrieved. An important advantage of the method is that the regularisation constants need not be given by the interpreter; the best combination of these constants is found automatically, within a reasonable range. Furthermore there is no need for a detailed initial model, a homogeneous initial model suffices. The inversion results can readily be interpreted in terms of structures in the subsurface.

4.6 Acknowledgments

M. Sambridge is thanked for supplying the code of his Neighborhood Algorithm.

Chapter 5

Structural inversion of seismic travel-times and joint structural inversion: an outlook

In this chapter, consisting of two parts, we first investigate the possibility of using the LP-based formalism for implicit structural inversion of seismic travel-times. Synthetic cross-well first arrival times, corrupted with Gaussian noise, are used for this study. In the 2.5-D examples presented, the subsurface is parameterised with homogeneous, rectangular cells. The parameters to be inverted for are the values of the slowness of each cell. The formalism can be extended to the full 3-D case. Unlike the inversion for density values in Chapter 2, inversion of first arrival travel-time data for slowness values is a non-linear problem; we attempt to solve it with an iterative approach.

Cross-well tomography is based on the lateral (ray-path) connection between sources and receivers. As a result, solutions suffer most from a lack of lateral resolution. Gravity inversion results suffer predominantly from a lack of vertical resolution. The combination of cross-well tomography and gravity inversion, i.e. a joint inversion of seismic and gravity data, may therefore be promising. In the second part, the potential improvement by performing a joint inversion is briefly illustrated.

In the outlook section, possible extensions to the joint inversion strategy are mentioned, which might improve the results and might also allow for inversion for more than two lithologies.

5.1 Part 1: Iterative seismic travel-times inversion for structure

5.1.1 Introduction

Measurement of the propagation speed of seismic waves - henceforth referred to as seismic velocity - are important in geophysics, in fields as varied as search for oil/gas and global seismology. This property is used both for diagnostic purposes and for further use in the imaging of the sub-surface; for an introductory overview, see Sheriff and Geldart (1995). In this context, *traveltime tomography* groups methods whereby surface- and/or borehole-measurements of first-arrival times or wave-forms are inverted to obtain the slowness distribution in the sub-surface. Much literature is available on the subject. Some examples dealing with paramerisation choices, resolution and optimal experimental design are Lévêque et al. (1993), Yamamoto et al. (1995), Spakman and Bijwaard (2001) and Curtis (2004).

The inversion of travel-times to obtain velocity information is, inherently, a non-linear inverse problem, requiring an iterative approach for its solution. A review of several approaches for calculating a model update in order to improve the datafit - an essential part of an iterative approach - was given by Spakman (1993). An interesting approach was described in Berryman (1989) and used in Berryman (1990); the author used Fermat's principle - the stationarity of the travel time with respect to the travel path - to classify the possible intermediate solutions in terms of their feasibility violations, and thus obtained a criterion for the optimum model update.

Musil et al. (2003) showed the use of discrete tomography based on mixed-integer Linear Programming, to invert seismic travel-times and to perform a joint inversion by addition of georadar data sets.

5.1.2 Theory

Seismic travel-times are dependent on the slowness distribution in the medium. In the high-frequency approximation (ray approximation), the energy of seismic waves can be taken to travel along ray-paths, the travel-time of the first arrival along a ray i being given by

$$d_i^t = \int_{\text{ray } i} s(\mathbf{r}) d\mathbf{r} \quad , \tag{5.1}$$

with $s(\mathbf{r})$ the slowness (the reciprocal of the velocity) field of the subsurface. The path integral is evaluated along the trajectory of the i-th ray, which depends on $s(\mathbf{r})$. We use the superscript *t* for quantities related to seismic travel-times.

For a gridded model with M homogeneous cells, ray-bending occurs only at the cell boundaries and is described by Snell's law. Given a gridded slowness field, the ray trajectory between a source and a receiver can thus be calculated. Each ray consists of a sequence of straight-line segments, one for each cell the ray passes through. Applying

this discretisation to equation 5.1 yields

$$d_i^t = \sum_{j=1}^{M} G_{ij}^t \; s_j \quad , \tag{5.2}$$

with s_j the slowness of cell j and G_{ij}^t the length of the ray segment of ray i through cell j.

Both d_i^t and G_{ij}^t have been divided by the assumed measurement error σ_i^t, to both sides of the equation dimensionless. Writing this in matrix form we get

$$\mathbf{d^t} = \mathbf{G}_{est}^t \, \mathbf{s} \quad , \tag{5.3}$$

$\mathbf{s}$ being the M-dimensional column vector containing the slowness values of the cells, $\mathbf{d^t}$ the column vector containing the N^t first arrival times and $\mathbf{G}_{est}^t$ the $N^t \times M$ forward operator matrix. The inverse problem is to obtain $\mathbf{s}$ given $\mathbf{d^t}$.

We added *est* in subscript to the operator $\mathbf{G}^t$ to emphasise that the operator is, in general, an approximation and not exact; the ray paths being dependent on the unknown slowness values in the cells. The operator $\mathbf{G}_{est}^t$ depends on the unknowns and hence, the inverse problem of calculating the slowness values on a grid from travel-time data is inherently non-linear. An iterative procedure has therefore to be designed to attempt to solve the inverse problem: Given an initial estimate of the slowness field, rays can be calculated and the forward operator is formed. Linear inversion techniques can then be used to invert the travel-time data to obtain "better" slowness values for the cells. This result can be used to update the slowness field, which can then be used for the next iteration.

5.1.3 Method

Given a two lithology setting, and a priori information regarding the respective slownesses of these lithologies, we want to investigate whether the technique of Chapter 2 will be useful for this non-linear inversion. The slowness range should be bounded by the minimum slowness and the maximum slowness of the lithologies. In the case of a relatively fast anomaly, we will limit the slowness range for the solution, $\mathbf{s}$, to $\mathbf{s_a} \leq \mathbf{s} \leq \mathbf{s_b}$, with $\mathbf{s_a}$ the anomalous slowness values (the same for each cell) and $\mathbf{s_b}$ the background slowness values (the same for each cell). (Relations comparing vectors should be read as relations between their corresponding components.)

We use Linear Programming to maximise the objective function

$$\phi = -|\mathbf{Gm} - \mathbf{d}| + \alpha|\mathbf{m}| \quad , \tag{5.4}$$

while satisfying certain constraints on the parameters $\mathbf{m}$. $|\cdot|$ denotes the L_1-norm. The first term of the objective function is the L_1-norm of the data residuals. The second term is analogous to the minimum-norm regularisation in an L_2-norm based inversion, with α being a kind of regularisation constant. We do not want a minimum-norm solution, but in order to prevent cells with no ray-coverage from being wrongly assigned the anomalous lithology, we enforce a minimum anomaly solution.

For a fast seismic anomaly, a positive α minimises the total amount of anomaly and a negative value maximises it. For example, assume a background wave speed of 2 km/s (slowness=0.5 s/km) and an anomalous wave speed of 3 km/s (slowness=0.33 s/km) and a slowness range limited to $[0.33, 0.5]$. The term $\alpha|\mathbf{m}|$ is then maximal when each cell has a slowness of 0.5 s/km, i.e. is assigned to the background. For a slow anomaly, it is the other way around: we have not investigated such models. In our synthetic experiments, the anomalous region has a higher wave speed, so α needs to be positive. The value of the parameter α, controlling the relative importance the minimum anomaly criterion, needs to be properly chosen.

Constraints imposed on the solution are that the parameter values should be in the interval bounded by the anomalous slowness and the background slowness, including both extremes. The constraints make up the 'feasible set'; the set of parameter values for which all constraints are satisfied.

At each iteration, we solve for the total value of the slowness field, $\mathbf{s^k}$, and not for a percentual change. The inversion result of one iteration is generally used to construct a new (updated) medium for the next round of ray-tracing. A possible problem with such an approach is that wrong intermediate models may lead to strongly varying ray coverage from iteration to iteration. To prevent such matters, we implemented an approach in which the model for ray-tracing, $\hat{\mathbf{s}}$, is obtained as follows:

$$\hat{\mathbf{s}}^k = \begin{cases} \mathbf{s_b} + \frac{k}{K-2}(\mathbf{s}^k - \mathbf{s_b}) & \text{if} \quad 0 \leq k < K-2 \\ \mathbf{s}^k & \text{if} \quad K-2 \leq k < K \end{cases}, \tag{5.5}$$

k being the iteration number and $K(> 2)$ is the total number of iterations. The initial medium ($k = 0$) for ray tracing is the model in which all cells have the background slowness s_b. In the beginning stages, the updated slowness model resembles the initial slowness model. Later in the process, the updated model progressively resembles the previous inversion result. This algorithm seems intuitively reasonable because the homogeneous background medium has a more uniform ray coverage and more varied ray orientations than any model with an anomaly.

This strategy of starting the iterations with a homogeneous medium and applying conservative updates in the beginning will also limit error propagation and (hopefully) help convergence to an acceptable solution. In the last few iteration steps (we have chosen 2), the previous inversion result is used fully as the new slowness model; the initial homogeneous model not playing any role anymore. Without such a strategy, perfect solutions cannot be achieved for noise free data (read: good solutions for good data). It will, in any case, not harm when the solution is already converging to the proper model.

Note that each iteration solves for the *total* slowness field and not for a small change to be added to the previous result. The main reason for this strategy is that in linear programming, the variables need to be positive, hence an additive update retrieved using LP can only be positive. Using such an additive update - as is mostly done - would therefore prevent a decrease in the slowness field in an LP-based formalism.

To be able to objectively compare inversion results, we introduce the percentage model misfit:

$$\text{model misfit} = \frac{100\%}{M\delta m} \sum_{j}^{M} |m_j - m_j^{true}|, \tag{5.6}$$

with m_j^{true} being the parameter value of cell j of the true synthetic earth model, as, for instance, in Figure 5.1a. δm is the exact contrast between the characteristic end-values of the parameters ($|s_a - s_b|$) and m_j is the inversion result of the j-th parameter. With this definition, if all cells got assigned either the background value or the value of the anomalous lithology, the model misfit is simply the percentage of wrong cells.

The LP problem is solved with the implementation of the Simplex algorithm, as given by Press et al. (1992). Results with clearly defined structures are expected to emerge because the optimum of the objective function (Equation 5.4) is always located at a corner of the feasible set. This will often be the value corresponding to one of the lithologies. The algorithm developed by Ditmar (1995) is used for the calculation of the ray paths.

5.1.4 Results

To test the applicability of the method from Chapter 2 for seismic first arrival travel-time data, a simple slowness model was constructed; a "diapir-like" body (+30% faster) embedded in a homogeneous background medium ($s_a = 0.35$ s/km and $s_b = 0.45$ s/km). This model of 10×10 cells is shown in Figure 5.1a. The dimensions of the 100 cells are $100 \times 1000 \times 100$ m, 1000 m being the strike length in the y-direction (perpendicular to the plane of the figure). Seismic sources are situated in a borehole on the left-side and the receivers in a bore-hole on the right. There are 10 sources and 10 receivers, between which 100 rays are traced (Figure 5.1a). The size of the cells in the strike direction is not important here, as both sources and receivers are in the same vertical plane, but will be important later for the joint inversion.

The first arrival travel-time data along the rays were scaled with a measurement error, taken to be 1.87 ms, to make the data dimensionless (Figure 5.1b). The forward operator was also scaled accordingly. Zero mean, unit variance Gaussian noise was added to the scaled travel-time observations. The expected value for the L_2-norm data misfit for this noise model is equal to the number of measurements (=100), while for the L_1-norm, this is 81 (Equations 2.10 and 2.12).

Figure 5.2 shows the model misfit of inversion results for several values of α, which sets the relative strength of the minimum anomaly criterion in the objective function (Equation 5.4). To avoid non-linearity, the seismic forward operator ($\mathbf{G^t}$) was constructed, in this case, using the true rays, i.e. those shown in 5.1a. Figure 5.2 could, of course, not be made for a real data set, because the true rays will then be unknown, as would also be the case with $\mathbf{m^{true}}$, needed to calculate the model misfit. However, the data misfit (Figure 5.2) might be used to select a proper value for α. (Similar experiments have also been carried out using noise-free data; in such cases, with a small enough α, the model could be retrieved perfectly.)

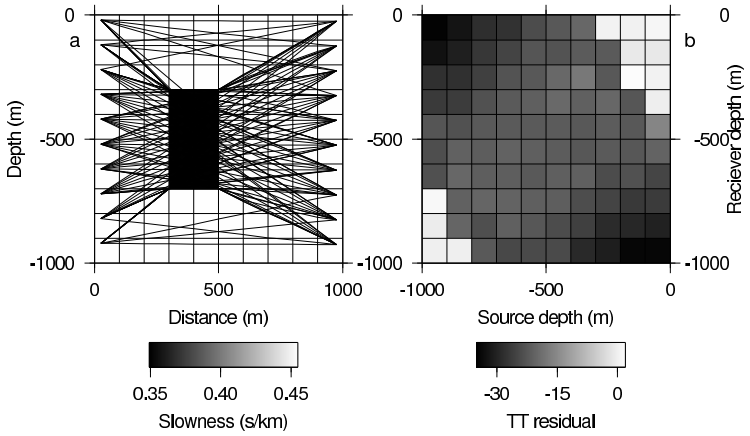

Figure 5.1: Setup for cross-well seismic tomography. The slowness model of (10×10) cells in panel a, has an anomalous zone with a velocity contrast of +30% ($s_a = 0.35$ s/km and $s_b = 0.45$ s/km). Seismic sources are situated in a borehole on the left-side (x=25 m, y=500 m) at depths of -20, -120, ... , -920 m. The borehole with receivers is on the right-side (x=925 m, y=500 m) with receivers at -25, -125, ... , -925 m depth. The ray paths from the sources to the receivers are shown. Panel b shows the scaled travel time anomaly; the differences in travel time between these rays and the rays through a homogeneous medium with background slowness, divided by a measurement error. Gaussian noise has been added to the travel time data after scaling ($\mathcal{S}/\mathcal{N}$-ratio=20). In this right plot, each row corresponds with one receiver, namely the receiver at the same depth. Each column corresponds with one source. For instance, the data in the left-most column are the travel time differences for the rays from the deepest source (-925 m).

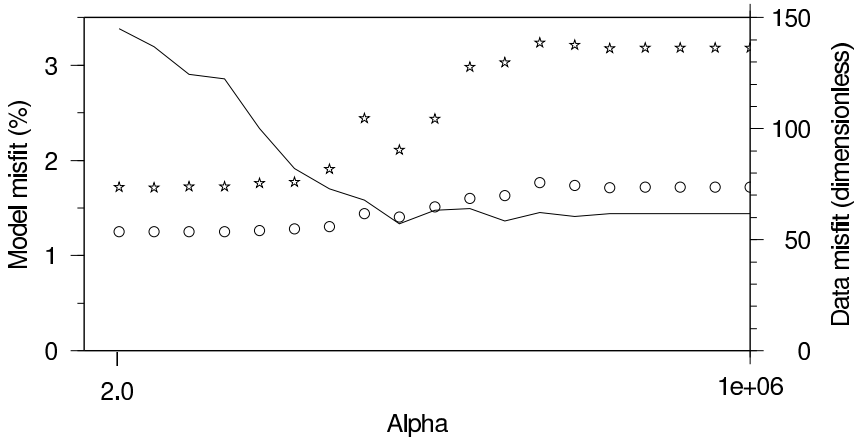

Figure 5.2: Results of the search for the best α using the true rays for the model in Figure 5.1a. The horizontal axis has a logarithmic scale. Each α is 2.09 times larger/smaller than its neighbor. The solid line is the model misfit of the inversion result (annotation on the left vertical axis). The open circles are the L_1-norm data misfit (dimensionless) of the data and the response of the inversion result (annotation on the right vertical axis). Similarly, the stars are the L_2-norm data misfit.

In a realistic case, the iterative scheme (Equation 5.5) should be used to invert the data. We have experimented with different values for K, the number of iterations; some of the resulting model misfits of 17 experiments ($K = 4, ..., 20$) are shown in Figure 5.3, using $\alpha = 1,000,000$.

The model misfit resulting after the very first iteration for all the experiments, thus using the straight rays through the homogeneous background medium, is about 8.3%. This is the left-most point in Figure 5.3, common to all the curves. Subsequent iterations improve the model misfit, which shows that ray bending has indeed to be taken into account. Furthermore, we see that using insufficient iterations leads to poor inversion results. In this experiment, the best final model misfit was found for $K = 15$, but using only 7 iterations also gave very good inversion results. For this last case $K = 7$, gradual emergence of the model through the iterative process is shown in Figure 5.5b-h. For each iteration, the retrieved models $\mathbf{s^k}$ are shown, together with the rays through $\hat{\mathbf{s}}^\mathbf{k}$, the model updated according to Equation 5.5. Due to the nature of the iterative scheme, the ray bending is more pronounced in the later iterations. Furthermore, the outline of the anomalous structure clearly becomes more apparent through the iterations.

After each iteration, the L_1-norm and L_2-norm datafit of the resulting model ($\mathbf{s^k}$) was calculated. The development of these two quantities with the iterations is shown in Figure 5.4.

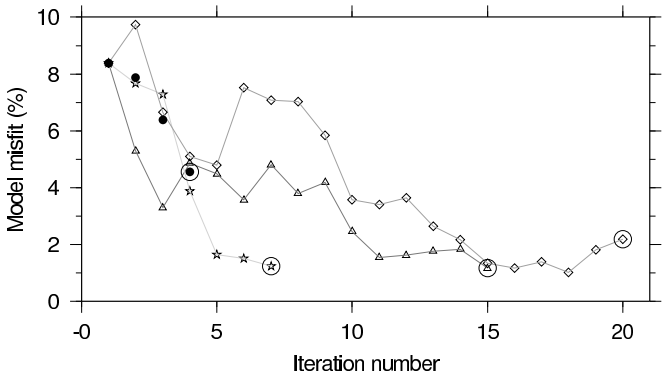

Figure 5.3: Examples of the model misfit of the inversion of the data shown in Figure 5.1 (using $\alpha = 1,000,000$), as a function of the number of iterations for 4 different values of K (Equation 5.5). The (common) left-most point is the inversion result using the (straight) rays through the homogeneous start model. The final results for each K are emphasised using open circles. The filled circles correspond to the inversion in 4 steps, the stars in 7 steps, the triangles in 15 steps and the diamonds in 20 steps.

5.1.5 Discussion

For the simple earth model, consisting of one anomalous body (Figure 5.1), the structural inversion of seismic first arrival travel-times works well. Even when noise is added to the observations, the inversion result is almost bi-modal and the true earth model can be accurately retrieved, both using the true rays and also using the iterative scheme. This occurs in spite of a rather high velocity contrast, which leads to significant ray-bending (and hence non-linearity in the inversion). (Compare, for example, the rays in the homogeneous earth model (Figure 5.5a) with the rays through the model in Figure 5.1a.) A suitable value for α, controlling the importance of the minimum anomaly criterion, must, however, be selected. In practice, the model misfit is unavailable and the selection should be made on basis of the data misfit and a priori knowledge of the noise characteristics.

The method was also tested on a more complicated earth model (Figure 5.6b) and the results (Figure 5.7d) were poor, probably due to inadequate ray-coverage. In the next part, the method is extended to perform a joint inversion of seismic data and gravity data, to explore possible improvement in the result with respect to inversion of the data sets individually.

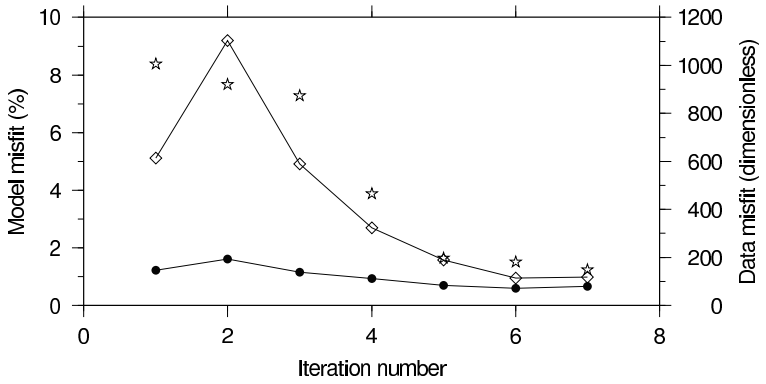

Figure 5.4: Characteristics of the intermediate models of Figure 5.5. The filled circles are the (dimensionless) L_1-norm data misfits (between predicted and observed data) corresponding to the intermediate inversion results (annotation on the right vertical axis). Similarly, the open diamonds are the corresponding L_2-norm data misfits. The stars are the model misfits of the inversion results (annotation on the left vertical axis) also shown by stars in Figure 5.3.

5.2 Part 2: Joint structural inversion

5.2.1 Introduction

Many authors have investigated joint inversion of different geophysical data types, e.g., Vozoff and Jupp (1975), Sasaki (1989), Lines et al. (1988), Haber and Oldenburg (1997), Bosch (1999) and Bosch (2001). The methods can - in general - be divided into two subgroups depending on whether or not the different data types are sensitive to the same rock property. An example of the first is joint inversion of EM-measurements and resistivity soundings, both sensitive to the conductive properties of the rocks. Joint inversion of magnetic measurements and gravity observations, which depend on the magnetic susceptibility and density, respectively, is an example of the second.

The second group further sub-divides into sequential (alternate) inversion and true joint inversion. The sequential technique works by first inverting one data type to get an image of the subsurface. Assuming that both rock properties vary at the same locations due to variation of lithologies, the second data set is inverted, using the first inversion result as the reference model, which the inversion result should resemble. After this second inversion, the first data set is inverted again, with the previous inversion result as the reference model. This procedure is continued until both data sets are fitted to some measure. See for instance Lines et al. (1988). A recent example is the work by O'Brien et al. (2005) in which sequential inversion of seismic reflection data and full tensor gravity gradiometry data was carried out to image the base of a large salt structure - to asses the extent of the K-2 oil field in the Gulf of Mexico.

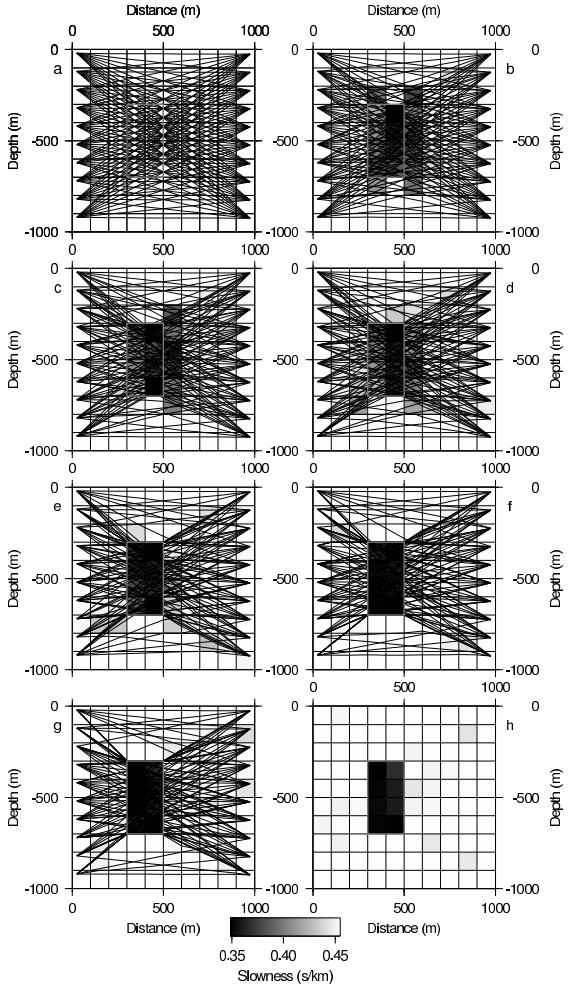

Figure 5.5: The intermediate inversion results of the inversion of the data in Figure 5.1, using $K = 7$ and $\alpha = 1,000,000$. In panel a, the starting model and the corresponding rays are shown. Panel b shows the inversion result after the first iteration. The plotted rays correspond to the medium updated according to Equation 5.5, i.e., the background medium plus 20% of the anomalous velocity contrast found. Similarly, the panels c through g each show the model ($\mathbf{s^k}$) resulting from the iteration. They also show the rays for the following iteration, traced through the updated model ($\hat{\mathbf{s}}^{\mathbf{k}}$). The model shown in panel h is the final inversion result. Anomalous slowness ($s_a = 0.35$ s/km) is given as black and the background slowness ($s_b = 0.45$ s/km) corresponds to white. In each panel, the outline of the fast anomaly is shown with a gray line.

The main drawback of the sequential inversion is that a limited resolution of the first data set could lead to a poor reference model for the second inversion, thus making the process potentially unstable without the guarantee of convergence. These problems can be overcome with a true joint inversion in which both physical properties are inverted for at the same time. In this approach, limited resolution obtainable from one data set could be compensated for by a suitably chosen second one. For this to happen the inversions should be coupled. Haber and Oldenburg (1997) did this by assuming a common structure for both parameter fields and replacing (mapping) one parameter by the other. Bosch (1999,2001) took another approach to couple the inversions; he calculated the likelihood of the combination of the physical property values, based on statistical information inferred from well-cores.

5.2.2 Approach

To improve on the results of separate inversions of different data types, they should be combined. We assume that the values of both rock properties in the sub-surface vary at the same locations, due to variations in lithology. Under this assumption, a parameter for one rock property can be replaced by a scaled version of the other. The scaling factor depends upon a physical relationship between the two properties, which can be either theoretically derived or determined by statistical analysis of empirical data (e.g. core measurements). For real lithologies, the relation between, for instance, slowness and density might not always hold for values 'in between' the lithologies. This might pose a problem for smooth inversions.

Structural inversion for two distinct lithologies, on the other hand, offers a very straightforward implementation for joint inversion, when the density and the velocity of each lithology is known. The undesirable 'in between' situations would, ideally, not occur, because most cells would be preferably assigned either the background or the anomalous lithology. The linear relation assumed between slowness and density needs thus to be valid only for both lithologies.

For the joint inversion of seismic travel-times and gravity data, we need to adopt a new notation for the formula in Chapter 2 regarding gravity: The linear equations (in matrix form) relating density contrasts with gravity observations (Equation 2.3) can be written as:

$$\mathbf{G^g}\boldsymbol{\rho} = \mathbf{d^g} \quad , \tag{5.7}$$

where $\mathbf{d^g}$ is the column vector with the N^g (dimensionless) gravity observations, $\mathbf{G^g}$ is the (scaled) $N_g \times M$ forward operator matrix and $\boldsymbol{\rho}$ is the column vector containing the density contrasts of the M cells. The superscript *g* is used for symbols related to gravity, to differentiate it from the second data type (travel times, superscript *t*).

For the joint inversion, we replace $\boldsymbol{\rho}$ (the densities contrasts) by $\mathbf{s}$ (the absolute slowness parameters), using the mapping (for a homogeneous background):

$$\begin{aligned} \rho_j &\rightarrow s_j \\ d_i^g &\rightarrow \mathcal{D}_i^g = \frac{\delta s}{\delta \rho} d_i^g + \sum_j G_{ij}^g s_b \quad , \end{aligned} \tag{5.8}$$

$\mathcal{D}_i^g$ being the adjusted and scaled gravity observations, δs is the slowness contrast of the two lithologies and $\delta\rho$ is their density contrast. The above scaling of the data ensures that $\mathbf{G^g s} = \boldsymbol{\mathcal{D}}^g$; we can now use the gravity forward operator on slowness parameters.

This allows for a reduction in the number of variables and is the implementation of the coupling between the two inversions. Using the above notation, and writing

$$\mathbf{G} = \begin{pmatrix} \mathbf{G}^t \\ \mathbf{G}^g \end{pmatrix} \quad \text{and} \quad \mathbf{d} = \begin{pmatrix} \mathbf{d}^t \\ \boldsymbol{\mathcal{D}}^g \end{pmatrix} \tag{5.9}$$

and calling the vector of unknown parameter values, $\mathbf{m}$, (slowness values in this case) we can put it all together in the well known compact notation:

$$\mathbf{Gm} = \mathbf{d} \quad . \tag{5.10}$$

This relationship is similar to the one in Chapter 2 and the same formalism for structural inversion can be used here too.

Berryman (1990) suggested the use of another constraint - that the sum of the observed values should be equal to the sum of the forward modeled data:

$$\sum_i d_i = \sum_{i,j} G_{ij} m_j. \tag{5.11}$$

The idea behind this constraint is that the noise in the observations tends to cancel out when summed (assuming that the noise has zero mean). We will call this the *constant data sum criterion* and have included it in the joint inversion scheme only for the gravity observations. This could improve the solution because the gravity operator is truly linear. This constraint has not been applied to the seismic travel-times, because the seismic forward operator is calculated given an estimate of the slowness model. When this 'estimated' operator is used to enforce the constant data sum criterion, it could do more harm than good.

5.2.3 Results

In Chapter 2 we discussed the need to know the depth to the top of a density anomaly for successful application of our method for structural inversion using gravity data. We stated that seismic data could provide such information. In the structural inversion of travel-time data above, for the simple earth model, good inversion results were obtained without the need for additional data. In Figure 5.6b, we present a more difficult situation, two

anomalous structures at different depth ($s_a = 0.35$ s/km) in a homogeneous background medium ($s_b = 0.45$ s/km), having a density contrast of -0.4 g/cc. The model consists of 20×15 cells of $50 \times 1000 \times 50$ m each (Strike length is 1000 m). Rays will be again traced between a series of sources in a bore-hole on the right and receivers in a bore-hole on the left.

Figure 5.6a shows the theoretical gravity values along the line y=500, at a height of 2 m above the surface, due to the two anomalies. The data has been divided by an assumed measurement error of 0.04 mGal. Zero mean, unit variance Gaussian noise has been added to the observations ($\mathcal{S}/\mathcal{N}$-ratio = 19) and are also shown. The forward operator $\mathbf{G^g}$ has been scaled accordingly.

The travel-time anomalies (signal) are shown in Figure 5.6c. They correspond to the travel-times along the (true) rays of Figure 5.6b minus the travel-times along (straight) rays through a homogeneous medium consisting of background slowness. The travel-times have been scaled with a measurement error of 1.875 ms ($\mathcal{S}/\mathcal{N}$-ratio of the travel-time signal = 20). The forward operator $\mathbf{G^t}$ has also been scaled accordingly.

Results for six inversion experiments are presented in Figure 5.7; the first two are inversions of only the gravity data, with and without an assumed minimum depth to the top. The data sum constraint was *not* applied here. The next two experiments are the inversions of the seismic data alone[1], using the true rays (not known in practice) and the iterative scheme, respectively. The last two experiments are joint inversions of both the seismic data set and the gravity data set, again using both the true rays and the iterative scheme.

In order to enable a better comparison of the results, given a fixed parameterisation, extra observations were added in the first 4 experiments mentioned above (Figure 5.7a-d). Otherwise, those inversions would solve for more parameters than there were observations; they would be highly under-determined problems. In the experiments 1 and 2, with gravity data alone, 200 extra gravity data values were inserted along the line of measurements given above, reducing the distance between the observations to 5 m. In experiments 3 and 4, using first arrival times, 15 extra seismic receivers were simulated in the borehole: The receiver spacing became 25 m. Thus increasing the total of number of seismic rays to 450.

For the experiments 5 and 6, the joint inversions, the density contrast ($\delta\rho = -0.4$ g/cm^3) and the seismic velocities (2.2 km/s and 2.86 km/s) were used to transform the density distribution to the slowness distribution (Equation 5.8). The seismic data and the - scaled and adjusted - gravity data were combined to perform the joint inversion.

The structural inversion of gravity data alone, without information regarding the depth of the top, is presented in Figure 5.7a, Additional gravity observations were added for this inversion, as mentioned above, to have the same amount of observations as parameters. The same data was then inverted using the extra constraint that the minimum depth to the top should be 150 m, which is the correct value for the anomalous body on the right. This result is shown in Figure 5.7b. The large structure on the right of the model is much better

[1] Due to the provisional implementation, it is necessary to include 2 gravity observations. The two leftmost observations were selected, because the model is the least sensitive to those.

resolved now. The smaller body, whose top was not provided, is still not imaged correctly. This result confirms our claim that the depth to the top should be known a priori for this method. The lateral positioning of both the structures is, however, reasonably.

The seismic travel-time data was inverted 'separately' as well (using only two distant gravity observations). The results are in Figure 5.7c and d. Again, extra receivers were added (now 30 receivers, with 25 m spacing) to have more observations than parameters. The inversion result using the true rays, and the optimal value for $\alpha(= 54898)$, is shown in Panel c. The small anomalous structure on the left is imaged adequately using the true rays. The depth of the structures is retrieved well. α was selected on basis of the model misfit. Several inversions were performed using the iterative scheme with different values for K. Panel d shows the best inversion result found (based on the model misfit). In these experiments, $\alpha = 1$. The iterative scheme performed poorly.

The two joint inversions of the 225 seismic travel-times and the 100 gravity observations are shown in Figure 5.7e (using the true rays) and f (using the iterative scheme). Both inversions show improvements, compared with the corresponding individual inversions. In particular, the joint inversion using the true rays is very structured and closely resembles the true model.

5.3 Discussion & outlook

For the model with two anomalous bodies (Figure 5.6b), the structural inversion of only travel-time data, suffers from the poor ray-coverage; there are some cells without rays and in the cells situated between the anomalies, there is not much variation in the ray-angles. The joint inversion of travel-time data and gravity data yields superior inversion results when using the true rays. The results presented using the iterative scheme do improve on the individual inversions, but the resulting model is still not satisfactory.

There are several possibilities for improvement worth exploring:

- Only a limited number of choices for the regularisation parameter, α and the number of iterations, K were used. A proper search for combinations of α and K could lead to better results. Further research is needed to find a method to select the proper values for α and K for a field data set, when the true model is not available for computation of the model misfit.

- The construction of the model for ray-tracing during the iterative process could be controlled more strictly. For a homogeneous background, one possibility could be to limit the size of the model update based on the change-of-length of the rays; straight rays have the shortest geometric ray-path. Putting limits on the change-of-length could help stabilise the solutions. Implementing this presents a difficulty, though; the travel-time along a ray is a smooth function of the slowness field, whereas the ray-path length could be a discontinuous function of the latter. Another option to ensure a smooth increase in ray-path length is scaling of the forward operator itself, but that would correspond to non-physical rays. Additional options

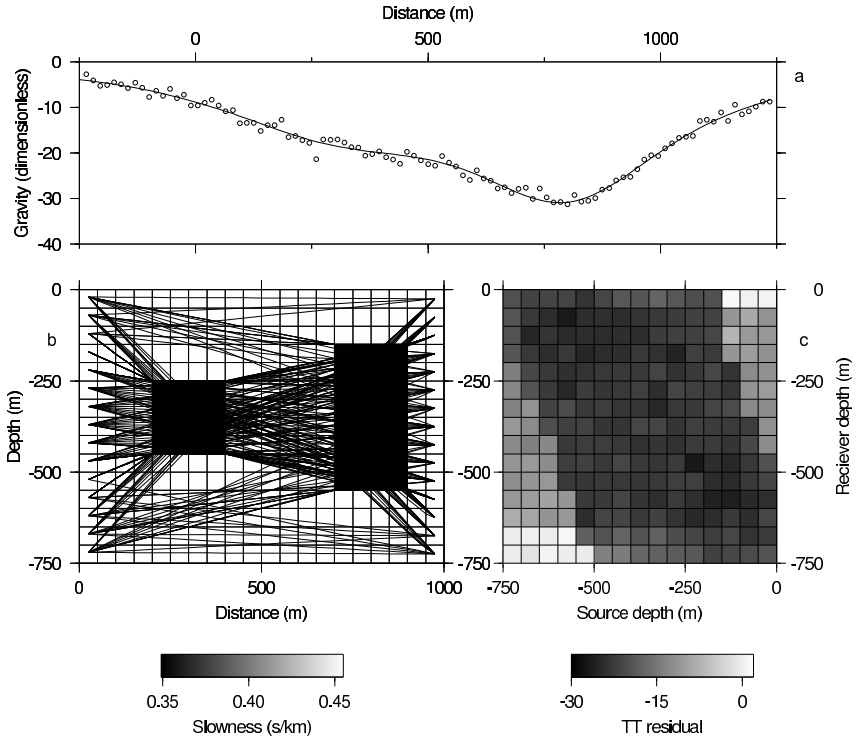

Figure 5.6: A synthetic model (20×15) cells, consisting of two lithologies. The slowness model is shown in b. This can be thought of as two salt structures (v=+30% and the density contrast, $d\rho = -0.4$g/cm^3) embedded in a sandstone medium ($v = 2.2$ km/s). The sources are on the left, spaced equally at 50 m depth intervals, starting at -20 m depth. The receivers are on the right, equally spaced from -25 m to -725 m depth. The 225 rays are shown and their (scaled) travel time residuals with a homogeneous background medium are shown in panel c. Each row corresponds to a receiver and each column to a source. Zero mean, unit variance Gaussian noise ($\mathcal{S}/\mathcal{N} = 20$) has been added to the data. In panel a, the gravity data, acquired at a height of 2 m above the surface, due to the anomalous mass (at the same location as the slowness anomaly), in dimensionless units. The solid line is the noise free data. The open circles are the observations corrupted with zero-mean, unit variance Gaussian noise ($\mathcal{S}/\mathcal{N} = 19$).

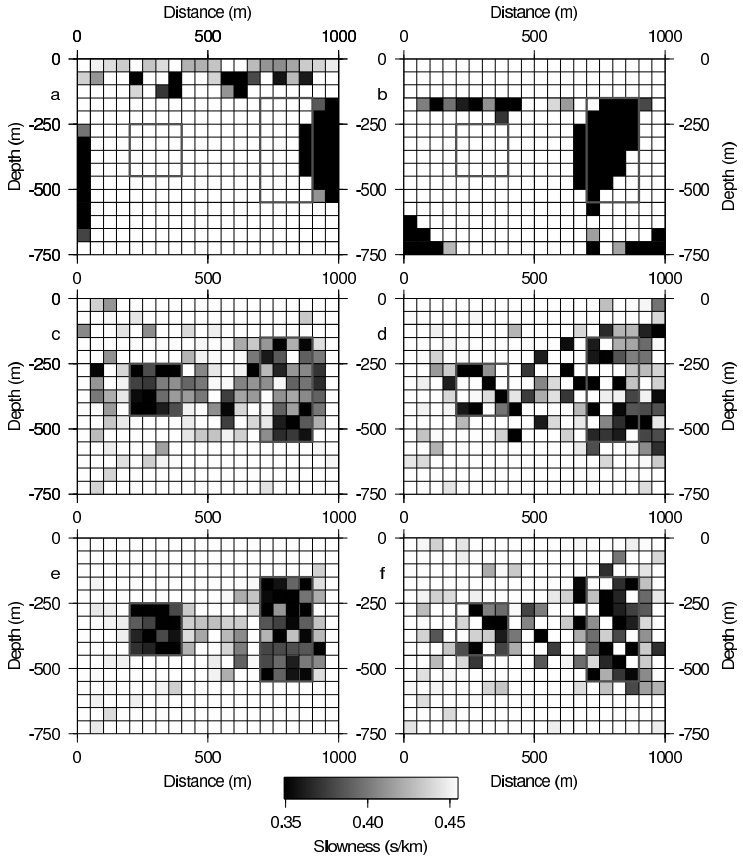

Figure 5.7: Separate and joint inversion results. In panel a, the inversion result using only gravity data (300 observations). (The density contrast is scaled to the slowness anomaly for plotting using the same gray-scale.) Panel b shows the inversion result of the same data as in panel a, but using a minimum depth to the top of 150 m. Panel c shows the inversion result using travel time data along 450 rays (15 sources and 30 receivers) and two distant gravity stations. The true rays were used to construct the seismic forward operator. This is the inversion result with $\alpha = 54898$, which was found to be optimal. In Panel d, the inversion result of the travel-time data used in c, but now using the iterative scheme with $K = 13$ and $\alpha = 1.0$. The joint inversion result of the data from Figure 5.6 (325 observations), using the true rays, is shown in panel e. For this joint inversion, $\alpha = 337$. The inversion result using the iterative scheme with $K = 9$ and $\alpha = 1.0$ is shown in panel f. In each panel, the outline of the two anomalous (fast) regions is given with the gray line.

of finding the size for the update for the ray-trace model are possible, e.g., Berryman (1990), who determined the number of rays, that had computed travel-times faster than the observed ones. The model update which minimised the number of such violations, was chosen.

- Extra constraints could be added for real gravity data; one could, for example, allow only for a band of values, constraining the predicted gravity data to be within a range, for instance, ± 3 standard deviations of the noise, around each observation. One advantage of this approach is that outliers would be detected, because these outliers will cause incompatible constraints and thus results in an empty feasible set. The outliers could subsequently be removed from the data set. Another advantage would be that the gravity data could then be 'better' fitted: The L_1-norm is much less sensitive to a large individual residual (possibly due to an erroneous outlier) than the L_2-norm. The extra constraints would act as the L_∞-norm for large residuals only. This combination of the L_1-norm and the L_∞-norm could provide the 'better' data misfit. This should only be applied to the gravity observations and not to the travel-times. The seismic forward operator being only an estimate could cause large residuals, which in turn could be falsely treated as outliers by the algorithm.

To continue the discussion in a broader sense, joint structural inversion could perhaps be implemented differently than described in this chapter; we will mention two possibilities. The first approach could be based on the inversion method in Chapter 4 with the L_2-norm. The inversions could now be combined by replacing one set of parameters with a scaled version of the other. As shown earlier, that method also produces bi-modal images of the subsurface, overcoming the non-physical combination of rock properties 'in between' the lithologies.

A totally different way of combining the separate inversions could be by not replacing (scaling) one rock property by the other, and thus enforcing a common structure, but by defining a likelihood, as a function of both physical properties. This likelihood function should describe how well a combination of the two properties satisfies the a priori information regarding the rock properties of the lithologies assumed present in the volume of interest. The likelihood should be maximal for the perfect combinations and low in between. The added advantage here would be that structural inversion for more than two lithologies would be possible. Of course, the disadvantage would be that each cell will have more than one parameter associated with it; one for each physical property inverted for. This would make the computations more expensive.

The joint likelihood function could be implemented into the formalism using the iterative approach in Chapter 4. The new function would replace the structural regularisation function ϕ (Equation 4.5) with parabolas with their minima at the most likely lithology (all parabolas will have their minimum at their respective physical property value for that lithology). Furthermore, the function t (Equation 4.8) should be replaced with a similar function of both physical properties; unity when the parameters are equal to the physical rock properties and decreasing to zero in between. An example of such a function is given in the gray-scale Figure 5.8 for 3 lithologies..

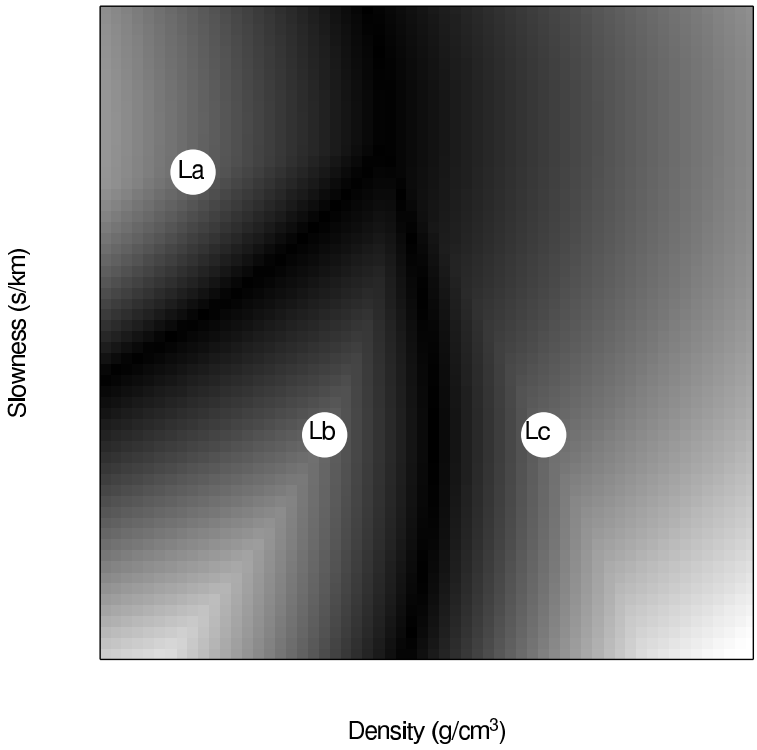

Figure 5.8: A possible weighting function to set the steepness of the structural regularisation parabolas for two physical rock properties, namely density and slowness. The three labels 'La', 'Lb' and 'Lc' mark the a priori information on the rock property values for three lithologies in the plot. Black corresponds to zero and the function is one at each lithology. The weighting function should depend on the statistical correlations between slowness and density for the lithologies and the consideration that the weighting should be small 'in between' lithologies.

The Linear Programming method might also be used for joint structural inversion for more than two lithologies. Instead of the joint likelihood function described above, the inversions could in this case be coupled by limiting the feasible set with extra constraints. For each cell, constraints like the ones shown in Figure 5.9, could be used. The lithologies (here three) are situated in the corners of the feasible set. The maximum of an objective function is always situated on a corner of the feasible set. So, if the other constraints do not interfere, the solution should yield only those combinations of the physical properties (e.g. slowness and density) of the lithologies present. Other constraints should of course be applied as well, to enable an absolute value objective function.

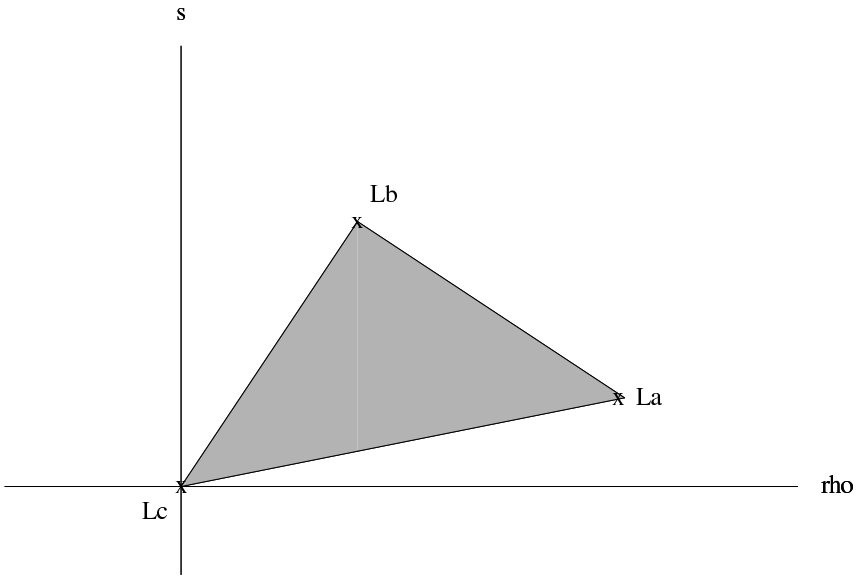

Figure 5.9: Constraints that make up a feasible set. The lines connecting the three lithologies ‘La’, ‘Lb’, ‘Lc’ are the extra constraints to couple a joint inversion using LP in the case of structural inversion for three lithologies. The gray area is the feasible set due to these constraints. Other lithologies can be added when their combination of slowness and density is outside of the feasible set spanned by the other lithologies. (The method for joint inversion with substitution of one property by the other can be seen as a feasible set consisting only of the line connecting two lithologies.)

Chapter 6

Summary

The aim of this PhD project was to develop a method for implicit structural inversion of geophysical data, using a gridded model. With the term *structural inversion* we mean an inversion to obtain an image of the (gridded) sub-surface, which can directly be interpreted in terms of sharp boundaries between different lithologies. *Implicit* indicates that we neither assume a fixed number of anomalies, nor enforce a fixed structure explicitly. We model the sub-surface with homogeneous, 2.5-D, rectangular prisms. The values of the physical property of the cells are the parameters inverted for. The structures in the model should be apparent from large contrasts in parameter values between neighboring cells. All the models investigated comprised of two lithologies, with one or more homogeneous anomalous regions embedded within a homogeneous background. In Chapter 2 we developed a method for such an implicit structural inversion, using Linear Programming. The latter is a mathematical technique to optimise a linear function of the parameters (the objective function) while satisfying linear constraints on them. These constraints define a high-dimensional subspace in which all constraints on the parameters are satisfied. This subspace is called the *feasible set* for the problem. Due to the linearity of the objective function, the optimum is always located in a corner of the feasible set. This is the property that makes Linear Programming suitable for structural inversion. When the contrast between rock properties for certain lithologies is known, that information can be used as constraints on the parameters. This forms a feasible set, for which the corners coincide with the rock property of one of the lithologies assumed present in the sub-surface. If other constraints do not interfere, each cell will ultimately be identifiable as belonging to one of the lithologies.

We used the sum of the absolute values of the data residuals (the L_1-norm), as the objective function. This is a non-linear function, but it can be implemented in the Linear Programming framework with the use of extra (dummy) variables. The method was tested on synthetic gravity data. The results compared favorably with a traditional inversion using truncated Singular Values. The new method was able to reconstruct important features of the test models. As a part of the study, we were able to also develop a criterion for the optimal parameterisation of the region of interest. Although our LP-based method

is able to reduce the inherent ambiguity of gravity inversion significantly, it does have the limitation that the minimum depth to the top needs to be known a priori.

In Chapter 3, the method from Chapter 2 was extended to simultaneously invert for both a linear trend in the data and for the density contrasts of the cells. The method was used to invert a field gravity data set, which had been acquired by the industry over a basin, to infer the depth to basement. The results of this inversion were robust, especially given that the aim was to obtain information regarding depressions in an otherwise dominantly sub-horizontal structural setting. The results are in agreement with the ones obtained by the industry using detailed forward modeling.

The same LP-based method was applied to invert synthetic seismic cross-well first arrival travel-times for the absolute inter-well slowness distribution in Chapter 5. This is a non-linear problem due to ray-bending in a heterogeneous medium. An iterative scheme is proposed for solving the problem. For a simple sub-surface model, the implicit structural inversion using the true rays, or using the iterative scheme, both produced very good models of the sub-surface. However, a regularisation parameter needs to be chosen properly. A more complicated model with poor ray-coverage, was difficult to retrieve using only seismic data, however, a joint inversion of seismic data (using the true rays) together with gravity data gave encouraging results.

Parallel to the LP-based approach using the L_1-norm, another method was developed in Chapter 4, this time using the L_2-norm of the data misfit. In an iterative scheme, a reference model was constructed, using information regarding the density contrasts, which the inversion result should resemble. The method needs 5 tuning constants; when proper (wide) ranges are supplied, a parameter search for the constants, yields their optimal values to be used for the inversion. One of the tuning parameter is dynamically increased during the iterations to ensure a result with clear structure. This search can be steered using the resulting data misfits. Experiments on synthetic gravity data gave good inversion results, even when the top of the anomaly was unknown.

Appendix A gives an illustration of the applicability of geophysical techniques in non-standard situations. The feasibility of using gravity measurements to assess the layered structure that develops with time inside a blast-furnace was explored. Detailed simulations yielded statistical information regarding the reliability of the thickness estimates that can be obtained from micro-gravity measurements. Such information, pertaining to different materials (densities) that form inside the blast-furnace near its base during its operation, can be used to set up a monitoring regime, and is - potentially - of great economic significance.

Bibliography

Acar, R. and Vogel, C. (1994). Analysis of total variation penalty methods. *Inverse problems*, 10:1217–1229.

Auken, E. and Christiansen, A. V. (2004). Layered and laterally constrained 2D inversion of resistivity data. *Geophysics*, 69(3):752–761.

Barbosa, V. C. F., Silva, J. a. B. C., and Medeiros, W. E. (1997). Gravity inversion of basement relief using approximate equality constraints on depths. *Geophysics*, 62(6):1745–1757.

Barbosa, V. C. F., Silva, J. a. B. C., and Medeiros, W. E. (2002). Practical applications of uniqueness theorems in gravimetry: Part II–Pragmatic incorporation of concrete geologic information. *Geophysics*, 67(3):795–800.

Berge, P. A., Berryman, J. G., Bonner, B. P., Roberts, J. J., and Wildenschild, D. (1997). Joint inversion of geophysical data for site characterization and restoration monitoring. Annual progress report for emsp, Lawrence Livermore National Laboratory.

Berryman, J. G. (1989). Germat's principle and nonlinear traveltime tomography. *Physical Review Letters*, 62(25):2953–2956.

Berryman, J. G. (1990). Stable iterative reconstruction algorithm for nonlinear traveltime tomography. *Inverse problems*, 6:21–42.

Bertete-Aguirre, H., Cherkaev, E., and Oristaglio, M. (2002). Non-smooth gravity problem with total variation penalization functional. *Geophysical Journal International*, 149:499–507.

Bosch, M. (1999). Lithologic tomography: From plural geophysical data to lithology estimation. *journal of geophysical research*, 104(B1):749–766.

Bosch, M. (2001). Joint inversion of gravity and magnetic data under lithologic constraints. *The leading Edge*, pages 877–881.

Brown, P. N. and Saad, Y. (1990). Hybrid krylov methods for nonlinear systems of equations. *SIAM J. sci. stat. comput.*, 11(3):450–481.

Brown, R. F. and Brown, B. W. (1992). *Finite Mathematics*. Ardsley House, Publishers, Inc.

Camacho, A. G., Montesinos, F. G., and Vieira, R. (2000). Gravity inversion by means of growing bodies. *Geophysics*, 65(1):95–101.

Cuer, M. and Bayer, R. (1980). Fortran routines for linear inverse problems. *Geophysics*, 45:1706–1719.

Curtis, A. (2004). Theory of model-based geophysical survey and experimental design. *The Leading Edge*, pages 997–1004.

de Groot-Hedlin, C. and Constable, S. (2004). Inversion of magnetotelluric data for 2D structure with sharp resistivity contrasts. *Geophysics*, 69(1):78–86.

Ditmar, P. (2002). Finding the shape of a local heterogeneity by means of a structural inversion with constraints. *Geophysical Prospecting*, 50:209–223.

Ditmar, P. G. (1995). An approach to 2-D ray tracing intended for non-linear seismic tomography. In *Proceedings of the Fourth International Symposium. Novosibirsk, Russia. M.M.Lavrent'ev (ed.). VSP*, pages 138–154, Utrecht, The Netherlands.

Esparza, F. J. and Gómez-Treviño, E. (1997). 1-D inversion of resistivity and induced polarization data for the least number of layers. *Geophysics*, 62:1724–1729.

Farquharson, C. and Oldenburg, D. (2003). Constructing piece-wise-constant models using general measures in non-linear, minimum-structure inversion algorithms. In *Extended abstracts,International symposium of the Society of Exploration Geophysicists of Japan, Tokyo, Japan*, pages 240–243.

Gidas, B. (1985). Nonstationary markov chains

and convergence of the annealing algorithm. *Journal of Statistical Physics*, 39:73–131.

Golub, G. H., Heath, M., and Wahba, G. (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223.

Gómez-Treviño, E., Esparza, F. J., and Méndez-Delgado, S. (2002). New theoretical and practical aspects of electromagnetic sounding at low induction numbers. *Geophysics*, 67(5):1441–1451.

Guillen, A. and Menichetti, V. (1984). Gravity and magnetic inversion with minimization of a specific functional. *Geophysics*, 49:1354–1360.

Guspi, F. (1993). Noniterative nonlinear gravity inversion. *Geophysics*, 58(7):935–940.

Haber, E. and Oldenburg, D. (1997). Joint inversion: a structural approach. *Inverse Problems*, 13:63–77.

Hansen, P. C. (1992). Analysis of discrete ill-posed problems by means of the l-curve. *SIAM Review*, 34(4):561–580.

Huestis, S. P. (1996). The use of linear programming in the construction of extremal solutions to linear inverse problems. *SIAM Review*, 38(3):496–506.

Jackson, D. D. (1972). Interpretation of inaccurate, insufficient and inconsistent data. *Geophysical Journal of the Royal Astronomical Society*, 28:97–109.

Jupp, D. and Vozoff, K. (1975). Stable iterative methods for the inversion of geophysical data. *Geophysical Journal of the Royal Astronomical Society*, 42:957–976.

Kilmer, M. E. and O'Leary, D. P. (2001). Choosing regularization parameters in iterative methods for ill-posed problems. *SIAM J. Matrix Anal. Appl.*, 22(4):1204–1221.

Kusche, J. and Klees, R. (2002). Regularization of gravity field estimation from satellite gravity gradients. *Journal of Geodesy*, 76:359–368.

Kusche, J. and Mayer-Guerr, T. (2001). On the iterative solution of ill-conditioned normal equations by the use of lanczos methods. In *Extended abstracts, IAG 2001 Scientific Assembly, Budapest, Hungary*.

Last, B. and Kubik, K. (1983). Compact gravity inversion. *Geophysics*, 48:713–721.

Lévêque, J.-J., Rivera, L., and Wittlinger, G. (1993). On the use of the checker-board test to assess the resolution of tomographic inversions. *Geophysical Journal International*, 115:313–318.

Li, Y. and Oldenburg, D. (1998). 3-D inversion of gravity data. *Geophysics*, 63:109–119.

Li, Y. and Oldenburg, D. (2000). Joint inversion of surface and three-component borehole magnetic data. *geophysics*, 65(2):540–552.

Lines, L., Schultz, A., and Treitel, S. (1988). Cooperative inversion of geophysical data. *Geophysics*, 53:8–20.

Matsu'ura, M. and Hirata, N. (1982). Generalized least-squares solutions to quasi-linear inverse problems with *a priori* information. *J. Phys. Earth*, 30:451–468.

O'Leary, D. P. (2001). Near-optimal parameters for tikhonov and other regularization methods. *SIAM J. Sci. Comput.*, 23(4):1161–1171.

Menke, M. (1984). *Geophysical data analysis: Discrete inverse theory*. Academic Press Inc.

Mirzaei, M. (1996). *Inversion of potential field data*. PhD thesis, Utrecht University. Geologica Ultraiectina, Vol 143.

Mirzaei, M., Bredewout, J., and Snieder, R. (1995). Gravity data inversion using the subspace method. In Gottlieb, J. and DuChateau, P., editors, *Parameter identification and inverse problem in Hydrology, Geology and Ecology*. Kluwer Academic Publishers.

Mundim, K. C., Lemaire, T. J., and Bassrei, A. (1998). Optimization of non-linear gravity models through generalized simulated annealing. *Physica A*, 252:405–416.

Musil, M., Maurer, H., and Green, A. (2003). Discrete tomography and joint inversion for loosely connected or unconnected physical properties: application to crosshole seismic and georadar data sets. *Geophysical Journal International*, 153:389–402.

Nabighian, M., Ander, M., Grauch, V., Hanson, R., LaFefr, T., Li, Y., Pearson, W., Peirce, J., Phillips, J., and Ruder, M. (2005). Historical

development of the gravity method in exploration. *Geophysics*, 70:63ND – 89ND.

Nagihara, S. and Hall, S. A. (2001). Three dimensional gravity inversion using simulated annealing: Constraints on the diapiric roots of allochthonous salt structures. *Geophysics*, 66(5):1438–1449.

Nagy, D. (1966). The gravitational attraction of a right rectangular prism. *Geophysics*, 31:362–371.

O'Brien, J., Rodriguez, A., Sixta, D., Davies, M. A., and Houghton, P. (2005). Resolving the K-2 salt structure in the gulf of mexico, an integrated approach using prestack depth imaging and full tensor gravity gradiometry. *The leading Edge*, pages 404–409.

Oldenburg, D. and Ellis, R. (1993). Efficient inversion of magnetotelluric data in two dimensions. *Physics of the earth and planetary interiors*, 81:177–200.

Oldenburg, D. W. and Ellis, R. (1991). Inversion of geophysical data using an approximate inverse mapping. *Geophysical Journal International*, 105:325–353.

Oldenburg, D. W., McGillivray, P. R., and Ellis, R. G. (1993). Generalized subspace methods for large-scale inverse problems. *Geophysical Journal International*, 114:12–20.

Parker, R. L. (1974). Best bounds on density and depth from gravity data. *Geophysics*, 39(5):644–649.

Parker, R. L. (1975). The theory of ideal bodies for gravity inversion. *Geophysical J. R. astr. Soc.*, 42:315–334.

Pilkington, M. and Crossley, D. (1986). Determination of crustal interface topography from potential fields. *Geophysics*, 51(6):1277–1284.

Portniaguine, O. and Zhdanov, M. S. (1999). Focusing geophysical inversion images. *Geophysics*, 64:874–887.

Portniaguine, O. and Zhdanov, M. S. (2002). 3-D magnetic inversion with data compression and image focusing. *Geophysics*, 67:1532–1541.

Press, W., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. (1992). *Numerical Recipes in C. The Art of Scientific Computing, 2nd. edition*. Cambridge University Press.

Rao, D. B. and Prakash, M. J. (1997). Fortran-77 computer programs for 2d and 2.5D analysis of gravity anomalies with variable density contrast. *Bollettino di Geofisica Teorica ed Applicata*, 38:187–209.

René, R. M. (1999). Three-dimensional "shape-of-anomaly" inversion of gravity and magnetic fields. In *Extended abstracts,69th Annual meeting International of the Society of Exploration Geophysicists*, pages 417–420.

Richard, V., Bayer, R., and Cuer, M. (1984). An attempt to formulate well-posed questions in gravity: Application of linear inverse techniques to mining exploration. *Geophysics*, 49(10):1781–1793.

Sabatier, P. C. (1977a). Positivity constraints in linear inverse problems –i. general theory. *Geophys. J. R. astr. Soc.*, 48:415–441.

Sabatier, P. C. (1977b). Positivity constraints in linear inverse problems –ii. applications. *Geophys. J. R. astr. Soc*, 48:443–459.

Safon, C., Vasseur, G., and Cuer, M. (1977). Some applications of linear programming to the inverse gravity problem. *Geophysics*, 42:1215–1229.

Sambridge, M. (1999). Geophysical inversion with a neighbourhood algorithm-1. searching a parameter space. *Geophys. J. Int*, 138(2):479–494.

Sasaki, Y. (1989). Two-dimensional joint inversion of magnetotelluric and dipole-dipole resistivity data. *Geophysics*, 54(2):254–262.

Sasaki, Y. (2001). Full 3-D inversion of electromagnetic data on pc. *Journal of Applied Geophysics*, 46:45–54.

Sheriff, R. and Geldart, L. (1995). *Exploration Seismology*. Cambridge University Press, 2nd edition.

Silva, J. a. B. C., Medeiros, W. E., and Barbosa, V. C. F. (2002). Practical applications of uniqueness theorems in gravimetry: Part I–Constructing sound interpretation methods. *Geophysics*, 67(3):788–794.

Smith, T., Hoversten, M., Gasperikova, E., and Morrison, F. (1999). Sharp boundary inver-

sion of 2D magnetotelluric data. *Geophysical Prospecting*, 47:469–486.

Spakman, W. (1993). *Seismic Tomography: Theory and practice*, chapter Iterative strategies for non-linear travel time tomography using global earthquake data, pages 190–226. Chapman & Hall, London.

Spakman, W. and Bijwaard, H. (2001). Optimization of cell parameterizations for tomographic inverse problems. *Pure appl. geophysics*, 158:1401–1423.

Tarantola, A. and Valette, B. (1982). Inverse problems = quest for information. *Journal of Geophysics*, 50:159–170.

Telford, W., Geldart, L., and Sheriff, R. (1990). *Applied Geophysics 2nd ed.*. Cambrigde University Press, 2 edition.

Tikhonov, A. N. and Arsenin, V. Y. (1977). *Solutions of ill-posed problems*. V.H.Winston and Sons, Washington, D.C.

Tsallis, A. and Stariolo, D. A. (1996). Generalized simulated annealing. *Physica A*, 233:395–406.

van Zon, A. T. and Roy Chowdhury, K. (2006). Structural inversion of gravity data using linear programming. *Geophysics, Accepted*.

Vernant, P., Masson, F., Bayer, R., and Paul, A. (2002). Sequential inversion of local earthquake traveltimes and gravity anomaly–the example of the western alps. *Geophysical Journal International*, 150:79–90.

Vozoff, K. and Jupp, D. B. (1975). Joint inversion of geophysical data. *Geophys. J. R. astr. Soc.*, 42:977–991.

Yamamoto, T., Nye, T., Kuru, M., and Sasaki, T. (1995). Super crosswell tomography for imaging the hydraulic structure of the earth. In *Extended abstracts, Annual meeting International of the Society of Exploration Geophysicists*, pages 450–452.

Zhdanov, M. S. and Hursan, G. (2000). 3-D electromagnetic inversion based on quasi-analytical approximation. *Inverse Problems*, 16:1297–1322.

Appendix A

Feasibility study for investigating the internal structure of a blast furnace with micro gravity measurements

A.1 Motivation

After several years of iron production in a blast furnace, its heat resistant internal layers are worn out. When these protecting layers get too thin, they need to be replaced. The down-time is expensive.

The base of a blast furnace is also protected with heat-resistant layers, which need to be replaced when they are worn out. However, sometimes there is a malfunction and the furnace is shut down temporarily. Each time a blast furnace is shut down, some iron (and possibly slag) solidifies at the bottom and will not become liquid again. A very tough, solid layer is formed then in the furnace. This needs to be removed before repairs can begin. To estimate the down-time of the blast furnace for repairs more accurately, a good estimate of the size of this remaining solid is needed. This estimate can have a major economic impact in the planning of the down-time.

[1] This chapter is a slightly modified version of the final technical report of a contract-research, delivered to the blast furnace division of the Corusgroup, IJmuiden/The Netherlands, and is included here with their permission. Authors: Tim van Zon and Kabir Roy Chowdhury. Applied geophysics, Utrecht University

A.2 Introduction

We were invited by Corus to join a brainstorm session in which we discussed possible techniques like gravity, acoustic, thermal images and seismics, to estimate the structures inside the blast furnace, with measurements carried out outside the steel casing.

Figure A.2 shows a cross section of the residue found during the last repair of blast-furnace HO-6, which illustrates the nature of the problem. The solid body has several layers with a large density difference. The thicknesses of these layers are of interest. The situation inside another blastfurnace - which may be shut down in the future - is expected to be similar although not the same.

Acoustics and seismics might be able to estimate the thicknesses of the layers. A layer can be detected with seismics if the product of density times (sound)velocity changes between layers. Part of the wave energy is then reflected back. However, the construction of the blast furnace already has several transitions. It is not clear whether enough energy can be transmitted to the remaining solid, to get clear reflections.

Gravity profiling might be worthwhile too, because of the large density contrasts found in HO6. Given an estimate of the densities of the layers, a line of gravity measurements can be used to get the thicknesses of the layers. It was decided to focus on the gravity profiling in this study.

Here, we discuss 3 possible scenarios for gravity measurements. A schematic of the blast furnace and approximate locations of the lines is shown in Figure A.1. One line "measures" right underneath the vessel, a second line is "measured" just next to the furnace and the third line is vertical, "measured" next to the vessel, at several heights.

A.3 Research plan

The aim is to study whether gravity measurements can be used to estimate the volume of the solid material remaining in the furnace. This is to be done by simulating measurements along certain lines. The procedure that could possibly allow inferences is:

- Along a chosen line, carry out N gravity measurements (d_i) $(1 \leq i \leq N)$. These measurements need to be repeated several times during the years to track changes. A reference station (d_{ref}) also needs to be established. This should be far enough from the furnace and other objects with time-varying gravity fields, should be easily accessible and its position should be fixed with high precision in location and height.

- Use the blast furnace design to theoretically calculate its gravity effect, b_i, at the measurement locations.

- Perform real measurements d_i after the furnace is repaired and filled with coke and iron-ore, but before production starts. (It is important to note the height of the measuring apparatus, to be able to perform accurate repeat measurements.) The real measurements d_i should ideally be equal to the theoretical measurements b_i.

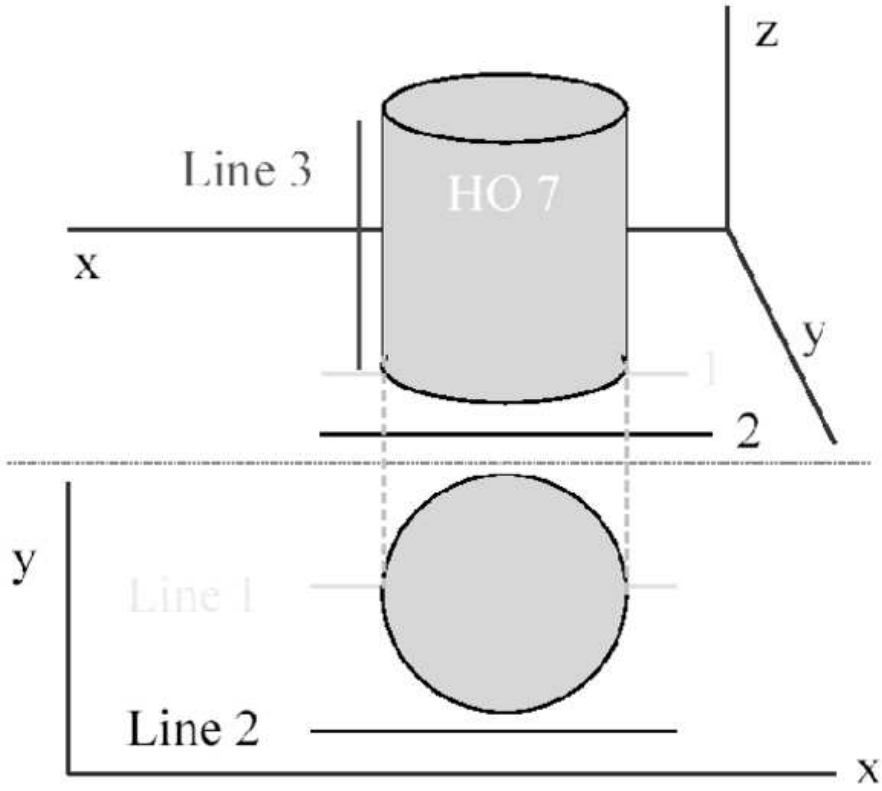

Figure A.1: Schematic overview of the location of possible lines of measurement for a gravity survey relative to the blastfurnace. The upper panel is in 3D perspective and the lower panel is the projection on an x-y-plane.

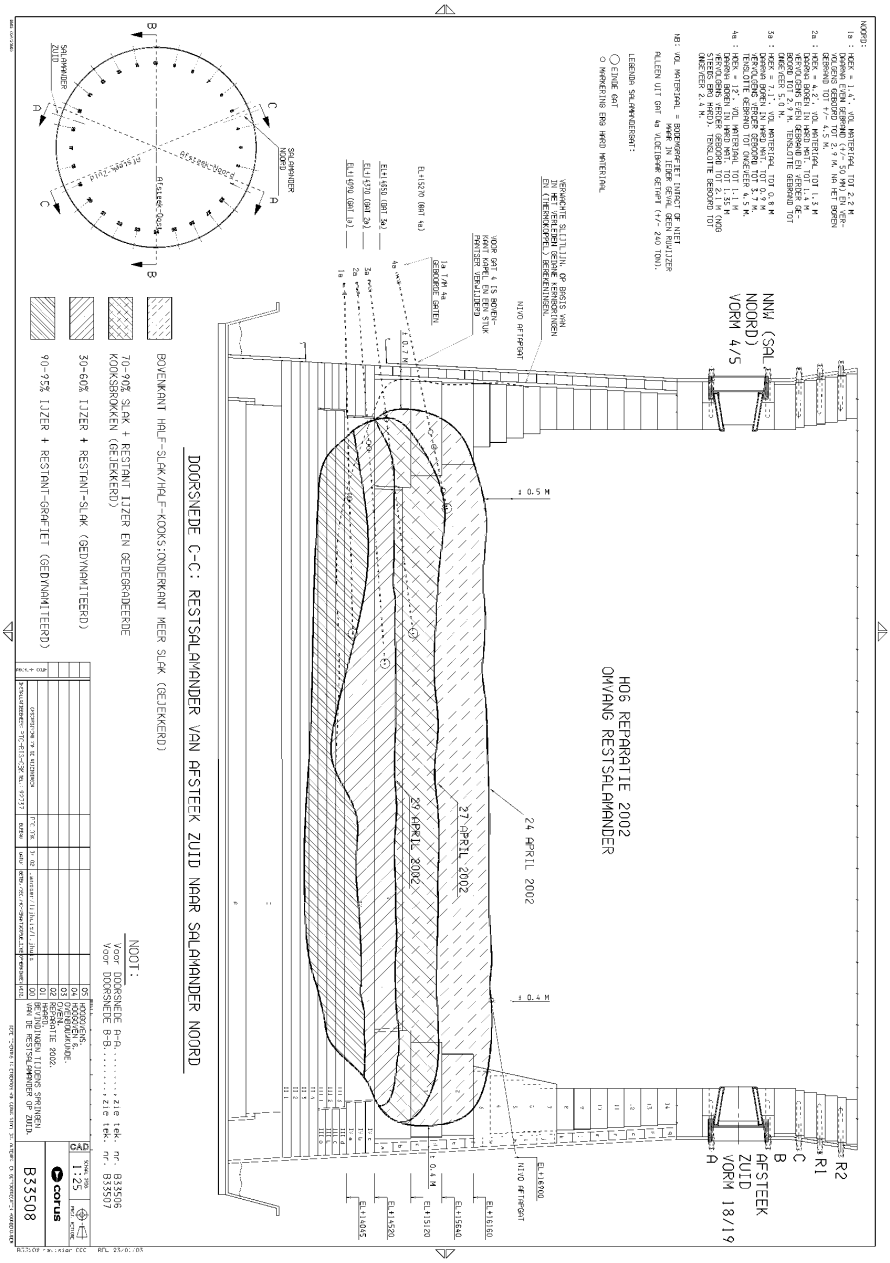

Figure A.2: The solid residue found in blast furnace 6 after shutdown. This is a design drawing and schematically, its internal "layering". (-B33508- by Corus)

Some differences could be present due to un-modeled effects from other parts of the factory and the earth below.

- The difference between the theoretical- and actual measurements is calculated and called *the correction*, c_i. Care must be taken to remove time-varying effects, t_i, due to tides etc. before calculating the corrections c_i. (Note that this step is not possible before the forth-coming repairs. For now, we will have to assume that $c_i = 0$. This is because no measurements were performed before production started.)

- What is of interest is the difference of repeat measurements, performed at the original locations and altitudes, with the reference station and after applying the corrections c_i. This difference, a_i, is clearly related to changes within the blast furnace since the last major maintenance. So in a formula:

$$a_i = d_i - d_{ref} - c_i \tag{A.1}$$

- The difference can then be used to invert for the amount of solid and liquid iron (remaining) in the blast-furnace, the number of intact graphite layers of the furnace floor, or any other measurable effect.

The study is divided in two parts. The first part of this study is to determine the minimum $\frac{\text{thickness}}{\text{change in density}}$ as a function of height in the furnace, that a disc with the diameter of the furnace needs to have, in order to be measurable at some outside location. The location of a layer within the blastfurnace also determines the type of construction material that is replaced. This determines the change in density.

The second part is to select locations which give the best resolution for layer thicknesses. A synthetic test is performed for a model consisting of 70 discs. (Instead of only calculating the effect of a single disc.) A grid-search is used to estimate the thicknesses of the 4 layers, to see whether the synthetic model can be retrieved. This is done in section A.5.

A.4 Part 1

A.4.1 Modeling the gravity effect of a disc.

Each distinct layer formed in the furnace can be modeled as a cylindrical disc of certain diameter, thickness density and location. Given a density difference and thickness, the disc can be used to calculate the difference in gravity when one material is replaced by the other (a_i). G_i, which gives the relation between density and the measured gravity, depends on the relative position of the center of the disc and the measurement location.

There is a linear dependence between the density of a disc (of given diameter and given center) and its gravitational effect d_i at any location i. This is the multiplicative factor G_i. Hence, if solid iron replaces a layer of graphite at the bottom of the furnace,

then the change in gravity at station i is $G_i \cdot (\rho_{iron} - \rho_{graphite})$.

$$\begin{aligned} a_i &= d_i - b_i - c_i \\ &= G_i \cdot \rho_{iron} - G_i \cdot \rho_{graphite} - c_i \\ &= G_i \cdot (\rho_{iron} - \rho_{graphite}) - c_i \end{aligned} \tag{A.2}$$

Thus the anomaly is linearly dependent on the density difference between the two materials. (c_i is the correction for un-modeled parts of the factory and the earth.)

Table A.1 gives the density of the materials which are used to calculate the theoretical gravity effect of the blast-furnace (b_i). Furthermore, this table will give the density differences between materials used in the calculations.

material	density (g/cm^3 (= ton/m^3))
concrete	2.5
graphite	1.7
EG stampmassa	2.0
half-graphite	1.7
HH-kwaliteit	2.5
carbon (koolstof)	1.6
liquid slag	2.6
(liquid) iron	6.7
coke	0.5

Table A.1: Density of materials (Personal communication with Corus)

A.4.2 Density model of the blast-furnace

A possible compositional layer-model of the blastfurnace at a moment of no production is given in Table A.2. The thickness of each solid layer is an estimate, derived from the solid residue found in HO6 in 2002, as given in Figure A.2. The layer consisting of 50% coke 50% slag is supposed to have formed after shutdown. *We are not sure whether the same is true for the layer beneath (70-90 % slag and iron and coke).* However, this working model is (for now) used as a guideline to get an indication for the heights of the discs and the density differences.

The height of a disc is set at the height of the transition of two layers and the density difference between the two layers is taken as the density difference. Now one can calculate the minimum thickness that the disc needs to have to produce a measurable gravity anomaly.

A.4.3 Results

We would like to check the outcome of our algorithm with a known formula. For a thin disc of radius r, the g_z is the gravity acceleration in the z-direction at a height z above the

layer	material	density (g/cm^3)	thickness (cm)
6	100% coke (s) + 5% slag (l)+5% iron (l)	1.0	400
5	50% slag (l) + 50% coke (s)	1.5	100
4	50% iron (l) + 50% coke (s)	3.6	100
3	10% coke (s)+ 80% slag (s)+10% iron (s)	2.8	50
2	45% iron (s)+ 55% slag (s)	4.5	70
1	90% iron (s) + 10% graphite (s)	6.2	100
0	carbon (s) + graphite (s)	1.7	115

Table A.2: Layered model of the blast-furnace HO6 in a week without production. The average layer composition and density is given for each layer, together with a rough estimate of the thickness. (s) means solid and immovable, (l) means liquid (in the pores of the coke). Layer 6 is the top-most layer. (*Layer number 2 and 3 might have formed during the shutdown of the blast furnace and might not be present in an operational furnace.*)

center of the disc, the analytical solution is (Marion and Thornton, classical dynamics of particles and systems, p205):

$$g_z = 2\pi\delta\rho\mathcal{G}\left[\frac{z}{\sqrt{r^2+z^2}} - 1\right] \tag{A.3}$$

With $\mathcal{G}$ the universal gravity constant. Note that this formula is only valid for points on the symmetry axis, so above or below the center of a thin disc. This formula is used to check the modeled result at the symmetry axis.

Let's fill in the formula for the graphite-solid iron border (layers 0 and 1 in Table A.2): $\rho = d\rho = 6200 - 1700 = 4500$kg/m^3, $r = 7.0$ meter, $z = 1.15 + 0.5 \approx 1.65$ (the 0.5 is the distance from the apparatus to the bottom of the vessel), $\mathcal{G} = 6.67 \cdot 10^{-11}$(The sign of the formula is changed so to make it compatible with the coordinate system used in the synthetic calculations):

$$\begin{aligned} g_z &= -2 \cdot 3.14 \cdot 4500 \cdot 6.67 \cdot 10^{-11} \cdot \left[\frac{1.65}{\sqrt{1.65^2+7^2}} - 1\right] \\ &= -1.88 \cdot 10^{-6} \cdot [0.23 - 1] \\ &= 1.45 \cdot 10^{-6} \text{ N/kg} = 0.15 \text{ mGal} \end{aligned} \tag{A.4}$$

The synthetic anomaly of a one meter thick disc (not really thin!) is 0.144 mGal. This was calculated for a disc made up out of prisms/cubes. (The implemented software uses a formula that calculates the combined gravity due to many prisms.) The equivalence between the analytical solution and the synthetic anomaly indicates that the forward modeling works properly. The small difference between the numerically predicted data and the analytical value could be due to the thickness of the disc.

Graphite-iron-transition

First the most promising transition is investigated; the graphite-iron transition between layer 0 and layer 1, as given in Table A.2. This transition has the largest density difference, which is 4.5 g/cc. The disc used has a diameter of 14.0 m and has its center at (7,7) (in m). The thickness of the disc is 15 cm. and its height above the bottom of the blast-furnace is taken to be 1.15 m. The height of the measure-apparatus is taken to be 0.5 m below the blast-furnace. The total distance to the transition is thus 1.65 m .

Measurements were simulated along 5 lines of 100 measurements each. (Figure A.3) The lines are each separated by two meters. Note that in practice these measurements cannot be collected at some of these positions due to the concrete structures beneath the vessel.

These synthetic data have their maximum value of 0.024 mGal just at the symmetry axis at (7,7). The maximum of the measurement line just outside/next to the blast-furnace is only 0.0060 mGal. A more detailed plot is shown in Figure A.4.

During our first meeting we also discussed whether it would be useful to measure at several heights along the blast-furnace. One simulated measurement for this disc is shown in Figure A.5. This is line 3 in Figure A.1 and in Table A.4.

This vertical line of measurements has both positive and negative values. In this test, a layer with a low density is replaced with a layer with a high density. With the chosen axis, a station below the disc measures a force upward, towards the disc. This is defined as positive. A station above the disc is pulled down towards the disc and this results in a negative change. The amplitude of the change is small.

Iron-slag transition

The iron-slag transition is modeled as the transition between layer 4 and layer 5 from Table A.2. The density contrast is thus 2.1 g/cm^3 and it is set at 4.35 m above the bottom of the furnace. Some results are shown in Figure A.6. The maximum value of the anomaly is much smaller, due to the increased distance and reduced density contrast. The tails are broader.

Conclusions

- Underneath the vessel, the amplitude of the anomaly is maximal.
- Thicker discs give broader tails.
- Along the vertical measurement line the amplitudes for g_z are small.

A.5 Part 2: Grid searches

To investigate the resolving power of the measurements, we use a simple model of the lower part of the blast-furnace. This is used to generate a synthetic dataset to which

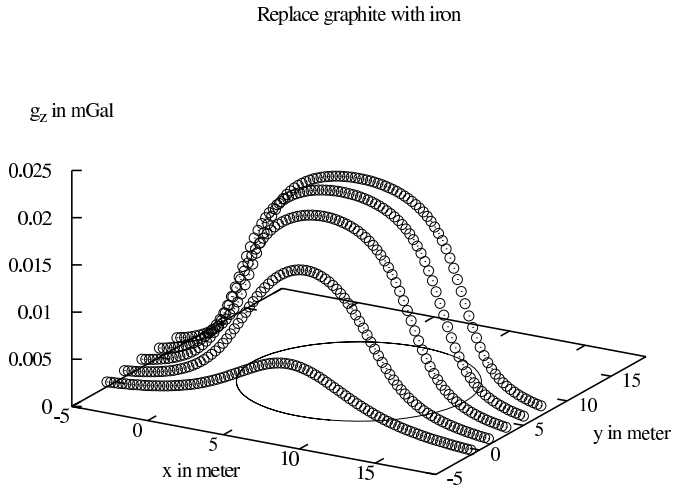

Figure A.3: The gravity (g_z) in mGal. (1 mGal = 10^{-5} Newton.) of a disc at 1.65 m above the measurement-plane. This is the change in gravity if a graphite layer is replaced with a layer of iron with the same thickness. The circle is the projection of the circumference of the modeled disc of 14 m in diameter. Its center is at (7,7). and the most outward line has $y = -1$, i.e. 1 m outside the furnace.

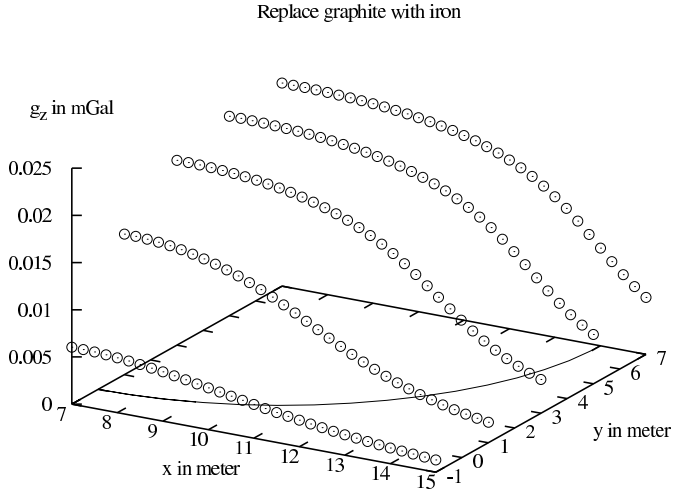

Figure A.4: The gravity (g_z) in mGal (1 mGal = 10^{-5} m/s^2). This is a zoom in of Figure A.3.

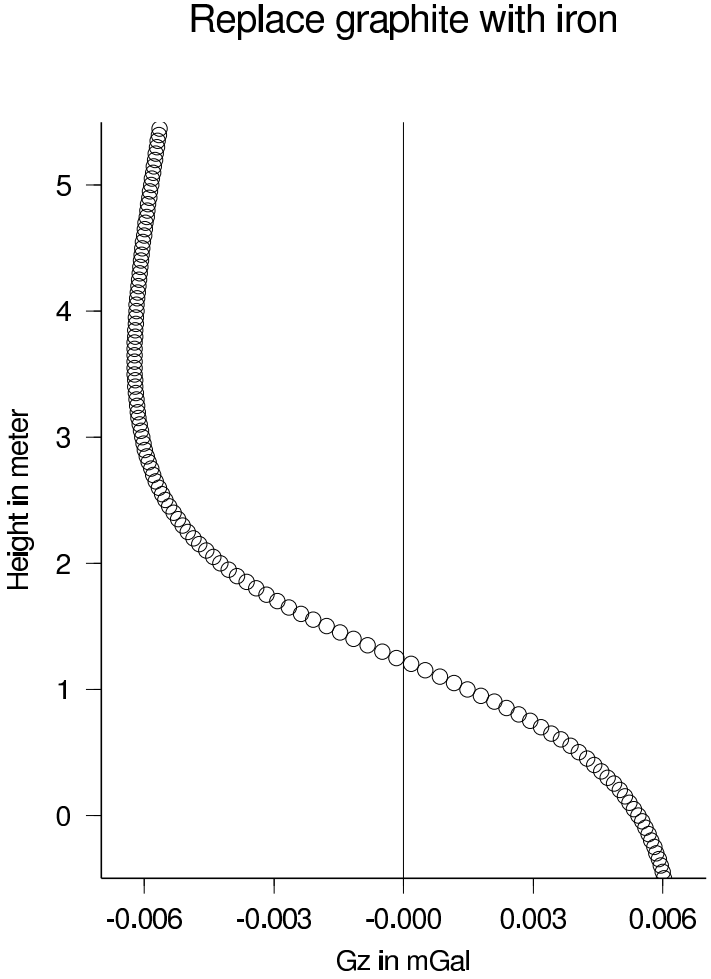

Figure A.5: The gravity (g_z) in mGal along the vertical line (line 2). (1 mGal = 10^{-5} m/s^2.) The height of measurement, relative to the bottom of the blast-furnace, is on the vertical axis. The vertical component of gravity (g_z) is on the horizontal axis. The lowest measurement location has coordinates (-0.5,-1) (in meter). i.e. half a meter outside the blastfurnace and one meter below the steel base. Note that the vertical force of gravity is zero at the middle of the disc. In this case this is at 1.15+0.075 = 1.225 m above the bottom of the blast-furnace.

Number	materials	ρ_{cur}	ρ_{con}	ρ_{model}
5	5% coke and ore as deployed	...	...	0.0
4	5% slag (l) and 5% iron (l) droplets in coke	1.0	0.5	0.5
3	50% slag (l) in coke replacing air in coke	1.8	0.5	1.3
2	50% iron (l) in coke replacing air in coke	3.8	0.5	3.3
1	90% iron (s) + 10% graphite replacing graphite	6.2	1.7	4.5
0	graphite replacing graphite and concrete	1.7	1.7	0.0

Table A.3: Overview of the materials and their density, as used for the model for the grid search. ρ_{cur} is the density in the current situation, ρ_{con} is the density of the original construction of the blast furnace in g/cc and ρ_{model} is the difference between the current density and the construction. It is the density used in the model. Note that the layers 2 and 3 as given in Table A.2 are not present in this (simplified) model. We think that these two layers might have formed during shutdown and for simplicity of the model, these were left out.

Gaussian noise of different levels can be added. These datasets can then be used to see whether the true model of the furnace can be retrieved. Furthermore, these test give some insight in possible deviations from the true model.

A.5.1 Model

Figure A.2 shows the wear of a blastfurnace and the residue found. A simplification of that situation will be used as our model, shown in Figure A.7, to test the possibility to discriminate several materials in the blastfurnace. This compositional data is also given in Table A.3. The test model consists of 5 cylindrical, horizontal layers. The density differences are shown with the color-scale. The differences in density with respect to the building specifications is plotted (because we are interested in the differences with a repaired state). The characteristics are tabulated in Table A.3. The diameter of each layer is 14 m.

In this study, it is assumed that corrections are made for all other materials. In a synthetic experiment, one can assign the density 0.0 g/cc to all unwanted sources. This is equivalent to not calculating their influence at all. For greater accuracy, one could model their gravity effect and correct the measured data.

Three important sources which should be corrected for are (1) the protective material at the wall of the furnace. (Gray in the Figure A.7) (2) the entire earth and (3) the complete blast furnace. In practice, the gravity should be measured repeatedly at a reference station, far enough away from the furnace (10 m would suffice). These measurements can be used to correct for meter drift and tidal influences. Using the gravity differences with respect to the reference value, together with the usual corrections for elevation, should eliminate the gravity effect of the earth.

Name	# stations	x	y	z	step-size	comment
Line 1	50	[-5.4, 19.6]	7	-0.5	dx = 0.50	horizontal, below
Line 2	50	[-5.4, 19.6]	-1	-0.5	dx = 0.50	horizontal, along
Line 3	50	7	-1	[-0.5, 7]	dz = 0.15	vertical, on the side

Table A.4: The three lines for the simulated measurements.

Measurement locations

For the sensitivity test, 3 lines of measurements are tested. These are a vertical line next to the furnace and 2 horizontal lines, as in Figure A.3. They are tabulated in Table A.4. Each line consists of 50 measurement locations.

Several synthetic datasets are collected, with different levels of Gaussian noise added to the measurements. For each of the three lines, three experiments were simulated, each with increasing amount of noise, so in total 9 experiments were performed. In each experiment, 20 datasets were generated. This is to reduce the influence of the noise. Otherwise an unlucky draw of noise could lead to wrong conclusions. However, the use of 20 datasets per experiment does *under*-estimate the possible variations, due to averaging.

In Table A.5, the nine experiments are tabulated. The standard deviation (σ) of the Gaussian noise added is given in micro Gal (μGal). The distribution is given by:

$$\frac{1}{\sigma\sqrt{2\pi}} e^{\frac{(n-\bar{n})^2}{2\sigma^2}} \tag{A.5}$$

with $\bar{n}$ the average value of the noise, which is zero. The signal to noise ratio ($\mathcal{S}/\mathcal{N}$) is

$$\mathcal{S}/\mathcal{N} = \sqrt{\frac{\sum_i^N d_i^2}{\sum_i^N n_i^2}} \tag{A.6}$$

with d_i the i'th measurement, n_i the noise on the i'th measurement and N is the number of measurements per line, so 50 in this study. For each of the 20 datasets of each experiment, the $\mathcal{S}/\mathcal{N}$ is calculated per dataset and the average value per experiment are shown in rightmost column of Table A.5. Each noise-level for all three lines was done with the same seed for the noise generator, so the differences is $\mathcal{S}/\mathcal{N}$ reflect the differences in the total signal of a line.

A.5.2 Grid search

The 'inversion' of the data is done with a grid search. The parameters of interest, to be determined from the data, are the heights of the transitions of the layers. Other interesting and important parameters are the densities of the layers, but these are assumed known for the present. A further assumption is that all the layers are present and in the same vertical order as tabulated.

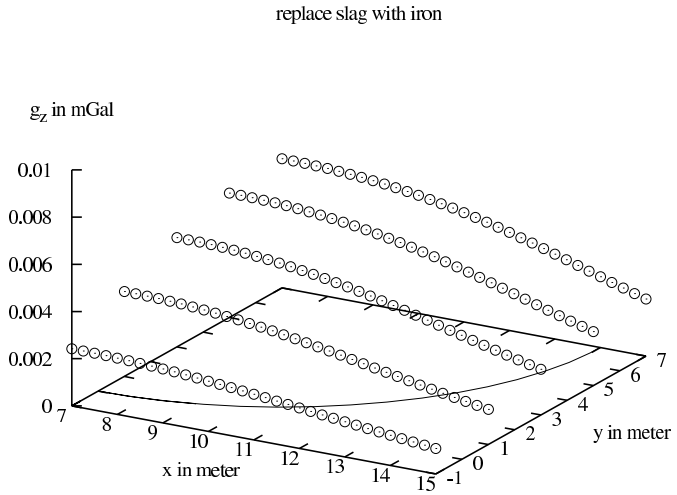

Figure A.6: The vertical force of gravity (g_z) in mGal for a disc of iron replacing a slag. The disc is 15 cm thick and is at 4.35 m. The density contrast is 2.1 g/cm^3. The quarter circle is the projection of the outline of the disc on the measurement plane.

Line	noise in μGal.	$\mathcal{S}/\mathcal{N}$
Line 1	10	22.1
Line 2	10	7.2
Line 3	10	9.1
Line 1	20	10.9
Line 2	20	3.6
Line 3	20	4.5
Line 1	40	5.6
Line 2	40	1.8
Line 3	40	2.3

Table A.5: Noise levels for the 3 lines in Table A.4.

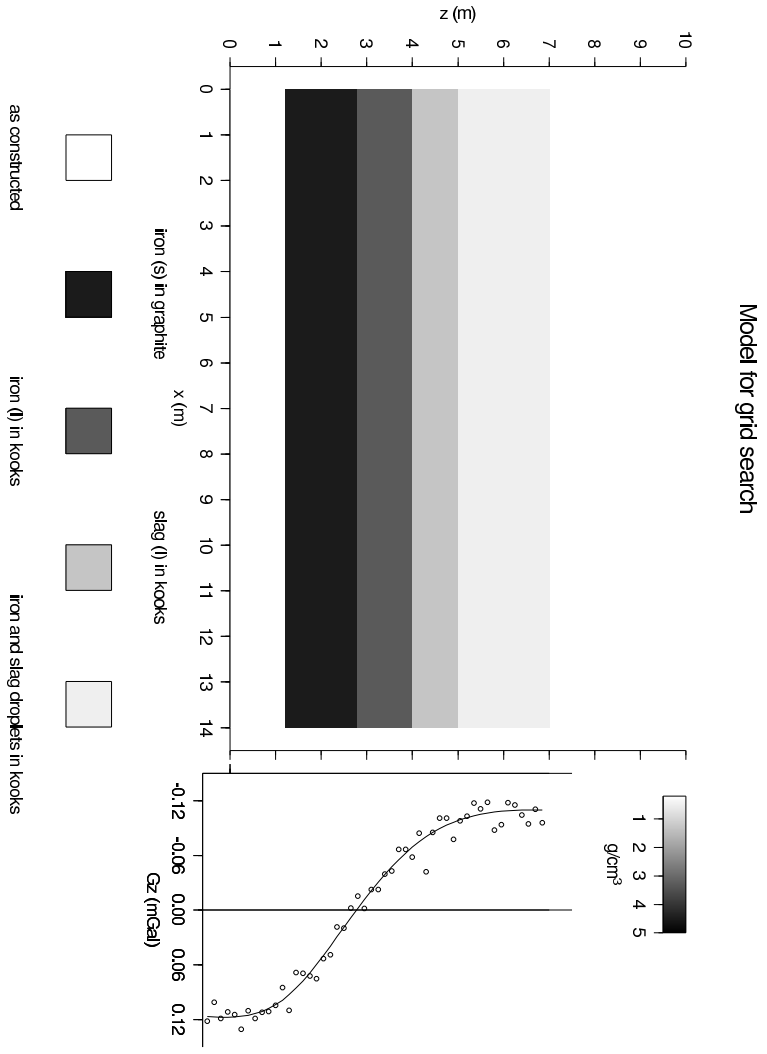

Figure A.7: A layered, cylinder symmetric model of the lower part of the blast furnace. The color-scale indicates the density difference between the original construction as given in a design drawing, and the current state. For instance, solid iron (with some graphite) replaces graphite at the brownish colors. Gray means that no changes have occurred. The synthetic measurements of line 3 with Gaussian noise ($\sigma = 0.01$ mGal) is shown on the right. The vertical axis of this plot is drawn at the correct (x,y)-position with respect to the coordinates of the vessel. The z-coordinates at which the measurements were simulated are also shown at the correct heights, with respect to the vessel.

Line	noise in μGal.	$\mathcal{S}/\mathcal{N}$	sd 1 (cm)	sd 2 (cm)	sd 3 (cm)	sd 4 (cm)
Line 1	10	22.1	23	82	56	81
Line 2	10	7.2	60	97	75	86
Line 3	10	9.14	17	73	30	65
Line 1	20	10.9	43	103	96	87
Line 2	20	3.55	79	11	99	90
Line 3	20	4.5	26	83	38	78
Line 1	40	5.55	55	120	124	104
Line 2	40	1.8	88	121	113	96
Line 3	40	2.3	40	88	61	83

Table A.6: Standard deviations for the 4 transitions.

Given the order of the layers and their densities, the lower 7 m of the furnace is divided in 70 discs, each 1 decimeter thick. Then all combinations of the 5 densities are given to the 70 discs. (Given the order and the number of layers, there are about 12 million possible configurations for these 70 discs.) For each of these 12 million models, the L_2-norm data-misfit is calculated as follows:

$$\text{data-misfit} = \sum_{i}^{N} (d_i - \sum_{j} G_{ij} m_j)^2 \tag{A.7}$$

with G_{ij} the components of the matrix, which forms the forward operator and the m_j's are the densities of the discs. The d_i's are the 'observations' plus the added noise. So it is the square of the differences between the observations, d_i, and the response that the generated layers would give ($\sum_j G_{ij} m_j$).

Given the Gaussian noise distribution, there is an expected value for the data-misfit, which is equal to the number of measurements times the variance; $\hat{n} = N\sigma^2$. In the grid search I marked all solutions with a data-misfit smaller than the expected fit ($\hat{n}$) as good. These were stored.

The procedure can be summarised as follows: synthetic data is generated for the test model and noise is added to the data. Then a search is done for the heights of the transitions between layers. Heights which gave a good data-fit were stored.

A.5.3 Results of the grid search

For each noise-level and measure line, a set with good models was calculated. One experiment is shown in histogram form in Figure A.8. It shows how often a particular height for a transition gave a good data-fit. The most likely height seems to coincide well with the true height. The shape of the histograms of transition 1 and 3 look like a Gaussian distribution. To summarize this histogram, I calculated the mean and standard deviation of the histograms. The information of this histogram is given in the middle panel of Figure A.11.

The Figures A.9-A.11 show the average heights and standard deviations for the nine experiments. The upper panels have the lowest noise-level with $\sigma = 0.01$ mGal, the middle panels have $\sigma = 0.02$ mGal and the bottom panels have $\sigma = 0.04$ mGal. Note that the histograms of transition 2 and 4 are not Gaussian shaped; The average in combination with the standard deviation does not properly describe the shape of these distributions. However, the most likely models are close to the true values.

One last visualisation of the results of line 3 with moderate noise is shown in Figure A.12. This Figure has six panels, showing the correlations between all pairs of transitions. The color-scale gives the number of times a combination gave a good data-fit. Again, transitions 1 and 3 are more or less (2D) Gaussian, while transitions 3 and 4 are not Gaussian at all. The most likely models (darkest colors) do appear close to the true transition heights, given by the green dot.

A.5.4 Summary of the results

- Gaussian noise with a larger standard deviation σ gives broader distributions for good models. This can be due to
  - The expected value of the data-fit increases, so more models satisfy the 'good model' criterium. (More good models are found.)
  - Some large noise contributions could make larger differences in height possible.
- The above effect is less pronounced for the two upper transitions for line 2, the line next to the vessel. (Although the standard deviation is poor for all noise-levels). The broad tail of the gravity anomaly for this line, as seen in the test with one disc, probably causes this behavior.
- The two transitions with the largest density contrast, transition 1 and 3 have fairly small standard deviation and the average (and the most likely height) are close to the true model.
- Although the amplitude of the vertical component of the gravity of line 3 is much smaller than for line 1, line 3 has the smallest standard deviation for all noise-levels. The forward operator for line 3 is better conditioned. (not shown)
- Especially for transitions 2 and 4, the average height is a poor estimator of the true height. The most likely height is a much better estimator.

A.6 Discussion

Line 1 is situated partly underneath the blast furnace vessel. For safety reasons, personnel should not come there. In practice, Line 1 could be measured with the measuring apparatus on a cart. Modern gravimeters are full-automatic, including the leveling. (It needs

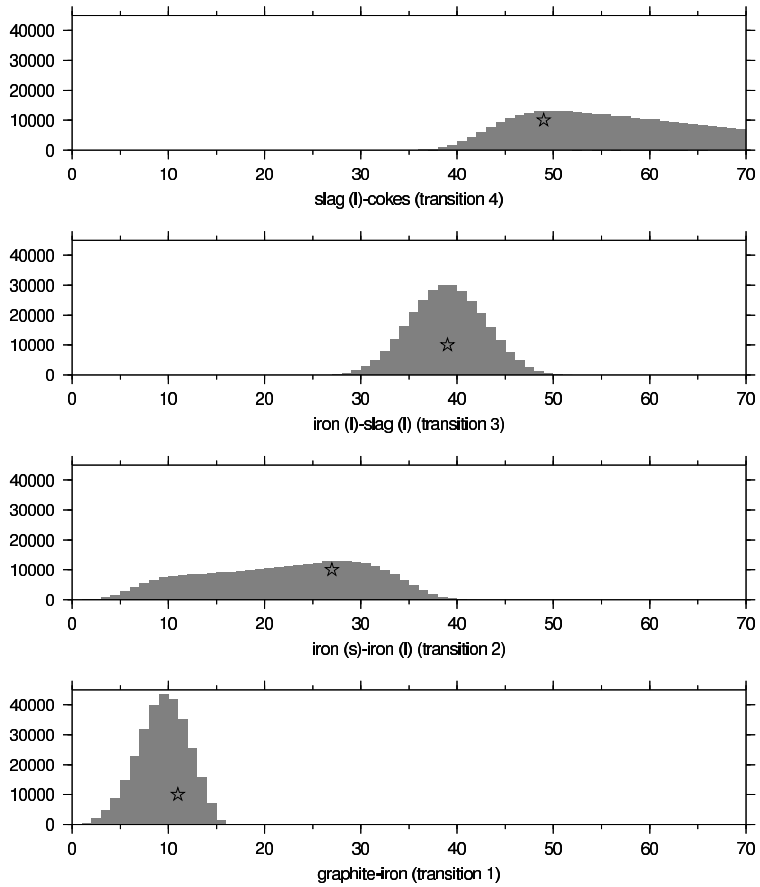

Figure A.8: Histogram of the good models for the grid search of data on line 3 and Gaussian noise with $\sigma = 0.02$ mGal. The figures show how often a tried height for the transition gave a good data-fit. The heights are in decimeter. The star indicates the true height of each transition of the model (Figure A.7).

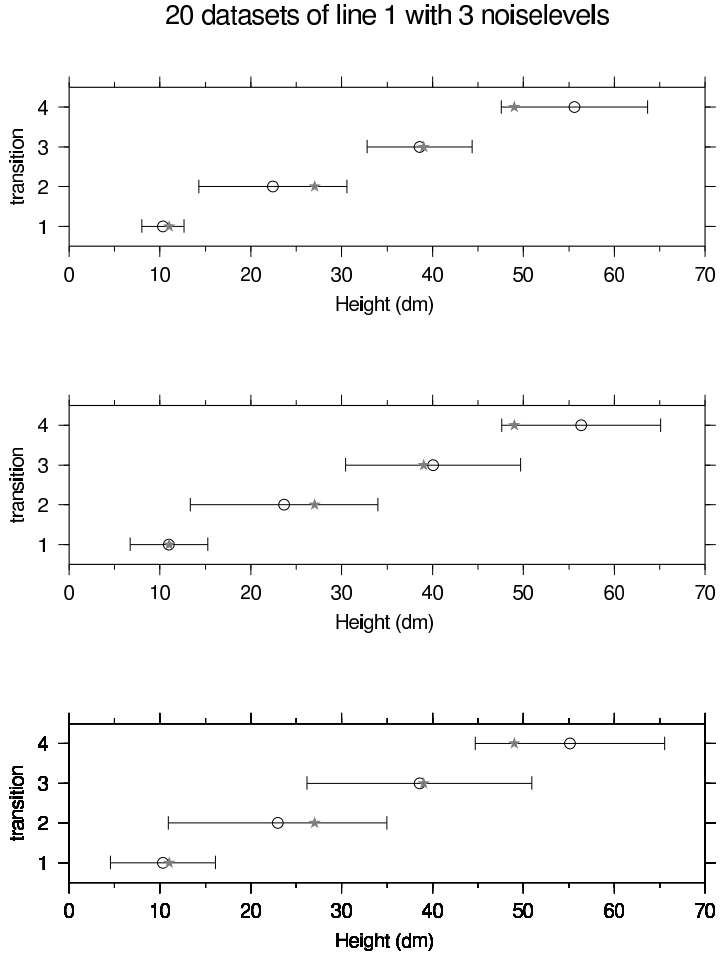

Figure A.9: The average heights of the 4 transitions in good models is given in decimeter above the bottom of the blast furnace, as obtained with data on line 1. The error-bars give the standard deviation of the distributions. The star denotes the true height of the transition in the test model. The experiment of the upper panel had Gaussian noise with $\sigma = 0.01$ mGal. The middle panel had noise of $\sigma = 0.02$ mGal and the bottom panel had $\sigma = 0.04$ mGal.

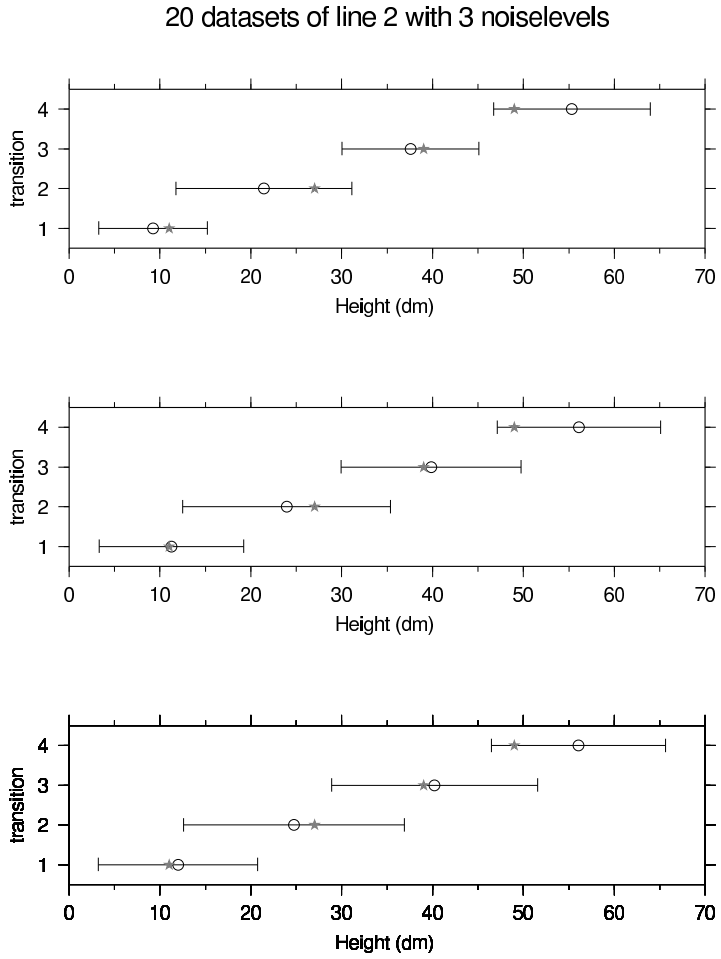

Figure A.10: The average heights of the 4 transitions in good models is given in decimeter above the bottom of the blast furnace, as obtained with data on line 2. The error-bars give the standard deviation of the distributions. The star denotes the true height of the transition in the test model. The experiment of the upper panel had Gaussian noise with $\sigma = 0.01$ mGal. The middle panel had noise of $\sigma = 0.02$ mGal and the bottom panel had $\sigma = 0.04$ mGal.

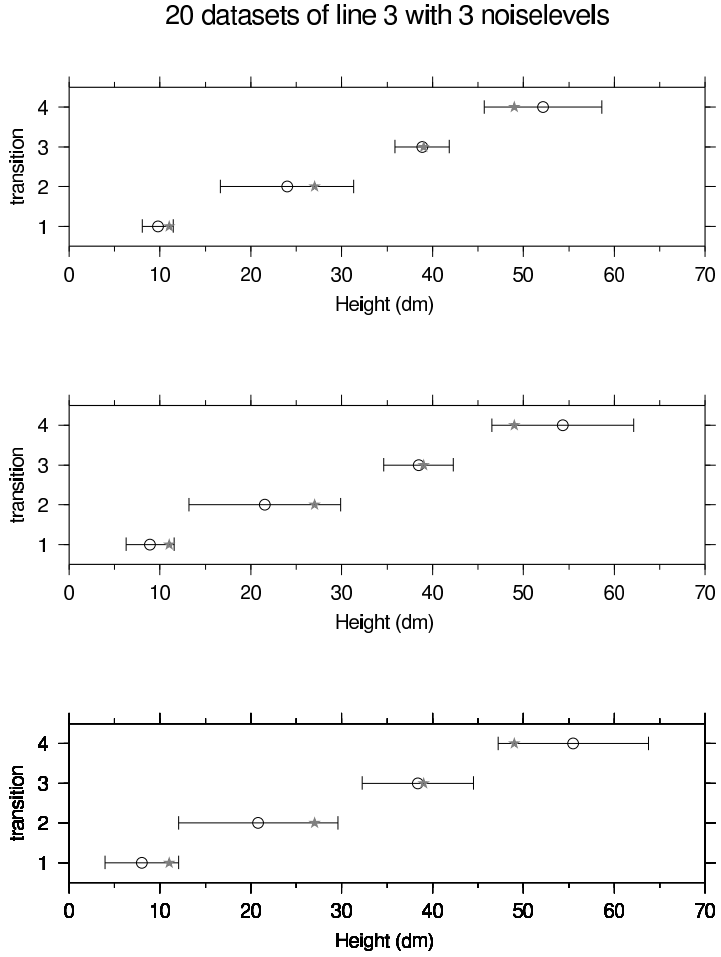

Figure A.11: The average heights of the 4 transitions in good models is given in decimeter above the bottom of the blast furnace, as obtained with data on line 3. The error-bars give the standard deviation of the distributions. The star denotes the true height of the transition in the test model. The experiment of the upper panel had Gaussian noise with $\sigma = 0.01$ mGal. The middle panel had noise of $\sigma = 0.02$ mGal and the bottom panel had $\sigma = 0.04$ mGal.

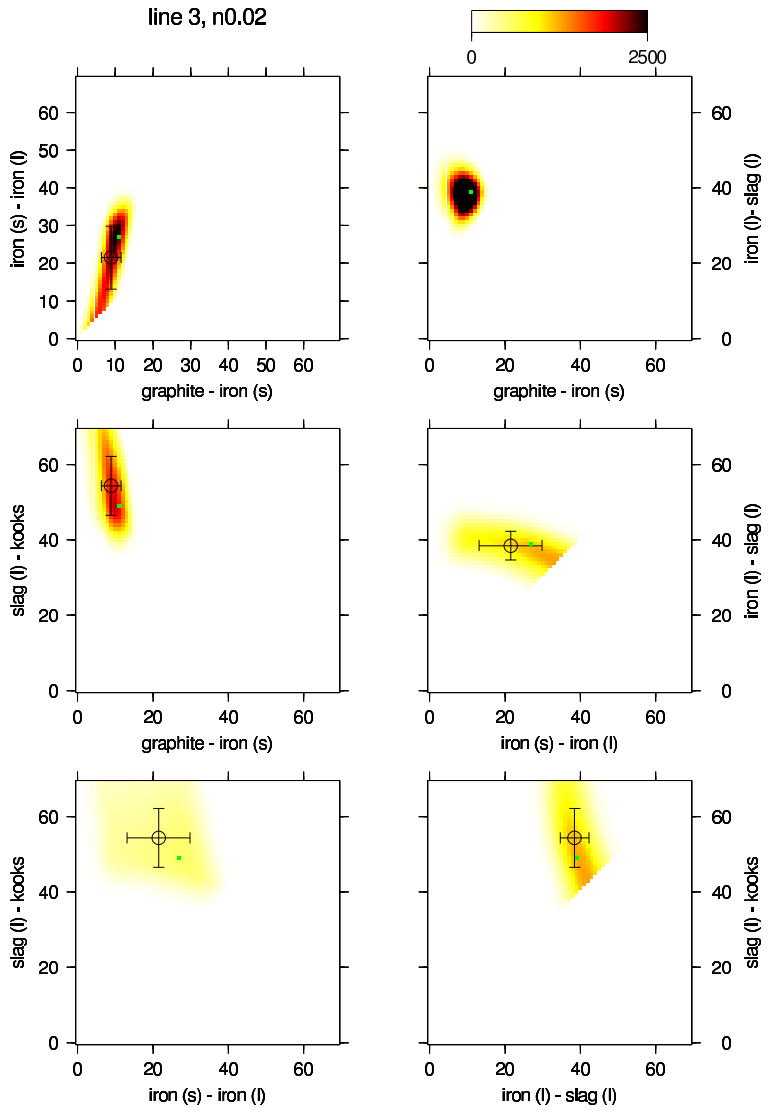

Figure A.12: The 6 cross-plots of the good models obtained with line 3 with the moderate noise with $\sigma = 0.02$ mGal. The colour-scale indicates how often a combination of the heights of a transition pair has a good data-fit. The average and standard deviation are shown with the circle and the error-bars. The true heights of the transitions are given by the green dot. Note that the lower right under the diagonal is not searched, because layer 4 is on top of layer 3, which is on top of layer 2, which is on top of layer 1. The darkest colours seem to correspond rather well to the true model.

to be level, in order to properly measure the vertical component of the gravity field.) The cart could be pulled underneath the furnace, in between the concrete supports.

Line 2 could also be measured with a cart, or a person can carry the equipment to the next position. Line 3 could for instance be measured with a movable platform, so that several heights can be reached.

This theoretical study and especially the grid search gives results that are probably more optimistic than a real life experiment.

- The density contrasts are important parameters and in the test, the correct values were used.
- Homogeneous layers were assumed, while in practice, there might be inhomogeneities, or a gradual change in densities.
- Only Gaussian noise was added to the synthetic data of the test model. For real measurements, there will be systematic, non-Gaussian errors due to small errors in the corrections for the data and influences from other sources. In this study, only the densities within the vessel contributed to the signal. In other words, I assumed perfect corrections for all other sources of gravity.

The final limitation for assessing the possible resolution is the amount of noise that a true dataset will have. It might be (much) larger than the test-cases I have performed. Although modern gravimeters are said to be repeatable up-to 0.01-0.02 mGal, I cannot estimate whether the vibrations near the blast furnace will affect the measurements. That depends on the frequency of the vibrations and the measuring apparatus.

A.7 Conclusions

- Under favorable noise conditions (Gaussian noise with $\sigma = 0.01$ mGal) and good a priori estimates of the density, a resolution of 20 cm can be obtained with measurements at small intervals along a vertical line, but only for large density contrast. This is probably the best gravity measurements can achieve.
- I recommend to gravity profiling if the resolution seems adequate to influence the decision making process by Corus. A test line could be measured to estimate the level of noise on the data.
- If the noise level is acceptable, for the future, I would recommend to perform a vertical line of measurements when the blast furnace is repaired and stocked and is (almost) ready to be used again. This is the perfect reference dataset to correct any measurements later on.
- If a reference dataset is measured, then measuring at regular time intervals (couple of times per year) could give information on buildup of solid iron at its bottom.
- If the noise-level is really low, maybe frequent measurements (several per day), could give information on the changes in the liquid iron level during production.

Samenvatting (Summary)

Een van de doelen van geofysica is het onderzoeken van de bodem met natuurkundige methoden. Een bekend voorbeeld is dat trillingen veroorzaakt door aardbevingen op verschillende plaatsen op de aarde met seismometers worden gemeten. Met deze data kan men iets zeggen over de opbouw van de Aarde. Geofysische metingen (seismiek, zwaartekracht, electromagnetisch) kunnen ook op veel kleinere schaal worden toegepast, bijvoorbeeld om olie of gas te vinden, of ijzererts. De typische diepte hiervoor is tot een paar kilometer. Grondradar kan werken voor de bovenste tientallen meters, tot zelfs maar een paar centimeter om de asfaltlaag van een weg te controleren! Door fysische metingen kunnen we dus een beeld vormen van de ondergrond, zonder te hoeven graven of boren.

Het doel van dit promotie onderzoek was het ontwikkelen van een methode voor impliciete, structurele inversie van geofysische data. Met *structurele inversie* bedoelen we dat het resulterende plaatje van de ondergrond makkelijk te interpreteren is in termen van structuren/gebieden met eenzelfde lithologie. (Lithologie is een kwalificatie van gesteente aan de hand van uiterlijke kenmerken en fysische eigenschappen, zoals de kleur, dichtheid, magnetische eigenschappen, etc.) Met het woord *impliciet* geven we aan dat de structuren niet expliciet een vaste vorm moeten hebben, noch dat er een vast aantal structuren aanwezig moet zijn. Een bolvormige structuur aannemen, of de ondergrond beschrijven met een aantal lagen, zijn voorbeelden van een expliciete structuur opleggen.

In plaats van expliciete vormen te gebruiken, delen we het gebied van de ondergrond waar we in geïnteresseerd zijn op in rechthoekige cellen. In 3 dimensies (3-D) hebben de cellen de vorm van een balk. We nemen aan dat overal in het volume van iedere afzonderlijke cel de fysische eigenschappen gelijk zijn en aan de hand van de gemeten data willen we die voor alle cellen berekenen. De getalwaarden van de gesteente eigenschap van iedere cel zijn de parameters voor onze inversie. Het inversie resultaat moet zijn structuur krijgen doordat de cellen (ongeveer de) waarden hebben die bij een bepaalde lithologie horen. De structuren van de ondergrond zijn dan dus opgebouwd uit clusters van cellen met ongeveer dezelfde fysische eigenschap en het verkregen plaatje is makkelijk te interpreten, zonder van te voren de vormen vast te hoeven leggen.

Bij geofysische inversie zijn er vaak veel verschillende oplossingen die de (beperkte hoeveelheid) meetgegevens even goed verklaren; de oplossing is *niet uniek*. Om toch een resultaat te krijgen moet er extra informatie gebruikt worden. Een veel gebruikt criterium is dat de oplossing *glad* moet zijn, dat wil zeggen dat (in het geval van inversie van zwaartekracht data) de gevonden dichtheid in iedere cel weinig mag verschillen van de dichtheid in zijn buurcellen. Door de gladste oplossing te kiezen, die de data nog goed fit, is het probleem uniek gemaakt. Het grote nadeel is dat de structuren in de ondergrond niet meer scherp omlijnd zijn, maar dat de grenzen zijn uitgemeerd. Dat maakt de interpretatie moeilijker. Wij veronderstellen in dit proefschrift als extra informatie dat er maar 2 soorten lithologieën aanwezig zijn in het stuk van de ondergrond waarvoor we inverteren. Verder gebruiken we dat we de dichtheden kennen, of het dichtheidsverschil. Voor inversie van seismische looptijden moeten we de golfvoortplantingssnelheiden in de media

kennen. Bij implicite structurele inversie wordt die oplossing gezocht waarvan de parameter waarden (dichtheden, snelheden) het dichtst bij de *a priori* waarden liggen, terwijl tegelijkertijd de metingen nog goed door het model worden verklaard. In **hoofdstuk 2** beschrijven we een methode voor impliciete structurele inversie gebaseerd op *Lineair Programmeren*. Dat is een wiskundige methode om het optimum van een lineaire functie te vinden, gegeven lineaire vergelijkingen die de oplossingsruimte beperken. Het is mogelijk om de a priori dichtheidsinformatie met dit soort vergelijkingen te beschrijven en om de L_1-norm van de datamisfit (de som van de absolute waarden tussen de metingen en de modelvoorspelling) te optimaliseren. Door een kenmerk van lineair programmeren zijn bijna alle gevonden waarden dichtbij de aangenomen waarden voor de lithologieën en zijn de resultaten dus makkelijk te interpreteren in termen van structuren. Het bleek echter dat er nog wat extra informatie nodig is: De minimale afstand tussen het oppervlakte van de Aarde en de bovenkant van een structuur. Deze methode was eerst getest op synthetische data en in **hoofdstuk 3** wordt een serie metingen gebruikt.

In **hoofdstuk 4** beschrijven we een tweede, iteratieve, methode. De L_2-norm van de datamisfit wordt geminimaliseerd, in combinatie met een gewogen term die er voor zorgt dat de nieuwe oplossing steeds meer op een ideaal referentie model gaat lijken. Dit referentie model en de mate van weging worden bepaald met behulp van het resultaat van de vorige iteratie en de a priori informatie over de dichtheden. Er moeten een aantal constanten gekozen worden. Door op een handige manier allerlei combinaties te proberen (met het *Neighborhood Algorithm*) en te kijken naar de uiteindelijk gevonden datamisfit, kan een goede combinatie van constanten worden gevonden. Deze methode is getest met synthetische data en zelfs zonder de minimale diepte van de structuren te kennen, waren de resultaten goed.

De methoden hierboven waren getest en toegepast op zwaartekrachtsdata. Het voordeel daarvan is dat de vergelijkingen lineair zijn. Seismische looptijden hangen op een niet-lineaire manier af van de golfvoortplantingssnelheden van de gesteenten. Een eerste aanzet om dit soort gegevens te inverteren met een iteratieve versie van de lineair programmeren methode is gepresenteerd in **hoofdstuk 5**. Als de ondergrond niet te complex is en het snelheidsverschil tussen de lithologieën niet te groot, is de methode toepasbaar. Een groot snelheidsverschil zorgt ervoor dat de berekende route die de seismische energie aflegt door de ondergrond sterk kan veranderen tussen 2 iteraties. Dit veroorzaakt problemen, die (deels) opgelost kunnen worden door in het begin van het iteratieve proces de verandering te beperken. Ook is onderzocht of de seismische- en zwaartekracht-inversie te combineren zijn. Deze gezamelijke inversie geeft betere resultaten.

Naast dit puur wetenschappelijke werk wordt in de **Appendix** een onderzoekje naar een mogelijke toepassing van geofysische methoden gegeven. Voor de ijzer producent *Corus* (vroeger *Hoogovens*) is numeriek onderzocht of veranderingen van een hoogoven, in opbouw en dichtheid van de inwendige materialen - door gebruik en slijtage - te meten zouden zijn met moderne zwaartekrachtmeters. Het lijkt erop dat de verandering in zwaartekracht met de tijd theoretisch wel te meten zou moeten zijn (gegeven de gevoeligheid van de meters in een ideale laboratorium opstelling), maar dat trillingen van de fabriek mogelijk de resolutie nadelig zullen beinvloeden. Ook lijkt het nodig om nauwkeurig te meten voordat een gerepareerde installatie in gebruik genomen wordt.

Dankwoord (Acknowledgments)

Dit proefschrift zou het licht niet gezien hebben zonder de hulp van velen. Ten eerste wil ik Kabir Roy-Chowdhury hartelijk danken voor zijn inzet, tijd, taal correcties en zeer makkelijke bereikbaarheid. Toen ik begon was Jaap Mondt in deeltijd aan de vakgroep verbonden, naast zijn baan bij Shell. Helaas moest hij zijn professorschap beëindigen, maar gelukkig kon hij toch nog tijd vinden - en veel geduld opbrengen - om betrokken te blijven. Ook Wim Spakman, die als aan de universiteit verbonden promotor een deel van de taken van Jaap overnam, heeft de nodige gaatjes in zijn drukke agenda kunnen maken. Pavel Ditmar supplied the basis of the research proposition. He always had most useful comments on my reports, containing figures with awful axis labels (if any) and inappropriate titles. Furthermore, he supplied the code for the seimic ray-tracer. Ook de leden van de leescommissie wil ik graag bedanken voor hun tijd en nuttig commentaar.

Mijn directe collega's, Karin, Robbert, Hanneke, Henk, Arie, Ueli, Vessa, Ebru, Mei, Ilaria, Wouter, Joe, Caroline, Tonie - en de anderen - en natuurlijk mijn kantoor-genoten Stefan en Dirk, dank ik voor de discussies, gezelligheid en roei-competities.

Mijn computer problemen werden snel en vakkundig door Mark van Alphen, Joop Hoofd of Theo van Zessen opgelost. Zo wilde ik eens een figuur maken met de legenda van de kleurschaal erboven. Hiervoor heeft Theo een ochtend moeten strijden om de nieuwste versie van het software pakket geïnstalleerd te krijgen. Jacqueline Landsheer was altijd behulpzaam met secretariële bezigheden.

Verder wil ik André ten Brinke, Libbe van Zwol en Eric van 't Hooft van Isodose Control b.v. danken voor de kans om mijn proefschrift af te kunnen maken nadat mijn contract met de universiteit afgelopen was; eerst door in te stemmen met een contract voor 4 dagen per week en later door me de gelegenheid te geven om het voltijds af te kunnen ronden. Super, zonder deze mogelijkheid zou het nu niet klaar zijn! Ook mijn andere IC collega's Martin, Michiel, Pier en Joeri dank ik voor hun nuttige uitleg, interesse en gezelligheid, op het werk en in de auto.

Maar bovenal wil ik mijn (schoon)familie danken voor hun steun en vertrouwen. Mary heeft me veel liefde en plezier gegeven; ik ben blij dat we elkaar aan het begin van mijn AIO-schap ontmoet hebben!

Tim

Curriculum vitae

21-6-1976	Geboren te Amsterdam.
1980 - 1988	Basisschool ‘de Vinkenbaan’ te Muiderberg.
1988 - 1994	VWO op het ‘St Vitus College’ te Bussum.
1994 - 1999	Natuurkunde aan de Universiteit van Amsterdam (UvA), met een afstudeerstage bij medische fysica op het Amsterdam Medisch Centrum (AMC).
2000 - 2004	AIO positie bij exploratie geofysica, departement Aardwetenschappen aan de Universiteit Utrecht.
2004	Junior universitair docent bij de Universiteit Utrecht voor een onderzoeksproject bij Corus.
2004 - 2005	Software engineer bij Isodose Control b.v. te Veenendaal.