Department of Mathematics and Statistics
College of Engineering

UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Summer Research Project

MARS Imaging - Computational Methods for Material Decomposition

Christopher J. Bateman

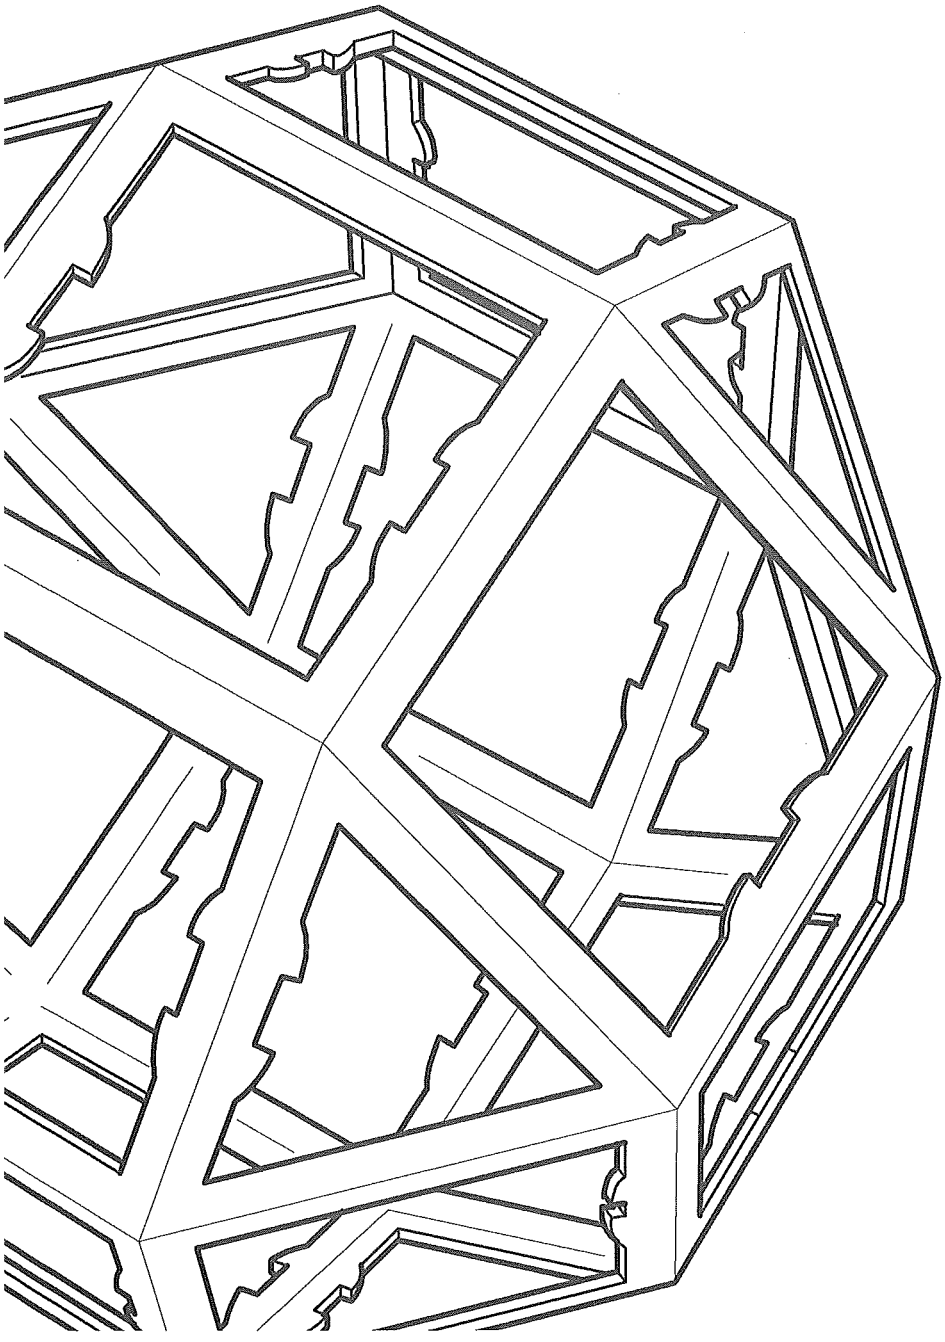

11/12

Department of Mathematics and Statistics, University of Canterbury, Private Bag 4800, Christchurch, New Zealand

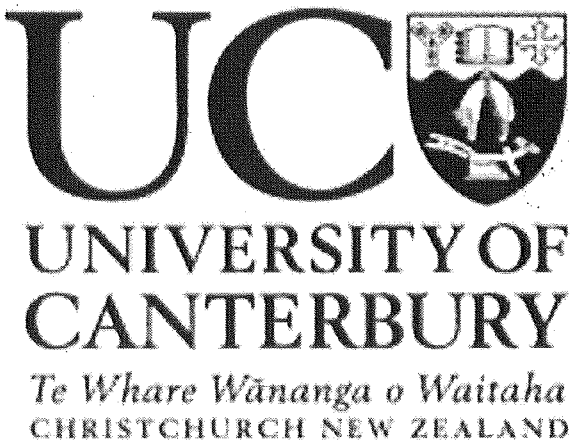

MARS Imaging - Computational Methods for Material Decomposition

Christopher J. Bateman
Summer Project 2011/2012

Supervisors: Peter Renaud, Rachael Tappenden, Philip Butler, Anthony Butler.

Abstract

The MARS (Medipix All Resolution System) research group is developing a CT (Computed Tomography) scanner which uses the Medipix range of energy discriminating photon counting detectors. These detectors have the potential to significantly improve the functionality of medical scanners and reduce patient radiation exposure. An image processing technique that benefits from this type of detector is material decomposition. The goal of material decomposition is to take data obtained at different energies and produce material density maps of each material in the scanned region.

This project contains the start of preliminary investigations into using compressive sensing as a method for material decomposition in the MARS Spectral CT Scanner system. L1 - Minimization methods from the YALL1 open source l1 optimization package [1] are currently being tested on phantom data representing 7 different medically relevant materials - Air, Water, Perspex, Calcium, Iron, Iodine and Fat. Calcium, Iron, Iodine and Fat were expressed in the phantom by the compounds Calcium Chloride (aq) [2M], Ferric Nitrate (aq) [0.4M], Iodine (aq) [0.02M] and Sunflower Oil.

More results are needed to draw firm conclusions about the viability of this method, however initial results on simulated data shown in this report are promising.

Contents

List of Figures

Chapter 1

Introduction

CT (computed tomography) scanners are used in medical applications to produce high spatial resolution images of internal tissues. A regular CT scan is a set of gray-scale images, similar to the common "x-ray", based on the attenuation properties of the materials that the x-rays pass through. Limited colour CT scans are possible with standard CT hardware if scans at two different x-ray energies are taken (Dual-Energy CT). A process called material decomposition is required to produce a colour CT image from a series of gray-scale scans. This determines the quantities of each material in every pixel in each scan by using the extra energy information provided in the multiple number of scans. The materials can then be colour-coded and shaded based on their densities to produce a colour CT image.

An ongoing issue with CT technology is the amount of radiation exposure delivered to the patient to obtain a quality image. The Medipix All Resolution System (MARS) team is one of the groups developing a CT system that will deliver a significantly reduced radiation dose and retain a high quality image. An advantage that the MARS Spectral CT Scanner has is that it uses the Medipix range of energy resolving photon counting detectors. In contrast to regular CT the Medipix detectors can measure the number of photons transmitted through an object and their energies. This allows more information to be gathered from the photons measured, therefore fewer photons need to be measured.

The purpose of this project is to begin preliminary studies on the application of compressive sensing for material decomposition. Compressive sensing is a technique that is used for solving under determined systems of equations. The main assumption for solving such a problem is that the solution is sparse [2] [3] [4] [5], which is true for material decomposition. Typically each voxel in the scan space contains a small number of dominant materials contributing to the attenuation properties of the voxel and a large number of trace materials with negligible contribution to the attenuation properties.

The data used is a set of CT scans taken with the MARS CT scanner of a phantom that is referred to as the Fat-CaFe phantom . The Fat-CaFe phantom is a Perspex cylinder with six holes drilled into it lengthways, each hole filled with a different material - Air, Water, Ferric Nitrate (aq) [0.4M],Calcium Chloride (aq) [2M], Iodine (aq) [0.02M] and Oil. Three different sample regions of this data set has been used, which cover all 7 materials. Prelimi-

nary testing is not yet complete however current progress shows that compressive sensing has potential to work for combinations of materials which fail in commonly used pseudo-inverse methods (for example Calcium, Iron and Oil).

This report contains a description of the mathematical formulation for material decomposition (Chapter 2), an account of a simple pseudo inverse method including conditions and common problems (Chapter 3), a brief introduction to the Basis Pursuit compressive sensing model and a description of the data used and the results obtained so far (Chapter 4).

Chapter 2

Basics of Material Decomposition

Material decomposition requires a set of images. Each of the h images is a scan of the same area in a different energy range (i.e. there are scans in h different energy ranges). Every pixel in each of these images represents the attenuation of the x-ray beam through one voxel of matter contained inside the scan space. We define the vector $\vec{P}_{lm}$ to be the "Pixel Vector" containing the numerical pixel values of each image for the pixel in the m^{th} column and the l^{th} row. A material decomposition on such a set of images treats each pixel vector $\vec{P}_{lm}$ separately. Therefore it can be seen that although the descriptions in the following sections are reasonably simple, with large images the computational cost will become a problem (This issue however is not looked at here).

This section gives a detailed account of how the material decomposition problem is formulated for a single voxel using pixel vectors.

2.1 Material Decomposition of a Single Voxel

The value of each pixel in each reconstructed image (before scaling) represents the total linear attenuation coefficient of each voxel of matter that each pixel represents over the energy range the image was taken at. Each voxel contains (possibly) different amounts of k different materials, therefore the total linear attenuation coefficient of a single voxel, μ_{total}, is the sum of the linear attenuation coefficients for each of the k materials in that voxel, μ_j $(j = 1, \ldots, k)$. This is represented by the equation

$$\mu_{total} = \mu_1 + \mu_2 + \cdots + \mu_k \tag{2.1}$$

These linear attenuation coefficients, μ_j $(j = 1, \ldots, k)$, depend on the density of each material they describe. For each μ_j there is a μ'_j that can be chosen to represent the linear attenuation coefficient of the respective material at a fixed density. Then weighting factors $w_1, \ldots, w_k$ can be used to rewrite equation 2.1 as

$$\mu_{total} = \mu'_1 w_1 + \mu'_2 w_2 + \cdots + \mu'_k w_k \tag{2.2}$$

where $\mu_j = \mu'_j w_j$ and the j^{th} weighting factor w_j depends only on the density of material j relative to the density implicitly described in the linear attenuation coefficient μ'_j.

Each set of weighting factors gives a description of the density for each material in each voxel. In a material decomposition $(w_j)_{lm}$ (the weighting factor of the j^{th} material for the pixel vector $\vec{P}_{lm}$) becomes the pixel value for the pixel in the m^{th} column and l^{th} row in the j^{th} material image. When the material decomposition is done for every pixel vector, a set of k material images are produced. These material images are effectively density distribution maps for each material.

The weighting factors can be obtained by writing the pixel vector in terms of equation 2.2. This produces a system of linear equations that can then be written in terms of matrices [6] [7].

$$\begin{pmatrix} \mu_{total}(E1) \\ \mu_{total}(E2) \\ \vdots \\ \mu_{total}(Eh) \end{pmatrix}_{lm} = \begin{pmatrix} \mu'_1(E1) & \mu'_2(E1) & \dots & \mu'_k(E1) \\ \mu'_1(E2) & \mu'_2(E2) & \dots & \mu'_k(E2) \\ \vdots & \vdots & \ddots & \vdots \\ \mu'_1(Eh) & \mu'_2(Eh) & \dots & \mu'_k(Eh) \end{pmatrix} \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_k \end{pmatrix}_{lm} \tag{2.3}$$

For simplicity equation 2.3 can be written symbolically as

$$(\vec{\mu_T})_{lm} = M(\vec{w})_{lm} \tag{2.4}$$

where M is a $h \times k$ matrix, $(\vec{\mu_T})_{lm}$ is a column vector with h elements and $(\vec{w})_{lm}$ is a column vector with k elements. $(\vec{\mu_T})_{lm}$ is the pixel vector $\vec{P}_{lm}$ from the set of images.

In a material decomposition $(\vec{w})_{lm}$ in Equation 2.4 needs to be calculated for every voxel.

Chapter 3

Conditions and Problems of Solving $(\vec{\mu_T})_{lm} = M(\vec{w})_{lm}$ Using the Pseudo Inverse.

The vector of weighting factors, $(\vec{w})_{lm}$, is commonly solved for by applying the pseudo inverse of M on the left of Equation 2.3 [6] [7] [8].

$$M^{+}(\vec{\mu_T})_{lm} = [(M^T M)^{-1} M^T](\vec{\mu_T})_{lm} = (\vec{w})_{lm} \tag{3.1}$$

This section describes the conditions imposed on material decomposition and common problems that arise when using the pseudo inverse.

3.1 Conditions for Material Decomposition

Material decomposition can be broken down into two types of problems based on the rank of the matrix M from equations 2.3 and 2.4:

- M satisfies the full rank condition (equation 2.3 contains as many linearly independent equations as the number of different materials that the voxel is being decomposed into). In this case M will either be a square matrix with full rank or will be over-determined non-square matrix with full rank. The pseudo inverse is a common method used for solving this class of problems. It is also implicit in this case that for equation 2.3 $h \geq k$.

- M is rank deficient. In this case techniques such as compressive sensing [9] must be used to obtain the weighting factors.

There are two major factors that will affect the rank of the matrix M:

(1) If $n \geq 3$ materials are solved for using three or more images at different energy ranges then to satisfy the full rank condition a minimum of $n - 2$ materials require unique K-edges in their attenuation spectra within the energy ranges used. More generally, for $n \geq 2$ materials the rank of M is at least 2, and will increase by 1 (unless the full rank condition has been met) for every material considered which contains a unique K-edge in its attenuation spectra within the energy ranges used [6] [7].

(2) The independent equations used can either come from images of different energy ranges or known physical relationships about the system scanned. The latter type of equations would be included in equation 2.3 if used. An example of such a relationship is the sum of the masses of each material in a voxel being equal to the total mass of the voxel, i.e. mass conservation [10].

3.2 Problems With the Pseudo Inverse Method

There are two major problems that are encountered when using the pseudo inverse method.

- The voxel contains at least one material which is not represented in the material matrix, M. This can result in materials being assigned negative densities and/or densities that are too large or too small.

- If two or more materials in the material matrix, M, have similar attenuation spectra over the investigated energy range then they may be treated linearly dependent by the pseudo inverse, lowering the effective rank of M.

Chapter 4

Preliminary investigations into Solving $(\vec{\mu_T})_{lm} = M(\vec{w})_{lm}$ Using Compressive Sensing.

Typically to solve a system of equations $(\vec{\mu_T})_{lm} = M(\vec{w})_{lm}$ the same number of linearly independent equations as unknowns is required. However if it is known that the solution is sparse (contains few non-zero components) then the system can be solved for exactly using fewer equations by searching specifically for the optimal sparse solution. The method applied in this project is L1-Minimization, also known as Basis Pursuit [2] [3] [4] [5]. The sparse reconstruction package used for this project is YALL1 and can be obtained from http://yall1.blogs.rice.edu/ [1].

4.1 Brief Summary of Compressive Sensing

Compressive sensing is a technique that can be used to solve an under determined system of equations for a sparse solution. When the solution vector, $(\vec{w})_{lm}$, is S sparse (contains S non-zero components) it can typically be reconstructed by $2S$ equations. A more general condition on the matrix M is that it must satisfy Equation 4.1, where M is an $h \times K$ matrix [2] [3] [4] [5]. A detailed analysis of further conditions is beyond the scope of this project.

$$h > |C|S \times log(K), \text{for a constant C.} \tag{4.1}$$

Basis Pursuit is the compressive sensing method focussed on in this project. Reconstructing a sparse solution from noise free data by Basis Pursuit is done by solving Equation 4.2. This will typically reconstruct $(\vec{w})_{lm}$ exactly.

$$\min \|(\vec{w})_{lm}\|_1 \text{ s.t. } (\vec{\mu_T})_{lm} = M(\vec{w})_{lm} \tag{4.2}$$

When the data used in Basis Pursuit is noisy $(\vec{\mu_T})_{lm} = M(\vec{w})_{lm}$ no longer has a certain equality, which is accounted for by using Equation 4.3. In this formulation λ is a positive parameter which controls the strength of the sparsity condition over minimizing the square error. Appropriate choice of λ will typically reconstruct an accurate solution despite the noisy data [2] [3] [4] [5].

$$\min \|M(\vec{w})_{lm} - (\vec{\mu_T})_{lm}\|_2^2 + \lambda\|(\vec{w})_{lm}\|_1 \tag{4.3}$$

4.2 Test Data

The test data used in this project was provided by the MARS research group. This data consists of six (438 × 438 pixel) CT scans of a phantom which is referred to as the "Fat CaFe" phantom. This phantom is made from an solid perspex cylinder with six internal holes drilled into it lengthwise. These holes are filled with different materials - Air, Water, Ferric Nitrate (aq) [0.4M], Calcium Chloride (aq) [2M], Iodine (aq) [0.02M] and Oil (Figure 4.2). These CT scans were taken using a 50 kVp x-ray tube in the MARS CT scanner at six

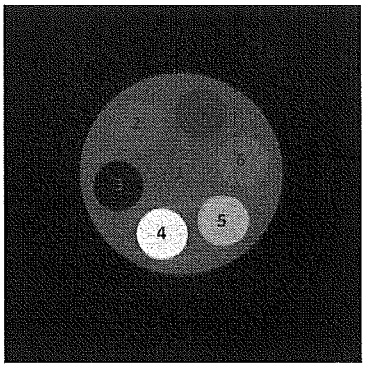

Figure 4.1: Scan of Fat CaFe phantom showing each material. (1) Perspex, (2) Water, (3) Air, (4) Calcium Chloride (aq) [2M], (5) Ferric Nitrate (aq) [0.4M], (6) Iodine (aq) [0.02M] and (7) Oil.

different energy ranges. These energy ranges are from a lower threshold up to the maximum x-ray energy, 50 keV. The lower thresholds are 9.8 keV, 15.1 keV, 20.4 keV, 25.6 keV, 30.9 keV and 36.2 keV. At each energy the CT scans have been reconstructed using filtered back projection and 30 adjacent scans have been averaged together to produce the images used in this project (each adjacent scan contains identical spatial and material information). This was done by the group working on pre-processing before the images were obtained for this project (Figure 4.2).

It should be noted that the images used for this project were not in units of linear attenuation. These images have gone through a series of scaling processes before reaching the material decomposition stage from the reconstruction program "Octopus". The program "Octopus" performs an automatic scaling on each newly reconstructed image. Following this the window and level for the grey-scale images are chosen. These images are then scaled to 16 bit TIFF files which have values ranging from 0 to 65335, which is the scaling of the images used for this project.

A synthetic set of data was also created for the preliminary testing stage. For each image in this set of data the values of each pixel representing the same material were averaged together. These values were then used to re-create the images in Figure 4.2.

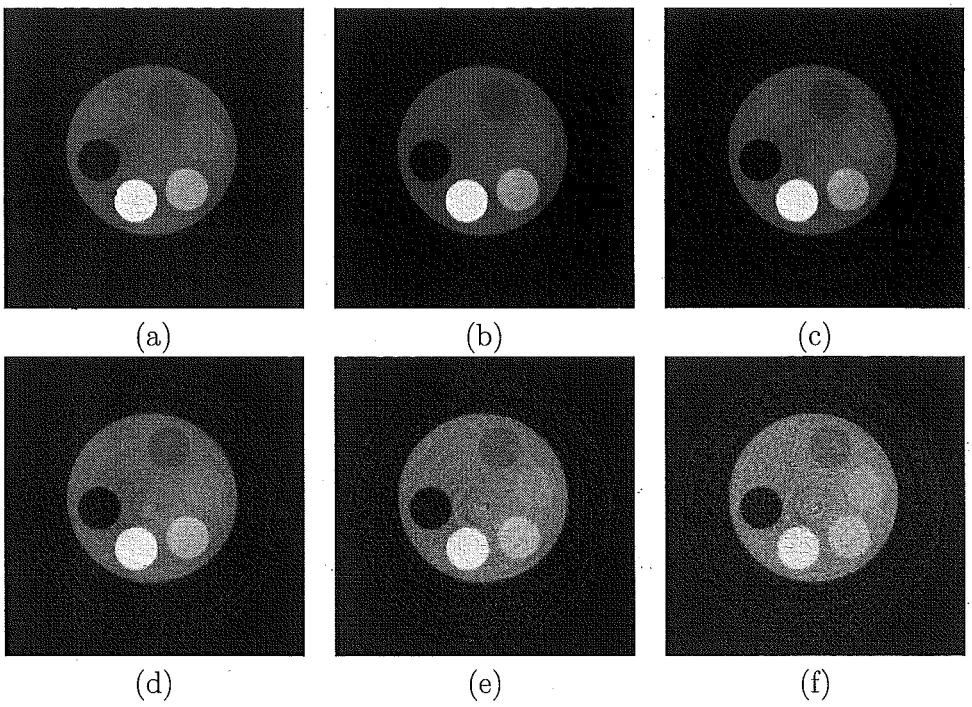

Figure 4.2: Data obtained from the MARS group. These represent scans taken at six different energy ranges: (a) 9.8 keV - 50 keV, (b) 15.1 keV - 50 keV, (c) 20.4 keV - 50 keV, (d) 25.6 keV - 50 keV, (e) 30.9 keV - 50 keV and (f) 36 keV - 50 keV.

4.3 Preliminary Investigation - Completed Trials

4.3.1 Specific Data Tested

Currently trials have only been performed on the synthetic set of data. Three sample regions of the data were chosen for testing, see Figure 4.3. The basic Basis Pursuit method has been

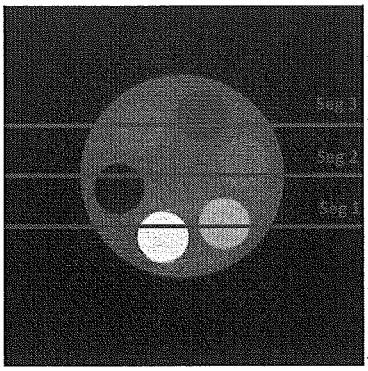

Figure 4.3: Sample regions of test data. Seg1 contains Calcium Chloride (aq) and Ferric Nitrate (aq) [0.4M], Seg2 contains Air and Iodine (aq), Seg3 contains Water and Oil. All three regions contain Perspex.

tested for these samples in both the TIFF scaling and Hounsfield Unit Scaling. The M matrix described in Chapter 2 in both cases were constructed from the average (mean) pixel

values for each material region in each image (The technicalities of attenuation units can be initially ignored for this scenario). Two different M matrices have been tested, one including all different materials in the scan (see Figure 4.1) and one including all materials except Perspex.

4.3.2 Results - Including Perspex

Figure 4.4 shows the material decomposition results for the three image segments shown in Figure 4.3 for TIFF scaled data and M matrix including Perspex. These show good decomposition in the material regions for - Calcium, Iron, Iodine, Oil and Air. The Algorithm failed to give a good reading in the Perspex and Water regions.

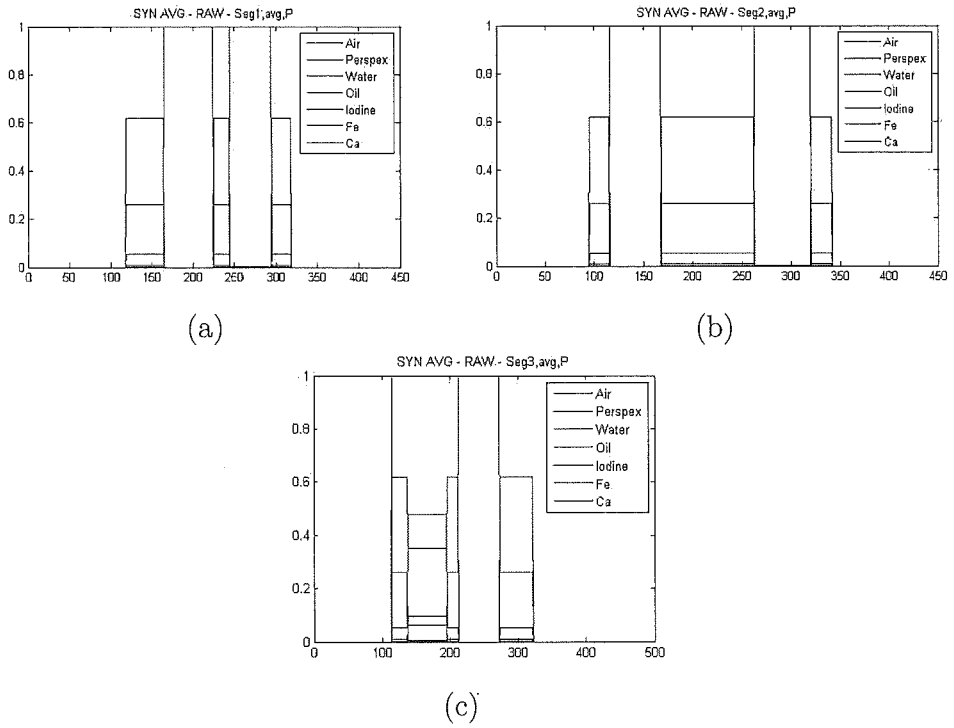

Figure 4.4: Material Decomposition on image segments shown in Figure 4.3 for TIFF scaled data and M matrix including Perspex. (a) Seg1 - Calcium and Iron region. (b) Seg2 - Air and Iodine region. (c) Seg3 - Water and Oil region. All regions also contain Perspex. Horizontal Axis: Pixel. Vertical Axis: Fractional Contribution

Figure 4.5 shows the material decomposition results for the three image segments shown in Figure 4.3 for Hounsfield Unit scaled data and M matrix including Perspex. These show good decomposition in the material regions for - Calcium, Iron, Iodine, Oil and Air. The Algorithm failed to give a good reading for the Perspex region. Water is given a zero value which is expected when using Hounsfield Units. It is also observed that in the Hounsfield Unit material images there is significant noise outside of the phantom. This is a consequence of the scaling providing the Basis Pursuit program with negative M matrix entries.

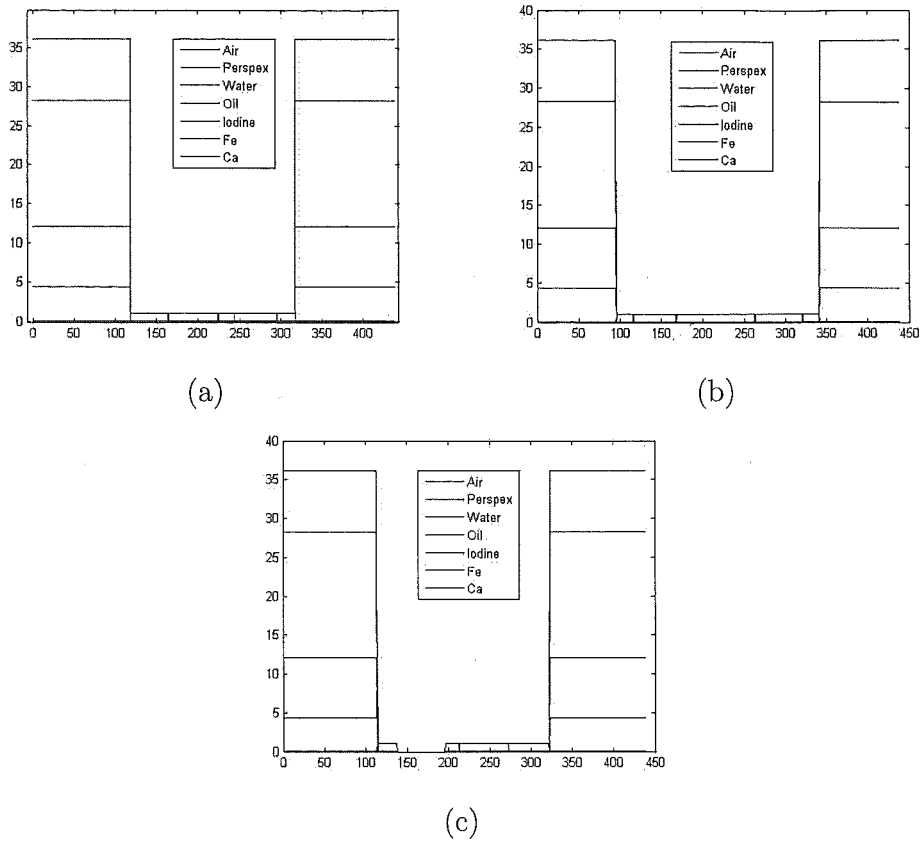

Figure 4.5: Material Decomposition on image segments shown in Figure 4.3 for Hounsfield Unit scaled data and M matrix including Perspex. (a) Seg1 - Calcium and Iron region. (b) Seg2 - Air and Iodine region. (c) Seg3 - Water and Oil region. All regions also contain Perspex. Horizontal Axis: Pixel. Vertical Axis: Fractional Contribution

4.3.3 Results - Excluding Perspex

Figure 4.6 shows the material decomposition results for the three image segments shown in Figure 4.3 for TIFF scaled data and M matrix excluding Perspex. These show good decomposition in the material regions for - Calcium, Iron, Iodine, Oil and Air. The Algorithm failed to give a good reading in the Perspex and Water regions. The cause of the Air spike on the outside has yet to be investigated.

Figure 4.7 shows the material decomposition results for the three image segments shown in Figure 4.3 for Hounsfield Unit scaled data and M matrix excluding Perspex. These show good decomposition in the material regions for - Calcium, Iron and Oil. The Algorithm failed to work on Seg2 for this scaling, the cause is still under investigation. Perspex regions were not given any material weighting under this scaling. Water is given a zero value which is expected when using Hounsfield Units. Again it is also observed that in the Hounsfield Unit material images there is significant noise outside of the phantom. This is a consequence of the scaling providing the Basis Pursuit program with negative M matrix entries.

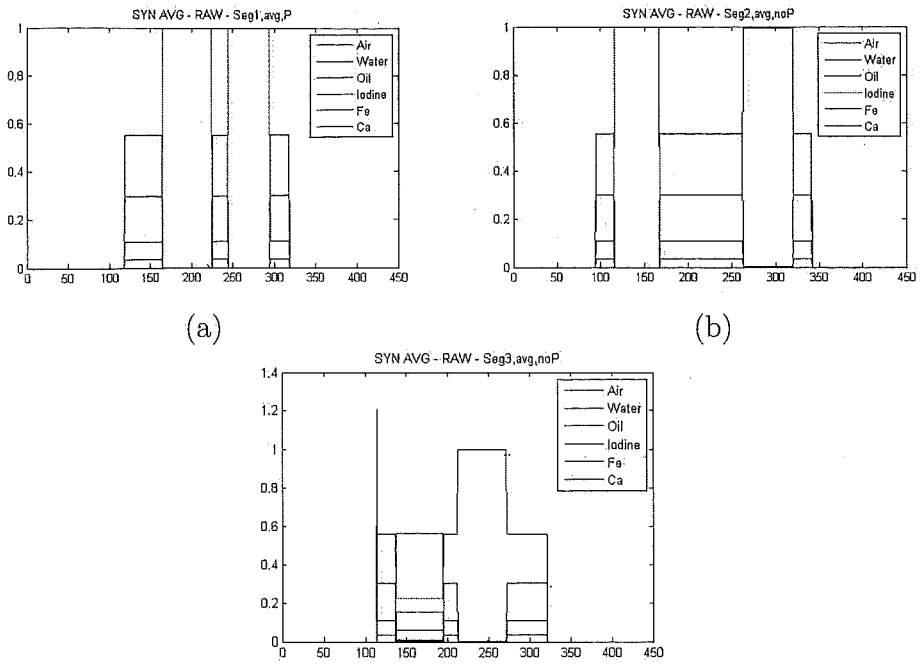

Figure 4.6: Material Decomposition on image segments shown in Figure 4.3 for TIFF scaled data and M matrix excluding Perspex. (a) Seg1 - Calcium and Iron region. (b) Seg2 - Air and Iodine region. (c) Seg3 - Water and Oil region. All regions also contain Perspex. Horizontal Axis: Pixel. Vertical Axis: Fractional Contribution

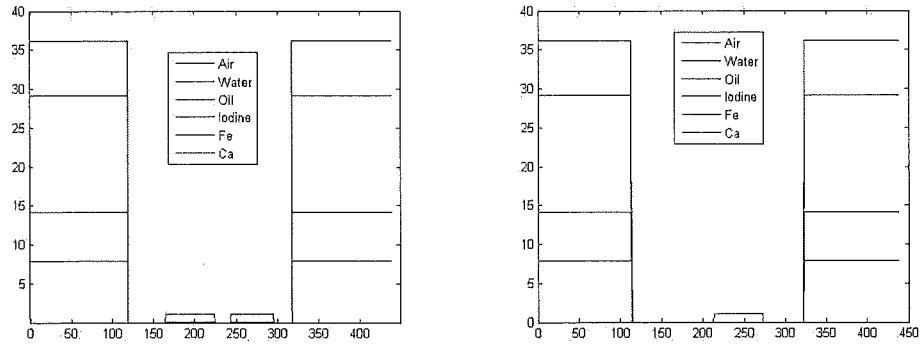

Figure 4.7: Material Decomposition on image segments shown in Figure 4.3 for Hounsfield Unit scaled data and M matrix excluding Perspex. (a) Seg1 - Calcium and Iron region. (b) Seg3 - Water and Oil region. All regions also contain Perspex. Horizontal Axis: Pixel. Vertical Axis: Fractional Contribution

Chapter 5

Discussion

This project is the start of an investigation into using compressive sensing as a material decomposition method for the MARS Spectral CT scanners. Currently only a portion of the first stage of preliminary testing has been completed. The results obtained so far show potential for compressive sensing to surpass other material decomposition methods which are currently used in Spectral CT.

The results of this project have successfully distinguished between 3 and 5 materials. A number of these materials (for example Calcium, Iron and Oil) are not distinguishable when using current material decomposition methods such as pseudo-inverse methods. It is important however to complete the preliminary testing before an inference is made for the application of compressive sensing to material decomposition.

Further preliminary testing will include:

- Testing on the corresponding non-idealized set of data.
- Noise sensitivity testing.
- Repeating tests in natural attenuation units (cm^2g^{-1} and/or cm^{-1}).

Full preliminary testing results will determine the content required for the next stage of testing and development for using compressive sensing for material decomposition.

Chapter 6

Conclusion

This report covers the first results for a preliminary investigation into using compressive sensing as a material decomposition method for the MARS Spectral CT scanners. No firm conclusions can be drawn until this stage of testing is completed. However compressive sensing does show potential to give accurate results for the cases where commonly used pseudo-inverse methods fail.

Acknowledgements

Thanks is required for my supervisors Peter Renaud, Rachael Tappenden, Philip Butler, Anthony Butler; Paul Ronaldson, Rafidah Zainon and Neils De Ruiter for providing data for this project; to the MARS research group and many members who offered advice; the University of Canterbury (UC) Summer Scholarship program for funding; the UC Department of Mathematics and Statistics for hosting this project; the UC Department of Physics and Astronomy for providing office space and a computer; and my partner Sandra Atkinson for many discussions on this research.

Bibliography

[1] Yin Zhang. User's guide for yall1: Your algorithms for l1 optimization. *CAAM Technical Report TR09-17, Department of Computational and Applied Mathematics, Rice University.*

[2] Yin Zhang. On theory of compressive sensing via l1-minimization: Simple derivations and extensions. *CAAM Technical Report TR08-11, Department of Computational and Applied Mathematics, Rice University*, 2008.

[3] Michael Wakin Justin Romberg. Compressed sensing: A tutorial. *IEEE Statistical Signal Processing Workshop*, 2007.

[4] Terence Tao. Compressed sensing, or:the equationax=b, revisited. *Mahler Lecture Series.*

[5] Stephen Wright. Sparse and structured optimization. *CSMP, Shanghai*, 2010.

[6] Daniel Niederlöhner Gisela Anton Markus Firsching, Jürgen Giersch. A method for stoichiometric material reconstruction with spectroscopic x-ray pixel detectors. *Nucl. Instr. and Meth.*, 546(1-2):125–130, 2005.

[7] T Michel G Anton Markus Firsching, D Niederlöhner. Quantitative material reconstruction in ct with spectroscopic x-ray pixel detectors - a simulation study. *Ieee Nuclear Science Symposium Conference Record*, 4(2):2257–2259, 2006.

[8] Alan J. Laub. Matrix analysis for scientists and engineers, siam. 2005.

[9] Ge Wang Hao Gao, Henyong Yu. True-color ct based on a prior rank, intensity and sparsity model (prism).

[10] Andrew N. Primak Cynthia H. McCollough Xin Liu, Lifeng Yu. Quantitative imaging of element composition and mass fraction using duel-energy ct: Three material decomposition. *Med Phys*, 36(5):1602–1609, 2009.